Electronics and Computer Science
Faculty of Physical Sciences and Engineering
University of Southampton

Jonathan NS Morrice

30th April 2013

SmartHeat: An Open Source Platform for Interactive Visualization of Sensor Data

Project supervisor: Dr. Enrico Costanza
Second examiner: Dr. Geoff Merrett

A project report submitted for the award of

BEng Electronic Engineering

Abstract

Visualisation of energy consumption data has shown to influence consumer behaviour and improve efficiency. Existing systems provide real time feedback and thus stimulate the user to save energy. While many current systems are based on electricity use, there is little research on visual temperature data and its proficiency.

This project tackled the key problem of visualising such data and emphasising where and when energy is wasted. SmartHeat enabled the user to easily understand and hence optimize the temperature development of a building. It was aimed particularly towards smart heating agents to simplify interaction, enhance user input and increase long-term efficiency. Core to the visualisation was putting data into context using a small, cheap set of sensors. Feedback from a user study showed how understanding was improved, with participants drawing good conclusions for their heating behaviour in the future.

Details of the background for this project, as well as technical design and implementation details can be found in this report. Furthermore this paper describes and evaluates the field trial performed. Finally the project as a whole is evaluated and the scope for further work presented.

Acknowledgements

I would like to thank my supervisor Enrico Costanza for the extensive support he provided throughout the entire project, making it a very enjoyable, challenging and rewarding experience. In addition I am very grateful to Evangelos Tolias for his support on interfacing to the FigureEnergy database. Finally many thanks go all participants of the field trial, whose feedback added a lot of value to this research.

Contents

1 Introduction

One of the major challenges the world currently faces is reducing energy consumption. This is due to excessive use of fossil fuels, which - despite damaging our environment - currently generate most of the energy we use. Technology provides opportunities to solve these issues by accessing renewable sources of energy and improving power efficiency, but can also assist humans in consuming more intelligently. Ubiquitous computing is particularly effective in domestic applications, which make up 39% of the overall energy consumption in the US and 28.8% in the EU [1]. Research has shown significant improvements in domestic energy efficiency when presenting visual feedback on electricity use to consumers. Likewise, this project explored the potential of visualisation systems for *temperature* data.

The primary goal of the project was to design an eco-feedback system that emphasized where and when energy was being wasted for space heating. Moreover, the project aimed to not only develop and implement the system, but also evaluate its effect on users. With a better understanding of their behaviour, users were likely to make better decisions on how to heat in the future. Understanding was achieved by adding context to temperature data. The context was provided by information from a small, cheap set of sensors. Ideally this showed whether a heating process was necessary (e.g. the house was cold and people were present) or wasteful (e.g. house was already warm). In addition, this research could help existing smart heating systems simplify user interaction, enhance user input and increase long-term efficiency. Giving consumers a better understanding of when and why they heat could help grasp system-generated schedules and make changes they add to it more valuable. Hence learning curves could be shortened and modern thermostats improved.

A wireless sensor node and a sophisticated web interface for users to interact with the data were developed. Designs were kept as simple and cheap as possible, to allow for implementation in the large open-source ubicomp community as future work. However, the project not only combined hard- and software development. The second goal was to evaluate the effect of the system on users in a field trial. SmartHeat was deployed in 6 different residencies for three consecutive days. Participants were interviewed before and after the trial to analyse how they had used the system and whether their understanding had improved.

The next chapter details the literature reviewed for this project, which focuses on existing heating control systems and visualisations for electricity consumption. This is followed by a chapter on the design details, implementation and testing of the wireless sensor node. Then different methods of processing the different sensor signals to receive the most useful information are explained. Another chapter details several visualisation approaches and their implementation, as well as reasoning of the final design chosen for testing. Finally the field trial is described and the project concluded as a whole, including a personal reflection and scope for future work.

2 Related Work

2.1 Heating control systems

2.1.1 Programmable Thermostats

Most heating systems around the world use thermostats as controllers. They sense the temperature of the system and regulate heat flow to keep the temperature near a certain value defined by the user. This is either done simply by turning the heating/cooling devices on and off, or by limiting the flow of a heat transfer fluid as required. More sophisticated systems use programmable thermostats, which allow the user to specify different temperatures for different times of the day. Starting as mechanical implementations in 1885 [2], these systems have become increasingly popular and are now used in most households in the UK. As technology has developed, programmable thermostats have been built using analogue electronics and more recently implemented as digital systems [2]. While their deployments have changed, functionality has hardly advanced.

Allowing the user to schedule heating is supposed to increase efficiency by adapting to his requirements. For example somebody who works during the day can choose only to heat in the evenings and at night when he is in the house and save resources during the day. However, studies have shown that programmable systems do not have the desired effect in many homes. In 2000 Nevius and Pigg showed that homes with programmable thermostats consumed roughly the same amount of energy that those with manual thermostats did [3]. Shiller summarised this and other studies as shown in Table 1 and claimed user confusion over their devices [4]. Contrary to these findings, RLW Analytics in 2007 showed that Energy Star[1] certified programmable thermostats achieved a 6.2% reduction in annual gas consumption [5]. They explained the difference to previous results to originate from newer, easier to program and hence more user-friendly thermostats.

More recent papers also see the user interface to be the main issue preventing people from using the true energy saving potential of their thermostats [6][7]. A survey in 2011 showed many users did not know at what times they were heating, how they could program their thermostat to change that or did not even know they had one [8]. Proposed solutions to the problem range from larger buttons on the devices to using a web interface or even voice recognition, to control the system [7].

---

[1] International standard for energy efficient consumer products

Organization	Investigators	Location/Year	Sample Size	Conclusions
Southern California Edison	Paul Reeves, Jeff Hirsch, Carlos Haiad	CA 2004	N/A	Energy savings depend on behaviour and can be + or -
Energy Center of Wisconsin	Monica Nevius, Scott Pigg	WI 1999	299 homes	No significant savings. PT's don't change behaviour.
Connecticut Natural Gas Corporation	David Cross, David Judd	CN 1996	100 homes	PT's cause significant behaviour change.
BPA/PNNL	Craig Conner	NW 2001	150 homes	No significant behaviour change/savings.
Florida Solar Energy Center	Danny Parker	FL 2000	150 homes	No savings, some increases.

Table 1: Summary of thermostat behaviour and energy saving studies [3]

2.1.2 Smart Thermostats

Modern smart thermostats attempt to solve these problems. Nest Labs released a thermostat in late 2011 that learns from user decisions and can also sense presence. Users manually adjust the temperature on the thermostat simply by turning it. The system learns user preferences from their decisions and builds a schedule accordingly. In addition, it does not heat when homes are not occupied and provides a web interface as well as a phone application with basic visualisation. Nest says their thermostats solve the problems complicated, inflexible thermostats have by building flexible schedules that adapt to the changing life of many users [9]. The ecobee smart thermostat is a similar product with a slightly more complex UI but more manual options.

Both devices have received good reviews and are promising options for the future. However, main issues are a high price at around $250 respectively, a learning process that takes time or a lot of adjustment and false presence detection.

2.2 Energy consumption visualisation

We have seen that HCI is the key to making heating control systems more efficient. Aside from technically being able to adjust the thermostat, users need to be willing to do so as well or they won't do it. A method that has shown to be effective in promoting energy savings in other applications is data visualisation. By visualising energy consumption data we feed it directly back to the user, increasing awareness, but more importantly an understanding of how energy is used. Understanding is important in order for users to apply their efforts and resources correctly [10]. A review in the year 2000 of 38 feedback studies carried out over a period of 25 years showed that feeding data back to consumers increased awareness and decreased consumption by about 10% [11].

With general consciousness increasing in recent years these figures have risen further. An experiment in 2005 where visual energy-monitoring systems were installed in 10 residential houses for 280 days showed total consumption was reduced by about 12% after installation of the system [12]. Figure 1 shows the interface residents were presented.

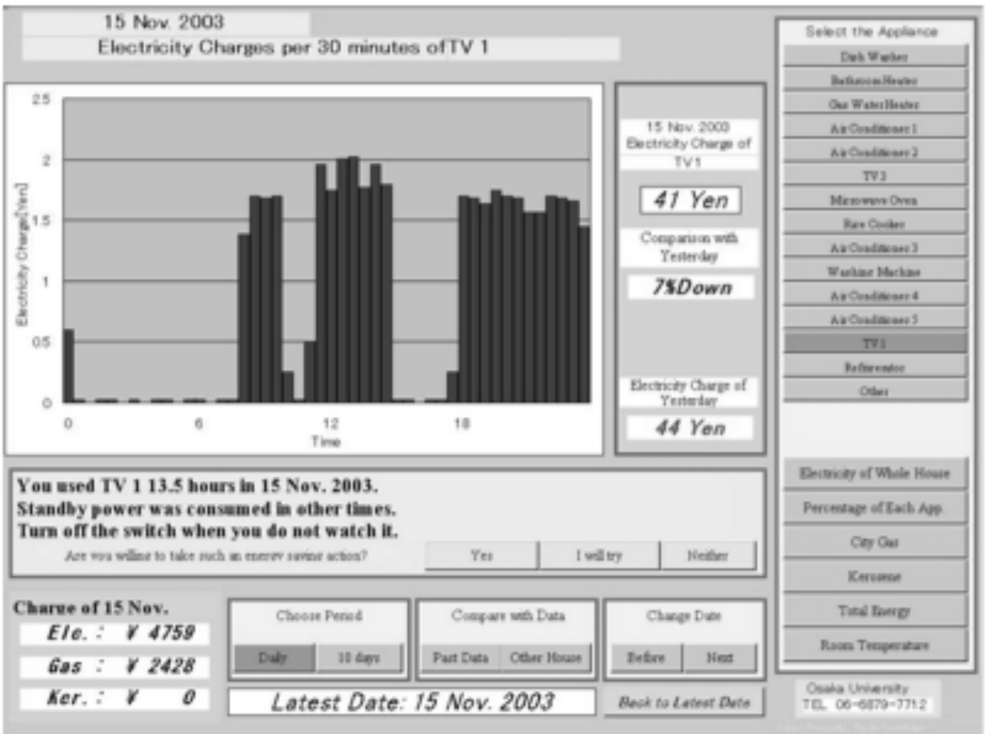

Figure 1: User interface of the ECOIS II [11]

Two years later a similar experiment conducted in student dormitories achieved a 32% reduction in electricity use [13]. This must be seen relative to the fact that it was designed as a competition and that it took place over a much shorter period of time. But it shows what an impact motivation has and moreover, also found that students presented with high resolution feedback were more efficient, reducing consumption by 55% as opposed to 31% by those with low resolution data.

FigureEnergy is a more recent project that engages users with their data beyond mere presentation by letting them assign context to data [10]. Users interact with the system using a web interface to manually assign certain peaks to their activities and the system then automatically annotates them in the future. A field study showed users that had used other electricity displays before discovered that some appliances consumed more than they expected, and that they became more aware of activities taking place in their household [10].

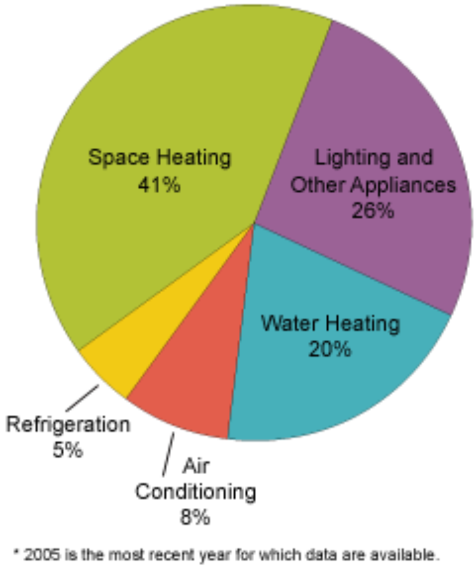

Figure 2: How Energy is used in homes in the US [16]

While most of the research in this area has focused on visualising electricity consumption, general theories on consumption feedback suggest that similar results can be expected for heating applications [14][15]. Moreover, a majority of 49% [16] of the average household energy consumption is used for heating (space and air conditioning). Hence users are more motivated to optimize their systems to cut not only their gas bills, but also electricity bills since air conditioning and some heaters run electrically. In the 2005 visualisation study, users were most interested in consumptions for space heating and some managed to cut down their heating demand by 20-45% by reducing heating durations.

A system deployed in 2011 also presented energy consumption in a larger context of data, similar to the approach taken in this project [17]. However, whilst achieving promising results regarding user awareness, the research only explored these possibilities in an office environment, not a domestic home. Moreover, individual nodes were created for each sensor rather than combining them all into one. Visualisation remained in form of a common line graph.

3 Wireless Sensor Node

3.1 Specification

3.1.1 Choice of Sensors

The visualisation system required data in order to be tested and evaluated. A wireless sensor node was needed to collect this data and transmit it to a central server that could be accessed by the web interface. Most sensors of interest were cheap and generally easy to implement or available in commercial systems. For this reason a wide range of sensors was chosen to discover each individual sensors value with respect to visualising temperature data.

Apart from the obvious temperature sensor a humidity sensor was also included. Possible benefits were detecting users cooking, open windows or other processes that affect the inside humidity and temperature.

Presence was an important factor to consider since previous studies had shown it to reduce consumption by 28% when considered in smart thermostats [18]. To detect this a PIR sensor commonly found in burglar alarms was used.

In addition, a light or brightness sensor was included to provide further presence information and possibly indicate sunshine on a building if placed near a window.

3.1.2 Technical Requirements

A hybrid approach was used to specify the sampling and transmission rates for the node. Generally, a higher transmission rate results in higher power consumption, while a lower rate is more efficient. The same applies for sampling rates, but for a sampling rate higher than the transmission rate temporary storage of sensor data is also required. It was chosen to only sample the presence sensor at a higher rate because it essentially detects motion, not presence. Occasional motion in the room means somebody is present and to detect that samples must be taken more frequently, unlike for temperate for example, which can be measured directly. Presence was specified to be sampled every 4 seconds and stored temporarily, while all other sensors were sampled every 2 minutes and then transmitted directly. Most domestic households have an existing Wi-Fi network and since an IP stack was required for server communication, these were specified for use in data transmission.

The sensor node had to be portable. This was important since positioning had a great influence on presence and light measurements. Moreover, this is the industry standard for smart home sensors, which was followed to prove that deployment in a domestic system was possible and realistic. Reliable collection of data over several weeks was also necessary. As a consequence, all electronics including sensors and communication modules had to consume low power to enable long term power supply from a battery.

3.1.3 Ease of Implementation

Due to the experimental set of sensors and interface to a specific server, a custom node was required, as no commercial product met this specification. However this was essentially a HCI and not an engineering project. Focus had to lie on the visualisation and not on the development of the sensor node. Hence a quick and easy implementation was essential to the design and an important part of the specification.

3.2 System Design

3.2.1 Component Selection

An Arduino development board was chosen for this project to simplify the required hard- and software while prototyping and shorten development time. The Arduino Uno runs on an AVR ATmega328 that can easily be programmed via USB and was the cheapest option from Arduino for this project. It provides all required power and reset circuitry and can be expanded with different shields for requirements like wireless communication in our case. Due to their simplicity and low price, Arduino boards have become very popular and there are many libraries available online that save a lot of coding. Many communities online work on open source Arduino projects, so using this platform also left that option for future work.

Since the system was specified to run on battery and to deliver data over several weeks, low power consumption was essential for all components. Size was less important, but fitting all components on the relatively small wireless shield prevented having to buy an extra prototyping shield. A through hole package or else an appropriate adapter were preferable for prototyping.

Specifically for temperature high accuracy was important, since it reduced the need to average data (see processing) and hence lose accuracy on top of that of the sensor. A typical error around ±0.1°C-±0.3°C was acceptable and realistic. The sample period for temperature was specified as 2 minutes. Most chips provide about 1Hz, which was more than sufficient.

Regarding humidity there were many sensors available that came with signal processors on-board to provide digital interfaces. They were based on capacitive, resistive or thermal conductivity sensors, which were also available separately. A standalone capacitive humidity sensor was chosen, since these were the most common in industrial and commercial applications [19]. The capacitance was to be measured using the Arduino by measuring the discharge time. This was the most cost effective and power efficient option, since no additional electronics were required and the error of ±2% RH (relative humidity) was small enough for this system.

