Detecting Total Building Occupancy for More Efficient Operation

Newsham, G. R.; Birt, B.

Publisher's version / Version de l'éditeur:
https://doi.org/10.4224/20374236
Research Report (National Research Council of Canada. Institute for Research in Construction), 2010-07-15

Detecting Total Building Occupancy for More Efficient Operation

RR-304

Newsham, G.R.; Birt, B.

July 15, 2010

Detecting Total Building Occupancy for More Efficient Operation

Guy R. Newsham, Benjamin J. Birt

National Research Council – Institute for Research in Construction

RR-304

July 15th, 2010

This research report was generated as part of the NRC-ICT Sector initiative on Sensor Networks for Commercial Buildings

Detecting Total Building Occupancy for More Efficient Operation

Abstract

Wireless sensors were installed in a three-storey building in eastern Ontario comprising laboratories and 81 individual work spaces. Contact closure sensors were placed on exterior doors, internal doors, and on the refrigerator door in the main break room, PIR motion sensors were placed in the main corridor on each floor, and a carbon-dioxide sensor was positioned in a circulation area. In addition, we collected data on the number of people who had logged in to the network on each day, network activity (bit transfer rates), and electrical energy use (total building, and chilling plant only). Data were recorded over the Summer, 2009. The data streams were clearly responsive to building occupancy at the whole building level (e.g. evenings, weekends, public holidays) and locally (e.g. doors near heavily-occupied meeting rooms). Further, we developed an ARIMAX model to forecast the power demand of the building in which an explicit measure of building occupancy level was a significant independent variable and increased the accuracy of the model. The results are promising, and suggest that further work on a larger and more typical office building would be beneficial. If building operators have a tool that can accurately forecast the energy use of their building several hours ahead they can better respond to utility price signals, and play a fuller role in the coming Smart Grid.

1.0 Introduction

In 2007 commercial and institutional buildings accounted for approximately 13% of the end use of all energy sources in Canada of which 25% was electricity [National Resources Canada, 2008]. Data from Ontario suggests that on-peak electrical demands are growing even faster than overall demand [Rowlands, 2008]. Concern is growing that additional electricity supply cannot be realised quickly enough to meet growing demand, particularly at peak times. As a potential solution, it is proposed that building power draw should be better tuned to the state of supply through more careful control and timing of loads in response to utility signals and other inputs; this may be one element of the larger concept of the Smart Grid [e.g. Gershenfeld et al., 2010].

A building's ability to reduce overall energy use, and to reduce demand at specific times may be substantial, depending on the systems in place and data available to inform decisions. Recent research demonstrates that "demand responsive" heating, ventilating and air conditioning (HVAC) and lighting system may reduce their power draw by 30%, or more, for periods of a few hours without major hardship to occupants [Piette et al., 2005; Newsham & Birt, 2010] The flexibility in future building energy use depends upon current operating conditions and projected future needs. Such forecasting and resulting action is rare in current practice. Tools that do exist are based on: previous experience, prior energy use, and weather conditions and forecasts; these are discussed in more detail in Section 3.4 below.

A factor that is not typically considered explicitly in energy forecasting is building occupancy. Occupants are a main driving factor behind commercial building energy use with their use of office equipment, lighting, plug loads and a need for suitable ventilation rates, thermal conditioning etc., thus including measures of occupancy in energy prediction models should improve them. Further to this, as people move away from the office and move to remote working locations (at home, office hubs etc.) to reduce daily travel time and improved work time flexibility, the efficient use of office space is critical [North, 2007]. For example, to assist in the efficient use of desk space, hotelling is becoming popular in some large offices [Freiman, 1994; Shimo, A., 2008]. Again, more detailed information on total occupancy would assist in more efficient space allocation. In certain buildings swipe card access can easily give building occupancy information. However in other buildings where such technology is not available, is there another economical way of determining how many people are in a building?

Information on occupancy may also lead to better and more energy-efficient indoor environment control. Localization; i.e., knowing where individual occupants are, and perhaps their preferences and needs, may lead to very fine-tuned delivery of lighting and HVAC [Zhen et al., 2008; Harle & Hopper, 2008; Wen et al., 2008; Tiller et al., 2009; Erickson et al., 2009]. The installation of sensors can be used to view social dynamics/habits in buildings that could be used for improved modeling and safety. For example, more than 200 motion sensors were installed in common areas across 2 floors and monitored over a year creating a large data set [Wren et al., 2007b]. This data has been analysed several ways to investigate relationships between soft drink consumption and motion [Wigdor et al., 2007], measuring the hidden social life of a building [Wren et al., 2007a] and visualising the history of living spaces [Ivanov et al, 2007]. A low cost network of sensors can provide powerful contextual information to building systems: improving efficiency of elevators, lighting, heating and cooling; enhancing safety and security [Wren et al., 2006].

This project investigated the use of several data streams, including several relatively cheap, off-the-shelf wireless sensors and receivers to measure various building occupant activities in order to develop metrics related to total building occupancy[1]; we also recorded total building energy use. We correlated occupancy information in a qualitative sense against various building events to demonstrate that monitoring occupancy could be used to detect (and react to) such events. In particular, and in a quantitative sense, we used metrics related to building occupancy as external inputs (independent predictor variables) in forecast models of building power draw. The question we wished to explore was whether including an occupancy metric as an explicit independent variable would provide additional useful information, and would therefore improve the model accuracy. Loveday & Craggs [1993] suggested that variance in their model of indoor temperature could be partially explained by variations in occupancy, and that forecasting accuracy could be improved by adding it as a variable in the model,

[1] We are also conducting a complementary project on linking occupant localization data to local delivery of lighting and HVAC, the results of this work will be reported separately.

and Neto & Fiorelli [2008] lamented the lack of occupancy data for use in their model of building energy use.

2.0 Methodology

A wireless sensor network monitoring several building characteristics was installed in a building located in eastern Ontario.   The building was a three-storey concrete structure, with metal studded dry walls and comprised a mixture of offices and laboratories. The laboratories included numerous pieces of industrial and electrical equipment, computer servers and metal cabinets. There were 81 individual work spaces (primarily private offices with some cubicles) in the building at the end of the study (some cubicles were added during the study).   Various sensors and receivers were installed throughout the study building to collect data on activities related to occupancy, and other relevant information. Network logins, network activity, and the wireless sensor network data were recorded between June 12th and September 30th 2009. Energy use data was also recorded during this period. The study period was chosen as a time when there is often the greatest variation in building occupancy, with occupants out of the office for summer holidays and conferences. Staffing levels typically start to return to normal at the end of August, when school resumes. A complete description of each of the data sources is in Section 2.1. The study protocol was approved by our organization's Research Ethics Board.

2.1 Data sources

2.1.1 Network logins

Logging in to the local computer network is necessary for all employees wanting access to the network and internet. Network logins were logged and maintained by the organization's System Support Unit (SSU) staff using Event Viewer in Microsoft Windows 2003 Server. For confidentiality, user identifiers were removed by SSU staff and replaced with a unique non-identifying number. The network logins log contained users from across the organization and was filtered to contain only those employees that were based in the study building. During the study period there was a total of 88 unique building occupants with network accounts, including full- and part-time staff, students and visiting staff. The vast majority of building occupants were on-site when they logged in, and did login when they arrived for work each morning. This data source does have errors associated with it. It did not include contractors and guests that were not using the network while in the building. Further, some people only logged on once per week (automatically re-authenticating when network was available again), while others logged on from remote locations.   Network logins from remote locations were included in the log with no markers distinguishing these from on-site logins. However, it could be reasonably assumed that those logging in to the network outside of 6:00 am and 6:00 pm were doing so remotely. The data file was filtered to remove any logins repeated within 5 seconds. This cleansed data gave a daily cumulative count of logins (multiple logins per user per day) and a daily cumulative count of unique users per day with at least one login. Note that the login data is limited to unique login events, subsequent logoffs were not recorded.

2.1.2 Network activity

Network activity was monitored by SSU by measuring the rate of data transfer between a network server and switch board. The average transfer rate over the preceding 30 minute interval was recorded over the study period. Two values were recorded, incoming transfer rate (data being sent to the server from a user's PC) and an outgoing transfer rate (data being sent from the server to a user's PC). The volume of data being transferred hourly from midnight was calculated daily for each transfer direction.

2.1.3 Building energy use

Two measures of electricity use were available: total building, and chilling plant only. These were recorded with separate PowerLogic ION7650 Power Meters manufactured by Schneider Electric. The power meters recorded several variables; we used the average power in preceding 15 minutes and cumulative energy consumed in our analysis.

2.1.4 Wireless sensors

Off-the-shelf wireless sensors and receivers were purchased to measure various parameters associated with in the building, including indicators of occupant activity. These were purchased from *Echoflex Solutions Inc.* (http://www.echoflexsolutions.com/) that use EnOcean technology and are part of the EnOcean Alliance (http://www.enocean-alliance.org/en/). Several sensors were available including magnetic contact closure sensors, motion sensors, illuminance sensors and temperature sensors. A carbon-dioxide sensor was connected to a transmitter via a prototyping board and included in the network. The wireless sensor network system recorded activity related to everyone that used the building, including full-time staff and visiting workers, and contractors and guests.

The sensors were wireless in terms of their method of communication, to transmit measured values back to the base station, some required wires for power supply. The motion sensors and carbon-dioxide (CO_2) sensor had a 110 V power source. The magnetic contact closure sensors, illuminance and temperature sensors were all powered with a small solar cell and capacitor. Data was transmitted in a small packet (telegram) upon activation using a 315 MHz signal. The 14-byte telegram was sent either upon activation or at a set period of time from its last activation, as determined by the "heartbeat" rate set on the transmitter. The heartbeat was essentially a telegram that simply indicated that the sensor was still working normally. Each telegram included a unique sensor serial identification number (4 bytes), 4 data channels (1 byte each) as well as telegram verification, status and sensor type identifiers. See the EnOcean Alliance website for a complete description of protocol used [EnOcean GmbH, 2008]. The transmitted signal strength was less than 10 mW and had a range of up to 30 m indoors (100 m free range) before the power of the signal was substantially dissipated and not detected by a receiver. The range depended upon the type of surrounding material, as observed by Jang & Healy [2010]. Metal cabinets, server rooms, bookshelves, concrete walls and experimental equipment had the greatest effect on signal propagation.

The illuminance and temperature sensors were used to monitor external weather conditions. The illuminance sensor was designed for internal daylight harvesting applications and had a maximum measurable illuminance of 510 lux. For external measurements we placed neutral density filters in front of the sensor raising the maximum illuminance to 130 klux. Both the illuminance and temperature sensors were placed in a waterproof container and placed on top of the building. The photosensor had an unobstructed view of the sky and the temperature sensor was shaded from direct sunlight (Figure 1). A comparison to a calibrated illuminance meter and thermometer was performed over several days to determine a correction factor for both sensors (Figure 2 and 3 respectively).

Figure 1. The Illuminance sensor and temperature sensor as installed on the rooftop. The Illuminance sensor is in the water proof container on the left. The temperature sensor is shaded from direct sunlight by the white covering on the right.

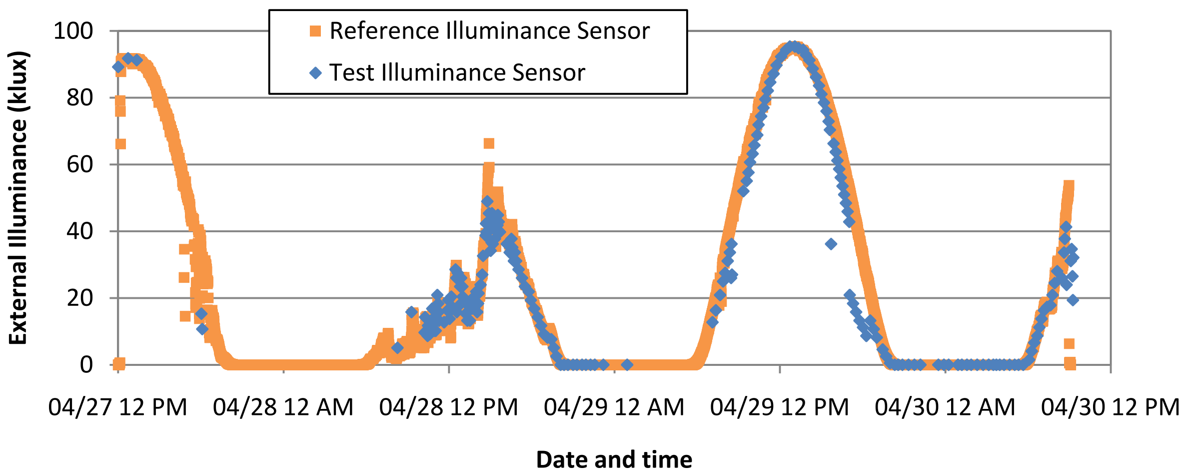

Figure 2. Calibration curve for the wireless Illuminance sensor (including correction factor) measuring external Illuminance for three days.

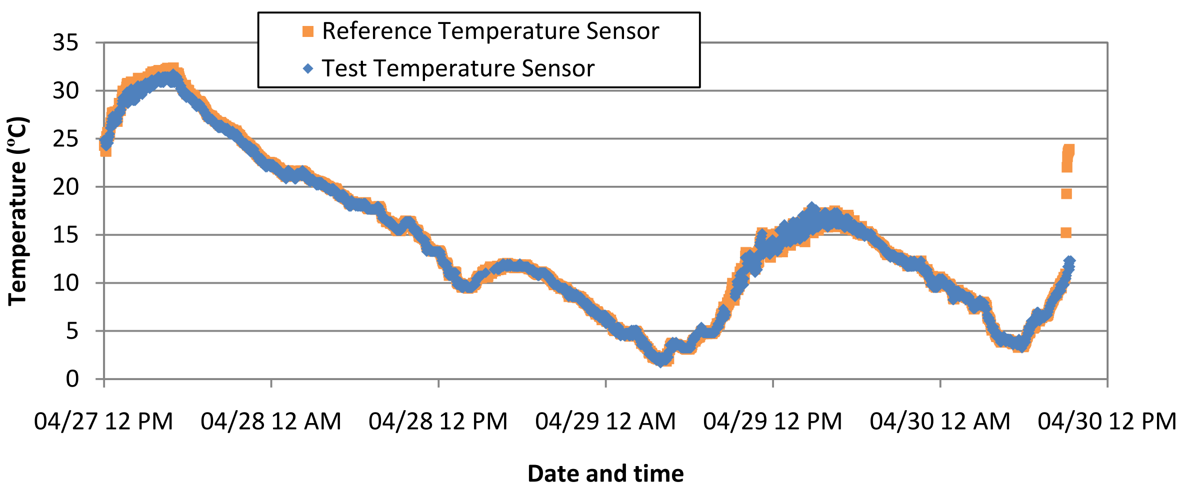

Figure 3. Calibration curve for the wireless temperature sensor measuring external temperatures over three days.

The CO_2 sensor (Figure 4) was a commercially available sensor (Vaisala GMW21) that was connected to a wireless transmitter. The output of the CO_2 sensor was 0 – 10V with the maximum corresponding to 2000 ppm, but the input of the transmitter was 0 – 2V only, with 256 increments. To ensure that we maximized the sensitivity of the device to the concentrations expected, we capped the output of the sensor at 4 V (800 ppm), and then stepped this down to 2V for transmission, meaning that the maximum transmitted value of 255 (2V) corresponded to 800 ppm. We initially planned to measure building CO_2 level in a common ventilation return duct, however this proved to be impractical. Large variations in the measured CO_2 levels were observed and there were difficulties in transmitting a signal from the confined spaces of the metal duct work. Another location that would measure general carbon-dioxide levels was chosen in a common circulation area on level three. The location was chosen such that it would not be biased by any group or person. It also serviced several communal areas. Calibration of the sensor took place in IRC's calibration lab. Calibration involved comparing the output of the wireless CO_2 sensor with that of 0 – 10 V output of the test sensor along with a reference sensor, both placed in a metal chamber. The chamber was dosed with nitrogen and CO_2 gas with a CO_2 concentration of 1000 ppm until a steady reading was observed (20 min). The chamber was then flushed with a pure nitrogen gas to achieve a zero point of CO_2 concentration. Since the final calibration point (1000 ppm) was beyond the digital output range of the wireless transmitter, it was compared to the test sensor at intermediate points during the flushing (see Figure 5). A linear line was fitted to this curve as seen in Figure 5 and used to convert the digital output to a meaningful CO_2 concentration.

