Vol.2017-UBI-55 No.1
Vol.2017-ASD-9 No.1
2017/8/24

Analysis of WiFi Router's Electric Power Consumption towards Determining Presence in an Office Environment

TIN PETROVIC[†1] KAZUYA ECHIGO [†1]
HIROYUKI MORIKAWA [†1]

Abstract: Electrical waste from using lights, AC and other plugged in devices while no one is in a room is noted as the biggest potential way to reduce electric consumption without affecting the end user. To achieve this, simple and cheap way to detect presence in a room is desirable. To this end we look into new ways to count the number of people in a space by looking at the power consumption of a WiFi router. We analyse if a router's power consumption can be correlated to the number of people in a room by setting two assumptions. First that an increased number of devices will results in a higher level of electrical consumption in a WiFi router. Second, that the aggregated network traffic from users is a stable and a good feature to use in prediction. We care out experiments in real world scenarios and showcase that it is possible to determine whether or not a room is occupied by looking at a WiFi's power consumption.

Keywords: Smart Plugs, Smart Grid, Presence detection

1. Introduction

Electric energy consumption of offices and homes take up between 25% to 30% of a country's overall electric consumption [1]-[3]. While the estimates might vary between countries, in most developed countries the residential and commercial sector take up around 70% of all electric consumption, making office and home electric control systems an important element to consider in the future of the smart grid. As shown by [4], building quality and energy efficiency is less important than occupancy. In this paper we consider a novel form of occupancy-based electric control system which used the electric power consumption of a WiFi router to determine the number of people in an office environment.

FOver the last 10 years, several approaches for context-based or activity-based energy management have been proposed focusing on lighting, heating ventilation and air conditioning (HVAC) or plug in loads. [5] uses a camera based occupancy monitoring for HVAC control. They achieve 80% occupancy estimation accuracy and calculate 14% energy savings, which are not significantly affected by the 20% error rate. Most systems [3], [6]-[11] use a passive infrared sensor (PIR) to monitor presence in real time. The authors in [7] use motion sensors and door sensors to infer occupancy with 88% accuracy, pointing out that for only 25 dollar’s worth of sensors it is possible to reduce the electrical energy consumption by 28% in HVAC. Other have looked at optimising the time delay of PIR sensors for electric lighting to reduce the time delay before light automatically turn themselves off and experimentally shown 25% savings. Finally, in [10] the authors look at savings from plug-in loads by implementing wireless sensor networks (similar to network of smart plugs), able to sense occupancy with PIR sensors and turn off devices when they are not in use. They show between 7% and 14.5% savings in real systems.

All systems show that an occupancy-based electric control system outperform other systems, making occupancy detection an interesting problem. The system should be cheap enough to allow widespread accessibility, simple enough for anyone to be able to install it and insure an acceptable level of privacy to improve adoption. While quite accurate, cameras can be relatively expensive, create installation challenges for wired systems and significant overhead for wireless systems. Also they are often perceived as an intrusion to personal privacy. PIR sensors are much cheaper and can be used in closed networks ensuring high levels of privacy. In [12] the authors use the already installed PIR sensor network to detect presence, while in most other situations the PIR sensors need to be installed by hand. From our experience, installing a PIR wireless sensor network can be somewhat problematic since line of sight and sensitivity were often quite difficult to adjust. Also, while we could easily deploy and manage up to 50 wireless sensor nodes for our experiment, deploying a wireless sensor network for presence detection and a separate network to control electrical loads was too costly and tedious for most users who are not so tech-savvy. This problem can be indirectly addressed by deploying PIR sensors directly in smart plugs. Unfortunately, this can create line-of-sight problems in many situations.

In this work we try to remove the line-of-sight and sensitivity problems of PIR sensors by instead focusing on the electric power consumption of a WiFi router to determine the number of people in an office environment. By doing so we try to create a power management system based only on smart plugs, but which still possess the ability to estimate the number of people in a room by considering parameters other than the overall electric power consumption. In this way we are able to more quickly react and reduce the number of false positive that occur due to wasteful power usage. We show that it is possible to accurately estimate the number of people in a room by monitoring the unique

†1 The University of Tokyo

characteristics of a WiFi routers electric power consumption and the fact that user behavior, on average, does not change significantly from day to day. In this way we hope to create a smartplug-based monitoring and control system which is also able to leverage occupancy data, without needing to rely on outside sensors.