For presence detection, PIR (passive infrared) sensors have proven sufficient for this type of application [18]. They support different ranges and angles but in contrast to other applications, losses in these areas were tolerable for price and ease of use in this project. The chosen sensor has a MCU on-board that filters noise and false triggers and has a range of 5m at 60° at a low price. Table 2 shows a summary of the sensors that were chosen for the prototype. If available, ICs with an I^2C slave interface were chosen. I^2C is a common bus system, which uses two bidirectional lines for communication, making it easy to add additional slave devices once the bus is implemented on the node.

Sensor	Manufacturer	Part Nr.	Interface	Details
Temperature	Microchip	MCP9808-E/MS	I^2C	±0.25°C accuracy (typical)
Humidity	Honeywell	HCH-1000-001	Analogue	Capacitance varies around 330pF depending on RH
Presence	Zilog	ZEPIR0AAS02MODG	UART	MCU as noise filter on board
Light	Silonex	NSL19M51	Analogue	LDR from 20MΩ (dark) to 20kΩ (light)

Table 2: Overview of sensors initially chosen for the prototype

Two communication protocols were considered for this project: Wi-Fi (IEEE 802.11) and ZigBee (IEEE 802.15.4). On the one hand, Wi-Fi is compatible with networks present in most homes and IP to upload data. On the other hand it uses a lot of power, is relatively expensive and requires the user to select a network and usually enter a passphrase. Meanwhile ZigBee is cheaper and has low power consumption, but requires an extra node connected to a computer or similar to communicate with a web server. Due to the increasing number of smart objects requiring IP connectivity in recent years, companies have developed low power Wi-Fi chips that are fully compatible with 802.11. The price is higher latency, slower response and lower data rates, tolerable for typical smart home applications like this. A study from early 2012 showed these chips could transmit sensor data for multiple year using two AA batteries [20]. The company that built the chip used in the paper is now called Roving Networks and belongs to Microchip.

Data was stored on the FigureEnergy development server, accessible via an API using HTTP POST and GET requests. This made it easy to access data from an Internet browser using JavaScript or similar, enabling deployment for this project at very low cost. The cheapest module from Roving Networks was chosen for this application. It has a ZigBee footprint compatible with the Arduino wireless shield and supports IP. Adequate libraries were also available, which made implementation short and simple. A block diagram of the whole system as it was finally implemented is shown in Figure 3. Note a design change regarding the humidity sensor, explained in more detail in section 3.3.2.

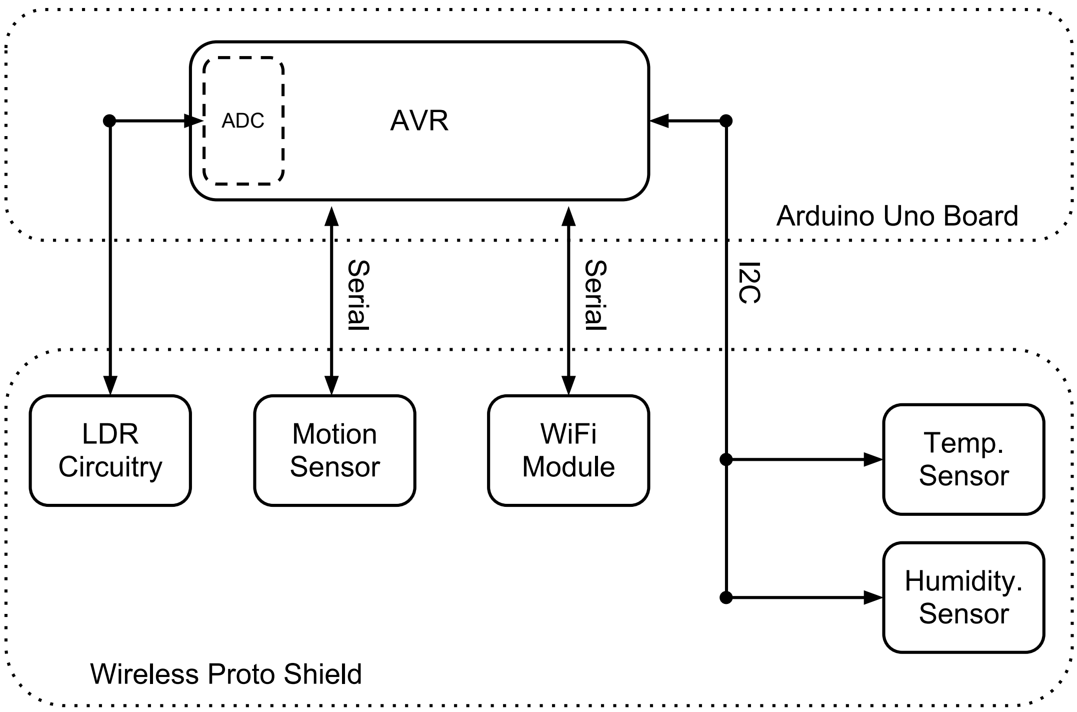

Figure 3: Overall Block Diagram

3.2.2 Software Design

The software running on the AVR was designed as a finite state machine (FSM) containing six states: INIT, READ, SEND, SLEEP and WAKE. The INIT state initialises the AVR, calibrates the motion sensor and configures the Wi-Fi module for a specific network. The READ state samples presence and increments an accumulator if true. Only if the accumulator was greater than a certain threshold value after 2 minutes, did the node transmit presence as true. This provided a realistic value of presence for this application because it filtered out very short occupancies that were not relevant to users. In SEND state, the Wi-Fi module connects to the server, samples all sensors and transmits them in a POST request respectively. In SLEEP state, the node simply waits 4 seconds. The WAKE state follows the SLEEP state and checks whether 2 minutes have passed, using a counter that is increments every 4 seconds in READ state. The state transition diagram is shown in Figure 4.

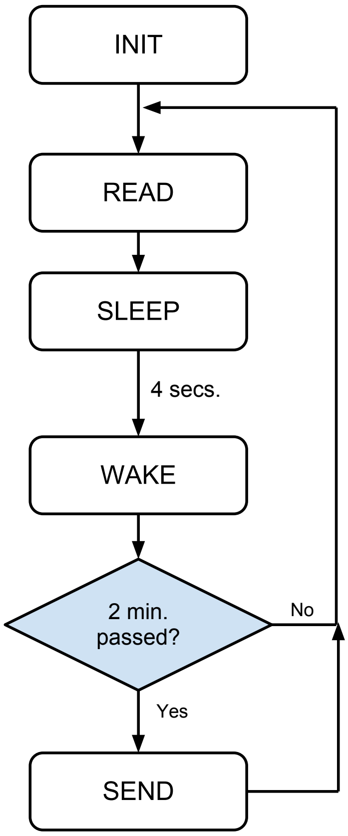

Figure 4: State Transition Diagram of the FSM implemented on the AVR

For optimal power consumption the Wi-Fi module is put into sleep mode unless it is transmitting or being configured. The PIR sensor and the AVR itself also have sleep modes that are active in the SLEEP state between presence readings.

A watchdog timer was used to perform a hardware reset of the AVR in case of a server connection timeout. This leads the system back into the INIT state so that all sensors and the Wi-Fi module are reconfigured. Since no fast error handling is required this was the safest option to avoid the node crashing and require a manual reset if Wi-Fi connectivity was lost or there were issues with the server.

In addition, it was very quick and easy to implement since the timer is simply started before a transmission and reset on success.

3.3 Implementation

3.3.1 Software Libraries

Software for the node was written in C and developed using the Arduino IDE, which uses the AVR-GCC compiler and 'avrdude' to upload to the AVR. Communication with the temperature sensor via I^2C and the motion sensor via UART was easy to implement using two official libraries from Arduino called 'Wire' and 'Serial'. For the Wi-Fi module a third party library based on the 'Serial' library was used called 'WiFlyHQ'. It is written specifically for the Roving Networks RN-XV module and provides basic functions to connect to a network and server, as well as send and receive packages. However, the library does not support sleep mode. Since this was essential to keeping power consumption low it was added to the library using C++. In addition functions to control the module’s three LEDs and to set some more specific authentication settings were added.

3.3.2 Debugging Tools and Testing

Debugging the sensors was straightforward, since debugging information could be sent from the AVR’s serial port to a USB port on a PC via the Arduino and viewed in a serial monitor. In addition a USB logic analyser from a company called Saleae was used. It supports I^2C and serial buses, decodes information as it is transmitted on the bus and displays it in hexadecimal or decimal form for example. However, debugging the Wi-Fi module was slightly more complex. Using a switch on the shield it could be connected to either the USB converter’s or the AVR’s serial port. Programming it via USB was useful to test settings, but to debug the code running on the AVR the USB logic analyser had to be used. Additionally another serial port was created using the 'SoftwareSerial' library to log extra information. This again had to be monitored using the extra logic analyser. The HTTP requests to the FigureEnergy API were tested using 'cURL', a command line tool for getting or sending files using URL syntax.

All sensors and the Wi-Fi module were tested with the AVR individually. After being tested successfully, sensors were included in the AVR and Wi-Fi system and transmission of their samples to the server tested. Sub systems (i.e. sensors) were added individually to a system that could upload dummy data successfully, until all sensors were included. One issue arose here regarding the humidity sensor. A capacitive sensor had been selected and implemented so that its discharge time could be measured as shown in Figure 5. Due to the very low capacity of 330pF that only varies within a range of about 100pF the discharge time is in the order of microseconds. Since the AVR only runs on a 16 MHz processor it could not be measured accurately. Even with a large resistor of 1MΩ the discharge time was still not long enough for accurate measurements. Different implementation using other circuitry was reviewed but ultimately a more expensive IC was purchased that could be added to the I^2C bus. This decision increased the cost of the node, but it saved a lot of time that would have been spent reviewing, implementing and testing circuitry that wasn't actually part of the research.

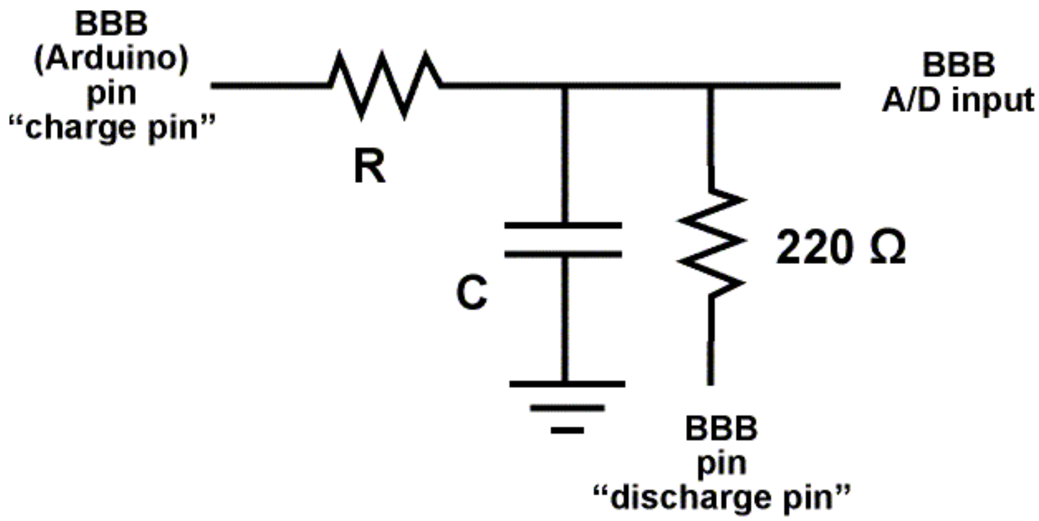

Figure 5: Initial circuit implemented to sample the capacitive humidity sensor [23]

Next stability of the node was tested. While uploads had been successful over a few hours in lab sessions, it failed a test over 24 hours. However, this was not due to a software error but caused by a lack of power. Results from a consumption test shown in Table 3 revealed that electronics on the Arduino board itself (excluding the AVR, Wi-Fi module and sensors) were drawing 75% of the overall current. This was more than expected and caused by the fact that these components could not be deactivated in the SLEEP state. By hacking the second MCU on the Arduino and deactivating components like the USB to serial converter for example, it may have been possible to save more power. Reducing the transmission frequency, storing data temporarily and transmitting larger amounts of data could have also saved power. However, a management decision was made to focus the project on visualisation and to differ from the specification by powering the node from a cable. Even when attached to a wall socket, the node was still flexible enough to be tested at several different locations and angles. Extension cables could be used if necessary.

Circuit	Current Drawn
Arduino Uno Board (AVR removed)	33 mA (measured)
Arduino Uno Board (with AVR in power down mode)	37 mA (measured)
Entire System in SLEEP State	44 mA (measured)
Wireless Shield with Sensors on-board	44 mA - 37 mA = 7 mA

Table 3: Power consumption of different parts of the node

Powered from a wall socket, the node uploaded data reliably over several days. When reviewing the data collected from different positions, the presence sensor was often either too sensitive, or indicated no presence at all. This was easily adjustable by adapting the threshold of motion samples over the transmission period of 2 minutes. In doing so the sensor was calibrated to detect presence of a user sitting at a desk. This was the most useful with respect to the field trial, since most homes have a room where people normally sit at a desk if they are in the room. Specific presence tests were performed, keeping a record of when somebody actually was at the desk. Measured times were very similar, but with a few slight disagreements when someone was in the room but not at the desk.

For the field trial, a second sensor node was constructed. However, while sensor circuitry was soldered onto the Arduino shield for the first node, it was constructed on a breadboard for the second. Circuitry was identical apart from an additional decoupling capacitor used in the breadboard implementation. Images of both nodes are shown in Figure 6.

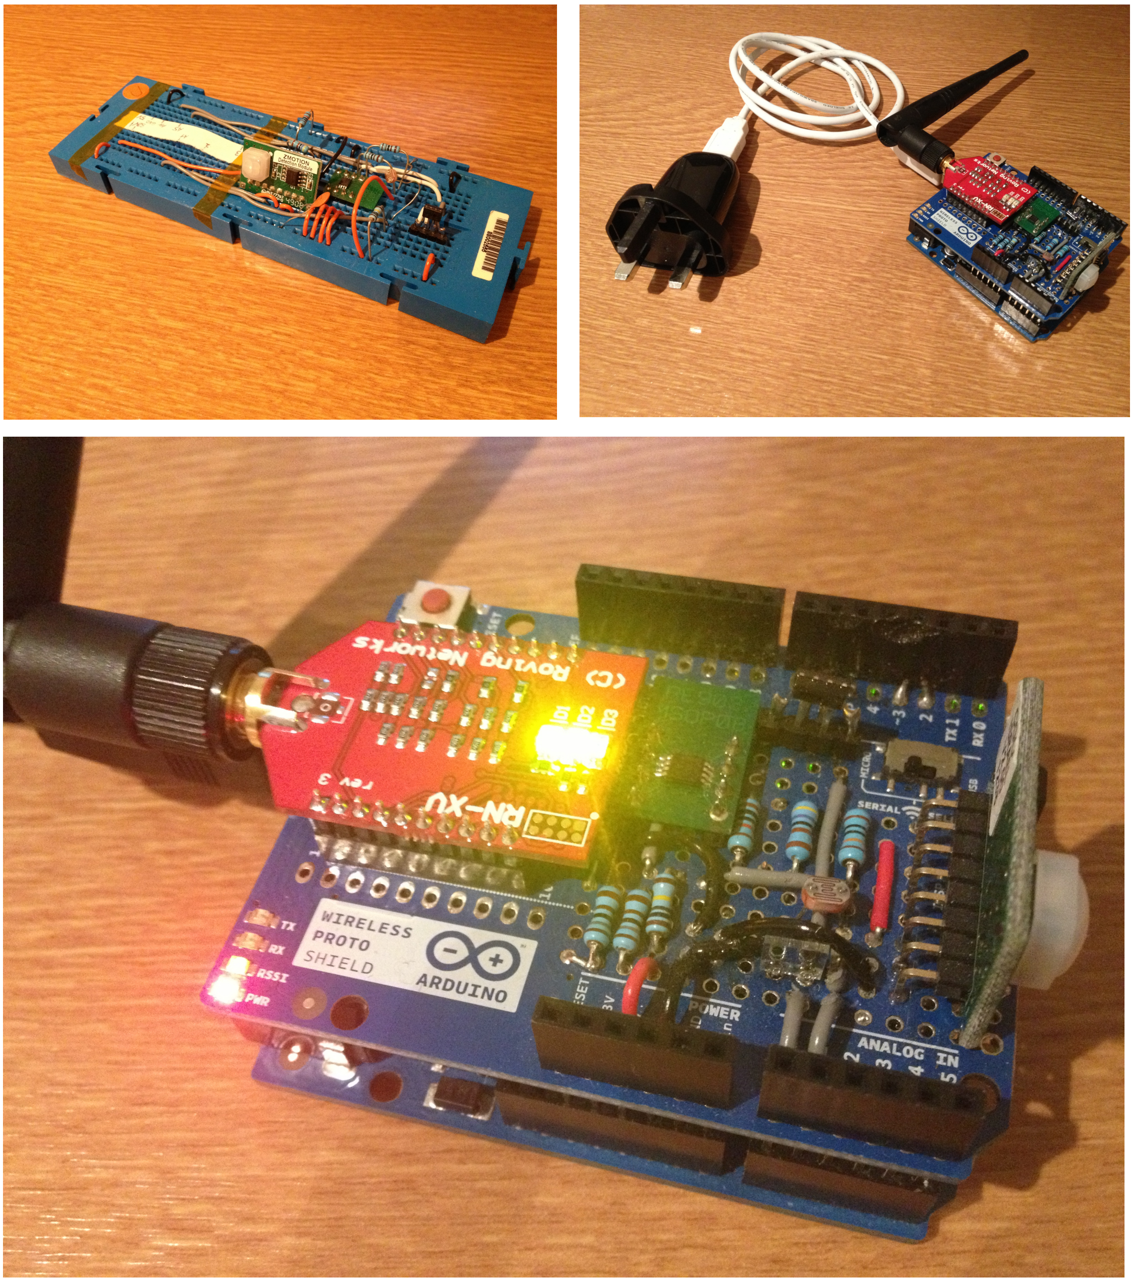

Figure 6: Implementation on the Arduino (right, bottom) and constructed on a breadboard (left)

