Software-only Occupancy Inference in a Workplace

Findings from a Field Trial

Mohit Jain, Vikas Chandan, Ashok Pon Kumar, Vijay Arya
IBM-Research

Sridhar R, Babitha Ramesh
IBM Real Estate Site Operations (RESO)

***Abstract*—HVAC and lighting loads contribute a significant fraction of total energy consumed in office buildings. These loads vary as a function of occupancy and therefore inferring occupancy is vital to optimizing energy efficiency within these buildings. This work presents evaluation and comparison results from a field trial conducted in a large office building, which involved estimating occupancy with the help of existing opportunistic context sources versus instrumented hardware sensors. Our results show that opportunistic sensing yielded an accuracy of 80% in comparison with expensive hardware sensors and may be used to continuously estimate fine-grained workplace occupancy in an inexpensive manner. Moreover the inferred occupancy information may also be used to identify anomalies in thermal management and space utilization within the building.**

Keywords—Energy Efficiency, Occupancy, Soft sensing.

I. Introduction

Global consumption of energy has been increasing rapidly. The commercial sector constitutes a significant share (18.6% in 2014 [15]) of energy consumption, and the energy consumption has grown by 69% between 1980 and 2009 [12]. Air conditioning and lighting loads (L-HVAC - Light, Heat, Ventilation and Air Conditioning) together accounts for 70% of all energy consumed in a typical office building [3]. To reduce the increasing energy cost, organizations have started taking measures to reduce their energy consumption [1], which also affects the environment positively.

Recently, researchers have proposed systems for making workplaces cost-effective yet comfortable and productive, such as mobile computing systems [4], occupancy-centric energy management systems [8, 7, 5, 6], measures for preventing resource wastage [10, 11], and feedback technologies for employees [9]. Majority of these research efforts focus on Information Technology (IT) companies, mainly due to two reasons: IT contributed ~$1 trillion to the US economy in 2011, *i.e.*, 7.1% of their GDP [13], and IT workplaces are early adopters of cutting edge technologies. In this regard, accurate occupancy information in commercial buildings can enable several useful applications such as energy management and dynamic seat allocation. Researchers have proposed several schemes [7, 5, 6], however most solutions rely on heavily instrumenting the office space with multiple sensors (including cameras, augmented PIRs, motion sensors) to monitor occupancy, and control L-HVAC. The cost of deployment and management of additional sensing infrastructure is significant.

In this paper, we propose and evaluate a method to use context-measurement software with existing Wi-Fi and computing infrastructure for occupancy detection in office space. We call such context sources *opportunistic context sources* or *soft-sensors*. Occupancy detection using such soft-sensors can be economical and easier to adopt, as it does not require installation of additional infrastructure. It is important to note that this approach does not preclude additional occupancy sensors but can incorporate data from hardware sensors in locations where the software sensors prove to be inadequate.

We present results obtained from the deployment of an opportunistic-sensing system in an IT office. The deployment involved an instrumented office space of 29,000m^2. To collect contextual data, a software application was installed on 122 participants' laptops. For collecting ground truth data (actual location of the employees), 83 participants carried an RFID (Radio Frequency Identification) tag as a wearable sensing device. 97 RFID readers were installed to receive signals from the RFID tags. In addition, to sense indoor workplace conditions to correlate occupancy with thermal and lighting comfort, each RFID reader box was equipped with these sensors - a temperature sensor, two to four motion detectors and an ambient light sensor. Thus in total 604 hardware sensors (83 RFID tags, 97 RFID readers, 97 temperature sensors, 97 light sensors, and 230 PIR motion detectors,) were deployed. Our results show that occupancy inferred using contextual data yielded an accuracy of around 80%.

The organization of the paper is as follows. In Section II, we present related work in the area of occupancy inference in office buildings. The design of the system used for occupancy inference using opportunistic context sources and for ground truth estimation is presented in Section III. Section IV provides results from the field study mentioned above and deduces appropriate inferences on resource utilization in the candidate environment. Finally, conclusions and avenues for future work are discussed in Section V.