Figure 4. The carbon dioxide sensor as installed on the ceiling of the 3rd floor corridor. The wireless transmitter component was hidden in the ceiling space.

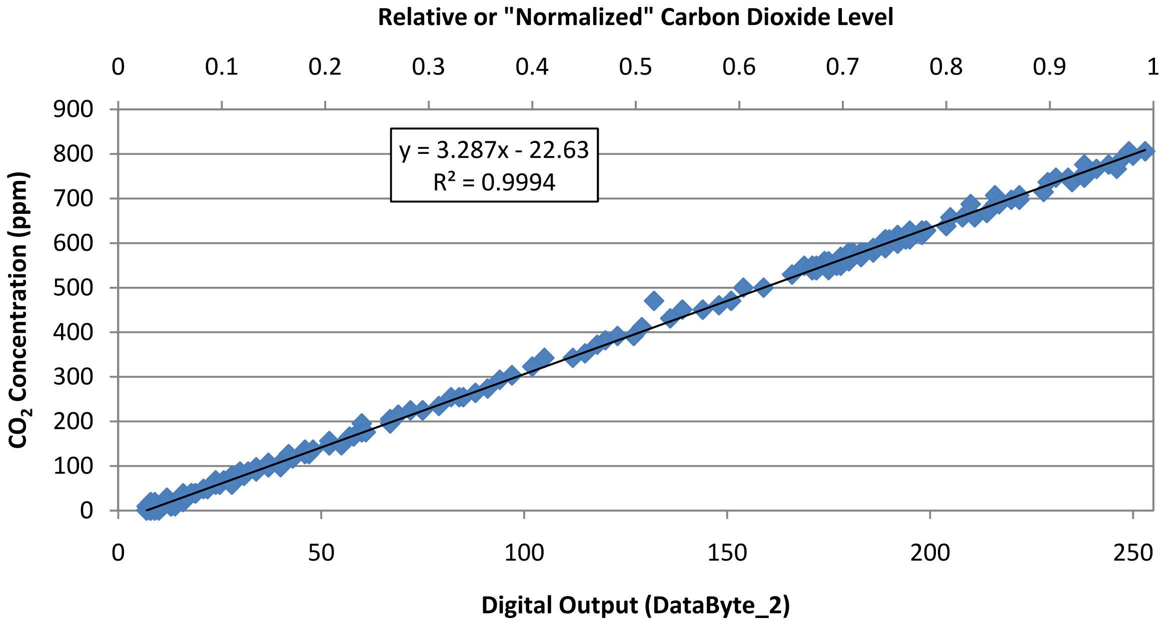

Figure 5. The calibration curve for the carbon dioxide sensor compared with a reference sensor. The relative carbon dioxide level as used previous to calibration is on the top horizontal axis.

Magnetic contact closure sensors (Figure 6) were fixed to the entrance doors of the building, internal doors along circulation routes, and a refrigerator. A telegram was sent upon each open and close of the door. A heartbeat telegram rate was set at 10 000 seconds. The heartbeat telegram of all sensors was a repeat of the last telegram sent. No identifying data bytes are used to differentiate a heartbeat telegram from an ordinary signal.

Figure 6. Magnetic contact closure sensor as installed in the study building on an internal door. The orange wire is the transmitter's antenna. The magnet is to the left of the sensor. The dark region on the sensor casing is the solar panel used for energy harvesting.

Passive infrared radiation (PIR) motion sensors (Figure 7) were used to count the number of times a person or a group of people moved through a certain region. The corridor locations were chosen not to be outside the entrance to a specific room. A telegram was transmitted on activation. A deactivation telegram was sent several seconds later. The time between the activation and deactivation telegram was dependent upon the length of time that a body remained in the field of view of the sensor. Heartbeat telegrams were transmitted every 80 seconds.

Figure 7. The PIR motion sensor as installed. The power supply is in the ceiling space.

2.2 Sensor locations

The goal of each interior sensor was to gain a measure of activity that would correlate with general building occupancy. For example, the number of times a door was used or the number of times people walked through a particular region were expected to correlate with building occupancy.

Each of the sensors was installed several weeks before the study period and monitored for reliability. Some sensor locations required the use of a repeater that re-transmitted the signal to the base station. The commercial system used was designed for wireless communications for building control systems (lighting, HVAC) within small areas/rooms. Setting the system up so that it provided reliable information across an entire building proved to be a challenge. Due to the building being composed of laboratories containing experimental equipment, metal cabinets/shelves and a server room, telegrams could not be transmitted through certain regions [Jang & Healy, 2010]. We used a signal strength meter to help in the placement of sensors. In addition, much time was spent on trial-and-error repositioning of transmitters and receivers. Sensors were installed in their proposed locations and test activations made, if the initial test activations were received then the location was monitored for several hours for reliability. Test activations were made periodically during this testing phase. Even after the testing phase, there were periods during the data collection period that the central receiver did not receive data from certain sensors consistently, requiring us to remove certain days from the data analysis for some variables. A building entrance door located on the north side of the building could not be consistently logged and was not included in the final network.

The locations of the sensors across the three floors are shown in Figures 8 to 10. The base station was located on level 3 and connected to a computer logging all incoming telegrams. Repeater C on level 3 serviced the magnetic contact closure sensors on the level 3 internal doors, refrigerator and the front internal and external doors. Repeater A on level 1 serviced the magnetic contact closure sensors fixed to the south door and level 1 internal doors and the motion sensor on level 1. Repeater B was used to attempt to get a signal from the north door, but was found to be unreliable due to the nature of construction around this area (i.e. reinforced glass, elevator shaft, server room etc.). A complete list of all sensors used with their location, identification number, sensor type and transmission path to the base station is in Table 1.

Table 1. Brief details of each of the sensor for the wireless network in the study building. North door was not included in final network due to unreliable signals.

Name	Sensor type	ID # (hex)	Floor	Transmission path
Motion sensor level 1	Motion	00010FAB	Level 1	Repeater A
Motion sensor level 2	Motion	00010EDF	Level 2	Direct
Motion sensor level 3	Motion	0000AD32	Level 3	Direct
Refrigerator door	Contact closure	00010F8B	Level 3	Repeater C
North door	Contact closure	000105D0	Level 1	Repeater B
South door	Contact closure	000105F4	Level 1	Repeater A
Front doors internal	Contact closure	00010623	Level 2	Repeater C
Front doors external	Contact closure	00010F62	Level 2	Repeater C
Level 1 doors	Contact closure	00010E57	Level 1	Repeater A
Level 3 doors	Contact closure	0001043E	Level 3	Repeater C
Carbon dioxide	Carbon dioxide	0000AD30	Level 3	Direct
Thermometer	Temperature	000105AB	Roof	Direct
Photosensor	Illuminance	00010E18	Roof	Direct

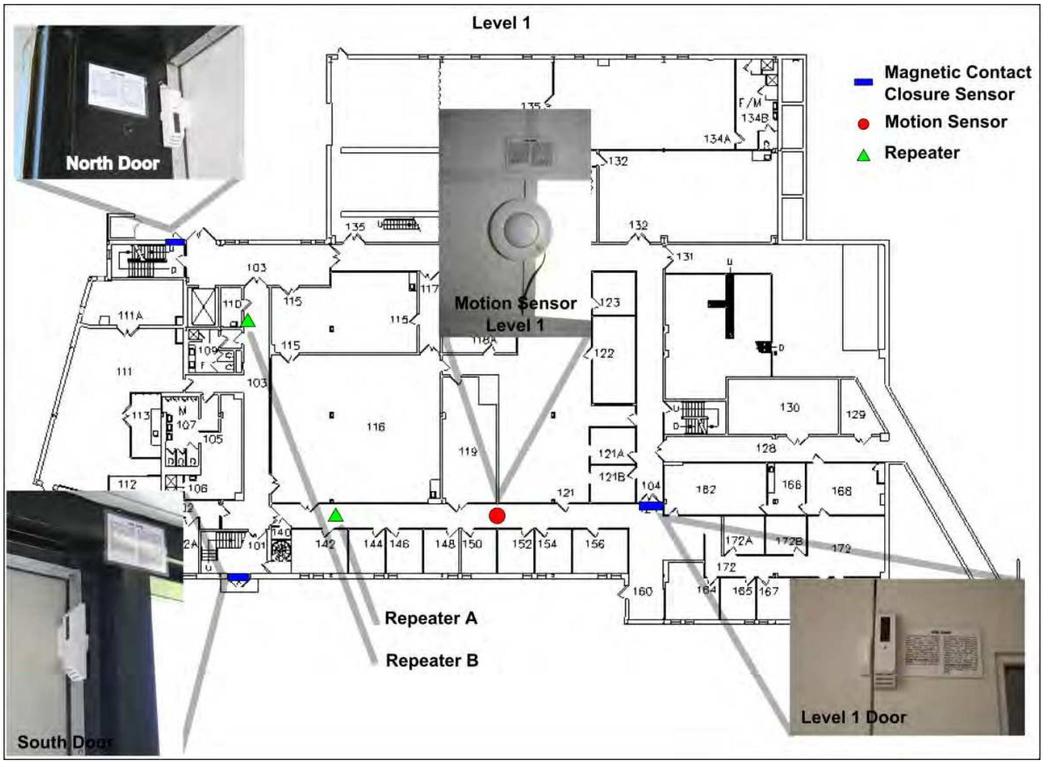

Figure 8. Level 1 of the study building

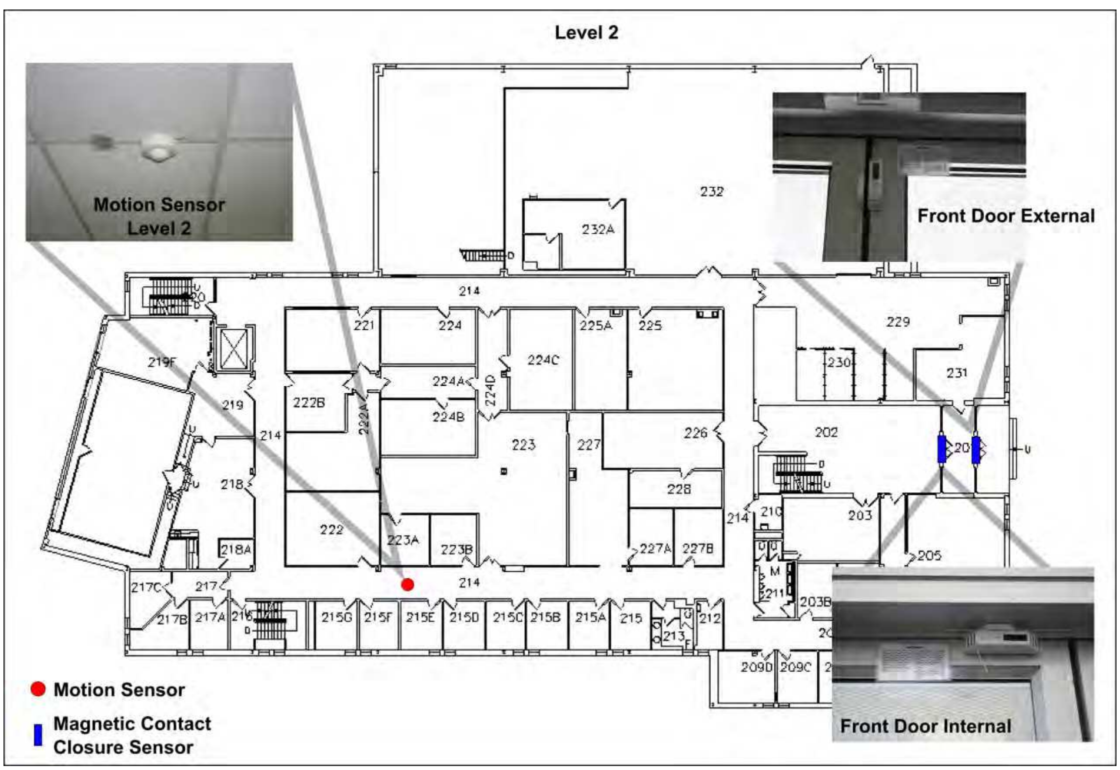

Figure 9. Level 2 of study building

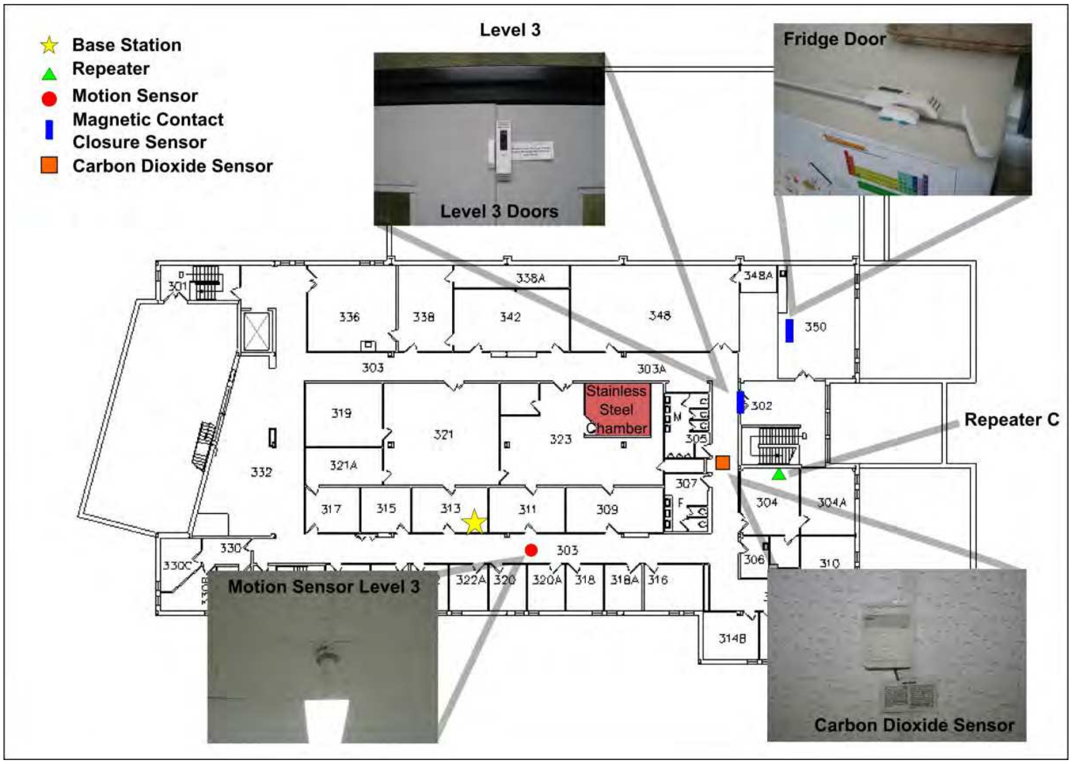

Figure 10. Level 3 of study building

2.3 Sensor data collection

Telegrams were logged at a base station on level 3. A program written in Agilent VEE was used to monitor and log the incoming telegrams with a timestamp. Preliminary data analysis was performed at this stage by converting the data bytes into meaningful units, or the sensors' status (opened/closed, active/inactive). Researchers inspected the data visually on a regular basis and any irregularities of missing data by sensors or repeaters not working were recorded and used to cleanse the data at the end of the collection period.

The telegram structure can be seen in the Echoflex documentation for each of the sensors (http://www.echoflexsolutions.com/) with a summary of the structure in Table 2. Equations 1 to 3 were used to convert the respective data byte transmitted by the carbon dioxide sensor, temperature and illuminance sensors respectively. For motion sensors, DataByte_1 has two values according to whether the sensor is active (2) or inactive (255). If the same value is received consecutively then it is considered as a heartbeat signal. For magnetic contact closures, DataByte_0 has four values according to whether the sensor is open (bit_0 = 1), closed (bit_0 = 0) or if a learn button is pressed in either situation (bit_3 = 0). If the same value is received consecutively then this is considered as a heartbeat signal.

