Lessons from large scale campus deployment

Rishiraj Adhikary
IIT Gandhinagar
rishiraj.a@iitgn.ac.in

Soham Pachpande
IIT Gandhinagar
pachpande.soham@iitgn.ac.in

Nipun Batra
IIT Gandhinagar
nipun.batra@iitgn.ac.in

ABSTRACT

Large scale campus deployments in the past have resulted in energy conservation measures, data validation, and software architectures. Inspired by the success and learnings from such previous deployments, we present our work on deployment involving sensing various aspect of campus sustainability like water, electricity, solar produce, air quality, and parking lot occupancy. Our full deployment spanned more than 171 days. We used 469 sensors, collecting a maximum of 190 MB of data daily. We discuss the deployment challenges and the learnings obtained from them. We address the data collection challenges by providing best practices measures and provide insights from the installation of wireless radio communication modules. Our deployment can act as a reconnaissance guide for campus deployment, especially in developing countries.

CCS CONCEPTS

• **Information systems**; • **Software and its engineering**;

KEYWORDS

deployment, sensor network

ACM Reference Format:
Rishiraj Adhikary, Soham Pachpande, and Nipun Batra. 2020. Lessons from large scale campus deployment. In *The 3rd International SenSys+BuildSys Workshop on Data: Acquisition to Analysis (DATA '20), November 16–19, 2020, Virtual Event, Japan.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3419016.3431490

1 INTRODUCTION

Large-scale sensing deployments in both the residential and commercial settings have been well-studied in the past. These have been primarily motivated by various application scenarios, including, but, not limited to: i) energy conservation [3, 4, 11, 13, 15–18, 20, 21], ii) sensor data validation and management [8, 12, 19, 22], iii) software architectures [3, 6, 9, 10], iv) occupancy estimation [7, 14].
Inspired by the success and learnings from such previous deployments, we present an extensive campus deployment (Figure 1) from our 400 acres campus in a developing country and offer various learnings and insights. The work involves sensing various aspects of campus sustainability, including i) water consumption, ii) electricity consumption, iii) air quality, iv) solar produce, and iv) Occupancy via Wi-Fi monitoring. These 469 sensors collect roughly 190 MB of data daily spread across 1211 data points.

The insights and learning presented herein can significantly reduce the cost in terms of time and money for deployments, where sustainability is the primary objective.

Dataset release: We have published our air quality and the electricity data in our project repository[1]. The data are stored as CSV files. We will release the other streams of data after validation. The same repository also contains the scripts used to collect the data.

2 DEPLOYMENT

Our efforts towards this work have involved sensing various aspects of campus sustainability, which includes:

(1) water consumption measured via 6 flow sensors.
(2) electricity consumption measured via 54 smart meters.
(3) air quality measured using 6 sensors
(4) solar produce monitored via 399 solar panels.
(5) WiFi monitoring for 295 access points.

These 469 sensors collect 190 MB of data daily spread across 1211 data points. Till date, we have collected more than 2.5 GB of data. The sampling frequency of water flow meter, electric meter and air quality sensor is 1 Hz. Each solar panel has a sampling frequency of 1/1200 Hz. Our deployment started in June 2019. We have continued monitoring air quality and WiFi till date (October 2020). We now discuss our deployment sensing structure.

Water Flow Sensor: We installed flow sensor i) to verify the water released by government agencies with the water received at the campus, and ii) to check the total water usage of the campus. To store the flow rate data in a server, we designed a circuit that would take the 4-20mA signal from the sensor output and convert it to a digital signal. We sampled the output every 1 second, stored in the local storage of Raspberry Pi (RPi) and also communicated periodically to a server via an HTTP POST request.

Electricity: The meter installation was done for two reasons: i) to correlate the water flow rate with apparent power delivered to the load, and ii) to understand the wear and tear of the pump as a function of time. These four smart electric meters gave output via RS-485 which was daisy-chained and connected to an RPi which sent the data as a CSV file every day to a local server beside storing it locally. The sampling frequency of these meters is 1 Hz. 50 other meters were already installed in different buildings of the campus prior to our deployment.