4 Data Processing

4.1 Initial Tests in Python

Prior to constructing a node that provided a real dataset, it was necessary to gain an understanding of residential temperature development. It was important to find out what information could be gained from the data and how it could be processed automatically. Data from an energy monitoring system called AlertMe positioned in my supervisor's house was used for this. The API provided data in a CSV format. Code to process, analyse and plot the data was written in python using the Matplotlib library amongst others. Additional data for outside temperatures was loaded from the Weather Underground (Wunderground) database using its JSON API. First plots as shown in Figure 7 simply displayed inside and outside temperature for a particular day.

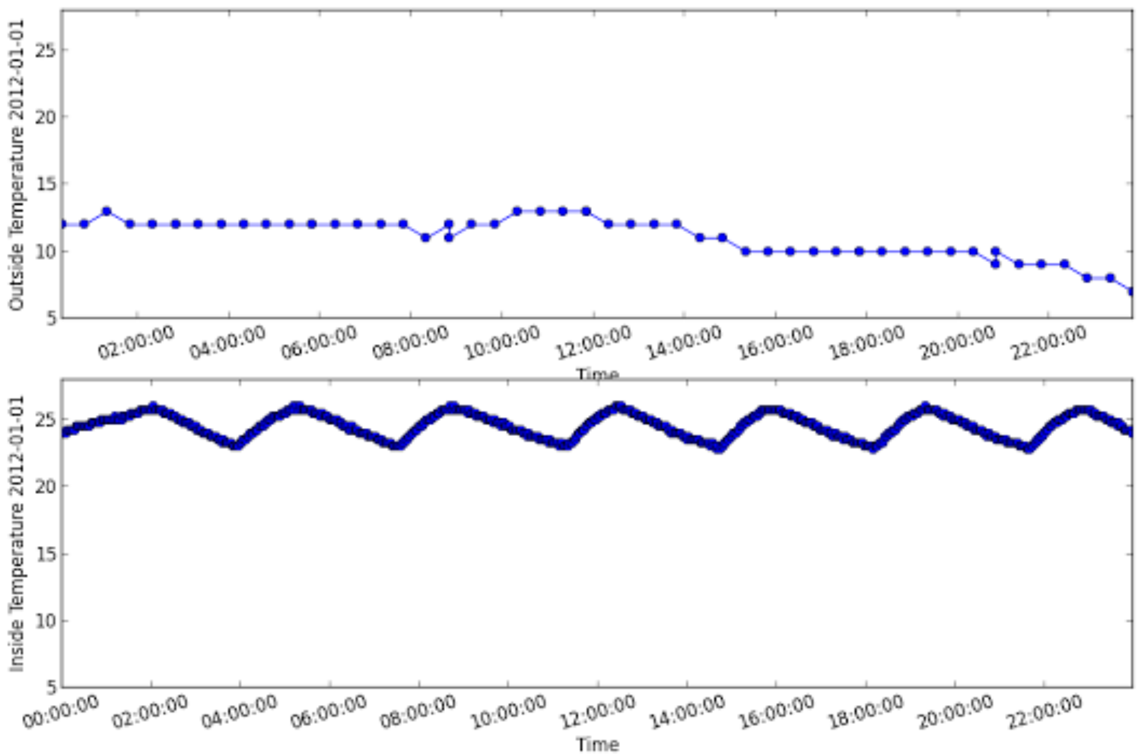

Figure 7: Initial Temperature Plots for 1/1/12

In a next step data was up-/subsampled to match the sample periods of both inside and outside temperature. Additionally, a moving average was used to 'smoothen' the graphs and reduce the effect of measurement and quantisation errors. Processing the data like this was necessary to produce a difference function and a realistic rate of change plot.

By simply using a threshold on the rate of temperature change, it could be detected when the heating had been turned on and count how many times this occurred. The threshold was set by inspection, but could be proportionally related to the outside temperate since houses warm up faster if it is cold outside and slower if it is warm. However, this depends on the heating characteristic of the individual building and how well it is insulated.

To determine the duration of a heating process, the length of a rise (i.e. how long the rate of change is above zero) was measured when a temperature change is above the threshold. This does not consider the time radiators take to heat up, since air temperature does not change instantly as soon as resources are used. However, radiators continue to stay warm after they have been switched off, so this roughly evens out the wasted energy at the start of every process.

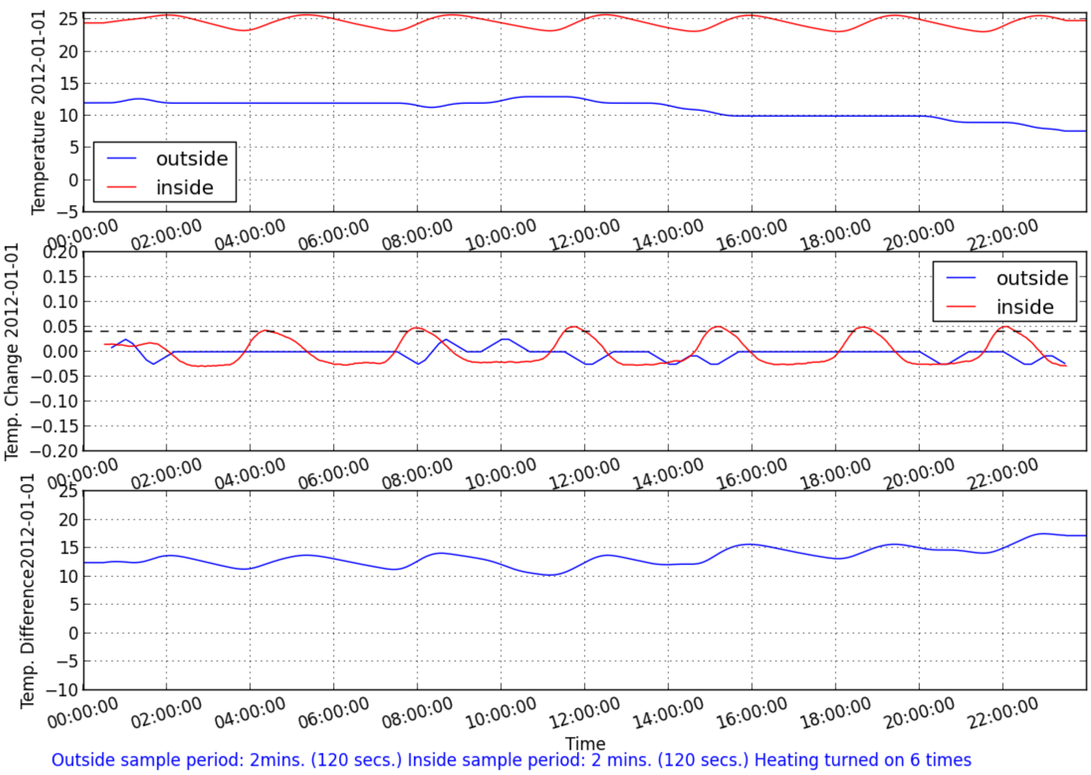

Figure 8: Latest Temperature Plot for 1/1/12

Figure 8 shows the final plots the tool produced. The most valuable discovery was that heating could be detected purely from the temperature data. By using a threshold on the rate of change, realistic numbers and durations of heating processes could be calculated for the given data set. This gives an idea of how much was consumed, all though some heating processes may be stronger than others if a thermostat or radiator allows different levels. Outside temperature proved less interesting on an hourly scale, but may be useful to calibrate heating detection from a daily mean value. A system that modelled the heating characteristics of a home by also taking the outside temperature into account was developed in 2011 [24]. This shows that automatic calibration is a realistic option for future work to improve the heat detection filter. The difference function correlated strongly with the inside temperature data, since it generally varied more than that outside, providing no relevant information.

4.2 Deployment for SmartHeat

When real data had been collected, parts of the data processing tested in Python were implemented for SmartHeat using JavaScript. The goal was to detect heating processes and measure their duration, like it had been done for the test data set. FigureEnergy already contained a processing module, which processed consumption data to be displayed clearly as a graph and hence averaged data differently depending on the zoom level. This was removed since it was not compatible with changes made to the UI and implementing the existing Python module would be quick either way.

Since the difference function had proven irrelevant, no resampling was necessary to synchronize samples between inside and outside data. However, when calculating the rate of change it was important to have regular time intervals between samples. Due to varying transmission times, data is not always uploaded at exact 2-minute intervals and transmission errors cause at least one missing sample. To avoid any issues in processing caused by irregular sampling periods, data was resampled at identical intervals, using a linear average between real samples.

After resampling, a moving average was applied to 'smoothen' the data. This assured a realistic rate of change without any short peaks due to minor measurement or quantisation errors. The width of the filter was set to 20 samples, which relates to 40 minutes. This value was found through trial and error, since it removed temperature inclines that clearly weren't the result of a heating process but kept those that were.

The rate of change was then calculated as the difference in temperature between two consecutive samples, since the time step was always 2 minutes. In processing a Boolean was added to the array that specified whether the rate was over the threshold (i.e. heating was on) or not. This made it easy to access from the UI. Additionally, the temperature for each sample was categorized as either too low (< 20°C), comfortable or too high(> 22°C). This was implemented using constant comparison values, but in future work these could perhaps be determined automatically from past behaviour, or manually by users themselves.

5 Visualisation

5.1 Implementation

5.1.1 Software Libraries

Similar to FigureEnergy, this project visualised data using the JavaScript library D3. It takes a data-driven approach to manipulating document object models in JavaScript and can create graphic, dynamic forms. Using HTML, CSS and in particular SVG models it provides powerful means of visualising data and is also open source, providing good documentation and examples online.

The Figure Energy system uses the free and open source web application framework Django. A tutorial was covered at the start of the project, but during implementation most of the existing code did not need to be altered significantly. The Django admin interface was used to add new sensors, channels and trial users to the system. It was also used to check that data had uploaded to the server correctly and likewise, that data shown in the UI was also correct. Since the database did not support presence and light, channels for these sensors had to be added. Additionally, the 'dataloader' module in FigureEnergy had to be altered to query data from four sensors and provide functions to export it to other modules.

Another JavaScript library FigureEnergy uses is jQuery. It is again free and open source and designed to simplify navigation through documents, selection of DOM objects and event handling for example. Unlike D3, it is only used in few parts of the FigureEnergy code, mainly for event handling. In SmartHeat the jQuery plugin 'tipsy' was used for the tooltips displayed when hovering over data.

5.1.2 Debugging Tools

All debugging of the UI was done in Google's browser 'Chrome', using the powerful web development tools it offers. Most importantly these include DOM object and style inspections, as well as a JavaScript console.

For deployment on the development server, a fabric script designed to deploy typical Django applications was used. It uploads all relevant files automatically, simplifying deployment to the execution of a single script. Generally an offline version was created for each different approach using hardcoded dummy data. This made changes to the code testable without having to run the fabric script first and fastened debugging significantly. Only when the core UI had been developed and was ready to be tested online with live data, was it deployed on the server.

5.2 Approaches

5.2.1 Calendar View

The calendar view was developed to become familiar with D3 using the test data set and the results of its analysis from python. It was created based on an example from the official D3 website that displayed CSV data arranged in columns by weeks and grouped by months and years [21]. Days were represented as squares and quantized to a colour scale depending on the responding value. The code was adapted to use processed temperature data exported in CSV format. A different colour scale was used that is commonly related to temperatures (red for hot, blue for cold). In addition, a wider range of colours was used to reduce the quantisation error in colouring and display data in more detail.

Figures 9 and 10 show visualised 'HeatCount' and 'HeatDuration' data. There is a strong correlation between the two, since duration increases with every additional heating process. However, especially in November and December a few days stand out in the second image that are not as noticeable in the first. This could be due to lower outside temperatures or open windows, causing the heating process to decelerate. Hence the user may have heated the same number of times for two days, but one will have a higher heating duration simply because it took longer to heat that day. Another possibility is that a thermostat was changed from programmable to manual mode and the user forgot to turn the heating off again. This results in few but very long heating processes.

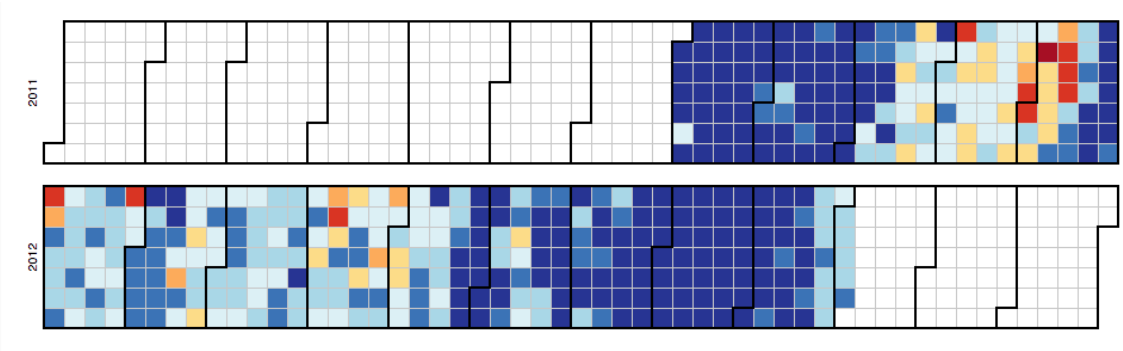

Figure 9: 'HeatCount' shows how often the heating was turned on

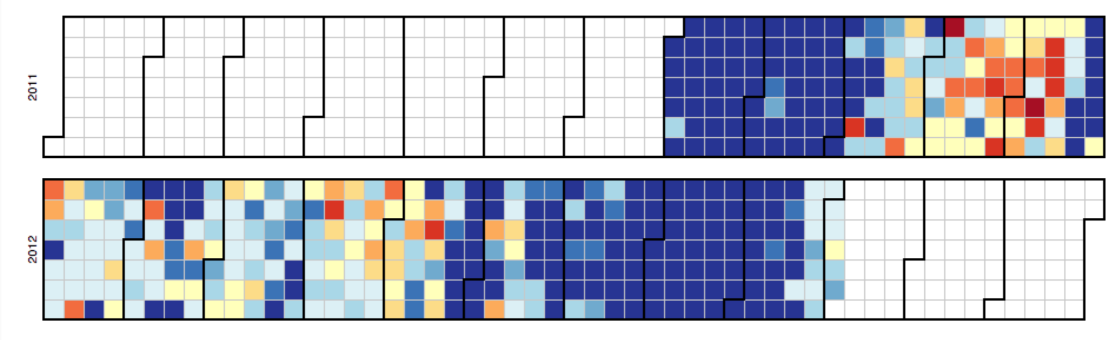

Figure 10: 'HeatDuration' shows how long the heating was turned on

A third image shown in Figure 11 was generated using the same code but for the mean outside temperature. Data was taken straight from the Wunderground database and not averaged locally. Comparing with the previous graphics, an outside temperature drop in late 2011 could be the cause for increased heating during that time. However, the days with the most heating do not always match the very cold days. This could be for several reasons such as good insulation, other factors (e.g. doors and windows) influencing inside temperature stronger or an inefficient change in thermostat setting perhaps.

