Poster Abstract: An In Situ System for Annotation of Home Energy Data

Sami Rollins
University of San Francisco
srollins@cs.usfca.edu

Nilanjan Banerjee
University of Maryland, Baltimore County
nilanb@umbc.edu

Lazeeb Choudhury
University of San Francisco
lfchoudhury@usfca.edu

David Lachut
University of Maryland, Baltimore County
dlachut1@umbc.edu

ABSTRACT

This work presents a system for collecting user activity annotations using in-home distributed energy monitoring combined with a novel algorithm that generates device-specific profiles used to identify potentially important changes in user context. In a five-week study of five homes, the system was able to generate profiles for 80% of the devices studied. Moreover, between four and five important changes in user context were identified when background loads were accurately distinguished.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous; D.2.8 [**Software Engineering**]: Metrics—*complexity measures, performance measures*

General Terms

Human Factors, Measurement

Keywords

energy, sustainability

1. INTRODUCTION

Home energy management is increasingly important and, though there are a plethora of tools for accessing energy consumption data, few provide concrete insights that can directly help users manage demand. Tools that allow users to connect their activities to their energy consumption habits have been shown to be a promising means of offering desired insights into a home or appliance energy profile [1], however solutions for collecting annotations from users can be error prone or intrusive. This work presents a system for collecting in situ annotations using a smart phone application coupled with an off-the-shelf home energy measurement infrastructure. Using only the information collected from a set of off-the-shelf energy meters, the system determines, in real time, whether the energy profile of a device suggests the user is performing an important energy consumption activity. If so, it sends a push notification to a user's smartphone requesting that he/she provide an annotation of his/her current context.

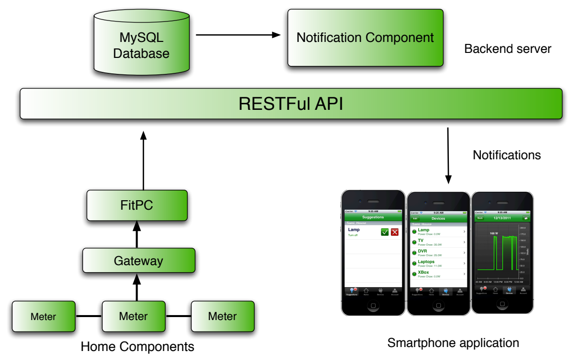

Figure 1: Green Homes architecture.

2. SYSTEM DESIGN

The Green Homes project is an ongoing effort to understand energy usage in homes, particularly those powered by renewable sources, and to encourage sustainability by developing automated mechanisms for matching energy demand with available supply. Figure 1 illustrates the architecture of the Green Homes system. The off-the-shelf energy meters communicate with a dual-radio gateway [2] in each home. Also in each home is a client component that polls the energy meters every 30 seconds and reports readings to a centralized server. Additionally, every 1 minute, the notification component on the server side executes the power profiling algorithm described below and, if appropriate, pushes a notification to the user's phone. Annotations entered by the user are then stored on the server for postprocessing.

The key challenge we address in the design of the notification system is deriving potentially important changes in user

BuildSys '13, Nov 14 - November 15 2013, Roma, Italy

Subject ID	Days of Data	Notifications/Day All	Inactive Days	Notifications/Day Active Only	Annotations/Day	Response Rate All/Interactive
1	35	22.1	2	23.4	2.9	6%/13%
2	33	12.8	9	16.2	1.2	3%/7%
3	35	3.2	8	4.1	2.0	33%
4	25	4.8	0	4.8	2.7	15%
5	34	3.1	4	3.5	1.9	43%

Table 1: The table overviews the data used in this analysis. Values reported as per day (/Day) are calculated as the mean across the days of the study. Inactive days are days in which no notifications were generated. The response rate is reported as the percentage of notifications that yield a response with a 10-minute window for notifications generated by all devices and notifications generated by devices not considered background loads, respectively.

context through analysis of raw power traces of individual devices. Our algorithm addresses the problem by attempting to identify when a device transitions from a low-power state to a higher-power state. This can mean capturing a transition from off to on, but it may also mean extracting a transition from a power-saving state to an active state. Our approach uses the DBSCAN clustering algorithm [3] to produce a unique profile for each device in our study. DBSCAN is a density-based algorithm that identifies clusters while excluding noise and is an ideal choice for this application as it does not require the number of clusters to be provided as input and it can be implemented very efficiently, particularly for one-dimensional data such as ours. Once the profile is generated, the final step of the algorithm is to evaluate whether a device is non-interactive and represents only background load for the user. To identify background loads, we apply a heuristic that will classify a device as non-interactive if in more than 80% of the hours for which data was reported for the device there was a change in power state for the device. Effectively, if a device transitions between power states in more than 19 hours of the day then the device is likely not manually controlled by the user and will not trigger notifications.

3. DEPLOYMENT AND FINDINGS

In order to understand whether the Green Homes notification component is an effective means of collecting contextual annotations for home energy data we report the results of a five-week deployment. To conduct the study, we enabled notifications for 5 of the 9 homes currently participating in the Green Homes project and collected data from the period between November 1, 2012 and December 6, 2012. Table 1 reports the number of days of data collected in each home as our deployment was rolled out incrementally. Participants were asked to provide feedback on their activities using our iPhone or Android application.

We report the two most salient findings of our study. First, we discovered that the profiling algorithm works for nearly 80% of devices, producing usable profiles for 31 of 39 devices in the study. For devices that were unsuccessfully profiled, we saw several scenarios. Several devices simply were not used during the study, or remained on with a consistent powerdraw throughout the experiment. This was the case, for example, with a DVR. Additionally, devices such as laptops produced a single cluster with a significant range as the power draw spikes when the device is plugged in but gradually decreases over time. This suggests that defaulting to a threshold-based approach would yield profiles for a broader range of devices.

Table 1 reports the mean number of daily notifications for each subject. There was considerable variation in the frequency with which the system requested annotations, with subjects 1 and 2 receiving significantly more notifications than the other subjects. In the homes of both subjects 1 and 2 a refrigerator was measured and accounts for a significant percentage of the daily notifications. In the home of subject 2, the refrigerator was identified as a non-interactive device on November 24 and, after that point, the number of notifications decreased to fewer than 10 per day. Moreover, excluding the refrigerator and space heater—another background load—from the notifications for subject 1, the average daily notifications drops to below 8. Across all homes, the mean number of notifications for interactive devices is 4.6, though this is skewed by a number of days when our subjects were traveling and no notifications were generated. Table 1 additionally shows, for each subject, the number of days when no notifications occurred—labeled Inactive Days—and the mean number of daily notifications excluding the inactive days when subjects were likely not home—labeled Active Only. This analysis demonstrates, first, that identifying background loads is a key component of maintaining a manageable number of potentially intrusive push notifications and, second, that in homes where we generate notifications for interactive devices only, our system is able to use device power profiles to identify 4–5 potentially important changes in user context each day.

4. REFERENCES

[1] Enrico Costanza, Sarvapali D. Ramchurn, and Nicholas R. Jennings. Understanding domestic energy consumption through interactive visualisation: a field study. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, UbiComp '12, pages 216–225, New York, NY, USA, 2012. ACM.

[2] `http://micasaverde.com/vera.php`. Micasaverde vera 2.

[3] Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proc. of 2nd International Conference on Knowledge Discovery and Data Mining (KDD-96)*, pages 226–231, 1996.