Pervasive and Mobile Computing 30 (2016) 132–150

Contents lists available at ScienceDirect

Pervasive and Mobile Computing

journal homepage: www.elsevier.com/locate/pmc

Fine-grained appliance usage and energy monitoring through mobile and power-line sensing

Nirmalya Roy [a,*], Nilavra Pathak [a], Archan Misra [b]

[a] Department of Information Systems, University of Maryland Baltimore County, United States

[b] School of Information Systems, Singapore Management University, Singapore

ARTICLE INFO

Article history:
Received 1 August 2015
Received in revised form 8 November 2015
Accepted 7 January 2016
Available online 22 January 2016

Keywords:
Energy
Plug loads
Green building
Mobile applications

ABSTRACT

To promote energy-efficient operations in residential and office buildings, non-intrusive load monitoring (NILM) techniques have been proposed to infer the fine-grained power consumption and usage patterns of appliances from power-line measurement data. Fine-grained monitoring of everyday appliances (such as toasters and coffee makers) can not only promote energy-efficient building operations, but also provide unique insights into the context and activities of individuals. Current building-level NILM techniques are unable to identify the consumption characteristics of relatively low-load appliances, whereas smart-plug based solutions incur significant deployment and maintenance costs. In this paper, we investigate an intermediate architecture, where smart circuit breakers provide measurements of aggregate power consumption at room (or section) level granularity. We then investigate techniques to identify the usage and energy consumption of individual appliances from such measurements. We first develop a novel correlation-based approach called CBPA to identify individual appliances based on both their unique transient and steady-state *power signatures*. While promising, CBPA fails when the set of candidate appliances is too large. To further improve the accuracy of *appliance level* usage estimation, we then propose a hybrid system called AARPA, which uses mobile sensing to first infer high-level activities of daily living (ADLs), and then uses knowledge of such ADLs to effectively reduce the set of candidate appliances that potentially contribute to the aggregate readings at any point. We evaluate two variants of this algorithm, and show, using real-life data traces gathered from 10 domestic users, that our fusion of mobile and power-line sensing is very promising: it identified all devices that were used in each data trace, and it identified the usage duration and energy consumption of low-load consumer appliances with ∼87% accuracy.

1. Introduction

There is widespread interest in developing solutions that provide knowledge of the *fine-grained* usage and power consumption of everyday appliances (such as coffee makers and televisions) in residential buildings. Such interest is primarily driven by recent interest in energy-efficient building operations, especially as empirical evidence suggests that empowering consumers with greater awareness of their energy consumption patterns can result in 5%–20% reduction in electricity usage [1,2]. However, we believe that, besides this energy-related benefit, the ability to precisely capture the

* Corresponding author.
E-mail addresses: nroy@umbc.edu (N. Roy), nilavra1@umbc.edu (N. Pathak), archanm@smu.edu.sg (A. Misra).

http://dx.doi.org/10.1016/j.pmcj.2016.01.003
1574-1192/

usage profile of everyday consumer appliances also provides insight into an individual's *context*, at a fine granularity that existing approaches (typically based on mobile sensing [3]) simply cannot provide. For example, while past approaches such as [4–6] can help classify activities such as "making dinner" or "watching TV", appliance monitoring can additionally indicate that the 'toaster was used today' (revealing details about the food items consumed) or 'the specific TV channel watched' [7].

This paper thus explores the technical feasibility of a vision where the sensing capabilities of body-worn pervasive devices are combined with the power-line sensing of appliance usage to provide significantly greater insight into the daily activities (formally called Activities of Daily Living or *ADLs*) of individuals, especially in their residential environments. While the empirical investigations carried out in this paper utilize smartphones (that may or may not be always carried around inside a home), an eventual embodiment will likely rely on wearable devices (e.g., smart-watches [8], smart-bracelets [9]) that are now gaining wider market acceptance and that a user will likely wear almost-continuously [10].