In this work our contributions are as follows:

- We analyse the WiFi router's electric power consumption characteristics to show that the electric power consumption is proportional to the amount of traffic, even for multiple access points, and that the number of high electric power consumption control messages is proportional to the number of devices, allowing for easy estimation of the number of connected devices
- We experimentally show that the user WiFi behaviour does not significantly change in periods on the scale of several minutes, which allows for good estimation of the number of users, when combined with knowledge of the electric power consumption when the access point is idle

The limitations and assumptions of this work are twofold. First, our experiment was carried out only in an office environment with up to 16 people which is often occupied by at least several people for long periods of time. Multiple access points electric power consumption is measured and there are multiple networks each individual user can connect to. While we assume that the same approach would work in a residential environment, since the problem would be much simpler, user behavior will be different and edge cases might significantly affect the accuracy. Secondly, to determine ground truth we used a camera system to monitor the exact number of people occupying the rooms. Since the sampling rate of the seven cameras was 10 minutes, we cannot claim real-time detection. We are still confident in the overall system accuracy since the number of people within the room changes gradually and slowly within a period of a single day.
This paper is structured as follows. In Section II we analyse a WiFi router's electric power consumption characteristics and define our assumptions. Section III looks at how the experiment was carried out and which features were selected for classification. Section IV experimentally shows the classification results and how they relate to our initial assumptions. In Section V we write the conclusion and discuss future work.

2. WiFi Electric Power Consumption Characteristics

In this section we briefly describe and test our hypothesis. To show the correlation between electric power consumption and data traffic for both single and multiple access points, we analyse which packets contribute to the increase in electric power consumption and briefly describe how the measurements were conducted.

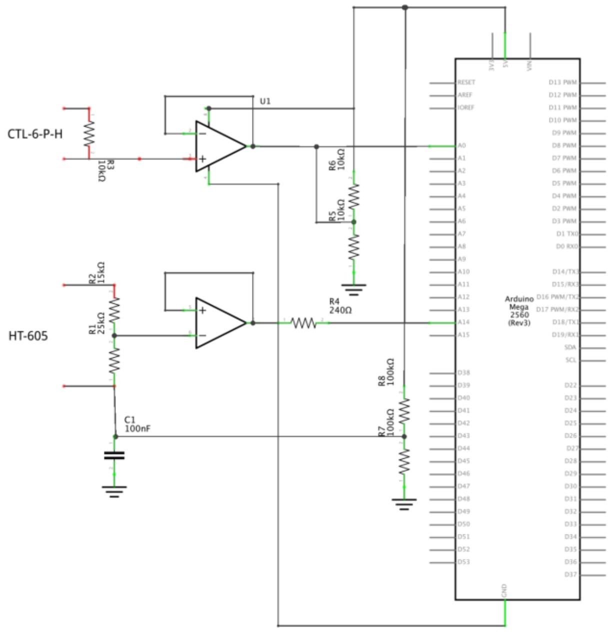

Fig. 1. Arduino-based wattmeter schematic

2.1 Hypothesis

To be able to monitor occupancy by using a WiFi router, we create and test two hypothesis. First is that the electric power consumption of a WiFi router will be proportional to the amount of traffic on the network. Specifically, that it is possible to discern the number of devices from the linear growth of the electric power consumption of a WiFi router.
Second is that the amount of data traffic will be proportional to the number of people in the room. While it is reasonable to expect that more people will produce more traffic in general, we need to confirm how consistent is user behaviour and how big is the impact of edge cases. This is covered later in Section III.

2.2 Measurement Setup

To measure the electric power consumption of a WiFi router we have constructed several simple Arduino-based wattmeters as shown in Fig. 1 that sample the voltage, current and wattage and combine them into one reading every 0.5 seconds. We have compared the results with 2 different commercial wattmeters and tested each part with an oscilloscope. There were no significant differences between our system and commercial wattmeter.
All the measurements were carried out on two domestic routers Buffalo Air Station WZR-HR-G301NH, Buffalo Air Station WZR-AGL300NH and an Apple Time Capsule A1470 router.