II. Related Work

This work is an extension of our work on using pre-existing opportunistic context sensors for saving energy in commercial buildings [7, 8]. In our previous work, we examined the possibility of using soft-sensors as a means to detect occupancy and using it to reduce energy consumption of a building. In this work, we extend the system to include multiple types of soft-sensors and present evaluation results from a real field trial conducted over 29,000m^2 of office space within a large building. Most other works have used extensive instrumentation to identify the occupancy of a location [14, 16]. Also, they are generally expensive [14] or possess limited accuracy [17]. Various researchers have also tried CO_2 based occupancy detection [18], sonar based methods [19], and camera based methods [20] but with varying degrees of success.

The deployments that involve CO2 based occupancy detection [18] are slow in detecting occupancy but might perform better compared to systems that rely on passive infrared based motion sensors in a less dynamic environment. The camera based sensors [20], in addition to being costly, are privacy intrusive and not preferred in many workplaces, even though it can provide better accuracy in even dynamic environments. Similarly, Sonar-based approaches [19] involve substantial cost.

A more plausible approach is to use optimization techniques to optimize the placement of cameras to reduce the cost and still predict the occupancy of the building. It is important to note that none of the techniques mentioned above use the context based sensors as we discuss in our paper.

III. SYSTEM DESIGN

A. Collecting Contextual Data

We developed a software application that collects the following contextual data from the laptops used by the employees: (i) meeting schedules from online calendar (every 2 hours), (ii) keyboard activity (every 5 seconds), (iii) Ethernet connectivity (every 60 seconds), (iv) Wi-Fi connectivity (every 20 seconds), (v) office VPN (Virtual Private Network) status (every 60 seconds), (vi) accelerometer readings (every 5 seconds), and (vii) battery state of charge (every 60 seconds). The collected data is stored locally and uploaded to a central server securely once every two hours for data analysis. The upload mechanism was designed to be delay and disruption tolerant, as we could not assume persistent connectivity to the organization's network.

The application was developed using the Microsoft .NET framework and runs as a background service on Microsoft Windows operating systems (supporting Windows XP and Windows 7 32/64 bit). The software collects data in the background without interacting with the user. Throughout the deployment, the application was installed on 122 Lenovo ThinkPad laptops.

B. Gathering Ground Truth

We gathered ground truth (actual location of employees) to measure the accuracy of locations inferred by the system. We opted for an IEEE 802.15.4 (low-rate wireless personal area networks standard) based active RFID for automated location tracking. Although the 2.4 GHz spectrum is crowded, we chose the 802.15.4 standard since the transceivers consume very little energy, support low-speed data rates and enable communication at ranges suitable for location tracking [29]. Moreover, since the standard mandates low-power transmission modes, the signals would not interfere with the 2.4 GHz Wi-Fi network installed at the site for business needs. As illustrated in Figure 1, a network of 97 active RFID readers distributed spatially across meeting rooms and the work area (cubicles), was deployed covering a total floor area of 29,000m^2.

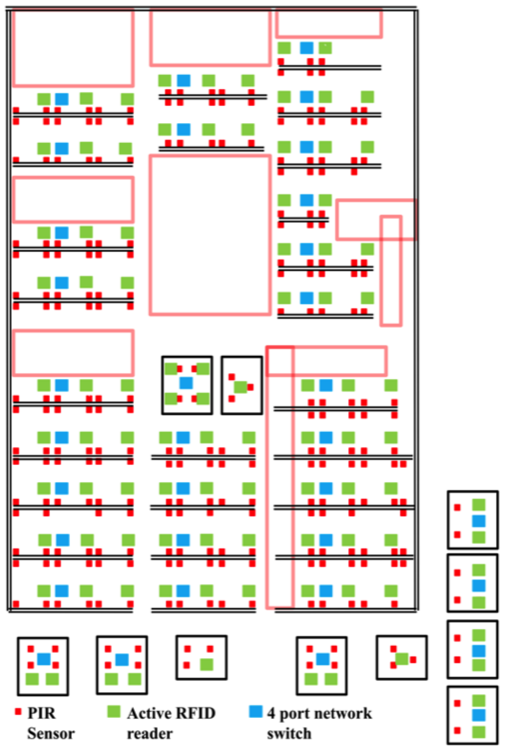

Figure 1. Layout of the office floor where the system was deployed, showing the locations of 230 PIR sensors, 97 RFID readers, 97 light and temperature sensors.