Table 2. Breakdown of the 14 bytes in a telegram of the Enocean protocol [EnOcean GmbH, 2008].

Byte Number	Name	Description
Byte 0	Sync_Byte_1	Constant (A5 hex)
Byte 1	Sync_Byte_2	Constant (5A hex)
Byte 2	Header	Header identifier and number of octets following
Byte 3	Org	Transmitter type
Byte 4	DataByte_3	Analogue data input 3 on transmitter
Byte 5	DataByte_2	Analogue data input 2 on transmitter
Byte 6	DataByte_1	Analogue data input 1 on transmitter
Byte 7	DataByte_0	Digital input on transmitter
Byte 8	ID_3	32 bit sensor ID
Byte 9	ID_2	
Byte 10	ID_1	
Byte 11	ID_0	
Byte 12	Status	Direct or repeated signal
Byte 13	Checksum	Last 8 least significant bits from addition of Bytes 2 to 12

$$\text{Carbon Dioxide Concentration (ppm)} = 3.287 \times \text{DataByte\_2} - 22.63 \quad (1)$$

$$\text{Temperature (°C)} = -\frac{80}{255} \times \text{DataByte\_1} + 60 \quad (2)$$

$\text{Illuminance (lux)} = \mathrm{k} \times \frac{510}{255} \times \text{DataByte\_2}$, where k is the empirically-derived calibration (filters transmission) constant, k = 255 (3)

2.4 Data cleansing

Over the study period it was noted that the reliability of some sensors varied daily. The variation in reliability led to some periods of an unknown number of missing telegrams. An error log was kept at the base station to record any potential problems and days to be removed from the final data set. A set of criteria was used to determine which days were removed for analysis purposes.

Motion sensors were the most reliable, likely attributable to a consistent signal strength due to being powered by a 110 V source. Due to the frequency of heartbeat telegrams, a minimum of 860 telegrams were expected in a day even if there were no activations. Therefore days with a total number of telegrams less than 860 were treated as erroneous and removed. For days that were close to this limit the days were individually inspected for missing periods of data on a daily profile plot.

The magnetic contact closure sensors were the least reliable in the installation. Depending upon their location telegrams often went astray. The transmitted signal strength and heartbeat rate was influenced by the amount of energy stored on the capacitor charged by the solar cell. As a result, a minimum number of daily telegrams could not be used for data cleansing. Individual days were visually inspected and removed from further analysis if there were atypical periods of inactivity. For example, the internal doors on level 3 were used regularly during business hours, therefore any inactivity during this time for more than one hour was uncharacteristic and the day was removed (unless an event was known that could account for the irregularity).

Network logins and network activity were also cleansed for erroneous days. Recorded events of network resets and logging faults eliminated days from network login data. The amount of network activity was extremely high (an order of magnitude greater) on several days near the end of the study period, these days were removed from the final analysis. The increased activity was due to a location remote from the study building, but sharing the same network, backing up data during these periods. Besides the obvious effect of greatly reduced activity on weekends there were other notable days or events that had an effect on one or all of the sensors in the network, Table 3 lists these events.

Table 3. Major events that occurred during sensor logging of study building

Date	Description
12th June 2009	Sensors and repeaters hardwired into building
24th June 2009	St Jean Baptiste Day (public holiday in Quebec)
30th June 2009	Computer Network was shut down for maintenance
1st July 2009	Canada Day (public holiday throughout Canada)
3rd August 2009	Civic Holiday (public holiday in Ontario)
12th August 2009	Extra cubicles on level 1 are starting to be used
12th - 17th August 2009	South car park closed
24th - 25th August 2009	Large meeting on level 3
27th August 2009	On-site BBQ – south door propped open at times
7th September 2009	Labour day (public holiday throughout Canada)
22nd - 24th September 2009	Large meeting on level 3

3.0 Results

3.1 Telegram reliability

It was known that some signals from the magnetic contact closure sensors were not received. Exactly how many telegrams went missing is unknown, but an estimate can be made by interpreting the incoming telegrams in a different way than the initial interpretation. Initially a count was only made when an open signal was followed by a consecutive close signal. Sometimes, a repeat of the previous signal was received during business hours. Initially the repeated signal was considered as a heartbeat (as described above), heartbeats should not occur during business hours, as doors were used more frequently than the heartbeat rate (1 per 10 000 seconds). A more liberal approach to calculate the number of times a door was used during business hours could be used. For example, if a closed signal was not received following an open event, and a further open signal was received a short time later, then originally this would have been considered as a heartbeat rather than as a new door opening event. Similarly should an open signal go astray, then two closed signals would be received consecutively, again, originally this would have been considered as a heartbeat and a new door opening event would have been missed. Assuming that repeated signals during business hours were actually the result of a missing telegram, then an estimate of the number of telegrams going missing per day could be made. Over the study period approximately 12.9% of the signals went missing (this is similar to the 90% packet delivery rate used as one benchmark by Jang & Healy [2010]. It is possible that both the open and close signals for a specific event were not received. If we assume that the chance of this happening is 12.9% x 12.9%, then this adds another 1.6% to the number of contact closure events not recorded.

Another estimate of missing telegrams comes from looking at the two doors of the main entry vestibule. It is assumed that the two doors (front door internal and external) are used within several seconds of each other when people are either entering or exiting the building. There were other ways out of the vestibule but they were very rarely used. An exact pairing of the internal and external door activations of the vestibule was not always evident in the data due to missing telegrams. Manually matching the internal and external door activations for the nearly 20 000 telegrams, we estimated that 12.0% and 15.8% of the total telegrams from the internal and external doors respectively were missing.

It is not clear if the pattern of missing telegrams was random or systematic. There were certainly dates and times when there were blocks of missing telegrams, but this can occur in random sequences too. We suspect that some missing telegrams were due to temporary physical obstructions or sources of interference.

3.2 General trends

We initially inspected descriptive representations of the data to verify expected trends. A weekly profile with weekends exhibiting much lower activity levels than weekdays was observed; activity levels in public holidays were similar to weekends (Figure 11). There was also reduced activity on Fridays for several measures with examples shown in Figure 12. Diurnal profiles of activity at regular times of the day were observed as expected. For example, an increase in activity at the beginning and end of day was seen on the external doors, and sensors that were located in paths leading to the lunch room (Figure 10, room 350) increased at break times (Figure 13).

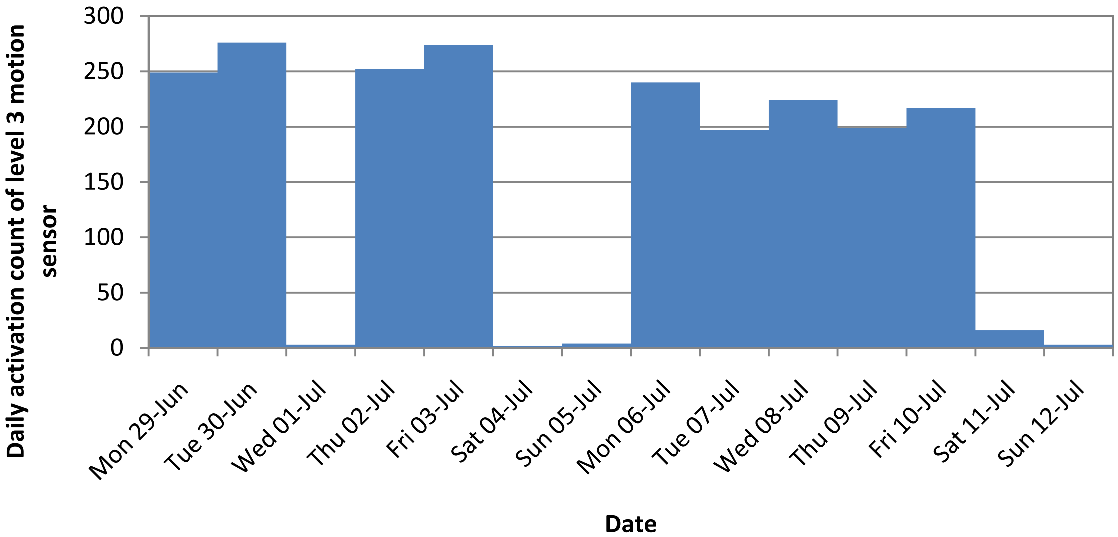

Figure 11. The daily total activation counts on the motion sensor on level 3 for an example two-week period showing weekends and public holidays (1st July) having a significantly lower number of activations over the day.

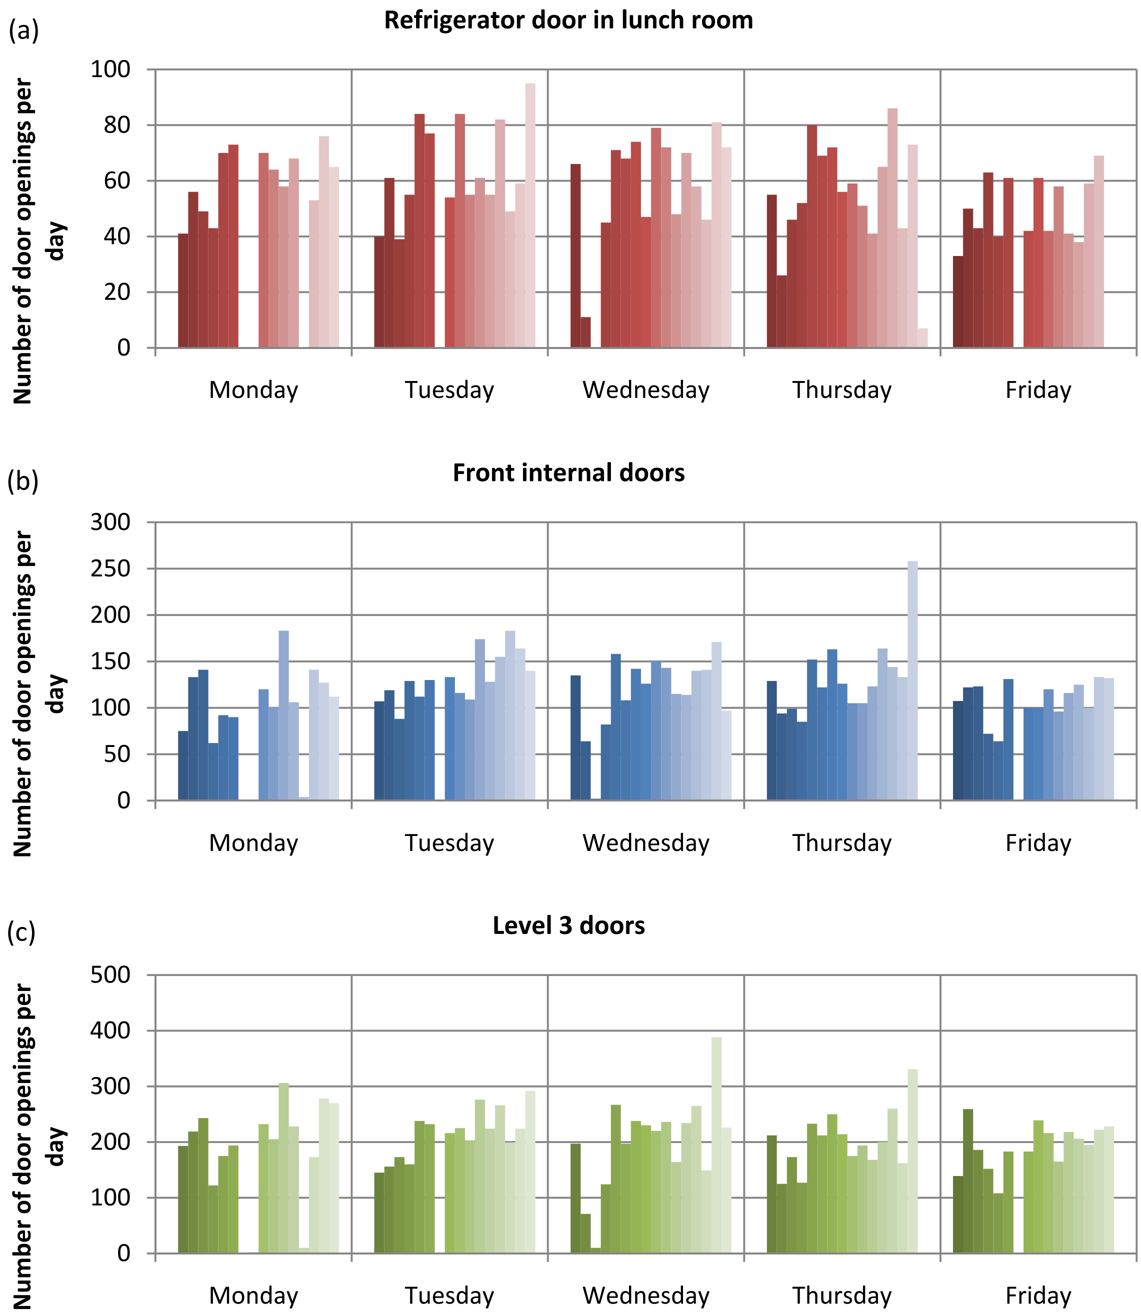

Figure 12. Examples of the sensors that showed a typically lower total number of daily activations on Friday: (a) refrigerator door, (b) front internal doors and (c) the level 3 doors. Days that have a total count of zero are either public holidays or were erroneous and removed from analysis. Data are shown by day of week, with consecutive bars showing values for that day on consecutive weeks.

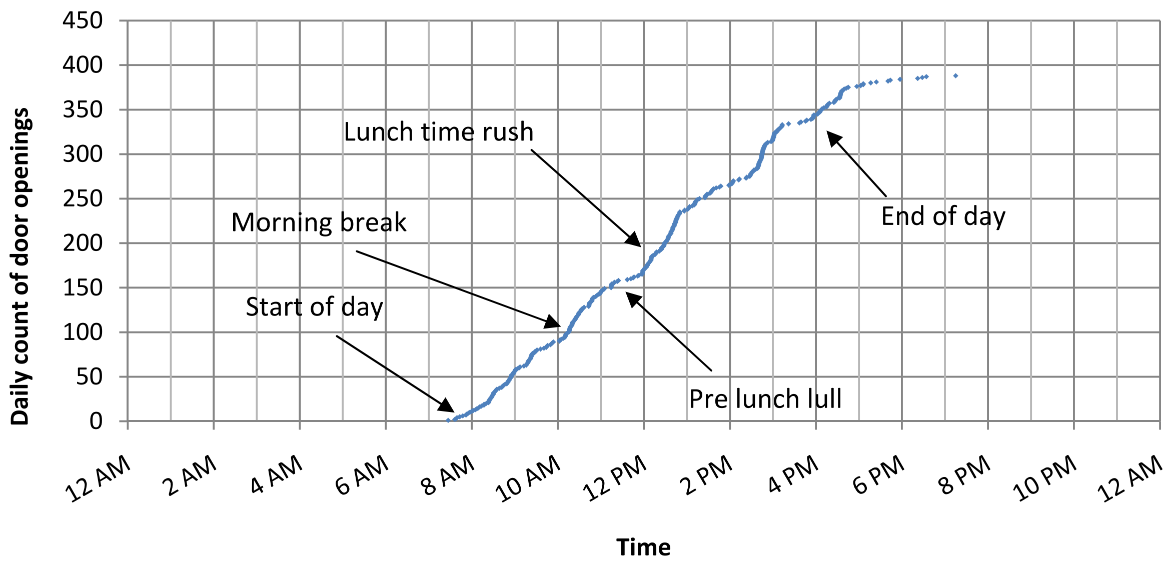

Figure 13. Example of a daily profile of the use of the internal doors (cumulative counts) on level 3. Increase in activity is seen at the start and end of the day, morning break and lunch times.

These general trends were also observed in the number of network logins (Figure 14). Figure 14(a) shows the total unique daily network logins that occurred between 6:00 am and 6:00 pm. The reduced Friday activity seen in the sensor network is also observable for network logins. Further to this, looking at the data from Fridays, there appears to be a decrease during the middle of the study period as expected for this time of year. Figure 14(b) shows that the majority of people are logged onto the network for the day by 10:00 am on weekdays (red lines). Unfortunately a measure of when a person logged out of the network was not obtained. The general trends were less obvious in network activity and building energy use. Figure 15 shows the daily energy use of the building excluding chiller loads. Due to the unique use of the building as a research facility, building energy use may not correlate with occupancy as strongly as in more conventional office buildings. For example, weekday loads were only approximately 25% more than weekend loads, suggesting the presence of high process loads.