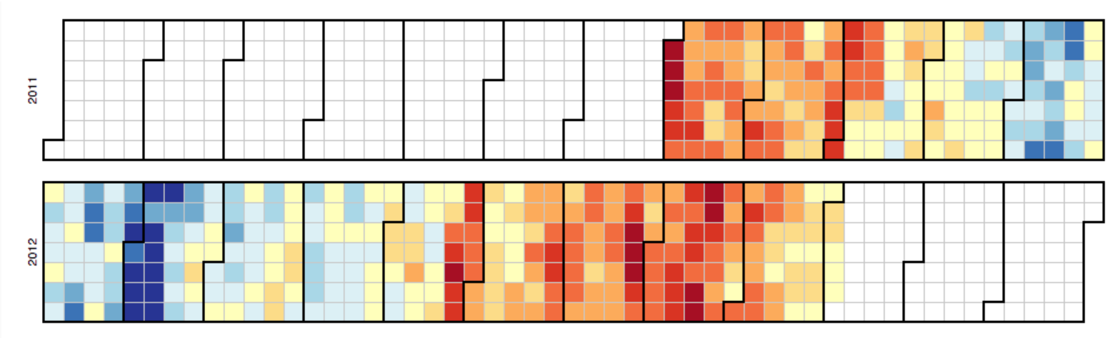

Figure 11: 'MeanTempOutside' shows what the mean temperature outside was each day

Initially developed just for learning purposes, the calendar view proved useful as a visualisation for long-term temperature development. Throughout a whole year the HeatDuration and HeatCount plots highlight the most interesting days, providing an overview of temperature data over a long period of time. However, it is restricted by the fact that only one value or colour can be shown per day, which makes it less valuable for analysing individual days.

5.2.2 Correlation Matrix

After construction of the node had been completed, the correlation matrix was the first attempt to visualise data from all sensors. Light data was processed further for this view to be either on or off. Reviewing the data had shown that room lighting gave a high constant value, while natural lighting through windows varied a lot during the day and was lower. A high threshold calibrated to positioning and lighting of the room could be used to determine the state.

An initial sketch of the layout idea is shown in Figure 12. Each square or element of the matrix simply represents how often a certain combination was true, e.g. ‘heating and light’ or ‘no heating or presence’. Rows and Columns for additional data could easily be added, for example for temperature ranking such as ‘too high’, ‘comfortable’ or ‘too low’. Hence, if the room were often too warm while the heating was on, the squares marked with an A would be highlighted. Likewise, if the room were often at a comfortable temperature when someone was present, the squares marked with a B would be highlighted.

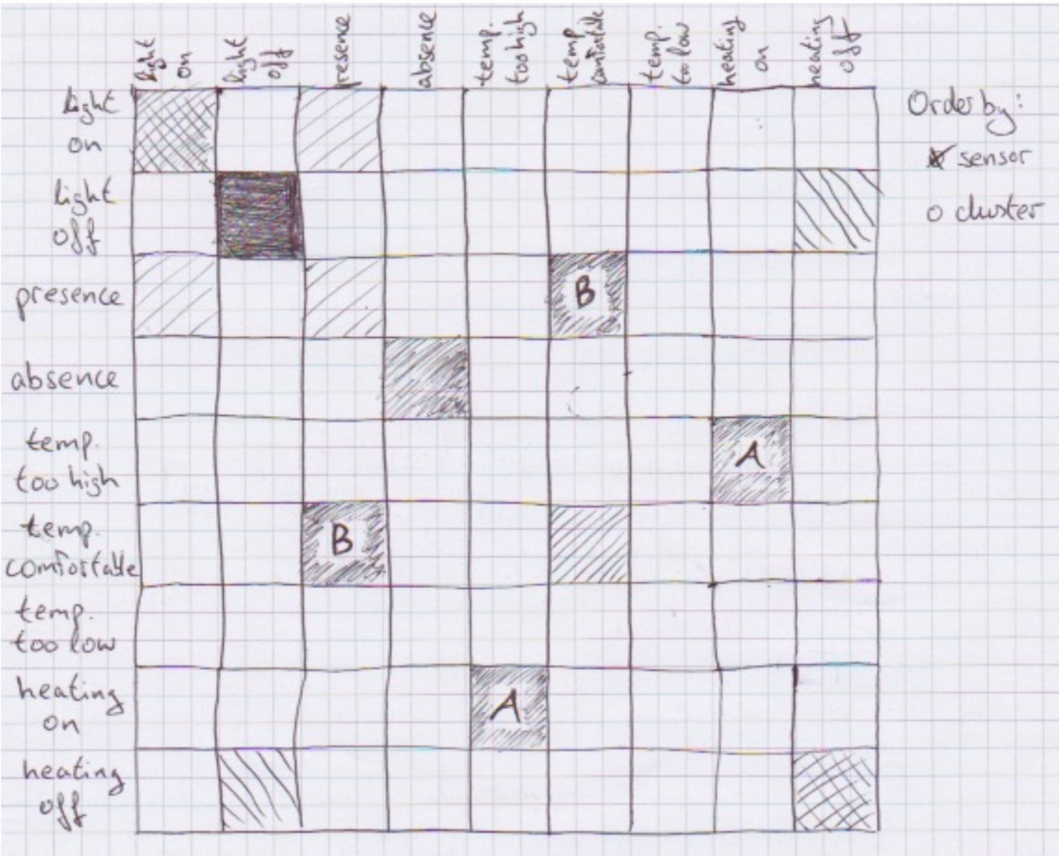

Figure 12: Initial sketch of the matrix view

Another idea was to value squares depending on something other than the frequency of occurrence of their state. For example the heating rows and columns could be removed and their correlation with all other states used to value each square, as shown in Figure 13. This visualises relationships between three sets of data, rather than just two. In that case, if the room were often heated when nobody was there and the light was off, the squares marked with a C would be highlighted red. Similarly, if the heating were never on when the room was too cold and someone was present, the squares marked with a D would be highlighted blue. Matrices could then be sorted by cluster, grouping states in which the room was often heated together.

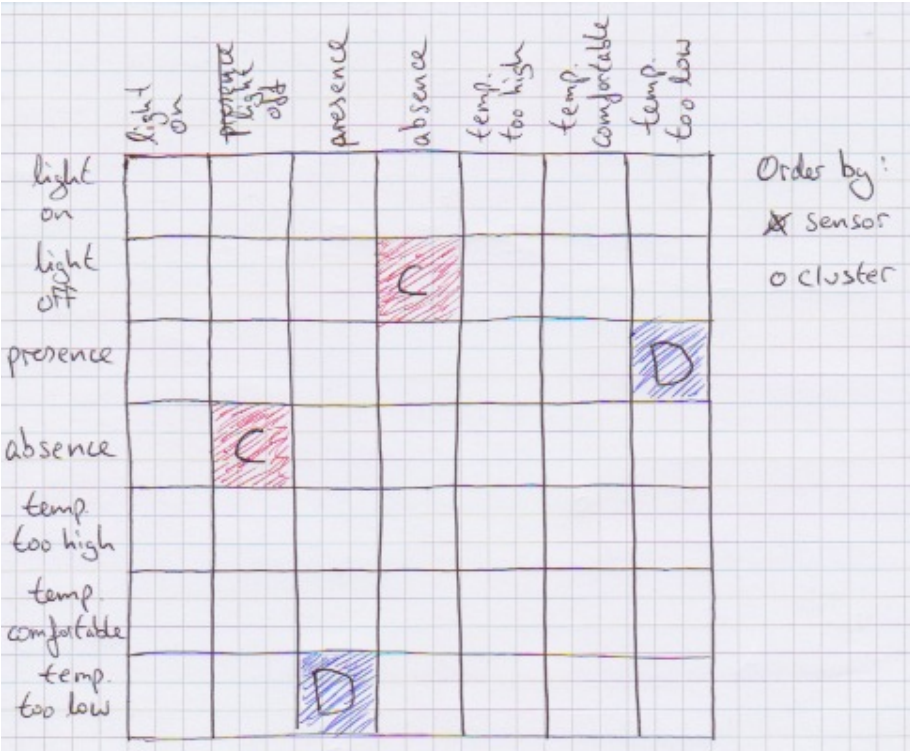

Figure 13: Second sketch of the matrix view

One advantage of this view is that data is not time dependant. For example, the data displayed could be from four different Mondays of the same month and hence give feedback on a specific weekday. It also gives a lot of emphasise to light and presence readings, and even displays a lot of information on the correlation between the two. However, relating heating behaviour to a certain state rather than a time makes it harder for users to identify when it occurred and may confuse some people. In addition, a lot of information is displayed that is not related to heating, such as correlation between light and presence. This distracts users from what they should really be looking out for.

A local prototype using a smaller set of dummy data was created for this view, shown in Figure 14. Due to its complexity it was decided that the correlation matrix was not suited for a field trial, which is why there is no server implementation with live data.

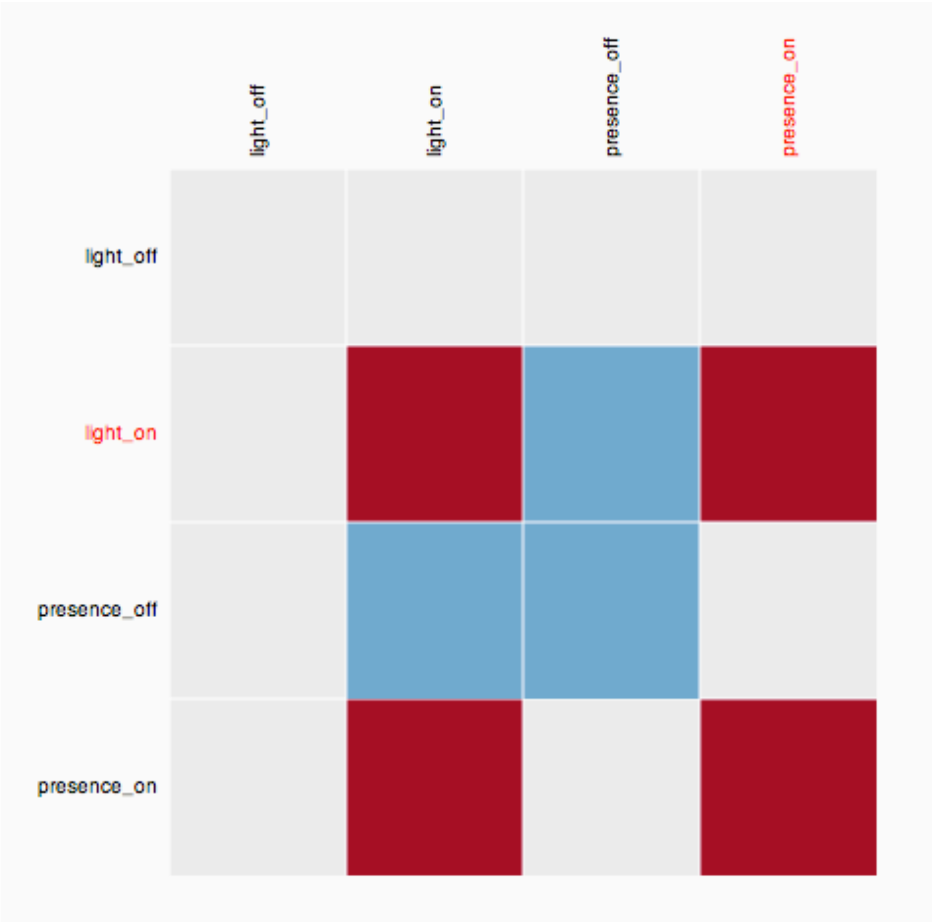

Figure 14: Offline version of the matrix view (red tags indicate hovering over a square)

5.2.3 Radial Sunburst

This view was inspired by a visualisation commonly used to display contents of a hard drive [22]. Similar to a pie chart, it is circular but data is shown in rings rather than triangles or 'slices'. In a first implementation shown in Figure 15, three colours indicated whether the temperature was too high, too low or in between. While the inner ring grouped all periods where this was the case, the outer ring was used to show them individually. By hovering over an area of a ring, the corresponding data on a classic line graph was highlighted.

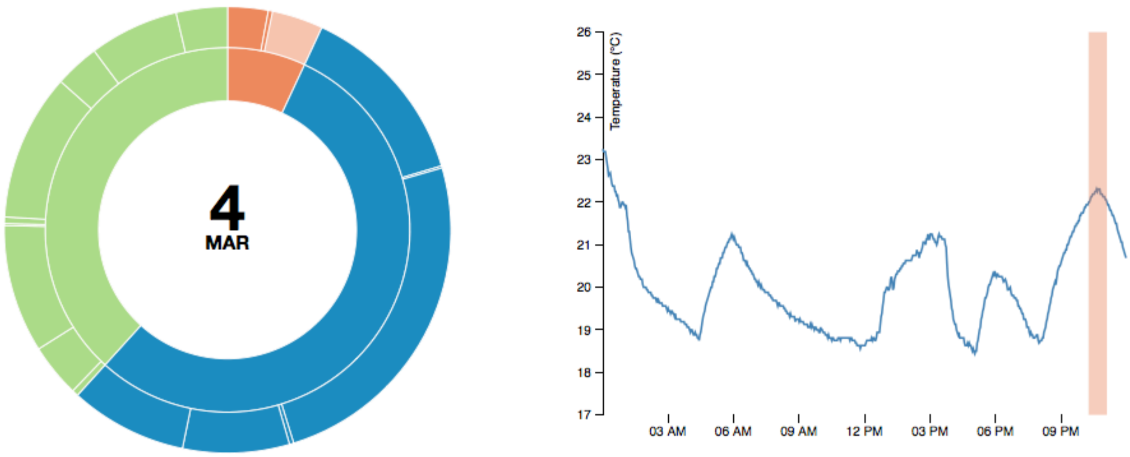

Figure 15: First Version of the Sunburst View showing three temperature categories

In a second version shown in Figure 16, the inner ring was altered to display whether the heating was on or off. The second ring still showed the temperature split into three categories as before, and the outer ring now added presence data to the visualisation. Moreover, the inner ring could be sorted by time or by group (the other rings would adapt accordingly).

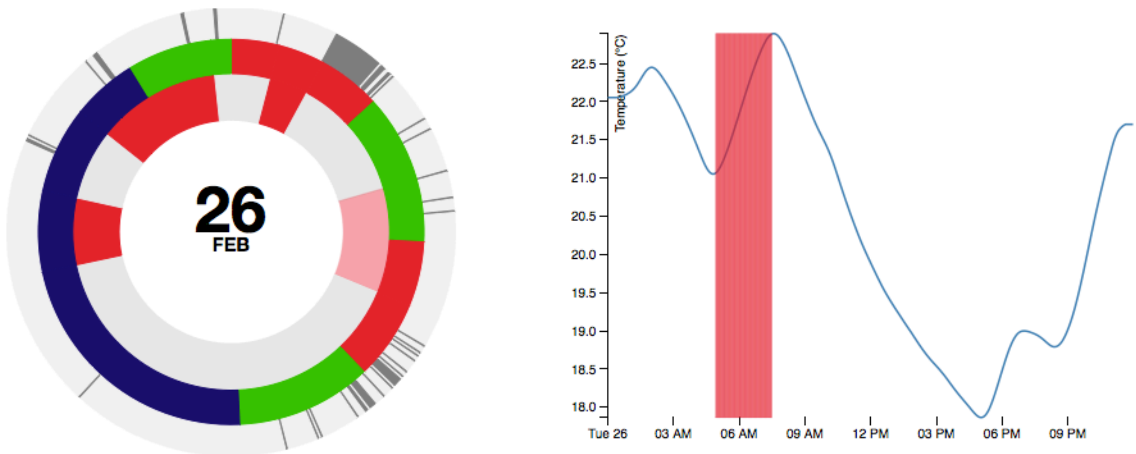

Figure 16: Second Version of the Sunburst View including heating and presence data

A third implementation used gradient colours for values over or under the thresholds in the temperature ring. Depending on the size of the integral between the temperature curve and the constant threshold value, a colour on a scale from green to red was chosen (green to blue for data below the threshold). Temperatures in the acceptable region remained green. This was a mixture of a full heatmap ring, where every temperature would be mapped to a different colour, and the 3-way rank based colouring used before. The implemented version is shown in Figure 17.