2.3 Impact of Data Traffic on Electric Power Consumption

In [13] the authors have extensively looked into the electric power consumption of WiFi routers using the 802.11g protocol and showed that it is possible to distinguish between data rates higher than 1Mbps and that the electric power consumption growth is linear. We were able to confirm this in Fig. 2 and expand on. As shown in Fig. 3, electric power consumption is proportional to data usage regardless of the router used. While different routers might have slightly different base power consumption, e.g. the Apple Time capsule's internal hard disk uses additional electricity, the relative jump between a passive

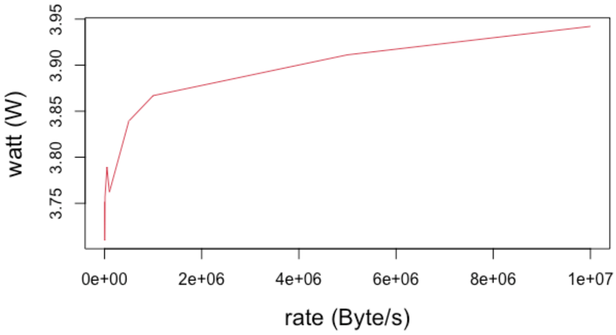

Fig. 2. Electric power consumption relative to the data rate

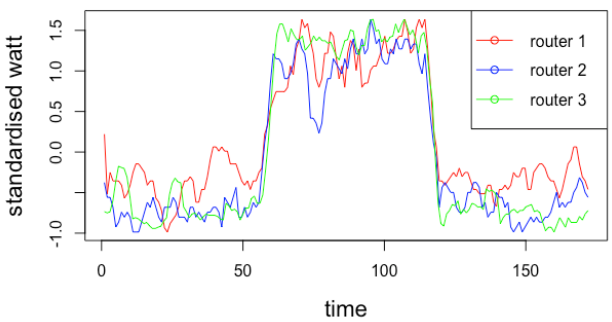

Fig. 3. Similar levels of electric power consumption regardless of the router

and active will be distinguishable. Also, in an office environment it likely to have multiple access points on the same network. For this reason, we have also looked into how the packet rate and electric power consumption are distributed between 2 routers. As we can see in Fig. 4, if the data rate is kept the same, the electric power consumption will still be evenly distributed over 2 routers depending on which access point is being used.

2.4 Packet Rates for Different Numbers of Devices

Next, we looked at the correlation between the packets type, size and electric power consumption. Since the electric power consumption changes proportionally to the amount of traffic, we still need some additional information to be able to distinguish between the number of devices. In Fig. 5, we can see a breakdown of packets depending on their size. While most of the time, the packets size does not seem to have any correlation to the number of devices, only packets which are 200 to 600 bits long grow proportionally to the number of devices. In Fig. 6, we can see that the number of packets length 320 to 639 increases linearly with the number of devices. Fig. 7 also shows the data rate increasing as well.

We further analysed the traffic with the Wireshark network sniffing software and isolated packets of this size refer to the beacon signal, probe response and quality of service (QoS) packets which are used to broadcast the network and signal to the access point that a device wants to communicate with it. This would indicate that the minimum electric power consumption

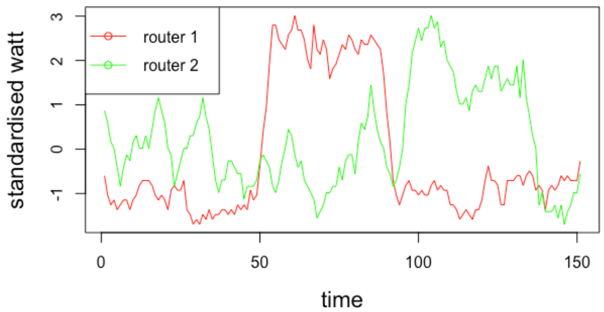

Fig. 4. Electric power consumption of 2 access points on the same network with mobile 10 Mbps link