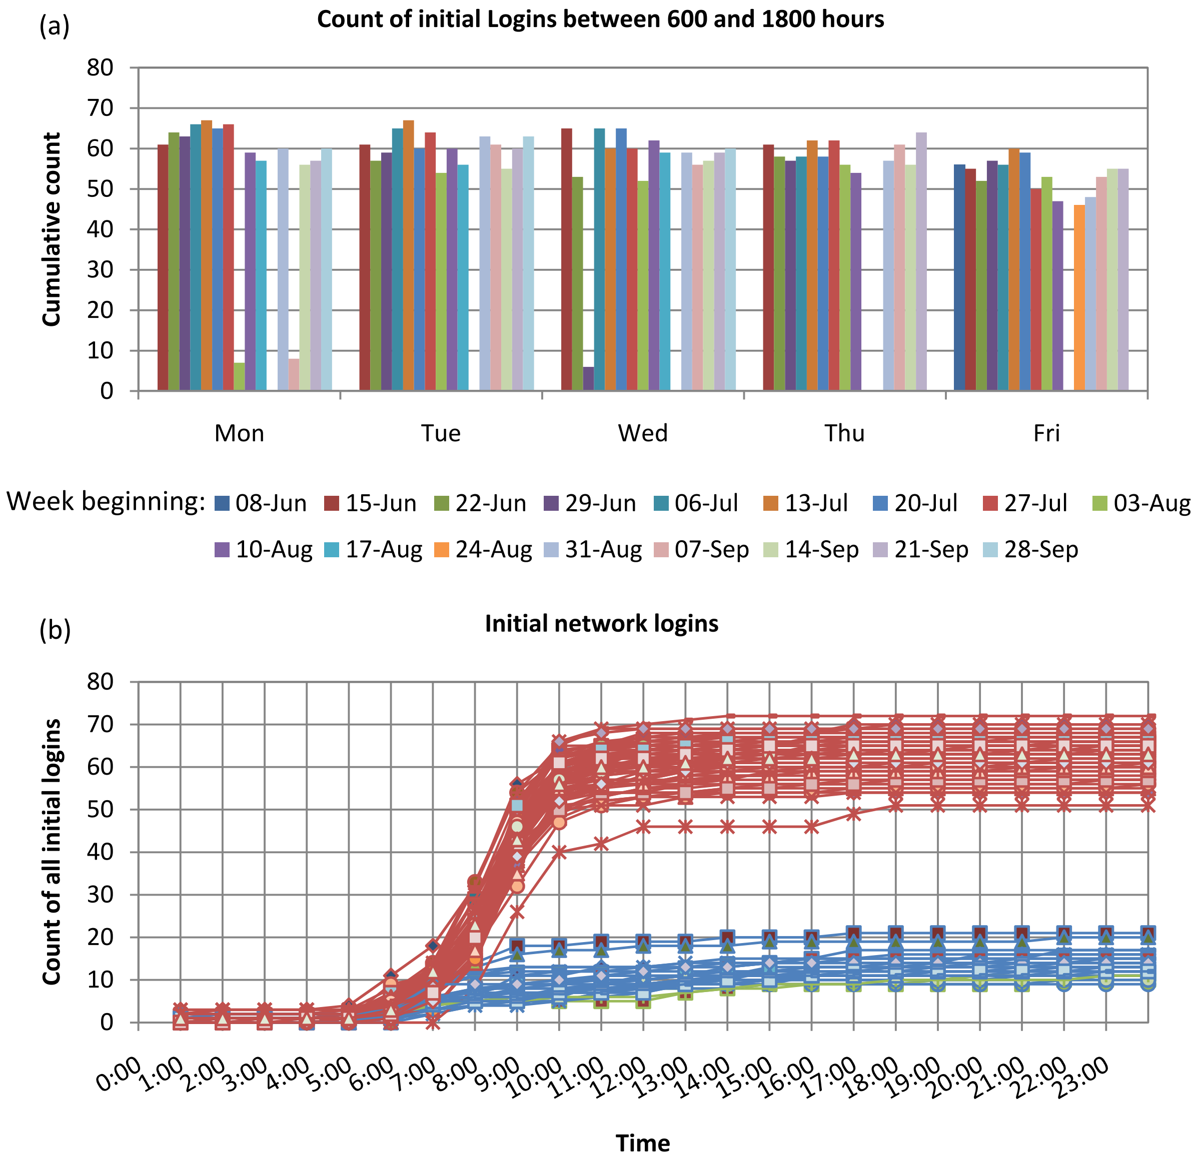

Figures 14(a) and 14(b). Figure 14(a) shows the total unique daily logins between 6:00 am and 6:00 pm. Missing columns were erroneous data days. Extremely low values on Monday and Wednesday were public holidays. Figure 14(b) shows the daily cumulative network login profile. Red lines are weekdays, blue lines are weekend days and green lines are public holidays.

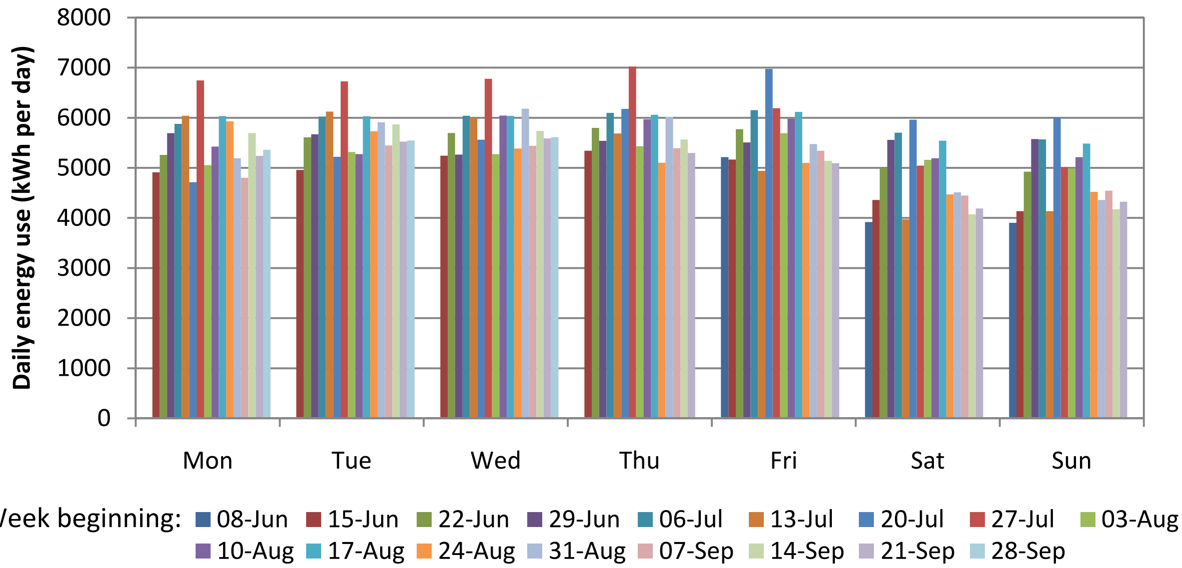

Figure 15. Daily energy use of the building excluding any chiller loads. The increase in load of approximately 80 occupants in the building on weekdays is around 25% of the weekend base load.

Several events (meetings, work activities) occurred during the study period that we expected to have an influence on certain sensors. Were these events, in fact, detected? Reduced activity was observed for public holidays in Ontario similar to weekends (Figure 11) and a slight reduction in activity for public holidays in neighbouring provinces (Figure 16). Other expected trends were an increase in activity for meetings and either an increase/decrease for other activities (Table 3) on some sensors. For example the commonly used car park was closed for a period of the study, requiring people to park in alternate lots and perhaps enter the building via a different door. This resulted in a reduction in use of the south door over the day (Figure 17(a)) and an increase in front door use (Figure 17(b)). An increase was only observed on one of these days, this could be due to people realising that there was another car park located to the north of the building and using the north door to enter the building instead. If the sensor on the north door had been working, one would have expected to see an increase in use of this door. Potentially the sum of all three doors would have remained constant throughout this period. Later in August we see a substantial increase in use of the front door (Figure 17(b)) coincident with a major meeting with partners from outside the host organization held in the main meeting room on level 3; external visitors would have used the front door to access the building. The same meeting room on level 3 was booked for a three day meeting in September, again with many external visitors. In Figure 18 an increase in activity is seen on a doors leading to this meeting room on two of these days. It is not known why door activity on the first day of the meeting was lower.

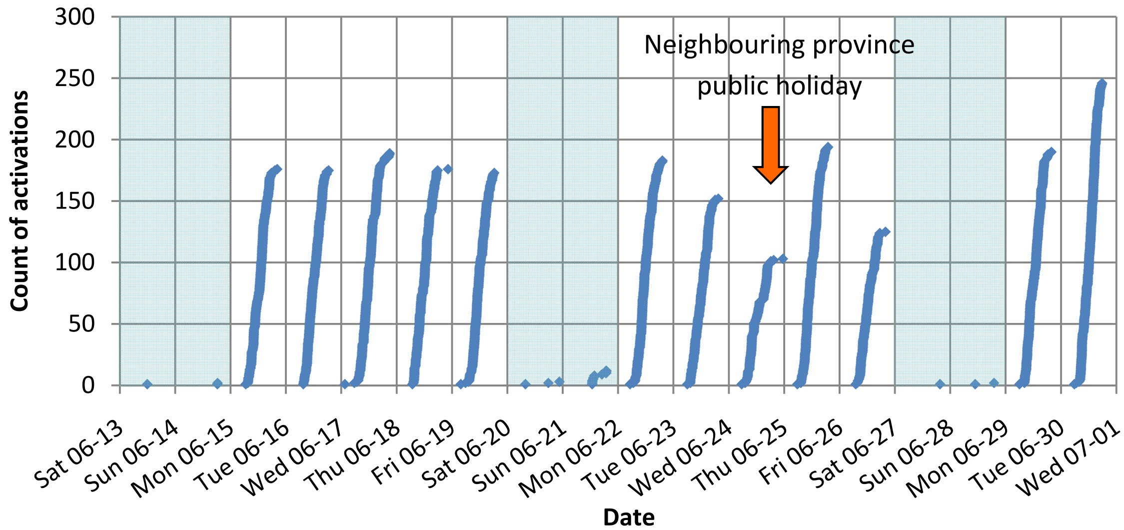

Figure 16. Cumulative motion sensor activity, per day, on level 2 during an example period.

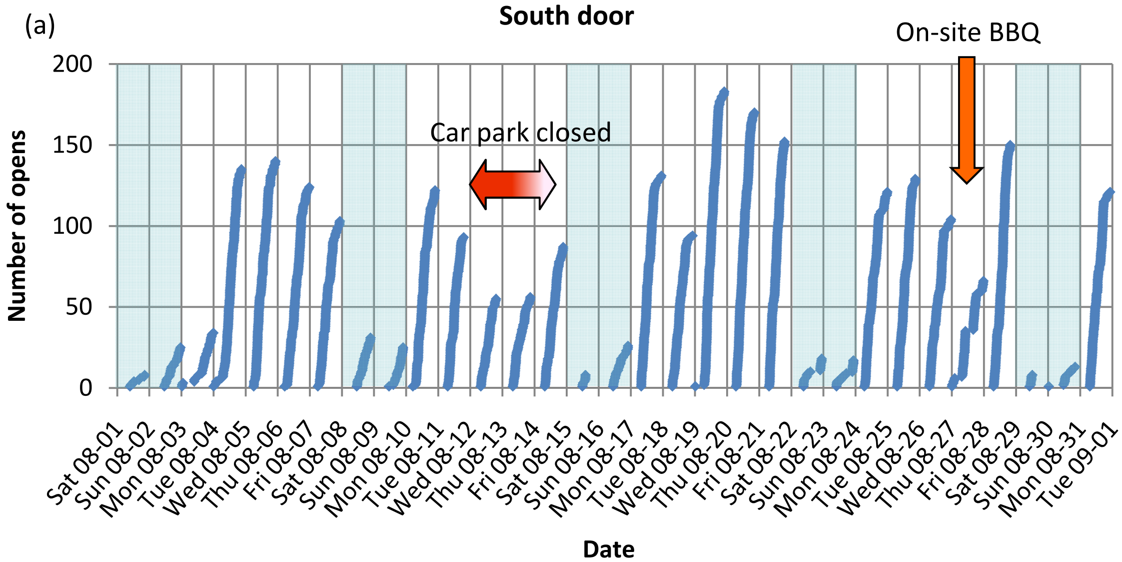

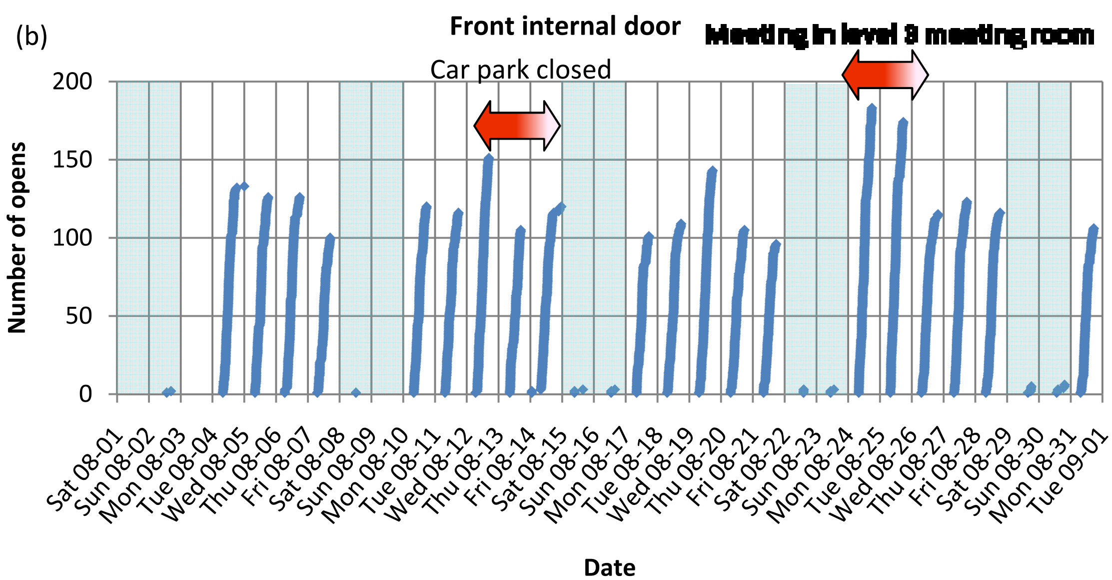

Figures 17(a) and (b). Cumulative door activity, per day, on the south door (a), and front door (b).

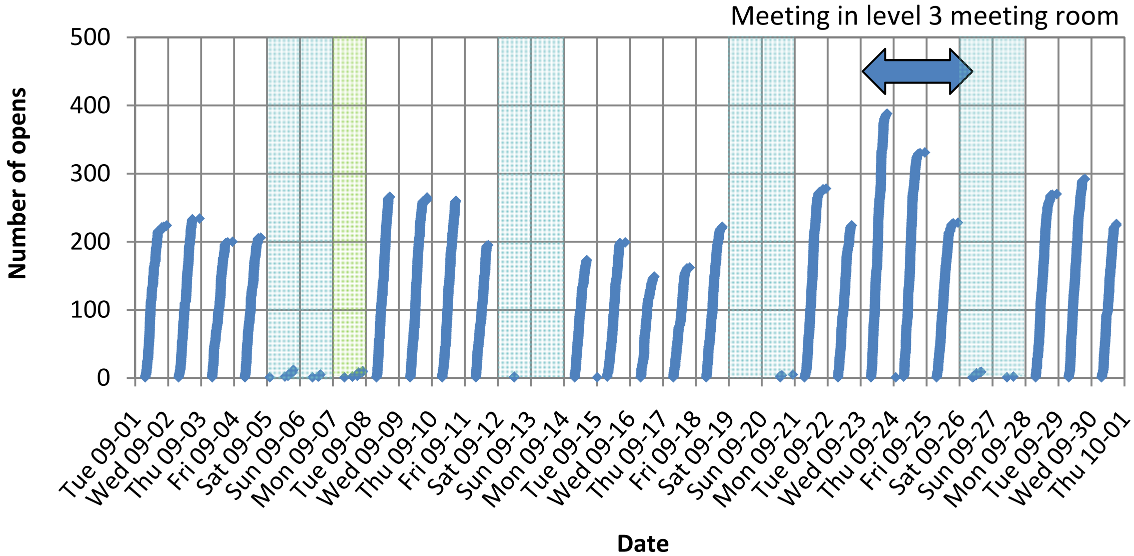

Figure 18. Cumulative door activity, per day, on the level 3 internal door.