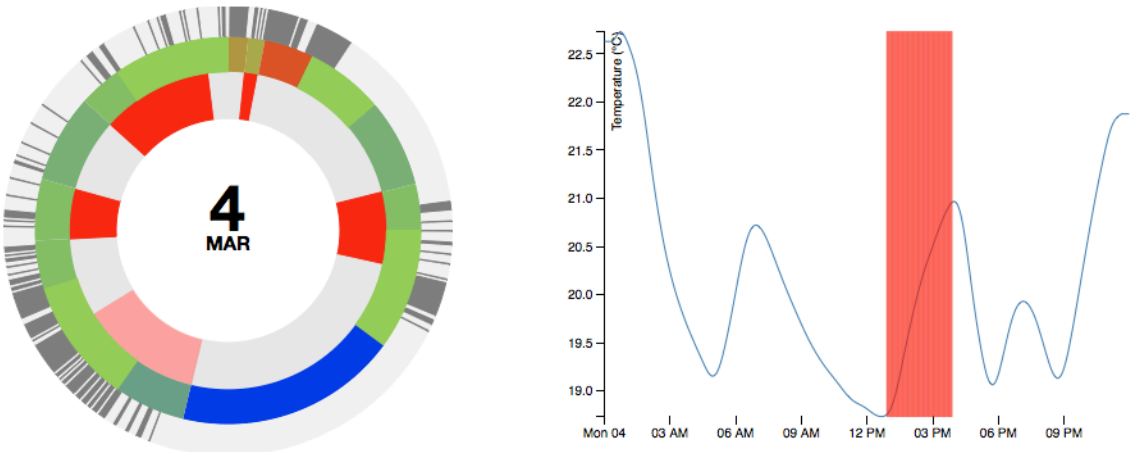

Figure 17: Third Version of the Sunburst View using a mixture of heatmap and categories

Finally a version with a full heatmap ring was implemented as shown in Figure 18. Advantages of the heatmap were that it highlighted drastic changes, as well as slow development in temperature. This made it possible to remove the line graph, since all temperature changes could now be seen in the ring. Additionally, tooltips were added when hovering over the graphic that displayed more detail on the data selected. This allowed users to read an exact temperature at a certain time for example, without using an axis as before.

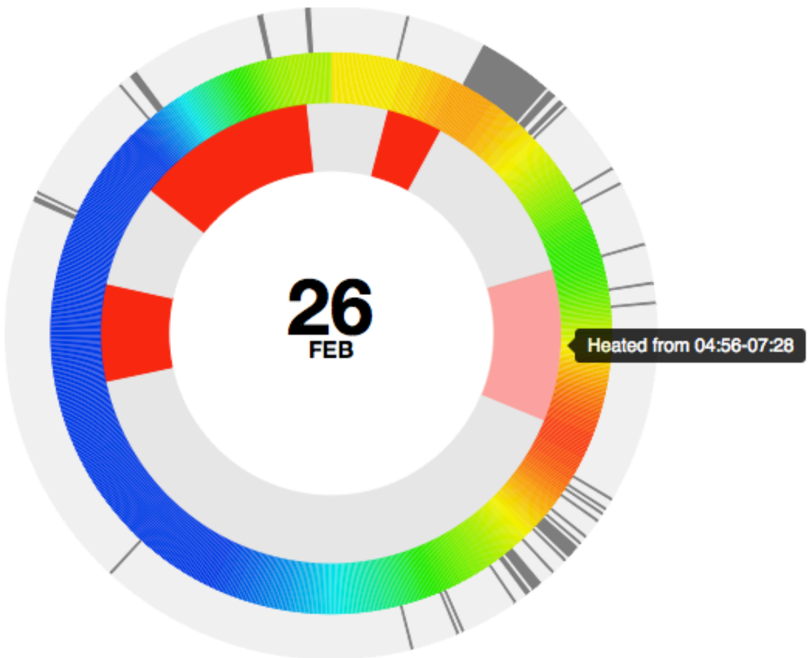

Figure 18: Final version of the Sunburst or Radial view, using a full heatmap

A key advantage of this view is that it highlights proportions well, similar to a pie chart which many users are familiar with already. If sorted by time, it is also similar to a clock, which again is a familiar concept to users. This would however limit the view to only display data from half a day (AM or PM). The ring concept also makes it easy to add data for light and humidity for example, if this was to be presented to users. In the latest form, the circular form representing the time domain is more of a disadvantage, since it can be misleading for users. This applies particularly to the fact that the start and end of a day link up at the top.

5.2.4 Linear Stack

The linear stack is very similar to the radial view, but instead of having a circular shape it is rectangular as can bee seen in Figure 19. Implementation was very simple, since most of the code from the radial view could be reused, the only real difference being the use of rectangles instead of rings. Functionality is identical.

Users are more familiar with time represented on a linear axis, which is the main advantage of this view over the radial one. A disadvantage is the loss clear proportions in a pie chart like manner, as seen in the radial view.

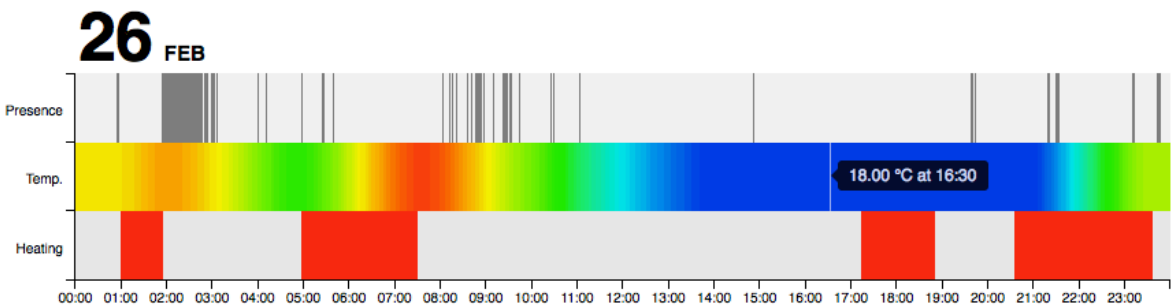

Figure 19: The linear stack view using a heatmap to display temperature

5.3 Choice of Final Design for Testing

The system implemented for the field trial combined the calendar view and the linear stack. While this was not necessarily the best option for deployment as a product, it was easiest for users to adapt to within a short period of time. The calendar view provides a quick overview of data over a long period of time and makes it easy to spot the most interesting days. This is why it was implemented as a browsing tool to provide an overview (using the heating duration of each day) and let users display data by clicking on a day of their choice. Days were then visualised in the linear stack view, which had been evaluated as the most intuitive to use due to its linear time axis.

The matrix view is the most complex and hence has a longer learning curve but could potentially highlight more interesting aspects of the data, especially over a longer period of time. Likewise, the radial sunburst may provide quicker understanding to experienced users or others who are regularly confronted with visual forms of data. For the purpose of this trial, these should not be anticipated.

Light and humidity data was excluded for the trial. Reviews of both sets of data showed that they provided little context to heating behaviour. Humidity data for example was very similar to temperature data, since air can hold more moisture at higher temperatures. As a consequence it did not add any value to the rest of the data. However, these sensors could still be applied in future work, as described in section 7.3.

The UI presented to users in the field trial is shown in Figure 20. Hovering over a day in the calendar view displays a tooltip showing how long the heating was on during that day. Moreover, days are coloured from white to red according to that value. Clicking a day loads the respective data into the linear stack, where another set of tooltips highlight the meaning of each individual rectangle when hovered over.

Figure 20: Overview of the UI used for the field trial

6 Field Trial

6.1 Goals and Protocol

The main goal of the field trial was to discover whether presenting consumers with visual feedback on their heating behaviour could improve their understanding and efficiency. Another aim was to test the SmartHeat system developed in this project, in particular to find the strengths and weaknesses of its UI.

To evaluate this, the wireless sensor node was deployed at a desk for 3 days. Participants received access to the SmartHeat interface and could view their activity online. In addition, interviews with each participant where held before and after deployments. These were designed to explore the participant's understanding of how the home was heated before the trial, and whether it had improved through use of the system.

6.2 Study Execution

The study took place between the 28th March and the 21st April 2013 in Hamburg, Germany. Two trials were performed simultaneously, since there were only two sensor nodes available. To provide three full days of data at least one day was left between trials for interviews, removal and installation of the nodes. Participants were all university students, male and friends of mine. There were 6 participants in total, 4 of which studied engineering related subjects. All residencies were apartments, which had a desk in the participant's bedroom where the node was deployed as shown in Figure 21.

From a technical point of view, the study was very successful. The node worked well with a variety of different Wi-Fi networks and submitted data reliably throughout the trial. One issue occurred with the node constructed on a breadboard, as it stopped transmitting data at one point when a power cable came loose. However, this was more due to the limitations a breadboard prototype has and the user ignoring instructions not to move the node, than a technical design fault. The user interface did not crash at any point and was always accessible to all participants.

6.3 Results and Evaluation

All interviews with participants were recorded to analyse them in more detail once the study had been completed. Statements were tagged to help group related information from different people. Tags used were data correctness, awareness, understanding, pattern recognition and suggestions. The following summarises information gained for each section, differentiating between understanding prior to and after the trial. Pattern recognition by users was included in section 6.3.4, as it confirmed comprehension.

Figure 21: Typical deployment of the sensor node at a desk during the field trial

6.3.1 Understanding prior to deployment

Initial interviews showed that all participants knew what type of thermostat they used and how it was configured. Four participants had thermostats directly on individual radiators that maintained one temperature in the room. The temperature could not be selected specifically, but on a scale from 1 to 5. The other two participants regulated their heating completely manually by adjusting their radiators directly. Naturally, those with thermostats changed the settings less frequently than those without. An interesting observation was that these two users already received feedback on the room temperature through alarm clocks. All participants said they changed the settings if they were uncomfortable, i.e. the room was too hot or too cold. Two said they also turn the heating down or off at night, and back on in the morning. One of those also turned the heating down if he left the apartment for several hours. Five participants believed they could improve the way they heated using a programmable thermostat and better isolation. Only two believed this was possible just through a change of behaviour.

6.3.2 Sensor Accuracy and Subjective Perceptions

Generally, data collected by the sensors was believed to be correct by participants. However, in three cases users thought the heating had been detected incorrectly. Two of them mentioned that the heating was on a low level throughout most of the day without showing so on the system. These had also noticed a very strong correlation between presence and heating. The third user said that some heating processes were wrong and that the heating had definitely been off at those times. In all three cases, the heating detection had not been calibrated ideally. For the first two cases the threshold was too low to detect the heating, even on a high level. It is most likely that the rooms heated up much faster with the participants present and that only these times exceeded the threshold, explaining the correlation between heating and presence. Users also felt that devices such as computer or TVs increased the temperature in the room, again adding to temperature rises during user presence. For the case where wrong processes were shown, the threshold was probably too low. This may have caused temperature rises through presence to exceed the threshold, even though the heating was off. Sunshine is most likely the reason for false heating detection with no user present. Although a few heating processes were incorrect or not shown in these three cases, a majority of the data presented to participants was still correct. They were aware which parts were incorrect, so their statements could still be considered for this trial.

6.3.3 Effects on Awareness

All participants said that using SmartHeat had increased their awareness of how they heated and how much they consumed. With two users only using the UI on the last day of the trial and four doing so at least once per day, frequency of interaction did not seem to change understanding or awareness. The people that used the UI the most were in general those present in the room the longest and moreover, those that manually adjusted their radiators. During the trial, more users turned off their heating at night than had said to do so previously.

6.3.4 Understanding and Discoveries

Having been instructed on how to use the system initially, all participants remembered how the UI worked after the trial and could describe all features. They had mainly used the time domain view and said the cluster view was confusing. One user had not discovered the percentages shown when hovering over field of the cluster view and said they could have been interesting, had he discovered them earlier. When asked at which points during the trial they had been wasting energy, all users pointed out at least one occasion. All users said that heating was wasted if no one was present. Three participants additionally pointed out points where the room was already hot when the heating was turned on. A selection of wasteful behaviour identified by users is shown in Figure 22.

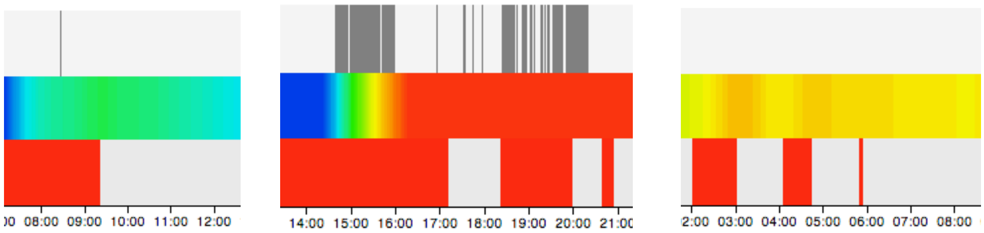

Figure 22: Users said heating was unnecessary when the room was warm or nobody was present

All three participants for whom heating was detected correctly found ways to improve their behaviour in the future. Two of them did not believe this was possible without upgrading the thermostat prior to using the system. One participant said (translated):

> *I can turn my heating off earlier at night because it's warmer than necessary when I go to bed. If I hadn't seen that, I wouldn't have thought of it.*

The visualisations this specific conclusion was drawn from are shown in Figure 23.

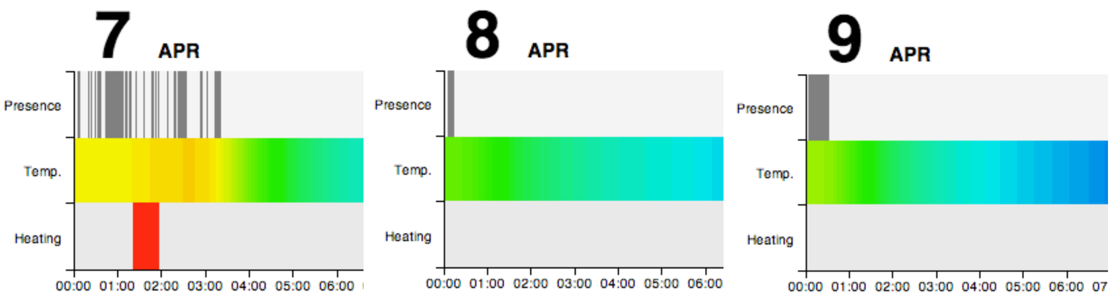

Figure 23: Extracts from a users data showing comfortable to warm temperature at night

Another user who controlled his radiator manually said:

> *When I heat, I need to stop when it is warm enough. A lot of the time I turn it off too late and it gets too hot because the radiator stays warm.*

Two users also spotted times when they opened a window, shown in Figure 24. The same users also related temperature rises to sunlight shining through windows. These were generally times when they were not present and had hence turned the heating off, but temperature still rose. A common behaviour amongst all users when presenting them with the UI was that they would link activities to certain times. They pointed out when they woke up, left the apartment or went to bed for example and this helped them understand their behaviour.

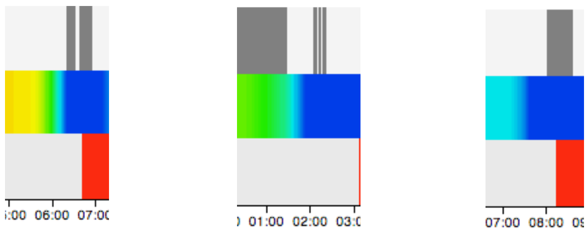

Figure 24: Users identified open windows from presence and sudden temperature drops

6.3.5 Suggestions from Participants

Users that could not conclude how to change their behaviour in the future claimed that 3 days were not sufficient for productive use of the system. They believed however, that a longer deployment would include more wasteful behaviour and hence highlight areas for improvement. A suggested UI improvement by two users was a customisable colour range for temperature, since different people have different comfort zones. With respect to implementing the system as a commercial product, participants mentioned that viewing data for multiple rooms of a house or apartment would be interesting. Some users also suggested including thermostat control in the UI. In particular, users were interested in controlling room temperature via a mobile phone application.