These readers have two communication interfaces: first, a RF transceiver to track the 0 dbm (approx. 5m range) beacon messages emitted once every 10 seconds by the active RFID tags worn by 83 participants; second, an embedded Ethernet module with a TCP/IP stack. These two interfaces provided a low-profile, low-cost, low-power network bridge for interconnecting the backend servers with tags and sensors (explained below). The readers receive the tag's ID and battery strength indicator as a beacon (broadcast) message from the tags. The RFID reader adds its own ID, a timestamp, and the Received Signal Strength Indicator (RSSI) to the received beacon message and forwards the augmented message to the server over Ethernet.

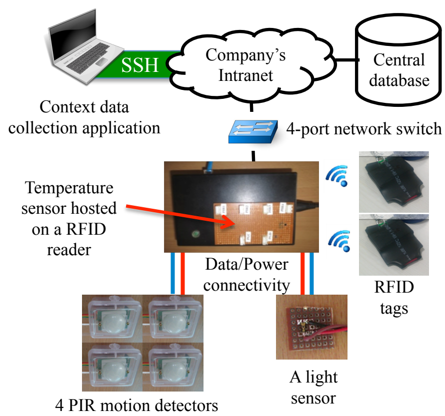

Figure 2. System architecture diagram

As multiple readers could pick a beacon message from an RFID tag, the RSSI value helps locate the nearest reader. The readers were positioned and oriented carefully so that their wireless reception was not impacted, and yet, the cables and sensors did not interfere with the employee's workspace.

C. *Sensing Indoor Conditions*

Each of the RFID readers was also equipped with these sensors (Figure 2): 2-4 motion detector, a temperature sensor, and an ambient light sensor.

Passive Infrared (PIR) Motion Detectors: These detectors are used to identify any movement in the cubicles and meeting rooms. Theses sensors were mounted under the desk and were oriented to detect leg movement. Since the purpose of the motion sensor was to detect occupancy at a desk, this helped in minimizing false negatives when there was really an employee at the seat and false positives when there was movement just outside the cubicle.

Temperature Sensor: Spot temperature was measured with an embedded temperature sensor (TC1046 from *Microchip*). TC1046 is a high precision temperature to voltage converter, which is directly interfaced to the Analog to Digital Converter (ADC) of the microcontroller. Temperature sensors were mounted on top of the readers.

Ambient Light Sensor: It was implemented using a photo-transistor (TEMT 6000 from *Vishay*), and was mounted on the cubicle frame and oriented towards the ceiling.

These additional sensors were included for identifying other potential use cases such as measuring the energy wasted in unoccupied areas. As shown in Figure 1, the RFID reader provided both network and electrical connectivity to the sensors.

IV. RESULTS

We present results from the pilot study along two dimensions: occupancy inference and resource utilization.

A. *Occupancy Inference*

We first provide results for occupancy inference using the RFID based system described in Section III, which was used for establishing ground truth occupancy. It was observed that out of the 83 volunteers who agreed to participate in the pilot

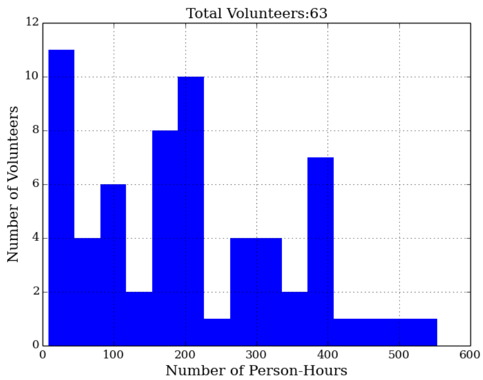

Figure 1. Distribution of number of participants with respect to person hours of occupancy data collected