3.3 Additional non-occupancy related information

Besides giving a potential measure of building occupancy, the sensors gave an insight in to how the doors operate and are used in the building. Five doors were monitored; two internal and three leading to the exterior of the building (two of which were on the main entry vestibule). All doors were fitted with an automatic door closer. The time it takes for the doors to close after being opened was calculated, and Figure 19 shows the distribution of the times. The average time for the doors to close was 5.6 (standard deviation 1.5) s for the level 3 door, 5.3 (s.d. 1.4) s for the level 1 door, 4.3 (s.d. 1.4) s for the south door, 4.8 (s.d. 1.3) s for the front internal door and 5.5 (s.d. 1.3) s for the front external door. Of the two doors at the main entry point to the building, the external door on average remained open longer. Assuming that the door closing mechanism was the same for both doors, the extended opening time was attributed to the negative pressurization of the building, which makes the external door slightly harder to open, as well as potentially restricting the closing speed of the door slightly. The level 1 door shows a distribution with two clear peaks, this might reflect a difference due to the direction of travel through the door: pulling the door open so that the gap is wide enough to pass through extends the total time the door is open compared to pushing the door open and passing through it. However, a similar effect on the other doors is not obvious.

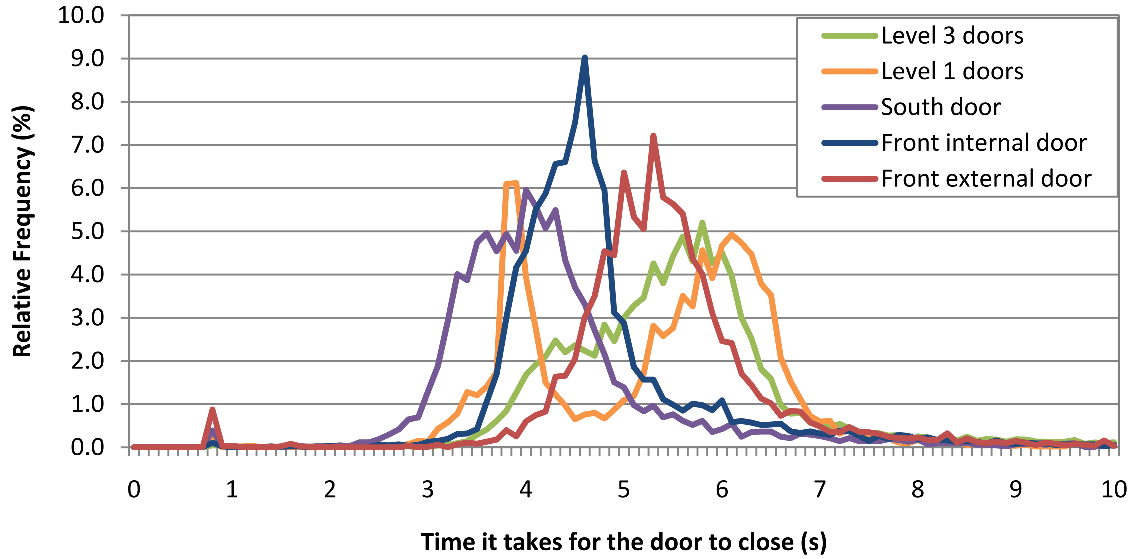

Figure 19. Curves showing the range in times that the 5 doors take to close upon opening. The direction of travel is unknown.

Of the double doors on the main exterior door vestibule, the sensor on the internal front door was more reliable than the sensor on the external front door, resulting in more signals being received from this sensor. Interpolating the results by removing heartbeats during office hours and matching each open and close of the external front door to that of the internal front door (as described above), an estimate of the time it takes for a person entering the building to open the external door first followed the internal door (negative number) was made, and vice versa on exiting of the building (positive number). Figure 20 shows a histogram of those entering and exiting the building. When entering the building (left hand side of the chart in Figure 20), it takes an average of 4.4 (s.d. 1.1) s to open the external door pass through the vestibule and open the internal door. On average there is 4.0 (s.d. 1.1) s between the two doors closing. Exiting the building is slightly quicker (right hand side of chart in Figure 20) at 3.2 (s.d. 1.2) s to consecutively open the two doors, a shorter period than entering the building due to the doors swinging in the direction of travel. Similarly to entering the building, the doors again close 4.1 (s.d. 0.8) s after each other. Time differences greater than 10 s were excluded from these analyses, as these times were assumed to be incorrectly matched.

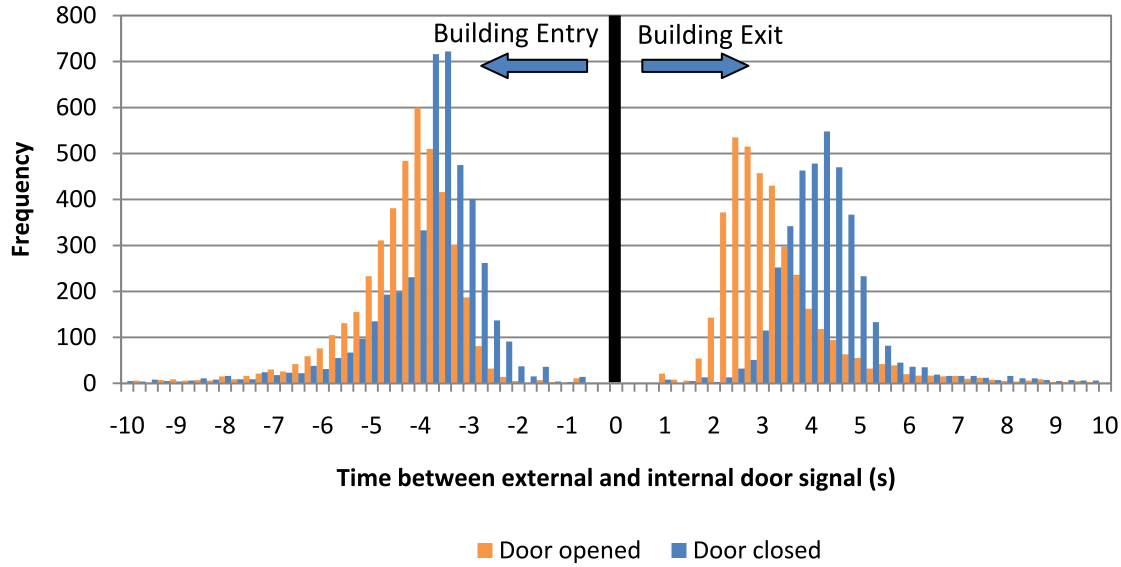

Figure 20. Matching the internal and external door signals of the main entry to the building, the time it takes from opening the first door to opening the second door of the vestibule was calculated (orange) as well as the time between the first door closing and the second door closing (blue). A negative time signifies that a person or group of people are entering the building. A positive time signifies that a person or group of people are exiting the building.

3.4 Correlations Between Different Metrics of Occupancy

Figure 21 shows the mean hourly values of the various occupancy-related metrics for which we collected data. The expected general patterns are apparent. All, except logins, show a general diurnal pattern, with some measures showing a decline over lunchtime. As described above, because we only have data on when people logged in, and not when they logged out, means that the login curve doesn't display a decline in the late afternoon.

Table 4 shows the simple correlations between the various metrics related to occupancy. Table 4(a) shows correlations derived from all data. In this case motion sensor data (for example) correlates better with the other metrics than logins does. However, this may be because of the invariability in login data in the afternoon, because we do not have data on subsequent logoffs. Table 4(b) shows correlations for data from midnight to noon (am period) only, when the login data is more likely to accurately represent the number of people currently logged in (as opposed to the number who have logged in, but may since have logged out). In this case we see the correlations between logins and the other metrics is much higher, and slightly higher than the correlations between motion sensor activations and other metrics. The intercorrelation between logins and motion sensor activations is very high.

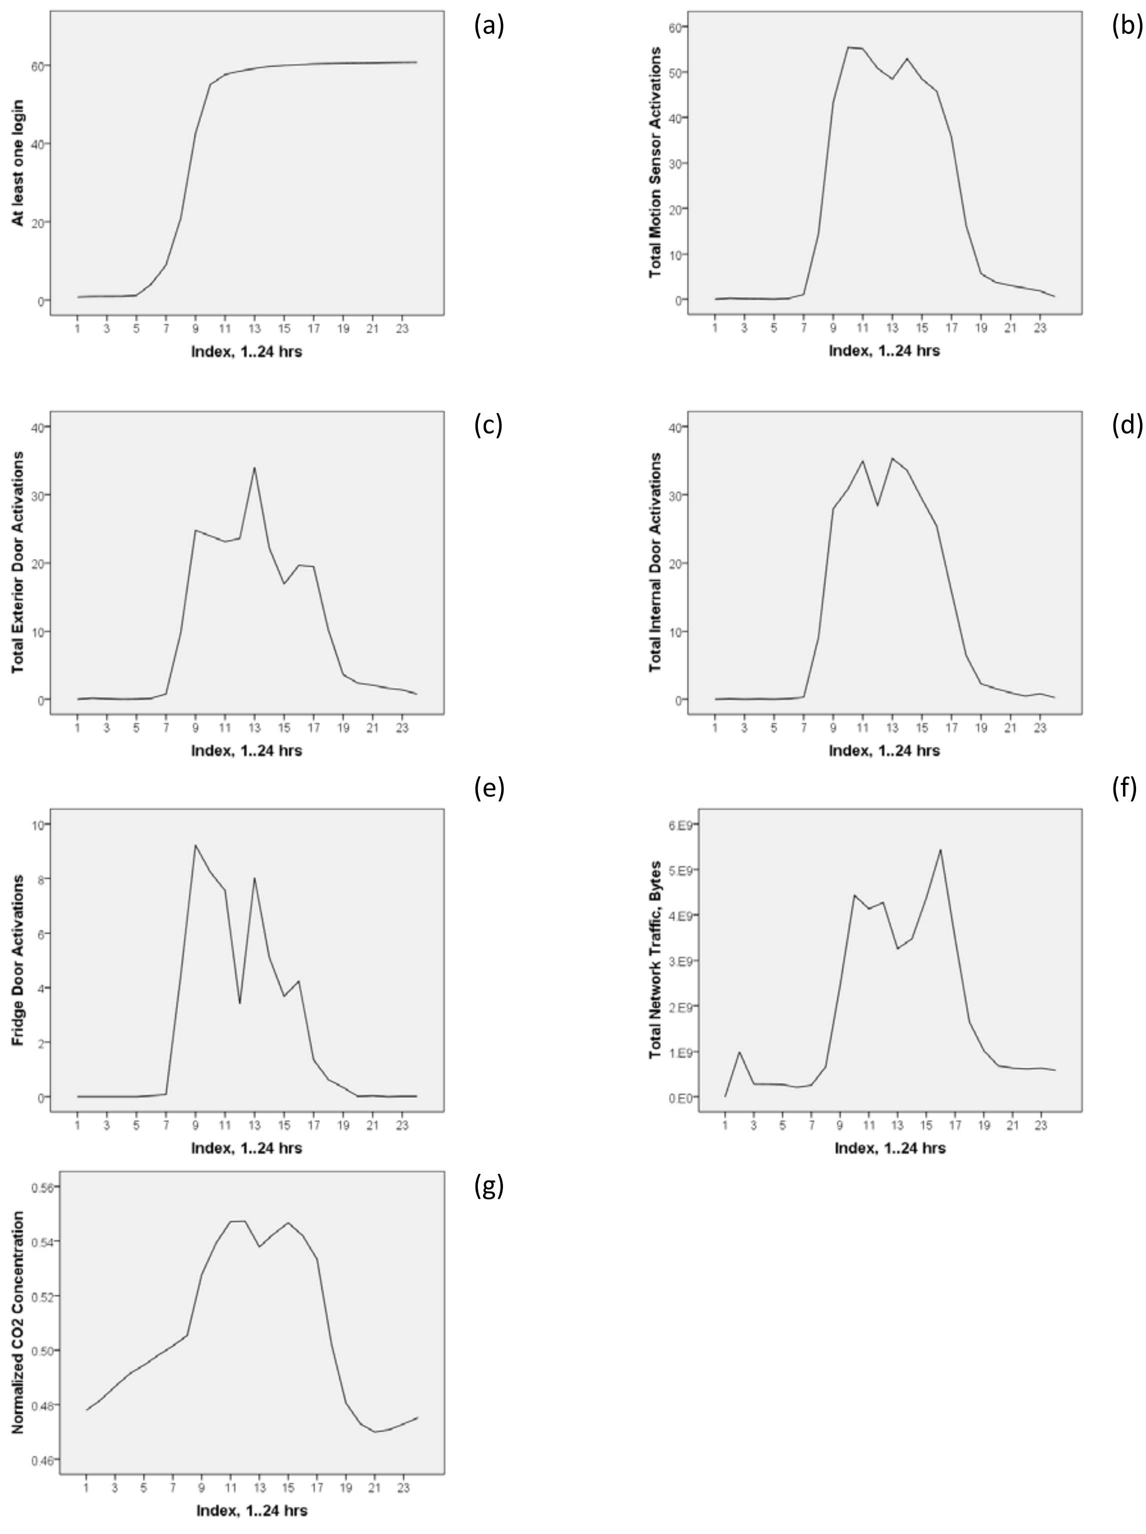

Figure 21. Mean hourly values of various occupancy-related metrics.

Table 4. Simple correlations between various metrics of occupancy over all days in the dataset (a) all hours; (b) For hours 1-12 only.

(a) All hours	At least one login	Total motion sensor activations	Total external door activations	Total internal door activations	Fridge door activations	Total network traffic	Carbon dioxide concent.
At least one login							
Total motion sensor activations	.506**						
Total external door activations	.521**	.867**					
Total internal door activations	.478**	.895**	.889**				
Fridge door activations	.315**	.699**	.744**	.801**			
Total network traffic	.337**	.525**	.435**	.472**	.342**		
Carbon dioxide concentration	.263**	.662**	.669**	.690**	.488**	.383**	

(b) AM hours only	At least one login	Total motion sensor activations	Total external door activations	Total internal door activations	Fridge door activations	Total network traffic	Carbon dioxide concent.
At least one login							
Total motion sensor activations	.946**						
Total external door activations	.925**	.911**					
Total internal door activations	.913**	.916**	.912**				
Fridge door activations	.742**	.739**	.790**	.823**			
Total network traffic	.533**	.530**	.476**	.512**	.371**		
Carbon dioxide concentration	.638**	.598**	.668**	.656**	.462**	.363**	

** Correlation is significant at the 0.01 level (2-tailed).

3.5 Energy Use Forecasts Using Time-Series Regression

3.5.1 Introduction

With the advent of the SmartGrid and utility tariffs that charge substantially more for electricity at times of peak demand, it is becoming increasingly desirable for facility managers to be able to forecast their facility's electricity use some hours into the future, and to take operations decisions based on this information. Such decisions might involve load shedding, pre-cooling, charging of ice storage, activation of local generation, or a variety of other actions to advantageously vary the load profile [e.g. Zhou et al., 2008; Neto & Fiorelli, 2008; Hoffman, 1998].

