Opportunistic Sensing for Efficient Energy Usage in Private Households

Wilhelm Kleiminger, Christian Beckel, and Silvia Santini

Institute for Pervasive Computing
ETH Zurich, 8092 Zurich, Switzerland
{kleiminger,beckel,santinis}@inf.ethz.ch

1 Introduction

Information and Communication Technology (ICT) recently became to be seen as an instrument with huge potential to improve energy efficiency in several sectors. For instance, recent studies have shown that modern sensing and communication systems can help in detecting inefficient use of household appliances [1]. Providing this information to the occupants of the households makes them aware of their electricity consumption footprint and often willing to change their utilization patterns [1]. Involving users in energy savings decisions, however, also has drawbacks. In particular, long-term user participation is hard to achieve unless durable incentives, like significant monetary savings, can be provided. Also, improper user actions may hamper the achievable energy savings [2]. In this context, the use of completely automated ICT solutions becomes necessary in order to guarantee long-term benefits. For example, systems that are able to estimate actual occupancy patterns of households can perform heating-control tasks efficiently. Such systems may enable significant energy savings, as heating represents a major source of energy consumption. In Switzerland, for instance, heating accounts for 70% of the total consumption in residential buildings [3].

Recently published studies show that the use of dynamically programmable thermostats may allow to save about one third of the total energy currently spent for heating [2]. However, the more residents are involved in actively programming the system, the lesser the actual savings will turn out to be [2]. To cope with these shortcomings, we are investigating strategies that can automatically adjust the heating levels depending on the actual occupancy state and activity patterns of the household. Indeed, the temperature at which an household must be heated varies if the occupants are away, at home and awake, or sleeping. To gather the information necessary to estimate such activity patterns we are building a system that opportunistically taps into existing sensor information such as GPS coordinates from residents' mobile phones, traces of connections to WiFi access points, or other data like readings from digital electricity meters. Our *smart heating system*, which we describe in more detail in the following sections, is a primary example of ICT-enabled solutions that can provide for significant energy savings.

2 Web-enabled sensors and actuators

The number of sensors deployed in urban or natural environments is constantly increasing [4]. The widespread use of smartphones adds to this abundance, as modern phones are equipped with several sensors, including GPS position sensors or near-field-communication (NFC) readers. Thanks to this large availability of devices, the possibility of unobtrusively gathering and sharing sensor data, usually referred to as *opportunistic sensing*, becomes possible [4]. This data can then be leveraged for several different applications, including urban and environmental monitoring.

The collection and sharing of sensor data, however, typically required the use of proprietary and application-specific protocols to interconnect the sensing devices. Bridging their data to the Internet in order to make it widely available then required ad-hoc gateways, whose use limited the extensibility and robustness of the application. On the other hand, the use of Internet protocols has traditionally been regarded as too heavyweight for being implemented directly on tiny sensors (or actuators), which typically do not comply with the necessary computational, communication and memory requirements [5]. More recently, however, the use of conventional Web protocols on resource-constrained devices is becoming feasible thus enabling a seamless integration of sensors and actuators into the Web [6]. For instance, an increasing number of devices offer service interfaces that rely on the REST (Representational State Transfer) [7] paradigm and encode data using open formats like JSON[1] (JavaScript Object Notation). The active research community working in this area also continuously pushes towards the development and adoption of *Open APIs*, i.e., application programming interfaces that allow to easily interconnect Web instances. This, in turn, opens new possibilities for the realization of mashup applications that can easily leverage, e.g., sensors developed for a security system (like an infrared sensor) to gather the occupancy information used by a smart thermostat. Last but not least, applications can dynamically incorporate available sensors thanks to discovery services for Web-enabled devices [8].

Our smart heating system will rely on available sensing devices and opportunistically gather their data. The fact that such devices can increasingly be accessed and controlled through standard Web protocols will significantly ease the deployment of our system in common households. Figure 1 schematically shows the connection of exemplary Web-enabled sensing and actuation sources.

3 Smart heating

Current programmable thermostats rely on user-supplied schedules to save energy when the home is unoccupied. When programming a thermostat, users are faced with two challenges. First, they need to predict their own occupancy schedule and, secondly, they need to consider the time it takes for the thermostat to arrive at the comfort temperature. If this ramp-up time is ignored, the home

[1] `http://www.json.org`

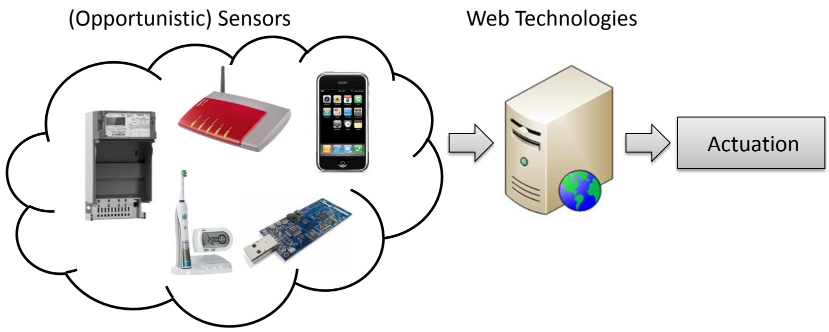

Fig. 1. Opportunistic Data Sources.

will not be at a comfortable temperature when the occupants return. Although the computation of the ramp-up time is taken over by some higher-end adaptive thermostats, for many lower-end thermostats, users still have to make this prediction themselves. [2] However, even with the adaptive variety, the user still has to continuously (re-)program the thermostat in order to accurately reflect her occupancy patterns. We believe that both these problems can be solved by a smart heating system that senses and predicts occupancy while keeping track of the operation of the thermostat.