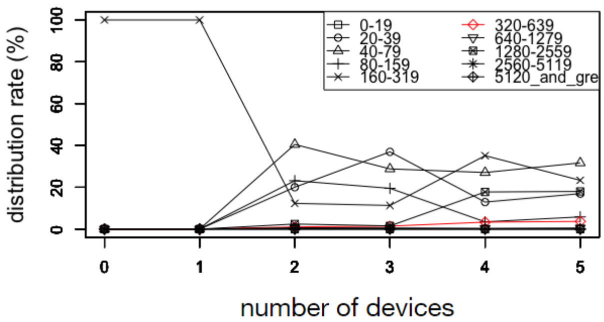

Fig. 5. Packet breakdown by packet size

should increase proportionally to the number of devices connected to the access point.

3. Presence Detection

In this section we describe the environment in which the experiment was conducted, how the data was gathered and classification results.

3.1 Experiment Environment

To obtain the data needed to test our second hypothesis that the data traffic is proportional to the number of people, we setup the following experiment. As shown in Fig. 8, we connected our wattmeters to three WiFi routers located in the room, labeled R1, R2 and R3. The access points R1 and R2 share the same network while R3 is a separate network. Everyone was instructed to continue using whichever network they preferred. We also connected a commercial smartplug to each desk to measure watt consumption as a secondary means of monitoring presence. The sampling rate of the desk smartplugs was 1 sample every 1 minute due to hardware limitations. A higher query frequency would often fill up the buffer of the smartplug which would end up requiring a manual hard reset. Lastly, cameras were placed around the rooms to determine ground truth. Each camera took 1 picture every 10 minutes to determine the number of people in each room. We also tried installing PIR sensors similar to other works [7], ultrasound crossing sensors and asking people to register upon entering. In the end we found that in our case, using the cameras gave the best and most consistent results for determining ground truth.

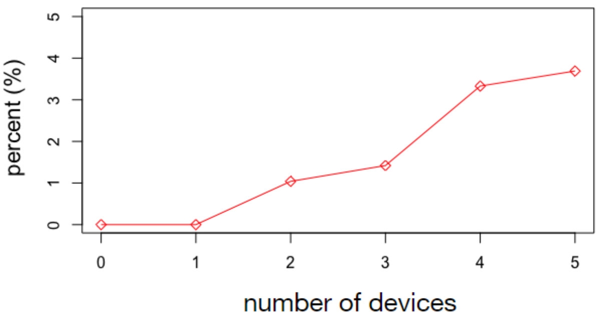

Fig. 6. Ratio of packets length 320-639 bits to overall packet number

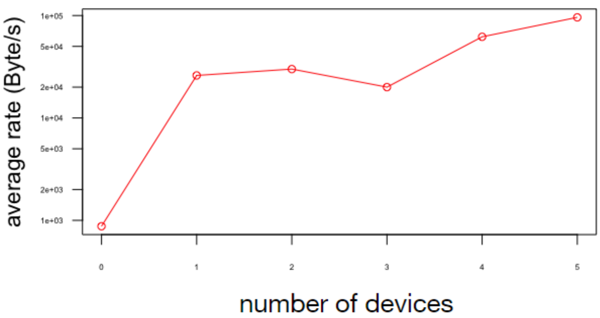

Fig. 7. Data rate of control packets depending on the number of devices

3.2 Feature Extraction

From each of the three wattmeters connected to the access points we extracted the wattage measurements. Next, we used a frame of 120 samples (one minute worth of samples) to calculate the minimum, average and maximum wattage value of a frame. This process was carried out for each consecutive 120 samples creating a form of low pass filter over the data set. The increments were done one sample at a time instead of of using a block of 120 sample at a time. A one minute frame was chosen as a compromise between classification accuracy and the minimum time required before the system could make a decision as shown in Fig. 9. This means that every data instance will contain the minimum, maximum and average value from the last 120 samples for each of the three routers. Those values are combined with the ground truth value taken from the cameras into a 10 dimensional vector to form the data table for later classification. Since the sampling speed of the cameras is much slower we extrapolate the number of people and assign an appropriate value to each data instance. Fortunately, the number of people coming and going out of the room is quite gradual without sudden changes, but it still limits our system to a an average 5 minute wait time before we can guarantee the number of people in a room. Since we show very high accuracy in the next section, we believe it would not significantly impact the accuracy if higher resolution ground truth data was available.