To date, relatively crude methods for prediction have been adopted in buildings, such as assuming the average load profile from previous days, or simple regression equations. In recent decades a new class of models has been developed that promises better prediction accuracy, and that has been successfully applied in other domains, particularly in the economic and business domains [Montgomery et al., 2008]. This class of models has the general name ARIMAX, which stands for Auto Regressive Integrated Moving Average with eXternal (or eXogenous) input, which is part of a broader class of time series analysis techniques[2]. The "integrated" part of the name refers to the fact that it is often required that one runs the analysis on the change in the dependent variable of interest (known as "differencing"), rather than the raw variable itself, to render the series stationary[3]. "Auto regressive" indicates that the forecasted value of the dependent variable may be predicted from prior, known, values of the dependent variable. "Moving average" indicates that the forecasted value of the dependent variable may be predicted from prior values of the error term in the prediction. "External input" refers to the optional use of independent predictor variables. The general notation for such a model is ARIMAX(p,d,q); if independent predictor variables are not employed then the notation is ARIMA only. The "p" indicates how far back in time one goes in using prior values of the variable of interest. For example, if the current value of a variable measured every hour is predicted using values of that variable from one and two hours ago (known as "lag 1" and "lag 2"), p=2. Similarly, q refers to how many lags in the error term are used and "d" indicates how many times one takes the difference of the dependent variable[4]. It is often the case that the variable of interest exhibits obvious periodic behaviour, generally referred to as "seasonal" behaviour. For example, building power use often displays a clear diurnal pattern; if one measures power hourly then there will be a seasonality of order 24. For modelling, one creates a new seasonal variable to reflect this variation, which is the current value of the dependent variable minus the value from one seasonal period ago. One can then apply differencing and lags to this variable and include these terms in the model. Thus the final general notation is ARIMAX(p,d,q)(P,D,Q)$_s$, where P, D, and Q have the same meaning as above, but now refer to the seasonal variable, and s is the order of seasonality with respect to the measurement interval[5].

The most general mathematical form of the ARIMAX model equation is as follows [UC, 2010]:

(Eq. 1)

$$(1-B)^d(1-B^s)^D Y_t = \mu + \Psi_i(B)X_{i,t} + \frac{\theta(B)\theta_s(B^s)}{\phi(B)\phi_s(B^s)}\, a_t$$

---

[2] Other advanced techniques that have also been applied to building-related applications include artificial neural networks (ANN) [Kawashima et al., 1995; Karatasou et al., 2006; Neto & Fiorelli, 2008].

[3] A "stationary" series is one in which the values vary around an unchanging mean, and the variance over time is constant. Stationary series are a requirement for ARIMA models.

[4] If no differencing is used, the model may be specified as ARMA(p,q) or ARIMAX(p,0,q), if there are no moving average components the model may be specified as AR(p,d) or ARIMAX(p,d,0), if there are no auto-regressive components the model may be specified as AM(d,q) or ARIMAX(0,d,q).

[5] A seasonal ARIMAX model is sometimes referred to as a SARIMAX model.

where,

Y_t is the dependent time series

$X_{i,t}$ is a set of *i* external predictor time series

a_t is a white noise time series representing random error, the values of this series are not known *a priori*, but are an outcome of the iterative parameter estimation methods used to generate the best-fitting model

t indexes time

μ is the mean of the series (=0 when series is differenced)

B is the backshift operator; i.e. $BY_t = Y_{t-1}$; $B^{12}Y_t = Y_{t-12}$; $BB^{12}Y_t = B^{13}Y_t$

$\phi(B)$ is the autoregressive operator, a polynomial of order p in the backshift operator:

$$\phi(B) = 1 - \phi_1 B - \cdots - \phi_p B^p$$

$\phi_s(B)$ is, similarly, the seasonal autoregressive operator, a polynomial of order P:

$$\phi_s(B^s) = 1 - \phi_{s,1} B^s - \cdots - \phi_{s,P} B^{sP}$$

$\theta(B)$ is the moving average operator, a polynomial of order q in the backshift operator:

$$\theta(B) = 1 - \theta_1 B - \cdots - \theta_q B^q$$

$\theta_s(B)$ is, similarly, the seasonal moving average operator, a polynomial of order Q:

$$\theta_s(B^s) = 1 - \theta_{s,1} B^s - \cdots - \theta_{s,Q} B^{sQ}$$

$\Psi_i(B)$ is a transfer function for the effect of $X_{i,t}$ on Y_t:

$$\Psi_i(B) = \frac{\omega_i(B)\omega_{s,i}(B^s)}{\delta_i(B)\delta_{s,i}(B^s)}(1-B)^{d_i}(1-B^s)^{D_i}B^{k_i}$$

$\delta_i(B)$ is the denominator polynomial in the backshift operator, for the *i*th predictor:

$$\delta_i(B) = 1 - \delta_{i,1} B - \cdots - \delta_{i,p_i} B^{p_i}$$

$\delta_{s,i}(B)$ is similarly, the denominator seasonal polynomial, for the *i*th predictor:

$$\delta_{s,i}(B) = 1 - \delta_{s,i,1} B - \cdots - \delta_{s,i,P_i} B^{sP_i}$$

$\omega_i(B)$ is the numerator polynomial in the backshift operator, for the *i*th predictor:

$$\omega_i(B) = \omega_{i,0} - \omega_{i,1} B - \cdots - \omega_{i,q_i} B^{q_i}$$

$\omega_{s,i}(B)$ is similarly, the numerator seasonal polynomial, for the *i*th predictor:

$$\omega_{s,i}(B) = \omega_{s,i,0} - \omega_{s,i,1}B - \cdots - \omega_{s,i,Q_i}B^{sQ_i}$$

k_i is the pure time delay for the effect of the *i*th predictor on the dependent variable (i.e. it may be the case that the predictor cannot have an effect for a certain number of time steps for basic physical reasons)

Note that the denominator polynomial plays a similar role for the predictor variable as the autoregressive operator plays for the dependent variable, and the numerator polynomial plays a similar role for the predictor variable as the moving average operator plays for the dependent variable.

There have been some attempts to apply ARIMAX models to building-related applications in the literature. Loveday & Craggs [1993], provide a more thorough mathematical introduction to the technique in the context of building issues (adding to the more general development found in Montgomery et al. [2008]), and describe its application to the modelling and forecasting of room temperature. Rios-Moreno et al. [2007] also modelled room temperature and included several external input variables[6]. Lowry et al. [2007] provide several excellent examples for modelling of water and fuel use in a variety of buildings. Hoffman [1998] uses such models to forecast and control the peak demand for electricity at a government complex.

In this report we are particularly interested in the value of metrics related to building occupancy as external inputs (independent predictor variables) in models of building power draw. Because ARIMA models use prior values of the dependent variable, in this case power use, and because we know that power use in a building is correlated with occupancy, the auto regressive and moving average components will implicitly carry the effect of occupancy. The question we wished to explore was whether including an occupancy metric as an explicit independent variable would provide additional useful information, and would therefore improve the model accuracy. Loveday & Craggs [1993] suggested that variance in their ARIMA model of indoor temperature could be partially explained by variations in occupancy, and that forecasting accuracy could be improved by adding it as a variable in the model. Neto & Fiorelli [2008] lament the lack of occupancy data for use in their ANN model of building energy use. Further, they conducted parametric studies with a building simulation tool based on physical principles (EnergyPlus) and observed that the sensitivity of energy use to changes in occupancy and occupancy-related loads (lighting and electrical equipment) was of a similar order to the sensitivity to external climate variables.

3.5.2 Method

3.5.2.1 Variables in the Model

In this study we apply ARIMAX to model and forecast total power draw of our study building, with metrics of total building occupancy as external inputs. All variables used were hourly values (derived

[6] Note, their model notation differs from the standard form

from the measurements of the various variables at shorter time scales). The dataset employed was for weekdays only (including three public holidays), because weekends were judged to be irrelevant for this building where weekend occupancy was virtually nil and our long term interest was in peak demand load control, which occurs exclusively on weekdays (Neto & Fiorelli [2008], found in applying ANN models that separate models for weekdays and weekends produced more accurate results than a single model with a day type variable).

As described above, we had data on total building electricity use, and also for the chiller separately. Initial analysis showed that chiller power draw displayed only a very small correlation with cooling degree days. This was very surprising, and caused us to question the simple relationship in this building between climate, occupancy and chiller power. Therefore, we decided to subtract chiller power from total building power, and to continue our analyses with power not associated with thermal load. Thus the power variable included lighting, office equipment, lab equipment, and other plug loads that were likely directly related to occupancy, and thus perhaps of more relevance to the goals of this study[7] (Lowry et al. [2007] suggested that non-weather related energy use would benefit from a separate analysis). Figure 22 shows the hourly time series for total building power minus chiller power, and Figure 23 shows the average hourly values. An initial analysis suggested that of the variables we collected that were indicative of total building occupancy; network logins and motion sensor counts were likely to be the most useful, thus we proceeded with these as potential external inputs to the models. Analysis showed that looking at the change of variables from hour-to-hour (one level of differencing) was more appropriate for modelling purposes than using absolute values, and Figure 24 shows average hourly values for these variables.

[7] An earlier, linear regression time series analysis suggested that using total building power including the chiller yielded similar final results and conclusions.

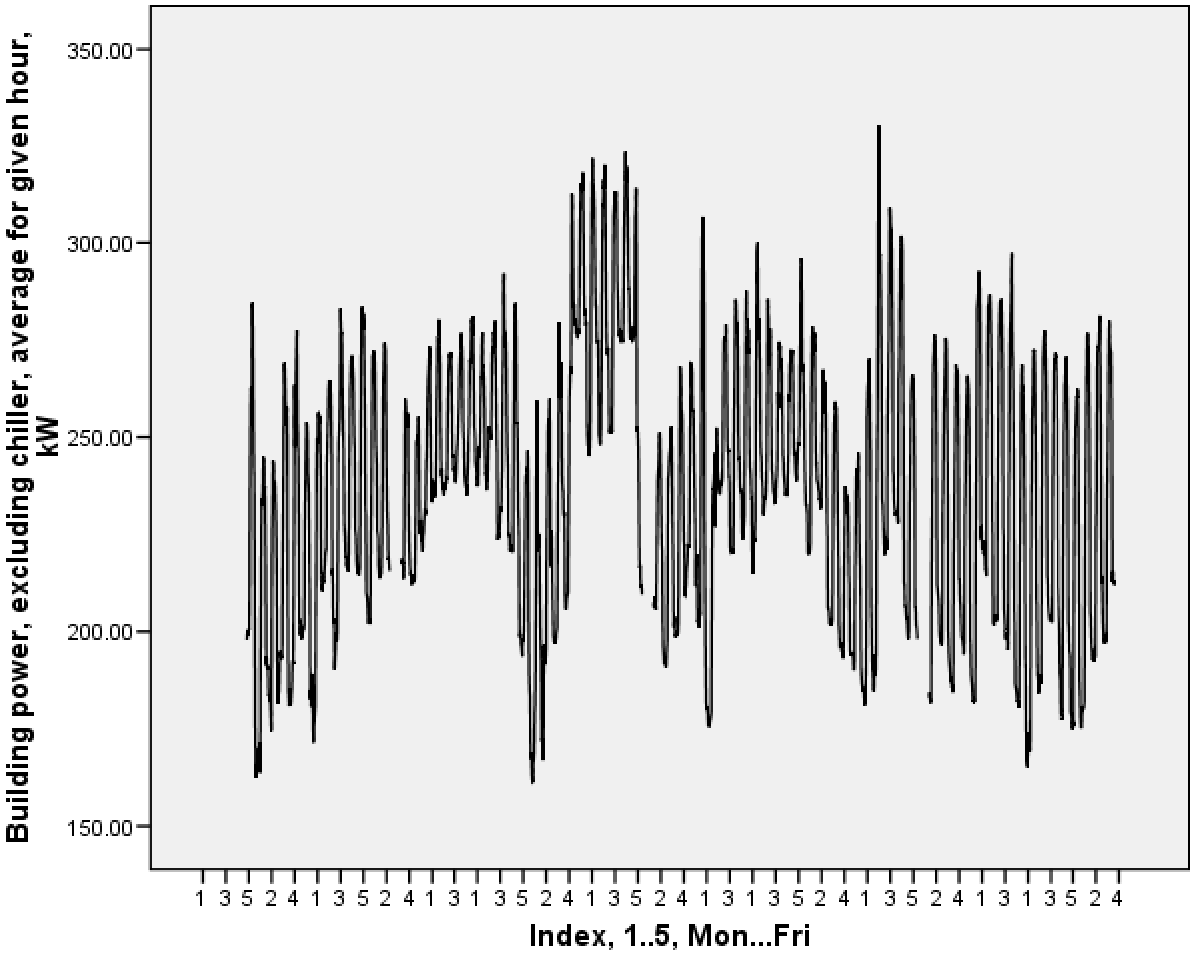

Figure 22. Hourly time series of total building power minus chiller.

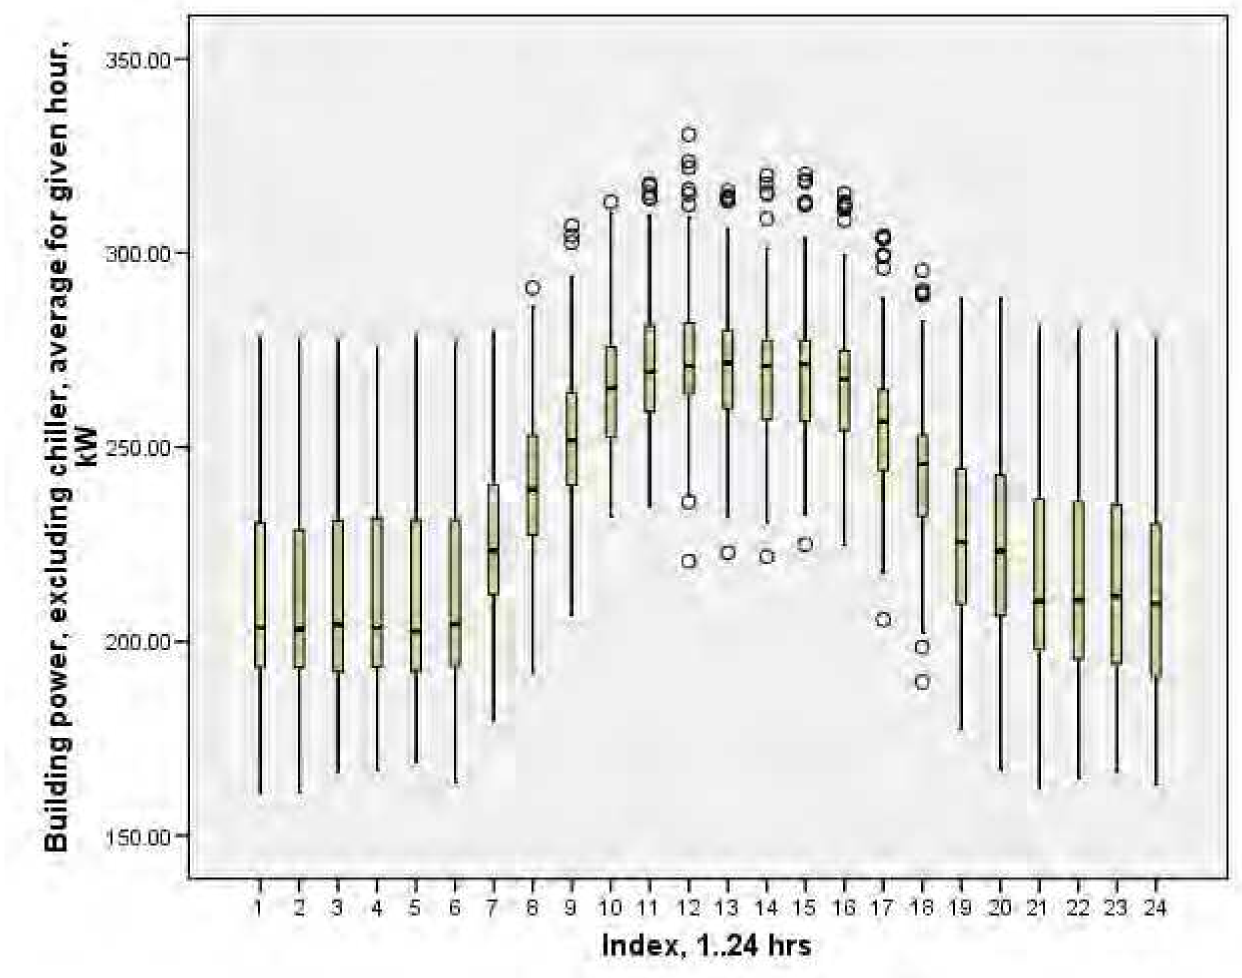

Figure 23. Average hourly values of total building power minus chiller. Length of box is the interquartile range (IQR); line in box is median; 'o' are outlier values more than 1.5 IQR from the end of the box; '*' are outlier values more than 3 IQR from the end of the box; whiskers show min. to max. range excluding outliers as defined above.