7 Reflection and Further Work

7.1 Project Management

7.1.1 Time Management

Appendix A shows three Gantt charts. An initial chart, another version updated as part of the interim report and a final version, showing all work completed for the project. Generally, Gantt charts were created regularly every couple of weeks to keep track of project progress and to ensure I stayed on schedule. This led to some very important management decisions, such as powering the node from a wall socket, or using a more expensive humidity sensor to save time. Nevertheless, more time was still spent developing the sensor node than initially planned. This left less time for design of the interface, but work was still completed in time to allow for the field trial.

The first main challenge I anticipated in the project was acquiring new programming languages, mainly for developing the UI. This is reflected in the Gantt charts, were some extra time was appointed for such learning processes. Secondly, I knew it would be difficult to produce a UI that was intuitive to use and would not require a long time to become familiar with. To prepare for this challenge, I allowed a lot of time for development of the visualisation.

Along with all Gantt charts, any relevant outcomes and progress of the project were posted on an online blog[2]. This made it easier to look back on work completed when writing reports. Moreover, the blog was a good medium for my supervisor to keep track of the project and a platform for communication.

7.1.2 Project Costing

A spreadsheet summarizing the cost of this project is shown in Appendix B. Most components were ordered in November, however some were left out due to budget limitations. Since the budget could only be extended once, I initially planned to wait with my request in case any additional components required were discovered whilst prototyping. Ultimately an extension was not necessary, since my supervisor supplied the missing components and I also reused some from previous work.

[2] The blog was still online when this report was submitted at 3ypjmorrice.blogspot.co.uk

7.1.3 Risk Assessment

The key risks were identified as those shown in Table 4. Each risk was scored on a scale from 1 to 25. Scores were calculated as the product of the potential loss through the risk and the likeliness of this occurring, scored from 1 to 5 respectively.

Risk	Loss /5	Probability /5	Risk /25	Solution
Loss of Source Code	5	2	10	Backups and version control using Git repositories[3] for project parts including code
Hardware Damage to the Prototype	4	2	8	Complete laboratory risk assessment, store prototypes securely, preferably in labs
No need to heat at time of field trial due to high outside temperatures	4	4	16	Use previously recorded data to evaluate effectiveness of the UI with users
Extra components required late in the project	3	3	9	Do not increase the budget until late in the project, so extra components can be included
Trial failure due to lack of power or WiFi	4	3	12	Advise participants not to unplug the node or turn off their WiFi, and monitor activity during the trial

Table 4: Risk analysis of the project relating loss, probability and risk scores

7.2 Review of Achievements

A robust wireless sensor node was developed and tested in multiple different environments. It successfully transmits temperature, humidity, light and motion data to the FigureEnergy database correctly over long periods of time. Most of the design uses open-source hard- and software, enabling replication for similar projects or by an online community for example.

Moreover, a filter was developed that detects heating processes reliably for a wide range of data using only temperature data. It is simple to implement and worked reliably in most cases tested, providing it was correctly calibrated.

[3] The repositories can be found at: https://bitbucket.org/morricj/smartheat_arduino
https://bitbucket.org/morricj/smartheat
https://bitbucket.org/morricj/heatcount
https://bitbucket.org/ecostanza/fe
https://bitbucket.org/ecostanza/fe_protected

Four different approaches to visualising the data collected were designed and implemented. Their strengths and weaknesses were evaluated respectively, showing that sacrifices had to be made either for ease of learning, or deeper value of the data presented.

Finally, a complete UI providing access to a wide range of data was designed and tested in a field trial. The entire system functioned reliably and remained stable throughout the trial. Users gave constructive feedback on strengths and weaknesses of the web interface. Half of the participants were able to draw at least one valuable conclusion from using the system for three days to reduce their heating consumption in the future.

7.3 Scope for Future Work

To deploy the system as a commercial product it will be essential to remove the Arduino board and design a standalone circuit around the AVR. Only this way will consumption be low enough to enable powering the node from a battery for several years. Moreover, there is room for improvement regarding presence detection. Ideas to try are detecting presence through sound using a cheap microphone, or electricity consumption, which is already supported by the FigureEnergy database. Finally a key addition to the hardware would be to interface to a thermostat, so that users can change settings straight from the UI.

For the software, it will be important to improve calibration of the heat detection. False detections were triggered through sunshine and users heating the room through their body temperature. These factors could be filtered out by detecting sunshine using the light sensor on-board, and processing presence more intelligently. Likewise, humidity may provide more relevant information if the temperature component was filtered out. Finally, since many users could link their activities to the data shown in the UI, event annotation similar to that in FigureEnergy may be a valuable addition to the system.

To determine whether the system can actually reduce energy consumption over a long period of time, a larger field trial is required. A user study similar to that in the FigureEnergy project would be appropriate, where the system was installed in 12 homes for two weeks. Another way of testing the system would be to share it with the Arduino and ubicomp communities online, by making code and circuitry available for free.

7.4 Conclusion

In this project a complete eco-feedback system was developed, including a reliable and robust sensor node, as well as an interactive web interface. A field study showed that visualising temperature data improved user understanding and in some cases, generated ideas to improve efficiency.

Following this research, the concept of reducing consumption from space heating through feedback of temperature data remains plausible. It has been shown that visualising heating times, temperature and presence provides sufficient context for users to detect where and when they waste energy. Moreover, it also led half the participants of the trial to ideas on how to behave more efficiently in the future. The system developed is robust enough for deployment in a larger trial, making this a technically feasible opportunity for the future. Longer deployments could show whether visualisation actually changes user behaviour to reduce long-term domestic energy consumption.

7.5 Critical Evaluation

This project primarily set out to build an eco-feedback system for temperature that emphasized where and when energy was being wasted for space heating in a domestic home. That goal was achieved, as described in the previous section. Another initial idea mentioned in the project brief was a control loop through an interface with a thermostat. However, it became clear quite early in the project that adding this would exceed the possibilities of a third year project. As a consequence it was removed as part of the progress report in December.

The second goal was to evaluate the effect the visualisation had on users. This was achieved in form of a field trial over 4 weeks. User understanding could be monitored through multiple interviews, that also helped determine the strengths and weaknesses of the UI developed.

Personally, I enjoyed being challenged in such a wide variety of tasks: Designing hardware to specification on the one hand, but being creative on the other by designing an innovative UI and finally, even interviewing users of the system. I learnt two new programming languages for the project and gained a lot of experience in web development, an area I was not familiar with before. On the whole I was particularly pleased with the progress of the project, since it was easy to get hooked on details. If I could complete the project again, I would try to spend more time on visualisations, since the node development took slightly longer than expected. I had more ideas on how to visualise data that I would like to have tested.

References

[1] P. Bertoldi and B. Atanasiu, "*Electricity consumption and efficiency trends in the enlarged European Union.*", Tech. rep., European Commission, Institute of Environment and Sustainability, 2007.

[2] Prothermostats (2006) [Online]. Available: http://www.prothermostats.com/history.php

[3] M.J. Nevius and S. Pigg, "Programmable Thermostats That Go Berserk: Taking a Social Perspective on Space Heating in Wisconsin", *Proceedings of the 2000 ACE Summer Study on Energy Efficiency in Buildings*, Washington D.C., 2000, pp.8.233-8.244.

[4] D. Shiller, "Programmable thermostat program proposal", U.S. EPA, 2006.

[5] RLW Analytics, "Validating the Impact of Programmable Thermostats", RLW Analytics for GasNetworks, Middletown, CT.

[6] A.K. Meier et al., "How People Actually Use Thermostats", Lawrence Berkeley National Laboratory and University of California Davis, Berkeley, CA, 2012.

[7] T. Peffer et al., "How people use thermostats in homes: A review", *Building and Environment*, vol. 46, pp. 2529-2541, May 2011.

[8] A.K. Meier et al., "Thermostat Interface and Usability: A Survey", Lawrence Berkeley National Laboratory, Berkeley, CA, Feb. 2011.

[9] Nest Labs (2012, October 19). *How is the Nest Learning Thermostat different from programmable thermostats?* [Online]. Available: http://support.nest.com/article/How-is-the-Nest-Learning-Thermostat-different-from-programmable-thermostats

[10] E. Costanza et al., "Understanding Domestic Energy Consumption through Interactive Visualisation: a Field Study", in *Ubicomp 2012, 14th ACM International Conference on Ubiquitous Computing*, Pittsburgh, PA, 2012, pp. 216-225.

[11] S. Darby, "Making it obvious: designing feedback into energy consumption", in *Proceedings of the 2nd International Conference on Energy Efficiency in Household Appliances and Lighting*, Naples, Italy, 2000.

[12] T. Ueno et al, "Tsuji Effectiveness of displaying energy consumption data in residential houses Analysis on how the residents respond", in *European Council for an Energy Efficient Economy (ECEEE 2005) Summer Study*, 2005, pp. 1289-1300.

[13] J. Petersen et al., “Dormitory residents reduce electricity consumption when exposed to real-time visual feedback and incentives“, *International Journal of Sustainability in Higher Education*, vol. 8, no.1, pp. 16-33, 2007.

[14] G. Fitzpatrick and G. Smith, “Technology-Enabled Feedback on Domestic Energy Consumption“, *Pervasive Computing, IEEE*, vol. 8, no. 1, pp. 37-44, 2009.

[15] B. Fogg, “Computers as Persuasive Social Actors“, in *Persuasive technology: Using Computers to Change What We Think and Do*, Morgan Kaufmann, 2002, ch. 5.

[16] U.S. Energy Information Administration, Residential Energy Consumption Survey 2005.

[17] M. Jahn et al., “EnergyPULSE: Tracking Sustainable Behavior in Office Environments”, *Proceedings of the 2nd International Conference on Energy-Efficient Computing and Networking*, 2011, pp. 87-96.

[18] J. Lu et al., “The Smart Thermostat: Using Occupancy Sensors to Save Energy in Homes”, in *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, 2010, pp. 211-224.

[19] D. Roveti (2001, July 1). *Choosing a Humidity Sensor: A Review of Three Technologies* [Online]. Available: http://www.sensorsmag.com/sensors/humidity-moisture/choosing-a-humidity-sensor-a-review-three-technologies-840

[20] S. Tozlu and M. Senel, “Battery lifetime performance of Wi-Fi enabled sensors”, in *Consumer Communications and Networking Conference (CCNC), 2012 IEEE* , vol., no., pp.429-433, 14-17 Jan. 2012.

[21] R. Wicklin and R. Allison (2009). *Visualizing Domestic Airline Traffic with SAS Software* [Online]. Available: http://stat-computing.org/dataexpo/2009/posters/wicklin-allison.pdf

[22] John Stasko (2013). *Sunburst Partition* [Online]. Available: http://bl.ocks.org/mbostock/4063423

[23] Arduino (2013) [Online]. Available: http://arduino.cc/en/Tutorial/CapacitanceMeter

[24] S. Ghosh et al., “Adaptive Home Heating Control Through Gaussian Process Prediction and Mathematical Programming.”, in, *Second International Workshop on Agent Technology for Energy Systems,* 2011, pp. 71-78.