Air Quality: We measure air quality using sensors placed on our campus and a moving bus. We are trying to understand how air

[1]https://github.com/sustainability-lab/DataCollection

DATA '20, November 16–19, 2020, Virtual Event, Japan

ACM ISBN 978-1-4503-8136-9/20/11…$15.00
https://doi.org/10.1145/3419016.3431490

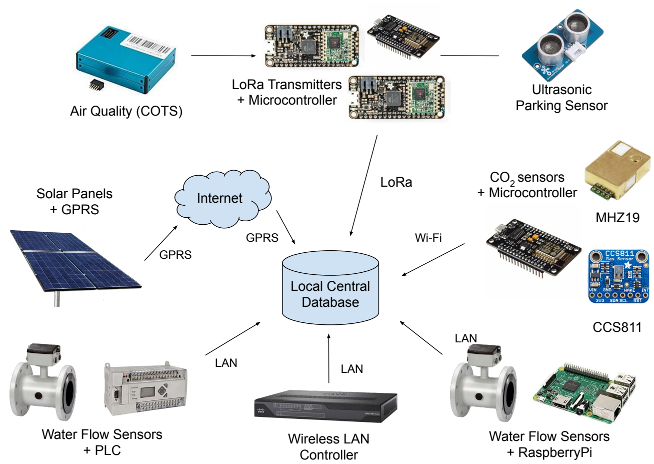

Figure 1: Our complete sensing architecture

quality affects human in different parts of the campus, viz. indoor classrooms, construction sites and dining hall. We placed a sensor inside the institute bus to quantify the change in air pollution exposure as the bus moves from the green campus to the city. We have put five sensors inside the campus (4 of them measure $PM_{2.5}$ and one measure indoor CO_2 levels), and another sensor inside the bus. Figure 2 shows the CO_2 sensor.

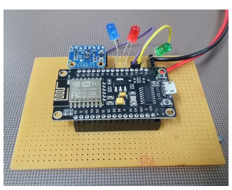

Figure 2: Circuit for monitoring indoor CO_2 and Organic Compound. It sends the data to a remote server via WiFi

Occupancy Measurement: Our goal was to determine occupancy in various zones in the campus using the existing infrastructure without any additional sensor deployment or intervention on devices of the residents. Occupancy measurement could lead to automatic grid balancing. WiFi is an integral part of any urban campus, and we show the application of WiFi logs to track occupancy to a respectable approximation based on previous work [2].

Solar Produce Solar panels are an integral part of our campus. The solar production information is processed by an inbuilt proprietary embedded system which sends the data over the GSM/GPRS to a remote server. Figure 3 shows the energy contribution of solar and grid for our campus.

3 LEARNINGS

3.1 Single board computers

We have two Raspberry Pi (RPi) at our water flow monitoring site. DC adapters power the RPi as well as a network switch to which RPi's ethernet is connected. We observed that one of the RPi's IP address was sometimes unresponsive, and at the same time, RPi failed to write data to our database, which resulted in the loss

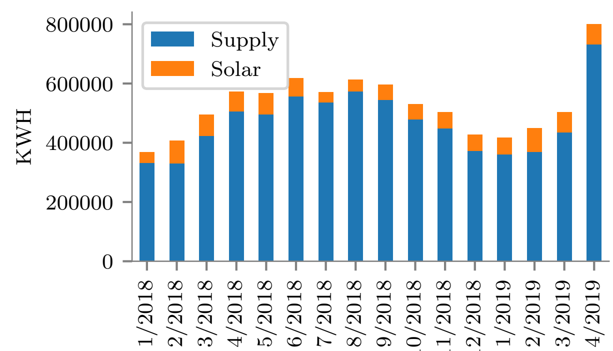

Figure 3: Figure shows the contribution of solar energy to the total power supply

of crucial data points. On-site verification revealed that the DC adapters connected to the RPi would often get disconnected due to physical interference from on-site staff who would use the same power points to recharge the battery of their smartphones. We later hardwired our adapter.

Learning 1: Ensure that a script in our remote server keeps pinging the RPi at a specific interval to see if it is online. Let the local server flag ping failures via an email to project investigators.

Different project investigators were involved in writing different data acquisition scripts. Some scripts would retrieve data from WiFi logs and some would retrieve data directly from sensors via RPi. We often used multiple such data for data analysis and found out that the timestamp information in these data was not consistent. Our WiFi-based occupancy scripts would store time information as Unix Epoch, which by default is in coordinated universal time format (UTC), other scripts running on RPi would store time information in the local timezone. It created an inconsistency where we had to change the time information into a uniform standard before analysing the data.