The most common technique for capturing appliance usage is Non-Intrusive Load Monitoring (NILM) [11], where a-priori knowledge of the *power signatures* of individual electrical devices is used to infer the activation and de-activation of such devices from an observation of the aggregate power consumption signal. Present NILM techniques however operate at two extremes. The most commonly deployed solution involves the use of Smart Energy Meters, installed outside the building, which provide only the *total* sum of power consumed by all the appliances belonging to that house or consumer [12]. Given such an aggregated observation, we can identify only the major loads in the house (e.g., the HVAC unit or the Washer/Dryer) [13], not the myriad "low-load" appliances (e.g., toasters, TVs, etc.). The other extreme involves the use of wireless "Smart Plugs" [14,15], inserted into individual electrical sockets, and thus providing power measurements at much finer, *individual appliance-level* granularity. The disadvantage, of course, is that this approach requires the provisioning of such smart plugs in every individual power outlet, resulting in fairly steep deployment challenges and increased monitoring cost.

Motivated by these limitations, our research attempts to answer two questions: (a) *is there a meaningful alternative, in between the extremes of total vs. individual power consumption monitoring, that helps us in our quest to infer individual appliance consumption patterns*?, and (b) *can the sensing capability of pervasive mobile/wearable devices be exploited to improve the accuracy of detecting such appliance usage, thereby providing finer-grained context monitoring?*

More specifically, we propose and utilize the *Smart Circuit Breaker* approach: here each room (or possibly multiple rooms) of a residential or office building is equipped with an instrumented circuit breaker (which we have implemented in ongoing work) which helps measure the aggregate room-level granularity [16]. Such circuit-breakers provide, for example, the total power consumed in the kitchen, without necessarily separating out the concurrent use of one or more appliances, such as the fridge, coffee maker, microwave or toaster. The challenge then is to devise an analytics technique that can help disaggregate such measurements to infer the use of even relatively low-power appliances.

In this paper, we first use real-life measurement studies to develop an enhanced Correlation-Based Power Analytics algorithm, called **CBPA**, that applies correlation over both macroscopic and microscopic power consumption features, to identify *the total usage duration*, and *the total energy consumption*, of individual devices, from such circuit-breaker level aggregated data. While CBPA helps to successfully disaggregate room-level power data into individual devices in *some* practical cases of interest, its accuracy diminishes if the candidate set of possible low-load devices becomes modestly large. Accordingly, we then explore a joint sensor fusion approach, that combines mobile plus power-line sensing data, to first obtain a smaller, *filtered* set of candidate appliances whose cumulative power consumption is reflected in the reading of the smart circuit breaker. We provide two different variants of this ADL-driven approach, called Activity-Aware Room-level Power Analytics (**AARPA**), one rule-based and the other probabilistically-weighted, and then use real-world usage traces to establish their efficacy. Our work thus establishes how a *joint fusion of mobile-sensing based ADL recognition and room-level power-line consumption data* can provide a practical solution that (a) helps capture the energy consumption characteristics of low-load, commonly-used domestic appliances and (b) provides useful additional context about the lifestyle habits and context of an individual.

Key Contributions:

- Our key contribution lies in the proposed AARPA approach, that jointly fuses the ADL recognition capabilities of pervasive mobile sensing with the practical power-line sensing offered by Smart Circuit Breakers, to reliably identify the *usage* and *energy consumption profile* of relatively low-load daily-use appliances (such as toasters, microwave ovens and treadmills). AARPA provides a practical way to achieve such fine-grained discrimination of appliance usage, without incurring the prohibitive operational overheads of per-device power monitoring via Smart Plugs.
- As a secondary contribution, we suggest CBPA, an enhanced power analytics algorithm, that utilizes correlation measures over long-term traces of both microscopic (*transient*) and macroscopic (*longer-term*) power consumption features, to identify the use of such low-load daily-use appliances. While CBPA's accuracy is not adequate when applied in isolation, it performs better than prior approaches [13,17], which are based solely on longer-term 'steady-state' power consumption features. More importantly, when used as part of AARPA, CBPA proves very useful in disambiguating appliance use.
- We evaluate the quantitative benefits of AARPA using real-life activity traces from 10 domestic users, collected over several weeks. Our results show that, given normal everyday patterns of domestic living, AARPA can provide very high accuracy in identifying device *usage* (AARPA identified 100% of all devices used in our studies), and significantly increase the accuracy (by approx. 45%) of estimating per-device *energy consumption*. This is achieved even under realistic limits of at-home activity recognition (accuracies of ≈90%). These results demonstrate the viability of the AARPA approach, for both finer-grained context/activity recognition and appliance-level energy consumption monitoring.