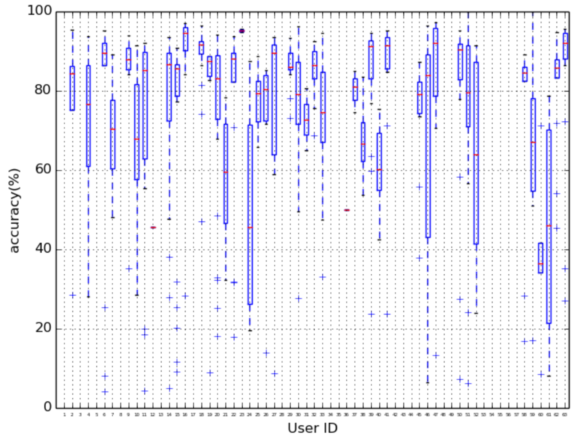

Figure 5. Box-plot showing occupancy inference accuracy for each participant

study, 63 carried a RFID tag, albeit for varying time durations. The spread of the number of volunteers over the number of person hours of data collected is depicted in Figure 3. It is observed that < 40 person hours of data was collected for around 12 volunteers, which is consistent with our experience that a few volunteers had returned the tags within a week. These challenges notwithstanding, a total of 12823 hours of ground truth data was collected over the duration of the pilot.

We now describe the methodology used to infer ground truth occupancy using the above RFID based system. Each location on the floor is represented using a pair of coordinates (X, Y). A RFID tag carried by a volunteer is detected by all RFID readers within its communication range. Each RFID reader records a triplet of reader ID, tag ID and corresponding RF signal strength at time intervals of 10 seconds. Next, for each tag and a given time instance, its location is set to be the (X, Y) coordinate of the reader which logs the highest signal strength for that tag. In this way, the instantaneous positions of all tags, and hence the participants carrying them are established. As an example, Figure 4 (top plot) shows the (X, Y) coordinates of one such volunteer established using this procedure, over the course of 9:30 hours to 20:00 hours during a chosen day. As observed, during certain durations, for example 13:00 – 13:30 hours, no data is available which suggests the employee has stepped out of the floor, in this example most likely for lunch.

With regard to opportunistic context sources, we utilized three main sources of contextual information – wifi, accelerometer and system activity. For each volunteer's laptop, we examine the strengths of the wifi signals from the various network routers/access points on the floor. We then identify the router relaying the strongest signals to that laptop, and declare the laptop's location as the (X, Y) coordinates of the identified router. Next, we examine the activities of keyboard and pointing devices (together referred to as system activity). We declare the laptop 'inactive' if no system activity was recorded for a duration longer than a threshold value of 10 minutes, and 'active' otherwise. Furthermore, we use accelerometer data to detect whether the laptop is being moved around, as indicated by a non-zero value. We then perform occupancy inference by