Learning 2: Ensure that the time format and time zone is consistent in all data collection end-points.

We use RPi extensively in our project, be it water data collection, electricity monitoring, or interfacing with communication (LORA) modules. Hence we made sure that certain necessary practices were followed in each of the RPi so that they become capable of research applications. We mention such **best practices** here:

1. Restart a script automatically after RPi reboots: Power outage, or grid failures can shut down the RPi. Therefore, the script running on RPi needs to restart after RPi reboots. Keeping the script path in the "rc.local" file of a Linux machine (RPi) ensures that the script starts itself after RPi reboots. It was a very significant practice concerning our project as we observed a power shutdown due to periodic grid maintenance every three weeks. It was also helpful in the smooth function of RPi after it has been subjected to physical interference like disconnection of the power adapter of the RPi from the powerpoint.

2. Update the time on RPi with a local NTP server: RPi lacks an onboard Real Time Clock and as such, updating the clock in RPi becomes difficult. We worked around this problem by adding an NTP server's IP address (inside our campus network) to the RPi.

3. Set a static IP to the RPi: We ensured that our RPi gets the same IP address from the Dynamic Host Control Protocol (DHCP) pool every time it reboots. This workaround did not work if the IP is already assigned to some other device.

4. Scripts to detect fault: It is not possible to monitor all sensor installations for failure at all times by personnel. Hence, an alternate mechanism to detect and report failures (via email/SMS) by scripts monitoring the database for updates by all the sensors will result in quicker fault detection.

5. Run a script after specific interval: We used RPi "cron" job scheduler to execute certain scripts after a specific interval of time. It is useful when we need a sensor data only at specific intervals, say five minutes.

3.2 External deployments

During our flow sensor installation in the water supply side infrastructure, there was a dire need for wired connectivity. We managed to get wired communication infrastructure after two weeks of getting approval from the campus development dean. Wired connectivity was also made available to us at the student dormitory terrace. As our deployment started to increase, getting wired connectivity in every location was challenging due to the following.

Time Constraints: A request for connectivity takes 1-2 weeks for approval. The campus development dean needs to approve the installation from the project or campus budget after which the technical officers have to hire contractors for the work of digging, conduit laying, and procurement of network switch and local area network cables. The entire process is resource and time expensive.

Cost: The process of getting connectivity adds to the project cost. In our case, we spent from 172 to 215 USD for each of the installations. It is a high cost when converted to our local currency.

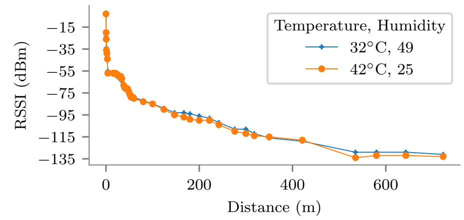

Figure 4: The range of Line of Sight communication via LORA is independent of temperature and humidity.

The first choice of wireless communication was GSM/GPRS. But during our air pollution sensor deployment, we realised that the GSM/GPRS connectivity on our campus was poor. Therefore, we decided to use LORA [5], a radio frequency-based communication system on our sensor deployment at the parking lot and also in all future sensor deployment. LORA is cost-effective and has served the IoT community very well [1, 5]. We conducted some indoor as well as outdoor experiments with LORA. An indoor investigation revealed that the Received Signal Strength (RSSI) was between -70 to -85 dBm for concrete whereas it was -20 dBm for glass. A hypothesis that RSSI does not change with the rise in temperature and humidity was proved correct in an outdoor experiment, the result of which is in Figure 4. For this outdoor experiment, the receiver LORA was taken to some extreme corner of our campus (Figure 5b). We observed a successful communication with the receiver, and the transmitter placed more than 800 meters away (Figure 5a) at the academic block of our campus without the line of sight. The primary objective for such experiment was to check the feasibility of LORA in our campus environment.

Learning 3: LORA can be used for a low bandwidth (suitable for many Internet of Things (IoT) applications) wireless communication ranging beyond 1-kilometer without a line of sight.