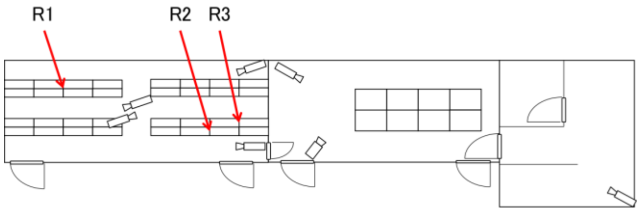

Fig. 8. Office layout

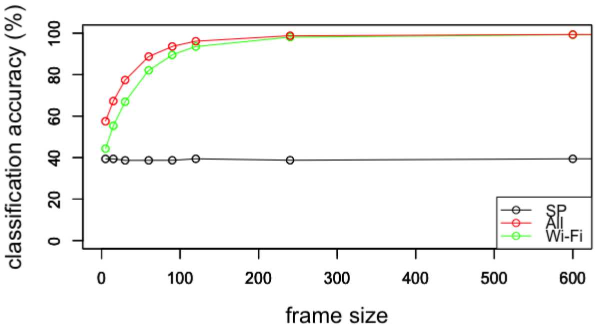

Fig. 9. Classification accuracy depending on the frame size

4. Results

For classification we used a random forest classifier [14] implemented in R programing language. The choice was made based on our previous result when dealing with data from electronic devices [15], good results from papers in similar fields [16] and recommendations from other papers [17]. 10 percent of data was used for training, 90 for testing. Also, without optimisation, the random forest was significantly better than the state vector machine and extreme learning machine implementations that we tested.

Also, we must point out that the classification accuracy for presence and accuracy based only on smart plugs, should be used as a naive reference. Whether or not someone was present at the time was determined by calculating whether a threshold value, calculated from the mean electric power consumption, was crossed. Since the sampling rate of the desk smart plugs is 1 minute and not 0.5 seconds, we use the table smartplug system only as a reference and not as a comparison.

4.1 Presence Detection

As shown in Table I and Fig. 10, we can quickly and easily classify whether there is anyone present in the room. Moreover, since we are comparing an empty room with an occupied room, the electric power consumption average plays the biggest role, as shown in Fig. 11. Despite the fact there is constant traffic from numerous devices, the traffic is constant and differentiable from active users.

TABLE I
PRESENCE CLASSIFICATION ACCURACY

Data used to classify	Frame length	Detection accuracy
Desk Smart Plug only	no frame	84.76%
WiFi router only	no frame	86.83%
Desk Smart Plug and WiFi router	no frame	87.46%
Router only	1 minute	99.34%
Desk Smart Plug and WiFi router	1 minute	99.34%