Appendix A: Gantt Charts

A.1 Initial Gantt Chart

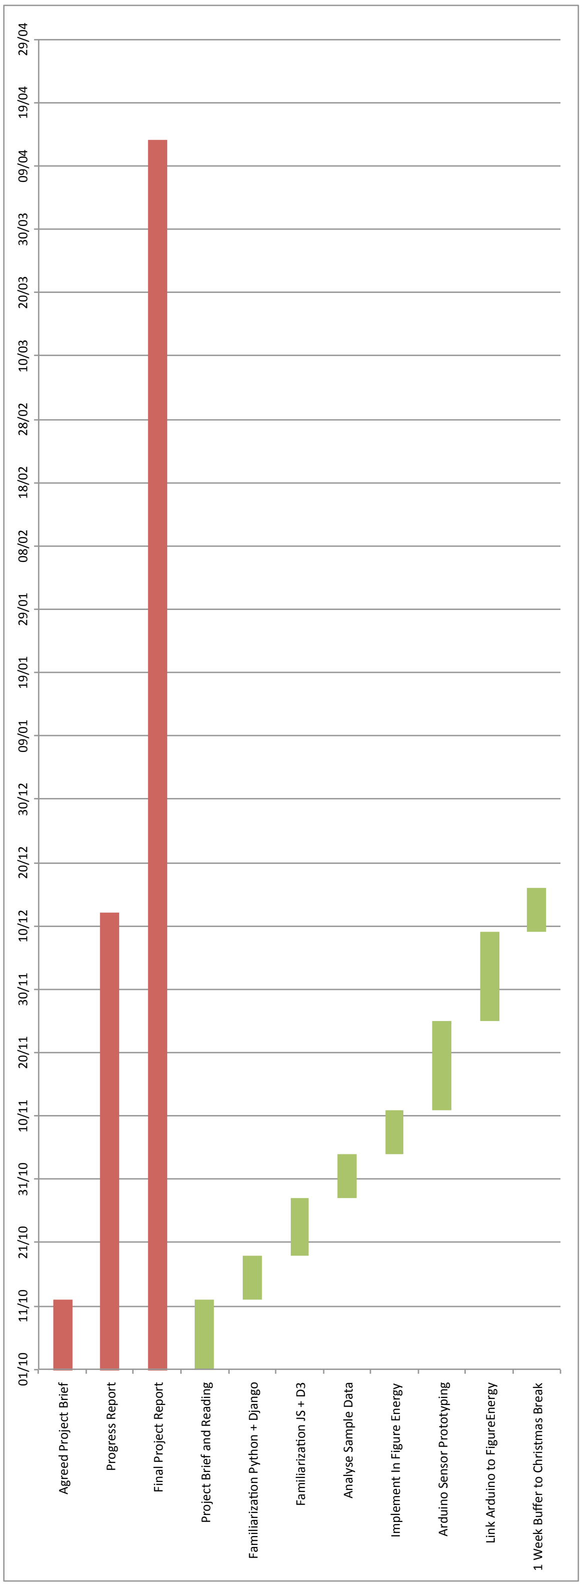

A.2 Gantt Chart at Interim Report Stage

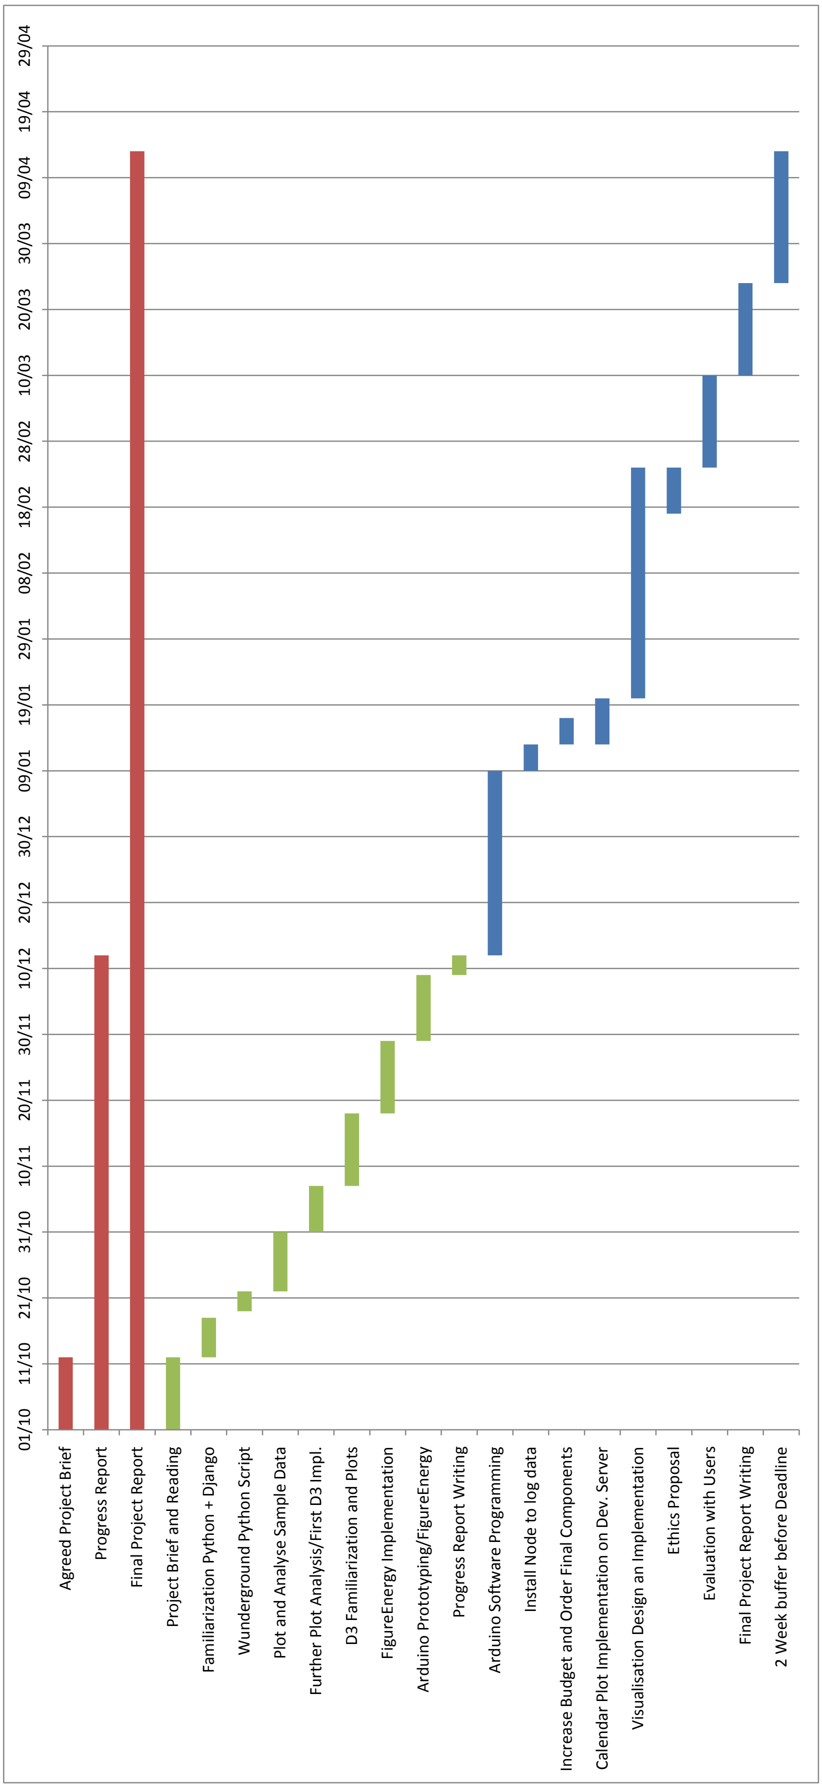

A.3 Final Gantt Chart

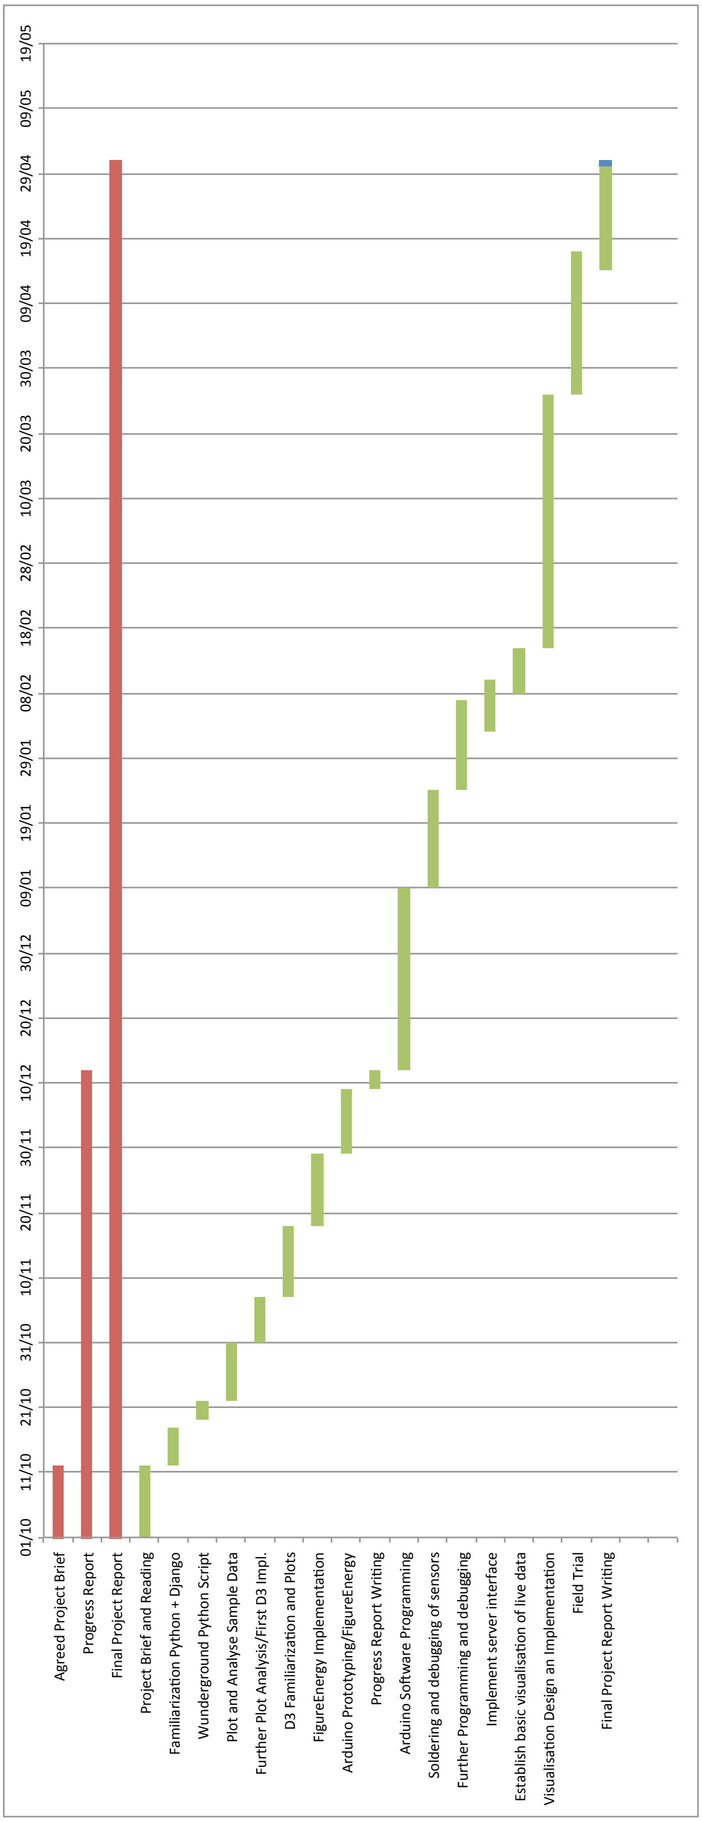

Appendix B: Budget Calculations

Store	Name	Stock Code	Price (£)	Quantity	Total excl. VAT	Total incl. VAT	Comments	Budget Remaining
								£ 80.00
OneCall	Arduino Uno	1848687	20.860	1	£ 20.86	£ 25.03	Ordered 26/11/12	£ 54.97
OneCall	LDR	3168335	0.590	2	£ 1.18	£ 1.42	Ordered 26/11/12	£ 53.55
OneCall	2.1mm Power Plug	1737256	0.390	1	£ 0.39	£ 0.47	Ordered 26/11/12	£ 53.08
OneCall	9V Battery Connector	1650675	0.600	1	£ 0.60	£ 0.72	Ordered 26/11/12	£ 52.36
RS	Arduino Wireless Proto Shield	748-5442	12.220	1	£ 12.22	£ 14.66	Ordered 26/11/12	£ 37.70
Digi-Key	ZMotion PIR Detection Module	269-2513-ND	4.48	1	£ 4.48	£ 5.38	Ordered 26/11/12 - Failed due to £50 min.	£ 37.70
Digi-Key	Microchip WiFi Module RN-XV	740-1044-ND	26.12	1	£ 26.12	£ 31.34	Ordered 26/11/12 - Failed due to £50 min.	£ 37.70
OneCall	Microchip WiFi Module RN-XV	2144012	22.410	1	£ 22.41	£ 26.89	Ordered 27/11/12	£ 10.81
OneCall	Microchip RN-SMA-4 Antenna	2143321	4.080	1	£ 4.08	£ 4.90	Ordered 27/11/12	£ 5.91
OneCall	Capacitive Humidity Sensor	SN36416	3.96	1	£ 3.96	£ 4.75	Ordered 27/11/12	£ 1.16

Appendix C: Project Brief

COMP3020 Individual Project: Project Brief

Jonathan N S Morrice
jnsm1g10@ecs.soton.ac.uk
BEng Electronic Engineering
Tutor: Dr. Maurits de Planque

Project Title: SmartHeat: An Open Source Platform for Interactive Visualization of Sensor Data

Project Supervisor: Dr. Enrico Constanza

Modern heating control systems use an intelligent agent to analyze the temperature development in a building and then adjust a thermostat to optimize consumption. We believe however, that by adding few, cheap sensors to these systems efficiency can be improved significantly. By putting temperature data into a larger context, we can improve the system's intelligence and visualize the heating process to enable much better user interaction.

This project tackles the key problem of visualizing such data and emphasizing where and when energy is wasted. Ultimately SmartHeat should enable the user to easily understand and hence optimize the temperature development of a building. It also aims to make the system more proactive by interfacing to thermostats and thus enabling it to put conclusions straight into practice.

Core of the attempt to visualize the data is answering questions we have about the data. This is done using additional sensors and hence putting the data in the right context. The outside temperature, as well as data from some additional sensors (e.g. humidity and occupancy) is taken into account. All data will be transmitted to a main server, where it will be analyzed to present to the user. In a next step more questions can then be answered by interacting directly with the user.

To close the loop a thermostat interface lets SmartHeat control the heating and improve efficiency either automatically or via the user. Since thermostats work at high voltages in the UK, mechanical hijacking is a more realistic and faster option for this project [1]. Prototyping will be done using Lego and possibly a 3D printed design. The code for the actual visualization will be based on the FigureEnergy project, which monitors electricity consumption in a similar fashion [2].

[1] Scott Davidoff, "Mechanical Hijacking: How Robots Can Accelerate Ubicomp Deployments", *UbiComp '11 Proceedings of the 13th international conference on Ubiquitous computing*, pp. 267-270, 2011.

[2] Enrico Constanza, "Understanding domestic energy consumption through interactive visualization: a field study", *UbiComp '12 Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, pp. 216-225, 2012.

Appendix D: Selected Screenshots from the User Study

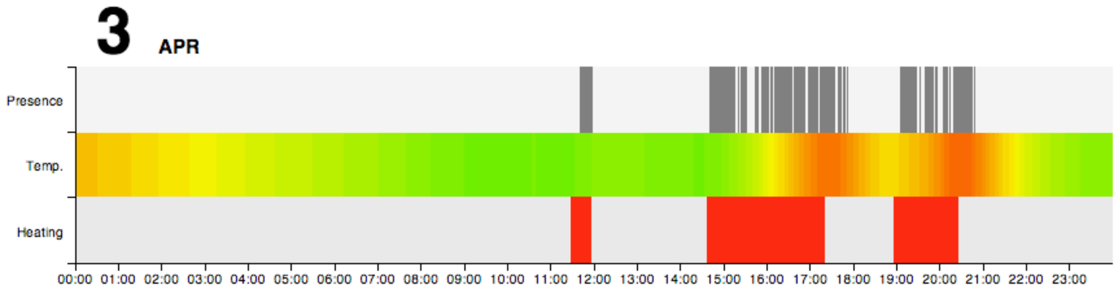

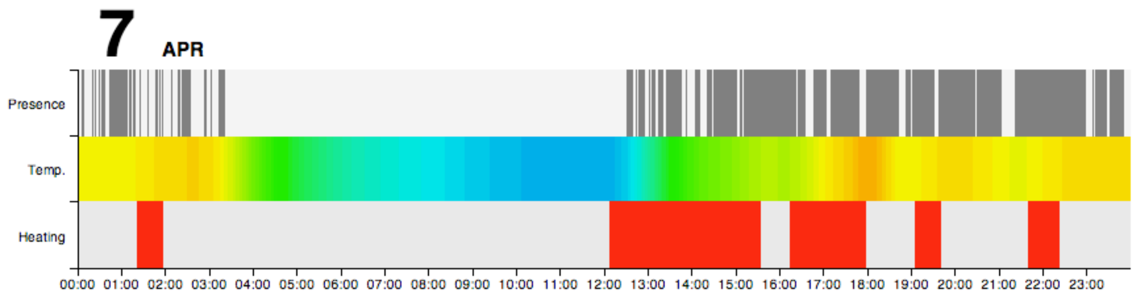

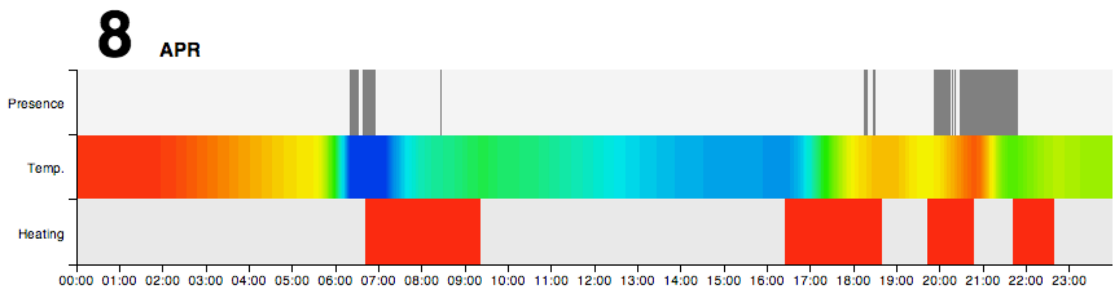

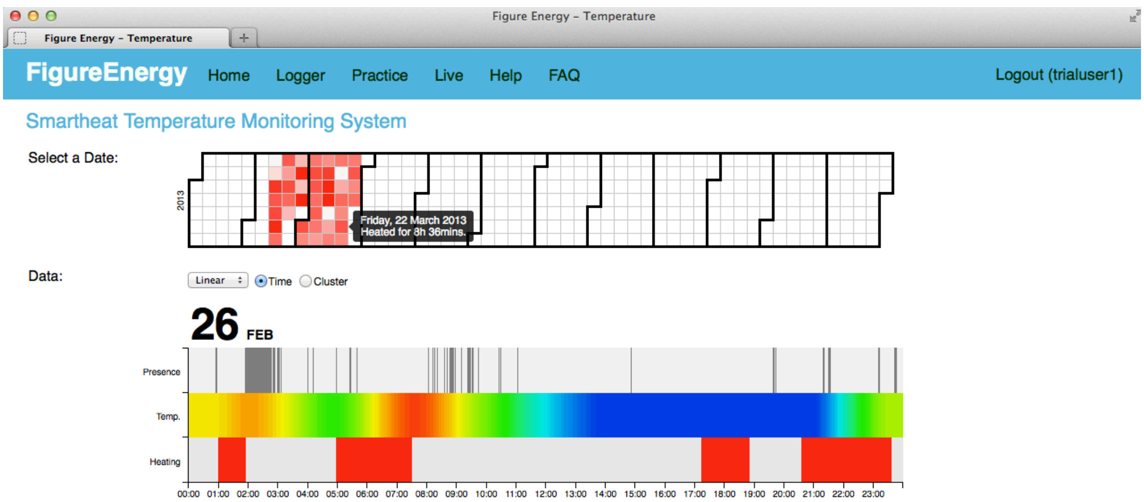

Appendix E: Ethics Documentation

E.1 Ethics Application Form

UNIVERSITY OF Southampton

ERGO application form – Ethics form

All mandatory fields are marked (M*). Applications without mandatory fields completed are likely to be rejected by reviewers. Other fields are marked "if applicable". Help text is provided, where appropriate, in italics after each question.