3.3 Proxy sensing

Occupancy Measurement Using WiFi: In Section 2 we explained the motivation of measuring occupancy via WiFi logs. Address duplication was an important challenge in estimating occupancy from WiFi logs. It was observed that a high proportion of campus members used multiple devices, which includes smartphones, laptops, and tablets. Thus, counting each device as a person would skew our estimations. We use the information about user names wherever available to remove duplication. Due to privacy concerns, we map every user and MAC address of every device of the user to a unique identification number. Whenever a client entry is detected in the logs, we first check whether a mapping between the user name of the client and mac address exists. If it exists, we use the existing unique identification number to add the entry in the database. If not, we allocate a new unique identification number and proceed with

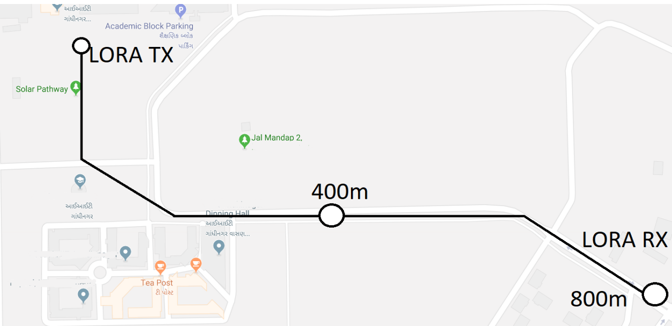

(a) LORA receiver path. Image source: Google Map

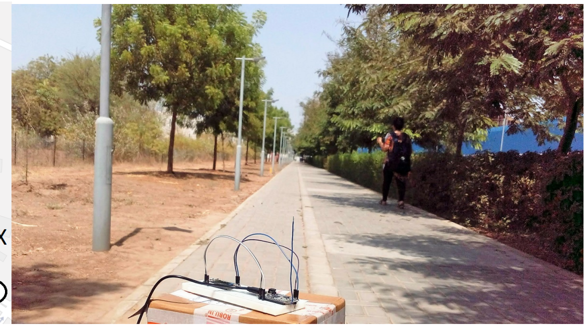

(b) A LORA receiver module

Figure 5: We obtained more than 1 KM range without a line of sight for LORA placed just above ground level.

this new mapping. We store these mappings in a secure database.

The WiFi network of our campus consists of one service set identifier (SSID) and one open network. Unique user identification and deduplication of WiFi counts is made using user credentials of SSID authenticated Active Directory of user name and password. The campus network relies on a single sign-on (SSO) method, which does not give details about users. The occupancy detection algorithm is shown in Figure 6a. The Wireless LAN Controller (WLC) provides details of MAC-address and associated access point name for every client visible to the network. The WLC also provides the username if the client is connected to the SSO network. The access point names are kept in the format: Area-Building-Floor-Room. This was done to identify the location of the access point within the campus. To verify the data given by WiFi log, we conducted an experiment across two different areas in our campus. Area 1 is an undergraduate workspace with frequent fluctuation of students count. Area 2 represents a lecture hall with a fixed number of students (21 and 34 for two days in an hour, respectively) in an interval. The ground truth compared with WiFi count is shown in Figure 6b.

Learning 4: Since several devices per user vary a lot, taking several connected WiFi devices inflates the occupancy measure. Addressing these duplication users is necessary.

Electric Meter as Proxy For Flow Meter: Flow sensors are difficult to install as we have to cut through the pipe to place them. Therefore, it made more sense to use a smart electric meter as a proxy for the flow meter. More apparent power delivered to the pump would mean more the flow across the pipe connected to the pump. We used the data from the electric meter and the flow sensor to correlate power and flow rate. In our academic campus, two freshwater pumps run alternatively when they pump out water to the dormitory block, but the same pumps run together when there is a need to transfer water from the freshwater reservoir to recycle water reservoir. We attempted to correlate the apparent power delivered to the load (water pump) and the instantaneous flow rate. Figure 7 shows the plot of meter data and flows data. The amplitude of the flow is highest when both the water pumps are running. Also, the presence of non-zero flow occurs when either or all of the water pumps are running. Hence, the use of electric meters as a proxy for flow meters is justified.

Learning 5: Smart electric meters concerning each electric pump can be used to measure the flow rate instead of flow sensors.