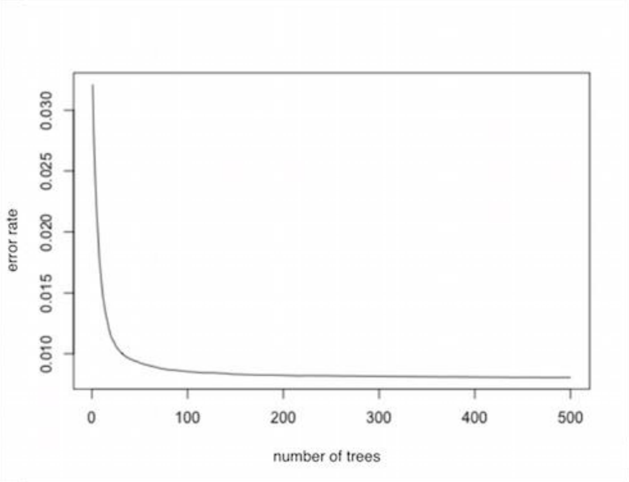

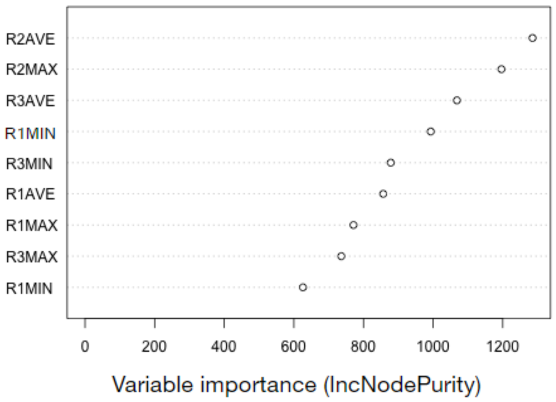

4.2 Occupancy Monitoring

In Table II we show that it is also possible to quickly classify the precise number of people in a room by just using the WiFi data, with the accuracy dropping to 93% while the conversion rate stayed the same as in the presence detection case. Also, the electric power consumption minimum data point is shown to be the most significant, as can be seen in Fig. 12. In Table III we can see the misclassification rate. Most of the time 5.37% the error will be off by one person, while 0.6% of the time the error will be by more than one person.

TABLE II
OCCUPANCY CLASSIFICATION ACCURACY

Data used to classify	Frame length	Detection accuracy
Desk Smart Plug only	no frame	39.69%
WiFi router only	no frame	44.54%
Desk Smart Plug and WiFi router	no frame	52.15%
Router only	1 minute	93.49%
Desk Smart Plug and WiFi router	1 minute	96.10%

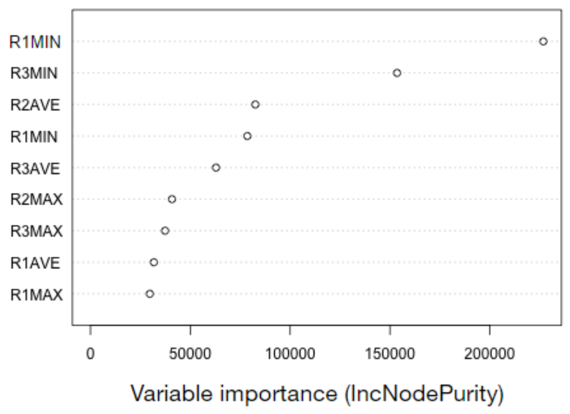

TABLE III
MISCLASSIFICATION RATE BY NUMBER OF PEOPLE

Classification offset	Percentage rate
1 person	5.37270%
2 person	0.48027%
3 person	0.08877%
4 person	0.01782%
5 person	0.00630%

This leads us to two conclusions. The office environment is static enough for us to be able to calculate the number of occupants just from the data traffic. While there are definitely edge cases where it is not possible to correctly classify, on average, when looking at a period of 10 minutes user behaviour does not significantly change and it can be used to predict the number of people present in a room. Secondly, since the minimum electric power consumption data point has been shown to be the most significant we can confirm that the beacon signal and probe response packets from additional devices are crucial for classifying the number of devices. As also shown in [13] these packet have the highest energy consumption. Since they are periodic and of constant size, they raise the overall electric power consumption of the router much more than the data packets, which allows for easy distinguishability of additional devices without relying on the assumption of the proportionally higher amount of traffic.

5. Conclusion

In this paper, for the first time we introduce the idea of using a WiFi router's electric power consumption in an occupancy-based energy control system. We hypothesize that both user behaviour is predictably stable and that the additional devices impact the electric power consumption more than data traffic. We are able to experimentally confirm both points and show that it is possible to determine the exact number of people in a room with very high accuracy. While we can confidently claim the system accuracy in one environment, we plan on gathering more measurement in different environments to be able to show the precision of our system in different environments.

Reference

[1] T. A. Nguyen and M. Aiello, "Energy intelligent buildings based on user activity: A survey," Energy and Buildings, vol. 56, pp. 244 – 257, 2013. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0378778812004537

[2] H. xiang Zhao and F. Magouls, "A review on the prediction of building energy consumption," Renewable and Sustainable Energy Reviews, vol. 16, no. 6, pp. 3586 – 3592, 2012. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1364032112001438

[3] G. R. Newsham and B. J. Birt, "Building-level occupancy data to improve arima-based electricity use forecasts," in Proceedings of the 2Nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building, ser. BuildSys '10. New York, NY, USA: ACM, 2010, pp. 13–18. [Online]. Available: http://doi.acm.org/10.1145/1878431.1878435

[4] J. Pan, R. Jain, S. Paul, T. Vu, A. Saifullah, and M. Sha, "An internet of things framework for smart energy in buildings: Designs, prototype, and experiments," IEEE Internet of Things Journal, vol. 2, no. 6, pp. 527–537, Dec 2015.