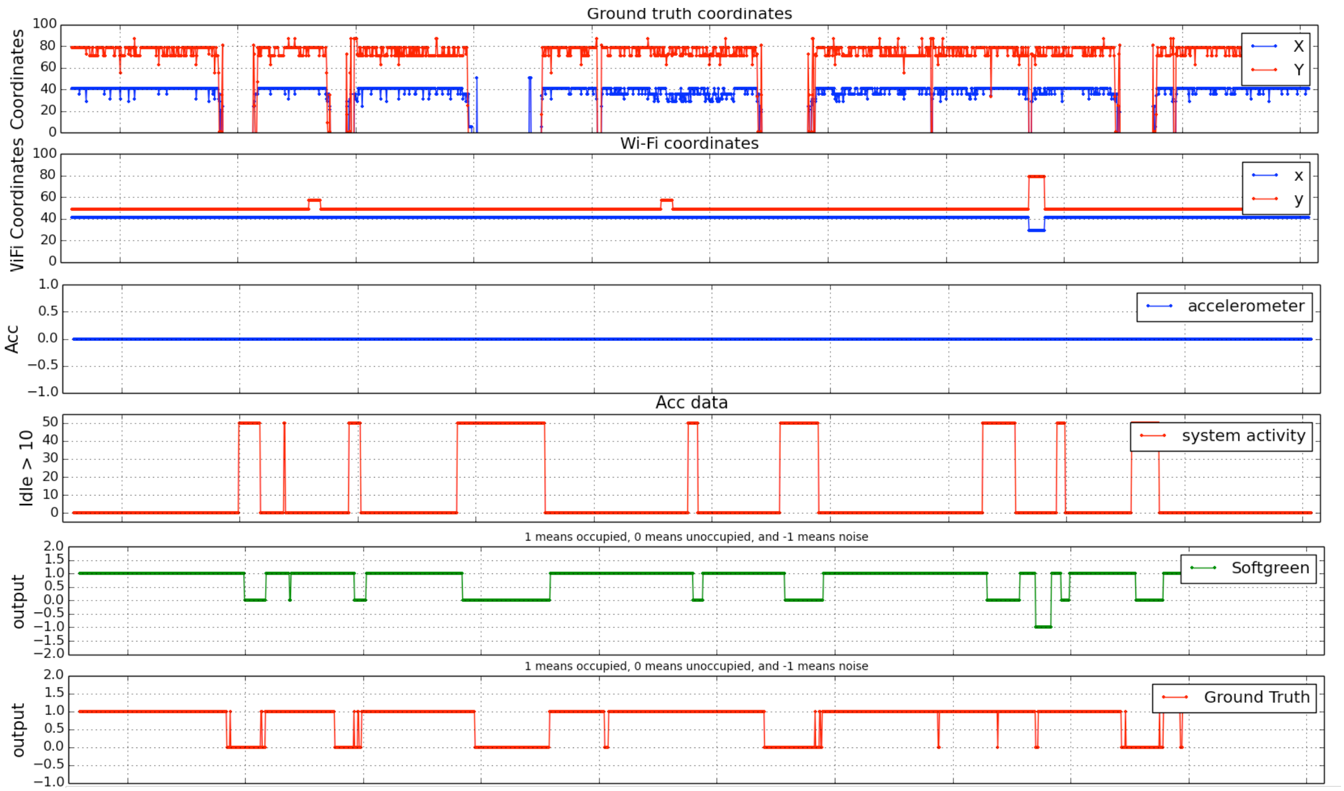

Figure 4. Illustration of occupancy inference procedure for one location using opportunistic context sources

comparing the above measurements from the various opportunistic context sources. For instance, if based on wifi, the laptop's location is estimated at coordinates (X, Y), and if the laptop is 'active', it is concluded that the location of the employee is also (X, Y). In other words the zone identified by (X, Y) is occupied. In this way, the occupancy in each zone of the floor can be estimated. Figure 4 illustrates this process for occupancy inference of one location, and a comparison with ground truth occupancy inferred using RFID tags.

Figure 5 shows, for each volunteer, a box plot of the percentage accuracy (percentage of correctly inferred instances in the day), plotted for the various days during the course of the pilot study. The edges of the box are the first and third quartile values. We conclude that for most volunteers, the mean percentage accuracy of inference (centroid of the box) is greater than 80%. This demonstrates that the proposed methodology for occupancy inference using opportunistic context sources has the ability to provide a high accuracy.

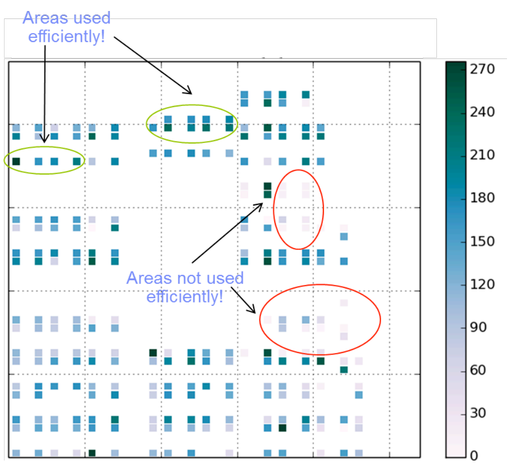

Figure 6. Total hours of seat usage observed between 9 AM and 10 PM over the course of one month.