3.4 Air Quality

In this section, we discuss the result of air quality monitoring from a sensor that was placed inside a kitchen mess. The monitor placed at the kitchen reported $PM_{2.5}$ value between 350 $\mu g/m^3$ to 400 $\mu g/m^3$ during cooking activity as shown in Figure 8. Even though the cooking occurs at the same time every day, certain days are more polluted compared to others. This can be explained by the difference in the cooked items (fried items would be more polluting), or, incorrect or no usage of the exhaust fan. An important caveat with monitoring cooking exposure is that our sensors can pick up humidity (or steam) as particulate matter. We ruled out this possibility in our experiments by monitoring the co-located humidity sensor. The 24-hour mean $PM_{2.5}$ value stands at 122 $\mu g/m^3$, nearly five times the WHO mandated air quality standard. The monitoring at the kitchen mess started on August 2019 and continues till date. We have shared the validated air quality data for the first three months in our project repository.

REFERENCES

[1] Aloÿs Augustin, Jiazi Yi, Thomas Clausen, and William Townsley. 2016. A study of LoRa: Long range & low power networks for the internet of things. *Sensors* 16, 9 (2016), 1466.

[2] Bharathan Balaji, Jian Xu, Anthony Nwokafor, Rajesh Gupta, and Yuvraj Agarwal. 2013. Sentinel: occupancy based HVAC actuation using existing WiFi infrastructure within commercial buildings. In *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems*. ACM, 17.

[3] Nipun Batra, Manoj Gulati, Amarjeet Singh, and Mani B Srivastava. 2013. It's Different: Insights into home energy consumption in India. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*. ACM.

[4] Alex Beltran, Varick L Erickson, and Alberto E Cerpa. 2013. Thermosense: Occupancy thermal based sensing for hvac control. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*. ACM, 1–8.

[5] Martin Bor, John Edward Vidler, and Utz Roedig. 2016. LoRa for the Internet of Things. (2016).

[6] Bradford Campbell and Prabal Dutta. 2014. An energy-harvesting sensor architecture and toolkit for building monitoring and event detection. In *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings*. ACM, 100–109.

[7] Dong Chen, Sean Barker, Adarsh Subbaswamy, David Irwin, and Prashant Shenoy. 2013. Non-intrusive occupancy monitoring using smart meters. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*.

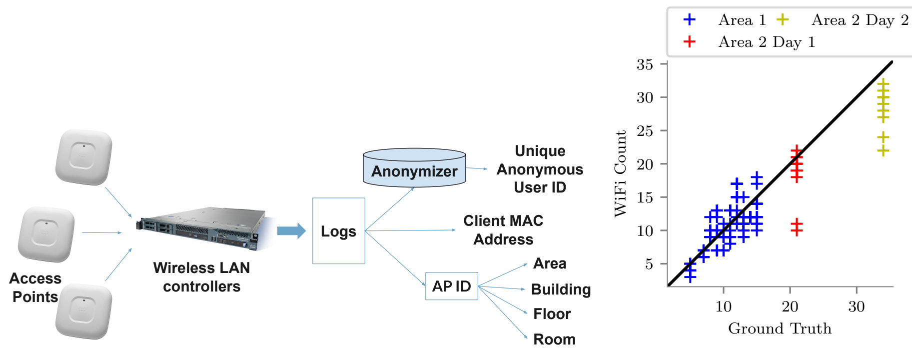

(a) Block diagram showing system of occupancy measurement

(b) WiFi based occupancy compared with ground truth

Figure 6: Occupancy measured via WiFi is highly correlated with true occupancy

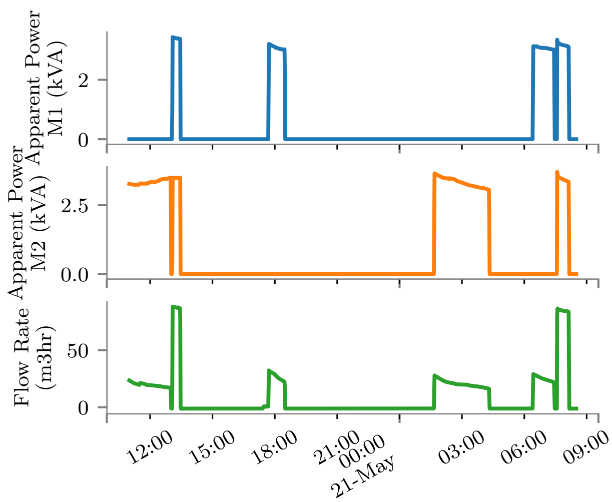

Figure 7: Water flow can be estimated from apparent power delivered to the load (Electric Motor)