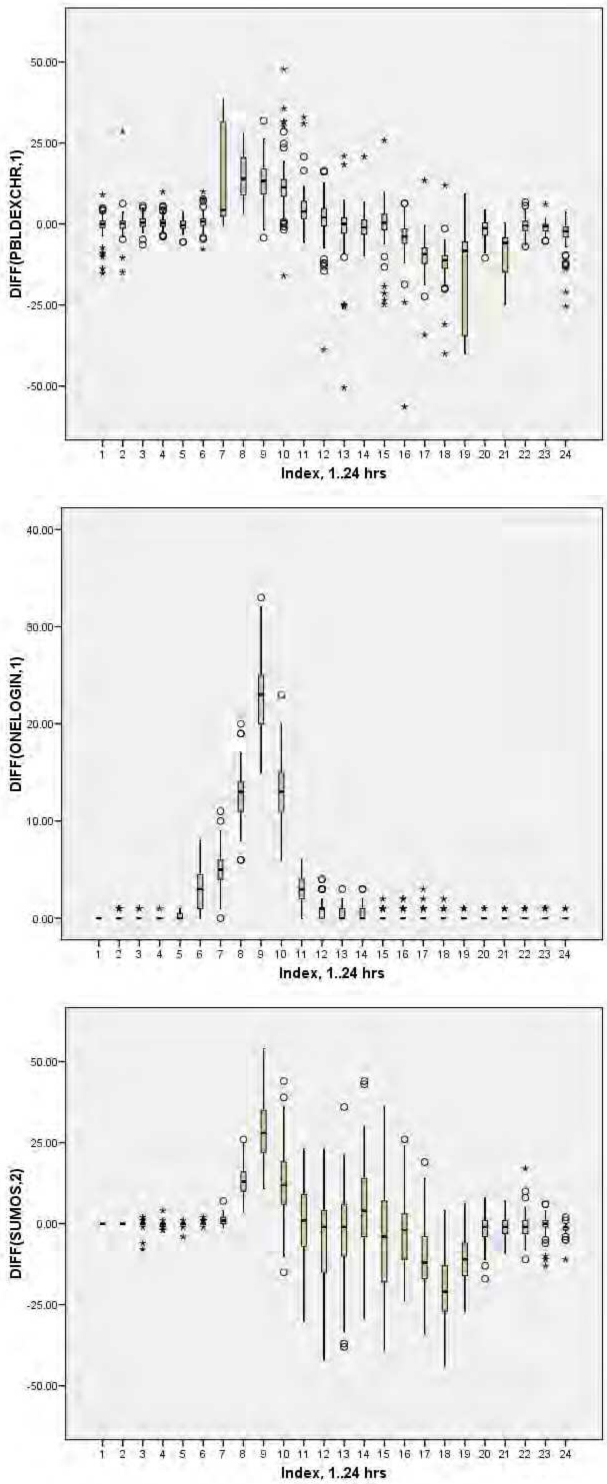

Figure 24. Average hourly values of the change in total building power minus chiller, unique network logins, and motion sensor activation (sum from all three sensors). Length of box is the interquartile range (IQR); line in box is median; 'o' are outlier values more than 1.5 IQR from the end of the box; '*' are outlier values more than 3 IQR from the end of the box; whiskers show min. to max. range excluding outliers as defined above.

Figure 23 shows that average power draw rise during the day compared to overnight was about 60-70 kW. Login data showed that there were typically 60-70 unique logins per day, suggesting a maximum building occupancy of the same magnitude, consistent with the 81 available workstations in the building. This also suggests that each occupant is responsible for about 1 kW of additional daytime load, this seems to be about the correct order of magnitude given the office equipment and lighting associated with each office and the laboratory loads that are initiated by the arrival of occupants. Figure 24 shows the expected rise (and fall) of building power draw coincident with the rise (and fall of indicators of occupancy). Recall that the login data is limited to unique login events, subsequent logoffs were not recorded. This means we have data related to occupancy as the building fills up, but not as it empties. This suggests that our login data might be only partially useful in predicting building power draw. This same limitation does not apply to the motion sensor data. On the other hand, we expect that login data, if coupled with logout data, would be a better measure of "ground truth" occupancy than motion sensor data.

3.5.2.2 Modelling Process

In general, we followed the model estimation and fitting procedure outlined in Montgomery et al. [2008]. All analyses were conducted using the Forecasting module in SPSS version 18. Some SPSS routines require complete data sets, whereas we had some gaps in our data for some variables due to imperfect data collection systems and subsequent data cleaning (as described above). In such cases we imputed the missing values with the mean of the non-missing values for that hour and day of the week. We had 79 days (with 24 hourly values each) of complete and continuous data for building power draw; of these 79 days, 5 complete days of network login data and 17 complete days of motion sensor data were missing and were imputed.

Data were available for some or all variables from 1 am on June 12th, 2009 (Week 1 Day 1 Hour 1) to midnight on September 30th, 2009 (Week 17 Day 3 Hour 24). Initial model exploration was conducted on the majority of the dataset (Week 2 Day 1 Hour 1 to Week 16 Day 1 Hour 7) to suggest the form of the best-fitting, parsimonious model. The dataset was then split into three for final model testing. The best-fitting model form from the prior step was applied to data from Week 2 Day 1 Hour 1 to Week 10 Day 5 Hour 24, to check that it still fit well, and was then applied to Week 11 Day 1 Hour 1 to Week 16 Day 5 Hour 24 to make sure it was robust. Finally, the model as derived from the majority of the sample was then used to forecast power draw for the immediate future hours (Week 16 Day 1 Hour 8 onwards) and compared to the actual power draw data for this same period; i.e., data that were not used in the derivation of the model.

Public holidays are not "events" in this model. Events in the terminology of such models are occurrences expected to have substantial, "hangover" effects on the next few days, or permanent effects. Inspection of the data showed that building power and occupancy returned to normal quickly on the days following public holidays[8]. Nevertheless, there were individual hours, or small sets of adjacent hours, that had unusual values, for reasons that we could not explain. Therefore, we did choose to model outliers in the analysis using SPSS routines (using Additive, Levelshift, Seasonaladditive and Localtrend options); this improved model fit markedly, even though this required substantially more degrees of freedom.

We included no delay factor in our models. We expected changes in occupancy metrics could show up in power draw within the one hour period of the data, occupants are likely to switch on office equipment and interior lights within the same hour they login and begin to trigger motion sensors.

3.5.3 Results

Inspection of the time series building power data suggested an obvious diurnal seasonality (order 24), and that the data should be differenced once. We then plotted the autocorrelation and partial autocorrelation functions (ACF and PACF) for the differenced and seasonally differenced data; the results are shown in Figure 25. The format of these plots suggests a moving average model of lag 2, with a seasonal lag, will be a good fit to the data. There are a few other significant lag terms indicated, but these are irregularly spaced and have no obvious theoretical basis. There are no higher-order seasonal lags that show strong significance. For example, if there was a strong weekly seasonality; i.e. all Mondays are like previous Mondays and different from Tuesdays, for example, then a significant and large peak at lag 120 would show up in the ACF (in this dataset we have 5-day (work) weeks, and 5 x 24 hrs = 120). Therefore, the following analyses adopted a single, diurnal seasonality.

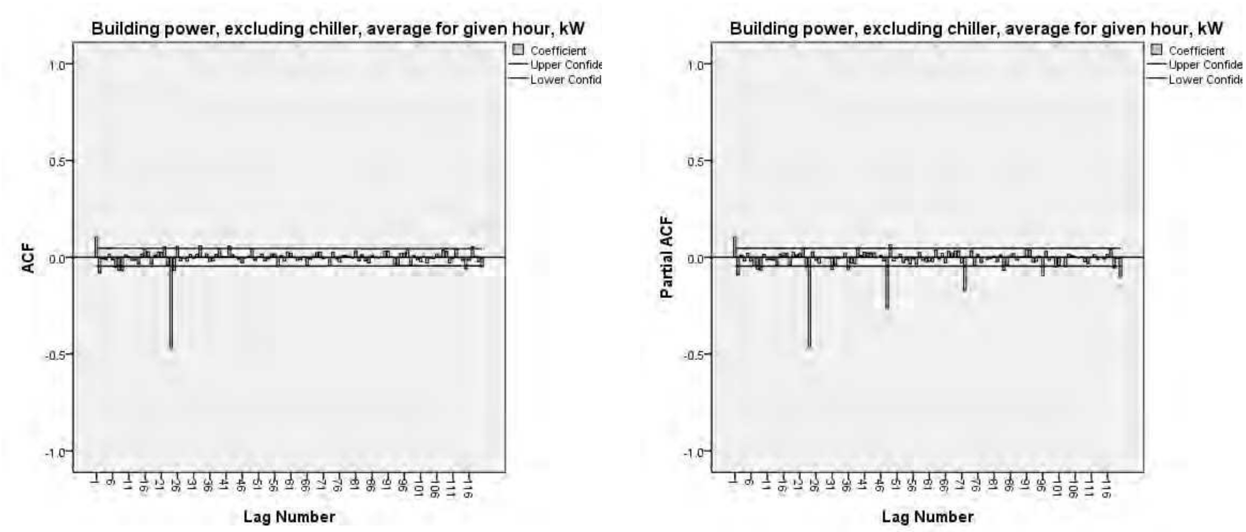

Figure 25. ACF and PACF for the differenced and seasonally differenced building power data.

---

[8] An event in the context of a building could include a permanent upgrade to a major building component, a change in tenant, or a change in utility tariffs, for example.

Initially we derived a model for building power on the majority of the dataset (Week 2 Day 1 Hour 1 to Week 16 Day 1 Hour 7) without occupancy-related or other predictors (i.e., an ARIMA model), with the goal of seeing if subsequently adding occupancy information improved the model. We did try cooling degree hours (base 18 ^{o}C, CDH18) for each hour (differenced) as a predictor in an ARIMAX model to check for any residual climate dependence; for example, there were several refrigerators in lunch areas and laboratories that might have been affected by external climate. CDH18 was significant in the model but actually worsened the model fit (defined by several parameters). We therefore decided on the pure ARIMA model as the base model for comparison to later models. The automatically-generated, best fit, model from SPSS Forecasting did not yield a lag 2 term, but did include a lag 7 term. However, this did not have any obvious physical explanation, and therefore to keep the model compact, simple, and parsimonious we chose to drop this term from the final base model; this choice had only a tiny effect on the overall model fit. Therefore, the final base model was ARIMA(0,1,1)(0,1,1)$_{24}$, and the general model of Eq. 1 thus simplifies to:

(Eq. 2)

$$(1-B)(1-B^{24})Y_t = (1-\theta_1 B)(1-\theta_{s,1}B^{24})\, a_t$$

The model parameters and fit statistics are shown in Table 5. Note that the model parameters included only moving average (MA) terms, and not autoregressive (AR) terms. Also note, that the procedure identified only 37 outliers from a total of 1687 hours, or 2.2% of values. Eq. 2 further expands as follows:

$$Y_t = Y_{t-1} + Y_{t-24} - Y_{t-25} - \theta_1 a_{t-1} - \theta_{s,1} a_{t-24} + \theta_1\theta_{s,1} a_{t-25} + a_t$$

This illustrates that the present value of building power depends on the value from 1, 24 and 25 hours in the past (due to differencing), and fractional values of the estimation error 1, 24 and 25 hours in the past (due to differencing in the moving average calculation), leaving the present value of the error, a_t, as the residual.

The fit statistics will serve as a guide to whether including explicit measures of occupancy as predictors improve the models. Note that the Stationary R-squared is a better measure of variance explained by the model than the simple R-squared when dealing with time-series data, and higher values indicate a better fit. RMSE (Root Mean Square Error), MAPE (Mean Absolute Percentage Error), MAE (Mean Absolute Error), MaxAPE (Maximum Absolute Percentage Error), MaxAE (Maximum Absolute Error) are all measures where lower values indicate better performance. Normalized BIC (Bayesian Information) accounts for the number of parameters used in the model, and may penalize non-parsimonious models; lower values indicate better model performance. This basic model seems to be a good fit to the data from which it was derived; the RMSE is 4.3 kW in values that vary in the 150-300 kW range, which translates into a mean absolute percentage error (MAPE) of 1.2%, and a maximum absolute percentage error (MaxAPE) of 8.8%.

Table 5. Model with no external predictors, Week 2 Day 1 Hour 1 to Week 16 Day 1 Hour 7. (a) shows model parameters, and (b) shows fit statistics.

(a) ARIMA Model Parameters					
		Estimate	SE	t	Sig.
MA (Power), θ_1	Lag 1	-.150	.025	-5.908	.000
MA (Power), Seasonal, $\theta_{s,1}$	Lag 1	.753	.017	43.377	.000

(b) Model Fit	
Fit Statistic	
Stationary R-squared	.679
R-squared	.984
RMSE	4.268
MAPE	1.244
MaxAPE	8.807
MAE	2.958
MaxAE	20.535
Normalized BIC	3.076

In the next step we added login data, again differenced and seasonally differenced, as a predictor in the model; this was done using the Transfer Function option in SPSS Forecasting. Logins were significant in the model. The final base model was ARIMAX(0,1,1)(0,1,1)$_{24}$, and the general model of Eq. 1 thus simplifies to:

(Eq. 3)

$$(1-B)(1-B^{24})Y_t = \omega_0(1-B)(1-B^{24})X_t + (1-\theta_1 B)(1-\theta_{s,1}B^{24})\,a_t$$

The model parameters and fit statistics shown are shown in Table 6; fit statistics were generally improved, albeit by relatively small amounts.

Table 6. Model with logins as predictor, Week 2 Day 1 Hour 1 to Week 16 Day 1 Hour 7. (a) shows model parameters, and (b) shows fit statistics.

(a) ARIMA Model Parameters					
		Estimate	SE	t	Sig.
MA (Power) , θ_1	Lag 1	-.150	.026	-5.856	.000
MA (Power), Seasonal, $\theta_{s,1}$	Lag 1	.757	.017	43.566	.000
Numerator (Logins), ω_0	Lag 0	.100	.034	2.968	.003

(b) Model Fit	
Fit Statistic	
Stationary R-squared	.702
R-squared	.985
RMSE	4.127
MAPE	1.217
MaxAPE	9.307
MAE	2.889
MaxAE	17.821
Normalized BIC	3.040

We tried adding motion sensor data as a predictor to the base model from Table 5, but this predictor was not statistically significant and did not improve the model.

Next, we explored model robustness by specifying the same model form in Table 5 to a split sample drawn from the same sample. Tables 7-8 demonstrate that the all three model parameters were maintained as significant in both split samples, and furthermore, the parameter estimates are similar.

Table 7. Model parameters with logins as predictor, Week 2 Day 1 Hour 1 to Week 10 Day 5 Hour 24.

ARIMA Model Parameters					
		Estimate	SE	t	Sig.
MA (Power) , θ_1	Lag 1	-.146	.032	-4.562	.000
MA (Power), Seasonal, $\theta_{s,1}$	Lag 1	.831	.020	41.326	.000
Numerator (Logins), ω_0	Lag 0	.133	.044	3.019	.003

Table 8. Model parameters with logins as predictor, Week 11 Day 1 Hour 1 to Week 16 Day 5 Hour 24.