B. Resource Utilization

Access to fine grained occupancy information can enable a holistic, occupancy centric profiling of resources in offices to provide a comfortable environment such as seats, lighting and air-conditioning. We performed a detailed analysis of resource utilization in the office floor under study, by making use of occupancy inference as presented earlier. For conciseness, we only report the most important findings in this paper. We create a seat-utilization map, which plots the total number of hours of seat usage for each cubicle location, for a specified time window. For example, the seat utilization map for one month is shown in Figure 6. Analysis of the map helps differentiate areas which are not being used efficiently. For instance, there are certain sections with less than 30 hours of occupancy over the month.

Next, we analyze the efficacy of thermal management in the office floor. For this, we divide the floor into 60 zones where each zone corresponds to an area of 64 to 100 square feet (containing 4 to 6 cubicles), which is served by an air conditioning supply vent. We then create a *thermal efficiency map*, which overlays temperature values recorded for each zone with its inferred occupancy. As an illustration, the *thermal efficiency map* at 2 PM on a particular day is shown in Figure 7, where, the radius of the circle in the top plot is proportional to the recorded temperature, whereas the square size in the bottom plot is proportional to the occupancy, with red representing an occupied zone. Analysis of this map helps identify zones, which are being managed inefficiently. For instance, we were able to find locations that were unoccupied, but the temperature was kept low.

The analysis illustrated above can be performed within the building management system to enable the building energy manager to take more informed energy conservation measures. For instance, after identification of sparsely occupied locations via the seat utilization map, the energy manager can permanently activate motion sensing based light control in these locations at all times to save lighting energy or optimally manage air conditioning vents and HVAC set points.

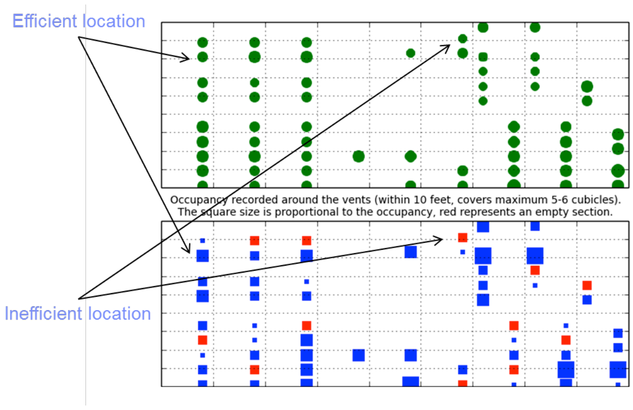

Figure 2. Efficiency map observed at 2 PM on a given day across various zones on the floor under study.

V. CONCLUSION AND FUTURE DIRECTIONS

Occupancy in office buildings continuously varies due to several internal and external factors such as weather, traffic conditions, vacations, and employee work schedules. HVAC and lighting loads, which constitute a significant fraction of the total energy consumed in office buildings, vary as a function of occupancy. Therefore continuously estimating fine-grained occupancy in the building can help improve energy efficiency, identify anomalies in thermal management, and improve thermal comfort and space utilization.

Inferring occupancy is challenging due to expensive hardware instrumentation and employee privacy issues. In this work, we showed that information from opportunistic context sources may be leveraged to estimate occupancy within offices workplaces. We presented evaluation results from a field trial conducted in a large IT office building involving several participants and compared the accuracy of occupancy inferred from opportunistic context sensors with occupancy from instrumented hardware sensors. Our results show that occupancy may be inferred with high accuracy (80%) using soft sensors. In future work, we plan to extend the field trial to cover the full office building, leverage information from other sources such as mobile phones, and consider the use of machine learning algorithms. We also plan to study the problem of occupancy forecast and its use in determining HVAC set points in an optimal manner.