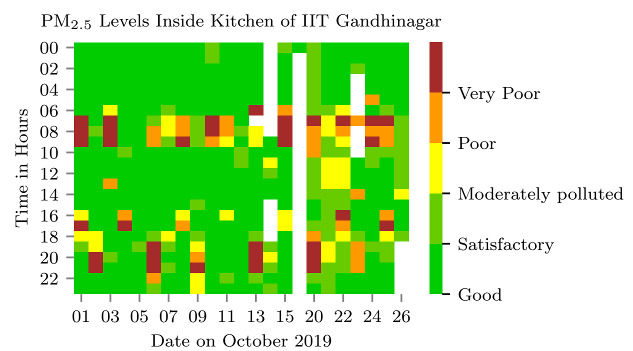

Figure 8: PM$_{2.5}$ levels inside the kitchen mess. White regions indicate missing data. The colour scale is as per Indian air quality standards.

[8] Christos Efstratiou, Ilias Leontiadis, Cecilia Mascolo, and Jon Crowcroft. 2010. A shared sensor network infrastructure. *SenSys' 10* (2010), 367–368.

[9] Frank Englert, Irina Diaconita, Andreas Reinhardt, Alaa Alhamoud, Richard Meister, Lucas Backert, and Ralf Steinmetz. 2013. Reduce the number of sensors: Sensing acoustic emissions to estimate appliance energy usage. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*.

[10] John Feminella, Devika Pisharoty, and Kamin Whitehouse. 2014. Piloteur: a lightweight platform for pilot studies of smart homes. In *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings*. ACM.

[11] Jereme Haack, Srinivas Katipamula, Bora Akyol, Brandon Carpenter, Kyle Monson, Nathan Tenney, Robert Lutes, Craig Allwardt, Poorva Sharma, and Timothy Kang. 2014. The transactional network platform. In *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings*. ACM, 208–209.

[12] Dezhi Hong, Jorge Ortiz, Kamin Whitehouse, and David Culler. 2013. Towards automatic spatial verification of sensor placement in buildings. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*.

[13] Zhichuan Huang, Ting Zhu, Yu Gu, David Irwin, Aditya Mishra, and Prashant Shenoy. 2014. Minimizing electricity costs by sharing energy in sustainable microgrids. In *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings*. ACM, 120–129.

[14] Wilhelm Kleiminger, Christian Beckel, Thorsten Staake, and Silvia Santini. 2013. Occupancy detection from electricity consumption data. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*. ACM, 1–8.

[15] Neil Klingensmith, Dale Willis, and Suman Banerjee. 2013. A distributed energy monitoring and analytics platform and its use cases. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*. ACM, 1–8.

[16] Balakrishnan Narayanaswamy, Bharathan Balaji, Rajesh Gupta, and Yuvraj Agarwal. 2014. Data driven investigation of faults in HVAC systems with model, cluster and compare (MCC). In *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings*. ACM, 50–59.

[17] Juhi Ranjan, Erin Griffiths, and Kamin Whitehouse. 2014. Discerning electrical and water usage by individuals in homes. In *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings*. ACM, 20–29.

[18] Alex Rogers, Siddhartha Ghosh, Reuben Wilcock, and Nicholas R Jennings. 2013. A scalable low-cost solution to provide personalised home heating advice to households. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*. ACM, 1–8.

[19] Anthony Rowe, Mario E Berges, Gaurav Bhatia, Ethan Goldman, Ragunathan Rajkumar, James H Garrett, José MF Moura, and Lucio Soibelman. 2011. Sensor Andrew: Large-scale campus-wide sensing and actuation. *IBM Journal of Research and Development* 55, 1.2 (2011), 6–1.

[20] Anika Schumann, Joern Ploennigs, and Bernard Gorman. 2014. Towards automating the deployment of energy saving approaches in buildings. In *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings*.

[21] Shailja Thakur, Manaswi Saha, Amarjeet Singh, and Yuvraj Agarwal. 2014. Wattshare: Detailed energy apportionment in shared living spaces within commercial buildings. In *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings*. ACM, 30–39.

[22] Thomas Weng, Anthony Nwokafor, and Yuvraj Agarwal. 2013. Buildingdepot 2.0: An integrated management system for building analysis and control. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*. ACM, 1–8.