1. APPLICANT DETAILS

1.1 (M*) Applicant name:	Jonathan Morrice
1.2 Supervisor (if applicable):	Enrico Costanza
1.3 Other researchers/collaborators (if applicable): *Name, address, email, telephone*	**Evangelos Tolias**

2. STUDY DETAILS

2.1 (M*) Title of study:	SmartHeat Field Trial
2.2 (M*) Type of study (*e.g. Undergraduate, Doctorate, Masters, Staff*):	Undergraduate
2.3 i) (M*) Proposed start date:	28/03/2013
2.3 ii) (M*) Proposed end date:	21/04/2013

2.4 (M*) What are the aims and objectives of this study?
The aim of this study is to evaluate SmartHeat, a system designed to help users understand how they heat their homes and (hopefully) help them in saving it.

2.5 (M*) Background to study (*a **brief** rationale for conducting the study*)**:**
Thermostats are used in most homes today to specify a heating schedule. However, they provide no sort of feedback to help the user understand how the temperature actually changes and how the schedule could be optimised to save energy. SmartHeat is a temperature visualisation system to help users understand the temperature development in their home and make it easier to optimise schedules.

2.6 (M*) Key research question (*Specify hypothesis if applicable*):
Can interactive visualizations help users make sense of their heat consumption data?

2.7 (M*) Study design (*Give a **brief** outline of basic study design*) *Outline what approach is being used, why certain methods have been chosen.*
The study design is a field trial. Participants will be asked to install the system at home and interact with it for 3 days. Interviews will be held at the start and the end.

3. SAMPLE AND SETTING

3.1 (M*) How are participants to be *approached*? *Give details of what you will do if recruitment is insufficient. If participants will be accessed through a third party (e.g. children accessed via a school) state if you have permission to contact them and **upload any letters of agreement to your submission in ERGO**.*
I will recruit through snow-ball sampling (similar to “word of mouth”), starting through personal contacts.

3.2 (M*) Who are the proposed sample and where are they from (e.g. fellow students, club members)? *List inclusion/exclusion criteria if applicable. NB The University does not condone the use of ‘blanket emails’ for contacting potential participants (i.e. fellow staff and/or students).* *It is usually advised to ensure groups of students/staff have given prior permission to be contacted in this way, or to use of a third party to pass on these requests. This is because there is a potential to take advantage of the access to ‘group emails’ and the relationship with colleagues and subordinates; we therefore generally do not support this method of approach.* *If this is the only way to access a chosen cohort, a reasonable compromise is to obtain explicit approval from the Faculty Ethics Committee (FEC) and also from a senior member of the Faculty in case of complaint.*
Most participants will be friends and/or fellow undergraduate students. Any other participants will be approached if referred to by previous participants.

3.3 (M*) Describe the relationship between researcher and sample (*Describe any relationship e.g. teacher, friend, boss, clinician, etc.*)
Most will be friends or colleagues, others members of the general public.

3.4 (M*) Describe how you will ensure that fully informed consent is being given: (*include how long participants have to decide whether to take part*)
Participants will sign the informed consent form before taking part in the study. After reading the form they will be given the chance to ask any questions and if they require more time to decide, then another meeting will be organized to install the node and hold the first interview.

4. RESEARCH PROCEDURES, INTERVENTIONS AND MEASUREMENTS

4.1 (M*) Give a brief account of the procedure as experienced by the participant *(Make clear who does what, how many times and in what order. Make clear the role of all assistants and collaborators. Make clear total demands made on participants, including time and travel). **Upload any copies of questionnaires and interview schedules to your submission in ERGO**.*
In order to participate a wireless sensor node must be installed in one room of the participants residence. Installation takes a few minutes and the node requires access to a WiFi network and a power source (wall socket) throughout the duration of the trial. The node logs temperature, humidity, brightness and presence data which are sampled every two minutes and transmitted to a development server. In addition, two interviews of about 20-30 minutes will be held before and after the trial.

UNIVERSITY OF Southampton

5. STUDY MANAGEMENT

5.1 (M*) State any potential for psychological or physical discomfort and/or distress?
I do not identify any such potential.

5.2 (M*) Explain how you intend to alleviate any psychological or physical discomfort and/or distress that may arise? (if applicable)
I do not identify any such potential.

5.3 Explain how you will care for any participants in 'special groups' *(i.e. those in a dependent relationship, vulnerable or lacking in mental capacity)* **(if applicable)?**
Not applicable.

5.4 Please give details of any payments or incentives being used to recruit participants (if applicable)?
Not applicable.

5.5 i) How will participant anonymity and/or data anonymity be maintained (if applicable)? *Two definitions of anonymity exist:* *i)* Unlinked anonymity - *Complete anonymity can only be promised if questionnaires or other requests for information are not targeted to, or received from, individuals using their name or address or any other identifiable characteristics. For example if questionnaires are sent out with no possible identifiers when returned, or if they are picked up by respondents in a public place, then anonymity can be claimed. Research methods using interviews cannot usually claim anonymity – unless using telephone interviews when participants dial in.* *ii)* Linked anonymity - *Using this method, complete anonymity cannot be promised because participants can be identified; their data may be coded so that participants are not identified by researchers, but the information provided to participants should indicate that they could be linked to their data.*
Linked anonymity. Data is stored on the server and is in no way associated to the users. It will be associated with an ID number to distinguish results from different participants from each other but no adresses, names or similar will be recorded.

5.5 ii) How will participant confidentiality be maintained (if applicable)? *Confidentiality is defined as the non-disclosure of research information except to another authorised person. Confidential information can be shared with those who are already party to it, and may also be disclosed where the person providing the information provides explicit consent.*
Any information that is obtained (either automatically or through questionnaires) in connection with this study and that can be identified with any participant will remain confidential and will be disclosed only with the permission of the participant or as required by law.

5.6 (M*) How will personal data and study results be stored securely during and after the study? *Researchers should be aware of, and compliant with, the Data Protection policy of the University. You must be able to demonstrate this in respect of handling, storage and retention of data.*

Personal data will be stored securely only on a secure server during trial and analysis. It will be deleted after the study has been completed.

5.7 (M*) Who will have access to these data?
Myself and the researchers have access to the data on the server.

N.B. – Before you upload this document to your ERGO submission remember to:

1. Complete ALL mandatory sections in this form

2. Upload any letters of agreement referred to in question 3.1 to your ERGO submission

3. Upload any interview schedules and copies of questionnaires referred to in question 4.1

E.2 Project Protocol

UNIVERSITY OF Southampton

Project Description (Protocol)

Study Title: SmartHeat Field Trial

Researcher: Jonathan Morrice

Funder: University of Southampton

Background
Thermostats are used in most homes today to specify a heating schedule. However, they provide no sort of feedback to help the user understand how the temperature actually changes and how the schedule could be optimised to save energy. SmartHeat is a temperature visualisation system to help users understand the temperature development in their home and make it easier to optimise schedules.

Method
To test whether presenting visualisations to users improves their understanding, the system will be installed in a participant's home for 3 days. In addition, interviews will be held before and after the deployment to test understanding.

Materials
The project involves evaluating the impact of heat management interfaces on users' understanding of their energy consumption and their resulting behavior.
The hardware used consists of:

1. A wireless sensor node to collect temperature, humidity, light and presence data.
2. A server that takes data from the node and provides its own rich web user interface to access it. This interface provides one or more interactive visualizations of the energy consumed by each user. Users are required to login to access their data through a username and password. The data is served to the browser through SSL.

The data stored by our system includes the temperature and heat development in users' homes, as well as presence and light information sampled at two-minute intervals. No other personal data is stored on our server, as detailed below personal information is decoupled and keys are stored securely and separately. For the interviews, a rough outline of the questions asked is included in the Appendix.

To ensure participants of the nodes functionality, a guided tour of the node's internals will be given prior to signing the consent form. Each sensor or any other component that may concern the participant will be clearly identified to prevent suspicion of any hidden functionality.

Participants
Since the study will be held partly over the Easter break, I will be holding a majority of the study in my hometown Hamburg in Germany and the latter part in Southampton. Most participants will be friends and/or fellow undergraduate students between 20 and 24 years of age. I will not choose any households with residents who are under the age of 18. For areas with more than one user or resident, consent will be gained from each individual. Any other participants will be approached if referred to by previous participants. I am looking for 10 participants that meet the following criteria:

- Use central heating
- Have a broadband internet connection constantly available through WiFi at home

- Have lived at their current address for at least 1 month

Procedure
An initial interview is held to determine basic information on the residence the test will be performed in (size, thermostat). In addition the users' heating behavior to date will be determined. After an initial interview is held, data will be logged and displayed for three days (72 hours). In a second interview, users will then be questioned on their experience and asked how it has improved understanding of their consumption and whether it will change their behavior in the future.

Day 1: - Installation of sensor node (10 min.)
- Initial interview (20-30 mins.)

Day 2: - no meeting, data continues to be collected on the server

Day 3: - Removal of sensor node (5 min.)
- Review of collected data and final interview (20 – 30 mins.)

Statistical analysis
Analysis will mainly focus on interview answers rather than the data collected (also it may be included if referred to by a participant). I am mainly interested in similarities and differences between the initial interview and the final interview, to see how answers have changed after using the system for 3 days. Since the trial will only include 10 participants this will be done manually.

Ethical issues
I don't identify any such issues.

Data protection and anonymity
Any information that is obtained (either automatically or through interviews) in connection with this study will remain confidential and will be disclosed only with their permission or as required by law. Data is stored only on a password-protected server and the personal computer used by the investigator and is in no way associated to the participants. Results will be associated with an ID number to be distinguishable results from other participants, but no addresses, names or similar will be recorded.
During any interviews, audio recordings may be made. These will be kept separately from the rest of the data, as identities may be inferred from voices or from interview content. The only individual with access to this recorded audio data and its association with the rest of the study data will be the investigator. The audio data will be stored only on the personal computer used by him. If participants wish not to be recorded or at any later date decide they do not want the recordings to be kept, they may request that I delete the data from the dataset.
Since the sensor node requires WiFi access, any passwords restricting access to the network must be programmed into the device by the investigator during the installation. The code containing the password will only exist during the installation and relevant lines will be deleted directly after programming of the node. Moreover, participants will be advised to change their password (if existent) for the duration of the study.

E.3 Participant Information Sheet

FPAS Participant Information Sheet

Study Title: SmartHeat Field Trial

Researcher: Jonathan Morrice **Ethics Reference Number**: 5707

Please read this information carefully before deciding to take part in this research. If you are happy to participate you will be asked to sign a consent form.

What is the research about?
You are asked to participate in a research study conducted by Jonathan Morrice, an undergraduate student working towards a BEng degree in Electronic Engineering from the Faculty of Physical and Applied Sciences at the University of Southampton. This research is part of my third year individual project that strongly contributes towards my aspired degree.

This study evaluates SmartHeat, a system for visualizing temperature development in a home, designed to help users reduce their energy consumption. I am interested in understanding how *effective* and *easy to use* SmartHeat is. The experiments are aimed at measuring the qualities of the system, not the abilities of the participants.

Why have I been chosen?
You were selected as a possible participant in this study because you expressed an interest participating.

What will happen to me if I take part?
SmartHeat, the system we are evaluating, is web-based and relies on a wireless sensor node, developed to send temperature, humidity, presence and brightness data to a central server. If you volunteer to participate in this study, I would ask you to do the following things:

1. Participate in an initial interview covering some basic questions on your residence and current heating behavior (20-30 mins.)
2. Let me install a wireless sensor node in your home, grant it access to your WiFi network and not interfere with its functioning (e.g. disconnecting it or turning it or your WiFi off) for the entire duration of the study (3 days)
3. Let me uninstall the system from your home at the end of the study (5 min.)
4. Participate in a second interview to determine the effects of the system on your understanding and future behavior (20-30 mins.)

During the course of the study your temperature, humidity, presence and brightness data will be transmitted to the server every 2 minutes. The node is about the size of a credit card and will be mounted in the corner of a bed or living room.

Are there any benefits in my taking part?
Since this is a 3rd year undergraduate project I cannot offer any compensation for your participation. However, you would be bringing valuable information to the project and any future research done in the area.

Are there any risks involved?
The experiment does not involve any risk other than those related to using standard household electronic equipment.

Will my participation be confidential?
Any information that is obtained (either automatically or through interviews) in connection with this study and that can be identified with you will remain confidential and will be disclosed only with your permission or as required by law.

19/3/2013 Version 1.0

The information collected during the experiment (data about your heating behavior, your usage of the SmartHeat system and answers in interviews) will be kept separately from your personal identity. The information will be collected and stored on password-protected servers and on the personal computers used by the investigator.
During any interviews, audio recordings may be made. These will be kept separately from the rest of the data, as your identity may be inferred from your voice or from the interviews content. The only individuals with access to this recorded audio data and its association with the rest of the study data will be the investigators of this study. The audio data will be stored on a password-protected server different from the one used for the rest of the data and on the personal computers used by the investigator.
The information that links your data with your audio recordings and your identity will be available only to the investigator on a local machine.
If you decide you do not want the data to be recorded, or if you decide at a later date that you do not want these recordings to be kept in the form specified above, you may request that we delete the data from our dataset.

What happens if I change my mind?
Your participation in this study is completely voluntary. If you choose to be in this study, you may withdraw from it at any time without penalty or consequences of any kind and you may request that any data collected be destroyed. The investigators may withdraw you from this research if circumstances arise which warrant doing so.

What happens if something goes wrong?
If you have any questions, concerns or complaints about the research, please contact:

Jonathan Morrice
+44 (0)7428 705 255, +49 (0)173 954 1562
jnsm1g10@ecs.soton.ac.uk

or:

Faculty of Physical and Applied Sciences Ethics Committee
ergopas@soton.ac.uk

or

The University of Southampton Research Governance Office
Dr Martina Prude
Head of Research Governance telephone:
+44 (0)23 8059 5058
mad4@soton.ac.uk

Signature of Research Subject
I have read and understood the informed consent form (insert date/version no.) and have had the opportunity to ask questions about the study.
I agree to take part in this research project and agree for my data to be used for the purpose of this study I understand my participation is voluntary and I may withdraw at any time without consequence

Name of Participant... Signature of Participant................................

Signature of Investigator
In my judgment the subject is voluntarily and knowingly giving informed consent and possesses the legal capacity to give informed consent to participate in this research study.

Name of Researcher... Signature of Researcher................................

Date.....................

19/3/2013 Version 1.0

E.4 Consent Form

UNIVERSITY OF Southampton

CONSENT FORM (Version 1.1)

Study title: SmartHeat Field Trial

Researcher name: Jonathan Morrice
Ethics reference number: 5707

Please initial the box(es) if you agree with the statement(s):

I have read and understood the information sheet (19/3/2013 Version 1) and have had the opportunity to ask questions about the study.	
I agree to take part in this research project and agree for my data to be used for the purpose of this study	
I understand my participation is voluntary and I may withdraw at any time without my legal rights being affected	

Data Protection
I understand that information collected about me during my participation in this study will be stored on a password-protected computer and that this information will only be used for the purpose of this study. All files containing any personal data will be made anonymous.

Name of participant (print name)……………………………………………………

Signature of participant……………………………………………………………..

Date……………………………………………………………………………………

25/3/2013 Version 1.1

E.5 Interview Structure

Interview Descriptions

Version 1

SmartHeat Field Trial

Investigator: Jonathan Morrice

Questions for initial interview

1. How do you regulate the temperature in your house? If you use a thermostat, what type is it?
2. What affects the way you use this controller?
3. How often do you do change the setting(s)?
4. Do you know how many hours a day your heating is turned on and/or at what times?
5. Do you ever find the temperature in your room uncomfortable (i.e. too high/low)?
6. If so, how do you react?
7. Do you try to save energy and if so, why?
8. If possible, how do you think your settings could be improved?

Questions for final interview

Participants will be shown the SmartHeat interface during the interview.

1. How often did you view data collected by the SmartHeat system?
2. Do you believe any of the recorded data is wrong?
3. When was the heating turned on?
4. When was the temperature too high/too low?
5. When were you in the room according to the system?
6. In your opinion, when were you wasting energy?
7. Based on this, could your thermostat settings be improved and if so in what way?
8. Did the system motivate you to change your settings at any point during the trial?