REFERENCES

[1] Energy and Climate Programs. IBM Corporate Responsibility Report, 2010. http://www.webcitation.org/6FJqjhHTF.

[2] Electricity Tariff. Karnataka Electricity Regulatory Commission, 2013. http://www.webcitation.org/6FJrHa2OH.

[3] US Whole Building Design Guide Sustainable Committee. http://www.wbdg.org.

[4] R. Want, B. N. Schilit, N. I. Adams, R. Gold, K. Petersen, D. Goldberg, J. R. Ellis, and M. Weiser. An overview of the parctab ubiquitous computing experiment. *Personal Communications, IEEE*, 2(6):28–43, 1995.

[5] Y. Agarwal, B. Balaji, R. Gupta, J. Lyles, M. Wei, and T. Weng. Occupancy-driven energy management for smart building automation. In *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building*, pages 1–6. ACM, 2010.

[6] B. Balaji, J. Xu, A. Nwokafor, R. Gupta, and Y. Agarwal. Sen- tinel: Occupancy based hvac actuation using existing wifi infrastruc- ture within commercial buildings. 2013.

[7] L. V. Thanayankizil, S. K. Ghai, D. Chakraborty, and D. P. Seetharam. Softgreen: Towards energy management of green office buildings with soft sensors. In Communication Systems and Networks (COMSNETS), 2012 Fourth International Conference on, pages 1–6. IEEE, 2012.

[8] S. Ghai, L. Thanayankizil, D. Seetharam, and D. Chakraborty. Occupancy detection in commercial buildings using opportunistic context sources. In Pervasive Computing and Communications Workshops, 2012 IEEE International Conference on, 2012.

[9] F. W. Siero, A. B. Bakker, G. B. Dekker, and M. T. van den Burg. Changing organizational energy consumption behaviour through comparative feedback. *Journal of Environmental Psychology*, 16(3):235–246, 1996.

[10] D. Foster, S. Lawson, J. Wardman, M. Blythe, and C. Linehan. Watts in it for me?: design implications for implementing effective en- ergy interventions in organisations. In *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems*, pages 2357–2366. ACM, 2012.

[11] M. Jain, A. Agrawal, S. K. Ghai, K. N. Truong, and D. P. Seetharam. We are not in the loop: resource wastage and conservation attitude of employees in indian workplace. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*, pages 687–696. ACM, 2013.

[12] US Department of Energy. Buildings energy data book. http://buildingsdatabook.eren.doe.gov/ChapterIntro3.aspx.

[13] R. J. Shapiro and A. Mathur. The contributions of information and communication technologies to American growth, productivity, jobs and prosperity. pages 1–31.

[14] Real time location system, Ubisense. http://ubisense.net.

[15] US Energy Information Administration: http://www.eia.gov/totalenergy/data/monthly/#consumption

[16] Dwyer Instruments. http://www.dwyer-inst.com.au/.

[17] J. Lu, T. Sookoor, V. Srinivasan, G. Gao, B. Holben, J. Stankovic, E. Field, and K. Whitehouse. The smart thermostat: using occupancy sensors to save energy in homes. In Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems, SenSys '10, pages 211–224, New York, NY, USA, 2010. ACM

[18] State-of-the-art review of CO2 demand controlled ventilation technology and application. U.S. Dept. of Commerce, National Institute of Standards and Technology. Gaithersburg, Md. : 2001.

[19] Tarzia, R. P. Dick, P. A. Dinda, and G. Memik. Sonar-based measurement of user presence and attention. In Proceedings of the 11th international conference on Ubiquitous computing, Ubicomp '09, pages 89–92, NY, USA, 2009. ACM

[20] S. P.V. L. Erickson and A. E. Cerpa. Occupancy based demand response hvac control strategy. In Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building, BuildSys '10, pages 7–12, New York, NY, USA, 2010. ACM