ARIMA Model Parameters					
		Estimate	SE	t	Sig.
MA (Power) , θ_1	Lag 1	-.193	.039	-4.981	.000
MA (Power), Seasonal, $\theta_{s,1}$	Lag 1	.650	.033	19.973	.000
Numerator (Logins), ω_0	Lag 0	.119	.047	2.510	.012

We now have confidence that the model is robust, and so returned to the model derived from the majority of data (Table 5) to conduct the final step, which was to use the model to forecast data into the future; i.e. to compare its forecasts to data from which the model was not derived. Forecasting was done using the methods in the SPSS Forecasting module, which takes the model equation and predicts future values. Note, that using our final model equation, which took the form shown in Eq. 3, to forecast the value of Y_{t+1} requires X_{t+1}. In this case, X_{t+1} is unknown *a priori*, and therefore it too must be forecast in some manner. This may be done with a separate ARIMA model for X alone [Montgomery et al., 2008].

We used the model in Table 5, calculated from data up to Week 16 Day 1 Hour 7 to forecast hourly values for building power for the remainder of the dataset, thus allowing us to compare the forecast to the actual recorded building power. Here we present the forecast for the remainder of Week 16 Day 1 (a Monday). Figure 26 shows the forecast made at Hour 7 for the remainder of the day, and the actual building power. For comparison we chose two simple forecasts that might commonly be invoked: the average for all weekdays in the sample up to Week 15 Day 5; and the values from Week 15 Day 1 (the previous Monday). Week 16 Day 1 was obviously unusual in some ways compared to the average of all previous weekdays, in that although the peak value was similar to the average day, the early morning and late evening values were much lower, resulting in a load shape of much greater amplitude than the average day. The previous Monday shares the same low values in the early morning (indicative, perhaps, of more loads being switched off over the previous weekend), but does not share the same low values at the end of the day. The ARIMAX model shows excellent fit to the data up to Hour 7, but remember that these data were used to create the model. After that, the model tends to under-predict the building power draw, predicting a peak load around 20 kW lower than the actual peak; the previous Monday is a similar amount over the actual peak value, whereas the average of all days is close to the actual peak. However, the model's under prediction provides a better fit to the actual data in the late evening than either of the two simple forecasts. As a metric for accuracy of forecast, we calculated the RMSE for Hours 8 to 24 for this day, these values are shown in Table 9. The ARIMAX model performs better than simply assuming the values from the previous Monday, but slightly worse that assuming the average values from all weekdays.

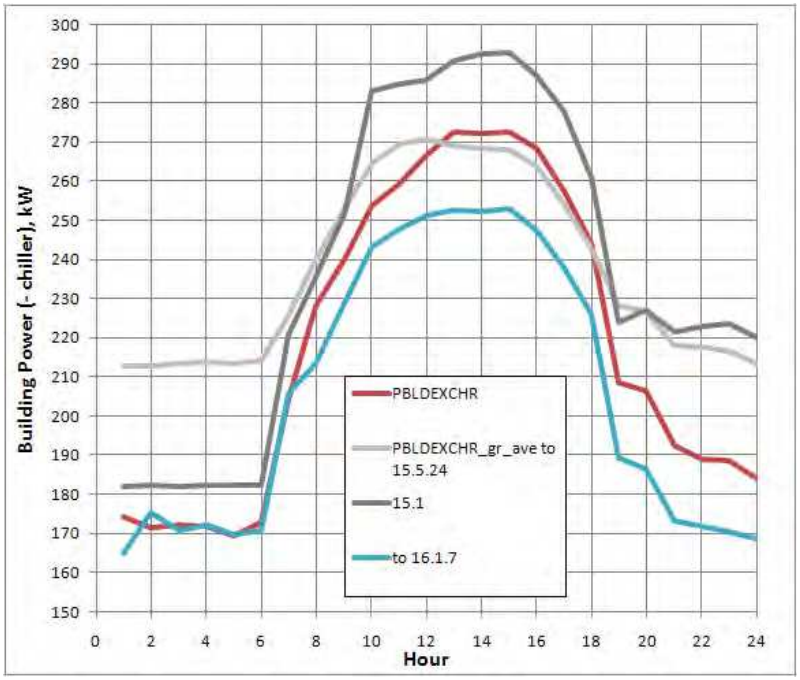

Figure 26. The building power forecast using the ARIMAX model for Week 16 Day 1 made at Hour 7 (blue line), compared to: the actual building power (red line); the average for all weekdays up to Week 15 Day 5 (light grey); and, the values from Week 15 Day 1 (the previous Monday) (dark grey).

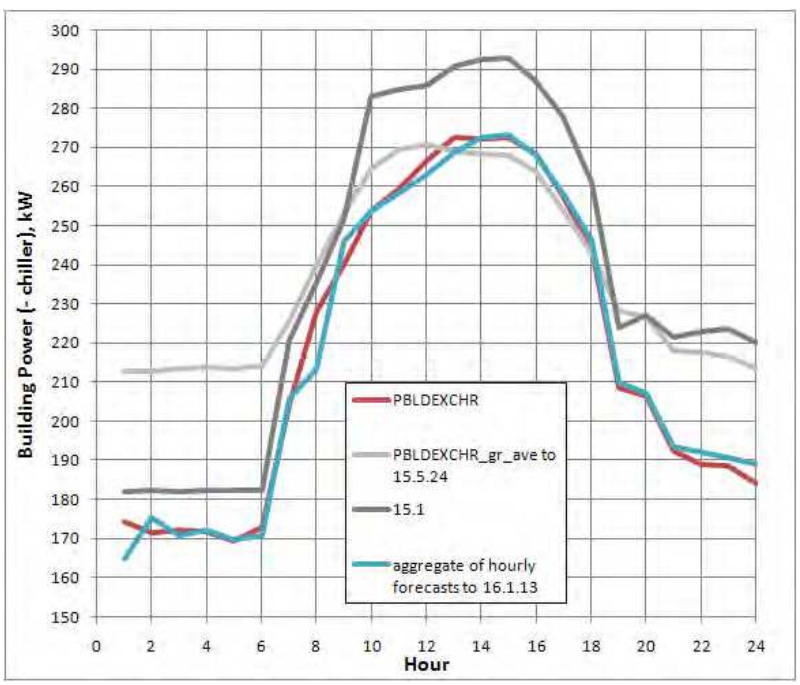

Figure 27. The building power forecast using the ARIMAX model for Week 16 Day 1 made initially at Hour 7 and then updated every hour to Hour 13 (blue line), compared to: the actual building power (red line); the average for all weekdays up to Week 15 Day 5 (light grey); and, the values from Week 15 Day 1 (the previous Monday) (dark grey).

However, it is common practice to update ARIMAX models as the new, real data become available. We recalculated the model at every hour [Kawashima et al., 1995] after Hour 7, and restricted ourselves to one-hour ahead forecasts; Figure 27 shows the aggregate one-hour ahead forecasts up to Hour 13, and the forecast out to Hour 24 with the model now updated with real data to that point. In this mode the model performs very well, with very accurate one-hour ahead forecasts through the morning to the peak in the early afternoon. Once "seeded" with real data to this peak, the forecast up to 11 hours ahead (at Hour 13 up to Hour 24) is also very good. The RMSE is now 73% lower than that from assuming the average of all previous weekdays.

Table 9. Accuracy of forecasts for Hours 8 – 24 on Week 16 Day 1, for various methods

	RMSE (Hours 8 – 24)
ARIMAX model (at 16.1.7)	17.4
ARIMAX model (aggr. to 16.1.13)	4.4
Average of all weekdays	16.3
Previous Monday	23.5

4.0 Discussion

The modelling and forecasting analysis with a measure of occupancy as an input is encouraging. Logins as an indicator of building occupancy was a significant external predictor in an ARIMAX model, and did improve the model fit. The improvement was small, but there are several reasons why this might be the case in this building, and why we could expect a larger effect in a more typical office building. First, the study building hosted a research organization, with substantial attached laboratories. This implies a high fraction of process loads for laboratory equipment, and thus a relatively lower fraction of loads related to the arrival and departure of occupants. This building was unusual in this regard, Hay & Rice [2009] provide building power and occupancy profiles for a university computer science building in a non-air conditioned week, which are remarkably similar to those in our study building. The amplitude in power draw between daytime and night-time, or weekday and weekend in this computer science building was similarly, relatively small. In a conventional office building one could expect a much greater fraction of loads tied to occupancy (e.g. office equipment, lighting), and therefore a greater sensitivity of building power to an occupancy-related independent variable. For example, Neto & Fiorelli [2008] provide a building power profile for a typical day for a university administration building in Brazil in summer, the peak power draw was similar to our study building, but the overnight power draw was only 20% of this peak. Second, logins was our most effective measure of occupancy in model development. However, we did not have logout data in this dataset. We expect that if logouts were also known, the effectiveness of this variable in the model would be improved.

Interestingly, motion sensor data did reflect the reduction in occupancy at the end of the day, but this variable did not improve our model. This was surprising, given the shape of the activity curve in Figure

21, and the fact that it reflected the later afternoon decline in occupancy, unlike login data. Perhaps this was because there was more variability in this parameter, or that it had less of a direct connection to occupancy than logins. It might also simply be an artefact of the modelling process. Also, recall that a relatively large number of days of data were missing for this variable and had to be imputed, this necessarily reduced the explanatory power of this variable. We expect that a more complete data set for this variable would have improved its contribution to the model, given the high correlation with login data in the morning, as shown in Table 5(b). None of the other occupancy measures (door events, CO_2 level) were effective in the model. Again, it would be interesting to explore whether such data streams were more effective in building power forecasting in a more conventional office building.

Separate from energy forecasting, the wireless sensor network we installed provided data that might be valuable to building operators (e.g. event detection), or those developing systems for buildings (e.g. door traverse time). This study has given us valuable experience in the issues associated with installing and commissioning sensor networks, and an understanding of the nature of the data that can be recovered. We recommend that a similar exercise be undertaken at a larger scale in a building more representative of a typical office building, with the expectation that potential uses for the data will be even more apparent.

Acknowledgements

We thank Loren Parfitt and Shawn Pedersen from Echoflex Solutions Inc for their assistance in setting up the network. Richard Laurin, Mario Laniel and David Fothergill (NRC) for provided valuable IT support, and Kevin Li (NRC) for helped us access the power meter data. Greg Nilsson (NRC) assisted with calibrating the carbon dioxide sensor. We are also grateful to Ruth Rayman of NRC's Information and Communications Technology Sector for financial and moral support of this research.

References

EnOcean GmbH. 2008. Standardization EnOcean communications profiles v1.04. http://www.enoceanalliance.org/en/enocean_standard/

Erickson, V.L.; Lin, Y.; Kamthe, A.; Brahme, R.; Surana, A.; Cerpa, A.E.; Sohn, M.D.; Narayanan, S. 2009. Energy efficient building environment control strategies using real-time occupancy measurements. Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings (BuildSys 2009, Berkeley, USA, part of SenSys 2009).

Freiman, Z. 1994. Hype vs. reality: The changing workplace. Progressive Architecture, March, 48-55,89,91.

Gershenfeld, N.; Samouhos, S.; Nordman, B. 2010. Intelligent infrastructure for energy efficiency. Science, 327 (Feb. 26th), 1086-1088.

Harle, R.K.; Hopper, A. 2008. The potential for location-aware power management. Proceedings of UbiComp, 10th International Conference on Ubiquitous Computing (Seoul, Korea), 302-311.

Hay, S.; Rice, A. 2009. The case for apportionment. Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings (BuildSys 2009, Berkeley, USA, part of SenSys 2009).

Hoffman, A.J. 1998. Peak demand control in commercial buildings with target peak adjustment based on load forecasting. Proceedings of the 1998 IEEE International Conference on Control Applications (Trieste, Italy), 1292-1296.

International Journal of Forecasting. 2008. Special Issue on Energy Forecasting, 24 (4), 561-792.

Ivanov, Y. A., Wren, C. R., Sorokin, A. and Kaur, I. 2007. Visualizing the history of living spaces. IEEE Transactions on Visualization and Computer graphics, 13(6), 1153-1160.

Jang, W.S.; Healy, W.M. 2010. Wireless sensor network performance metrics for building applications. Energy and Buildings, 42(6), 862-868.

Karatasou, S.; Santamouris, M.; Geros, V. 2006. Prediction of energy consumption in buildings with artificial intelligent techniques and chaos time series analysis. Proceedings of the International Workshop on Energy Performance and Environmental Quality of Buildings (Milos Island, Greece).

Kawashima, M.; Dorgan, C.E.; Mitchell, J.W. 1995. Hourly thermal load prediction for the next 24 hours by ARIMA, EWMA, LR, and an Artificial Neural Network. ASHRAE Transactions, 101 (1), 186-200.

Loveday, D.L.; Craggs, C. 1993. Stochastic modelling of temperatures for a full-scale occupied building zone subject to natural random influences. Applied Energy, 45, 295-312.

Lowry, G.; Bianeyin, F.U.; Shah, N. 2007. Seasonal autoregressive modeling of water and fuel consumptions in buildings. Applied Energy, 84, 542-552.

Montgomery, D.C.; Jennings, C.L.; Kulahci, M. 2008. Introduction to Time Series Analysis and Forecasting. Wiley Series in Probability and Statistics. John Wiley & Sons, Inc. (Hoboken, USA).

National Resources Canada. 2008. National energy use database. Retrieved 15th December 2009, from http://www.oee.nrcan.gc.ca/corporate/statistics/neud/dpa/handbook_tables.cfm?attr=0.

Neto, A.H.; Fiorelli, F.A.S. 2008. Comparison between detailed model simulation and artificial neural network for forecasting building energy consumption. Energy and Buildings, 40, 2169-2176.

North, S. 2007. Meeting the office occupancy challenge. Facilities Management

Newsham, G.R. Birt, B. 2010. Demand- responsive lighting: a field study. Leukos, 6 (3), 203-225.

Piette, M.A.; Watson, D.S.; Motegi, N.; Bourassa, N. 2005. Findings from the 2004 fully automated demand response tests in large facilities. Report for the PIER Demand Response Research Center. LBNL Report Number 58178. URL: http://drrc.lbl.gov/pubs/58178.pdf (last access: 2009-07-22).

Rios-Moreno, G.J.; Trejo-Perea, M.; Castaneda-Miranda, R.; Hernandez-Guzman, V.M.; Herrera-Ruiz, G. 2007. Modelling temperature in intelligent buildings by means of autoregressive models. Automation in Construction, 16, 713-722.

Rowlands, I. H. 2008. Demand response in Ontario: Exploring the issues, Prepared for the Independent Electricity System Operator (IESO), Toronto, Canada.

Shimo, A. 2008. Book your desk now - spaces are limited. Globe and Mail, January 21st. URL: http://www.theglobeandmail.com/life/article662192.ece (last access: 2010-04-14)

Tiller, D. K., Guo, X., Henze, G. P. and Waters, C. E. 2009. The application of sensor networks to lighting control. Leukos, 5(4), 313-325.

UC. 2010. Notation for ARIMA models. URL: http://www.uc.edu/sashtml/ets/chap30/sect13.htm (last accessed 2010-04-07).

Wen, Y.-J. and Agogino, A. M. (2008). Wireless networked lighting systems for optimizing energy savings and user satisfaction. Wireless Hive Networks Conference, Austin, Texas.

Wigdor, D., Ivanov, Y. and Wren, C. R. 2007. Soda pop zombies: soft drink consumption and motion. Workshop on Massive Datasets. Japan

Wren, C. R., Ivanov, Y. A., Kaur, I., Leigh, D. and Westhues, J. 2007a. SocialMotion: Measuring the hidden social life of a building. Location- and Context- Awareness. Berlin, Springer. 4718/2007.

Wren, C. R., Ivanov, Y. A., Leigh, D. and Westhues, J. 2007b. The MERL motion detector dataset. Workshop on Massive Datasets. Nagoya, Japan, 10-14

Wren, C. R., Minnen, D. C. and Rao, S. G. 2006. Similarity-based analysis for large networks of ultra-low resolution sensors. Pattern Recognition, 39, 1918-1931.

Zhen, Z.-N.; Jia, Q.-S.; Song, C.; Guan, X. 2008. An indoor localization algorithm for lighting control using RFID. Proceedings of the IEEE Energy 2030 Conference (Atlanta, USA).

Zhou, Q.; Wand, S.; Xu, X.; Xiao, F. 2008. A grey-box model of next-day building thermal load prediction for energy-efficient control. International Journal of Energy Research, 32, 1418-1431.