[5] V. L. Erickson, Y. Lin, A. Kamthe, R. Brahme, A. Surana, A. E. Cerpa, M. D. Sohn, and S. Narayanan, "Energy efficient building environment control strategies using real-time occupancy measurements," in Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings, ser. BuildSys '09. New York, NY, USA: ACM, 2009, pp. 19–24. [Online]. Available: http://doi.acm.org/10.1145/1810279.1810284

[6] A. Barbato, L. Borsani, A. Capone, and S. Melzi, "Home energy saving through a user profiling system based on wireless sensors,"in Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings, ser. BuildSys '09. New York, NY, USA: ACM, 2009, pp. 49–54. [Online]. Available: http://doi.acm.org/10.1145/1810279.1810291

[7] 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, ser. SenSys '10. New York, NY, USA: ACM, 2010, pp. 211–224. [Online]. Available: http://doi.acm.org/10.1145/1869983.1870005

[8] W. Kleiminger, C. Beckel, T. Staake, and S. Santini, "Occupancy detection from electricity consumption data," in Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings, ser. BuildSys'13. New York, NY, USA: ACM, 2013, pp. 10:1–10:8. [Online]. Available: http://doi.acm.org/10.1145/2528282.2528295

[9] 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, ser. BuildSys '10. New York, NY, USA: ACM, 2010, pp. 1–6. [Online]. Available: http://doi.acm.org/10.1145/1878431.1878433

[10] A. Marchiori and Q. Han, "Distributed wireless control for building energy management?" in Proceedings of the 2Nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building, ser. BuildSys '10. New York, NY, USA: ACM, 2010, pp. 37–42. [Online]. Available: http://doi.acm.org/10.1145/1878431.1878441

[11] K. Padmanabh, A. Malikarjuna, V, S. Sen, S. P. Katru, A. Kumar, S. P. C, S. K. Vuppala, and S. Paul, "isense: A wireless sensor network based conference room management system," in Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings, ser. BuildSys '09. New York, NY, USA: ACM, 2009, pp. 37–42. [Online]. Available: http://doi.acm.org/10.1145/1810279.1810288

[12] M. Milenkovic and O. Amft, "An opportunistic activity-sensing approach to save energy in office buildings," in Proceedings of the Fourth International Conference on Future Energy Systems, ser. e-Energy '13. New York, NY, USA: ACM, 2013, pp. 247–258. [Online]. Available: http://doi.acm.org/10.1145/2487166.2487194

[13] K. Gomez, R. Riggio, T. Rasheed, and F. Granelli, "Analysing the energy consumption behaviour of wifi networks," in 2011 IEEE Online Conference on Green Communications, Sept 2011, pp. 98–104.

[14] randomforest: Breiman and cutler's random forests for classification and regression." [Online]. Available: https://cran.rproject.org/web/packages/randomForest/index.html

[15] T. Petrovic and H. Morikawa, "Active sensing approach to electrical load classification by smart plug," in 2017 IEEE Power Energy Society Innovative Smart Grid Technologies Conference (ISGT), April 2017, pp.1–5, to be published.

[16] A. Reinhardt, P. Baumann, D. Burgstahler, M. Hollick, H. Chonov, M. Werner, and R. Steinmetz, "On the accuracy of appliance identification based on distributed load metering data," in Sustainable Internet and ICT for Sustainability (SustainIT), Oct. 2012.

[17] M. Fern´andez-Delgado, E. Cernadas, S. Barro, and D. Amorim, "Do we need hundreds of classifiers to solve real world classification problems?"Journal of Machine Learning Research, vol. 15, pp. 3133–3181, 2014. [Online]. Available: http://jmlr.org/papers/v15/delgado14a.html