3.1 Background

Our smart heating system will build and improve upon recent approaches that have been suggested to replace user-programmable thermostats. For instance, Gupta et al. developed a thermostat that retrieves the current position of a user using her smartphone's GPS sensor to infer households' occupancy [9]. When the occupant (and her smartphone) moves closer to home, the system starts preheating the flat. This approach works well for single-person households, but depends on a single source of information (the GPS sensor). A drained battery or a loss of connectivity causes the system to fail. Lu et al. showed a more fault-tolerant system by using traditional occupancy sensors like reed switches and infrared sensors [2]. The collected occupancy data is aggregated and used to build a Markov Model that predicts when the occupants are likely to return home. While this system is more tolerant to failures, it relies on the installation of additional sensors. Although wireless, battery-powered sensors might be installed by the occupant in a few minutes, they still suffer from the same reliability problems as the GPS-based thermostat. On the other hand, wired sensors offer much less flexibility and may cause high installation costs.

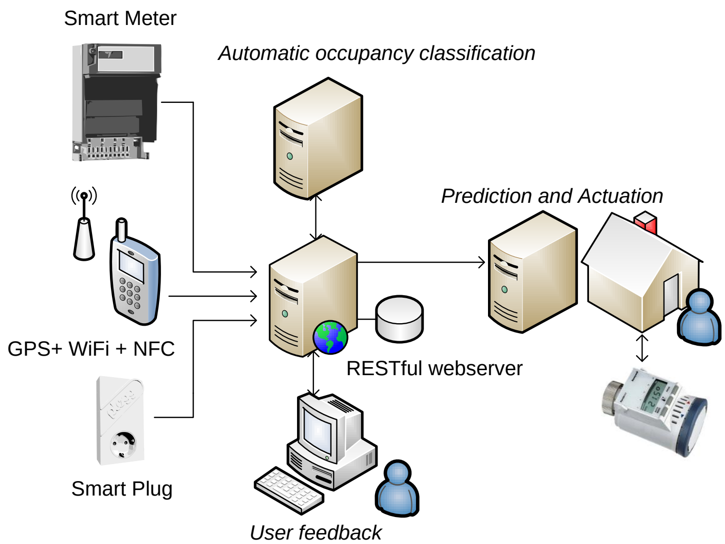

Fig. 2. Overview of opportunistic data sources and our proposed architecture.

3.2 An opportunistic approach

To cope with these problems we suggest that a smart heating control system should adapt to the available sensors opportunistically, as schematically depicted in Figure 2. If a GPS-equipped mobile phone is available, its current position may be included in the database used by the system. Furthermore, the system may also keep track of WiFi logins or of the use of specific household appliances, which can be detected by smart plugs or electricity meters. Smart electricity meters are particularly interesting for our system as they are intrinsically linked to the household and are becoming ubiquitous throughout Europe [1]. Electrical consumption data is harder to analyze than GPS coordinates. While the latter unambiguously signal presence or the current position of an occupant, the information contained in electricity consumption profiles is more fuzzy. Several devices such as fridges or wireless Internet routers are typically active even when the occupants are away or sleeping. To solve this potential ambiguity we can opportunistically leverage other data, for instance information about WiFi connections that reveal the presence or absence of given devices (and their owners) from the home. Additionally, we plan to leverage NFC technology, which is also increasingly becoming available on modern smartphones, to retrieve additional information about users' whereabouts and activities.

Besides the automatic occupancy classification, our system will also be able integrate user-feedback. By giving the users the ability to set occupancy schedules as guidelines for the smart heating system, we can make sure that we always

[2] http://www.energystar.gov/index.cfm?c=thermostats.pr_thermostats_guidelines

have a viable fall-back strategy. By registering user interaction with the system we can also fine-tune the algorithm used by our system to predict user schedules and thus maximize comfort.

3.3 Privacy implications

Obviously, the analysis of sensory information from various sources including a detailed breakdown of the energy consumption and smart phone localization data requires an architecture that adheres to strict privacy principles. When this data is used to compute schedules for the thermostat, there is no need for the data to be made available outside of the household. In addition, precaution has to be taken that the data is not being used for any purposes other than the smart heating system. Concerning privacy, we believe the contribution of our work to lie mainly in the sensibilization of users to the potential issues arising from the ubiquity of sensor information. By showing the amount of data needed to accurately infer users' occupancy schedules we aim to make users' aware of the significance of sharing data with social networks, utility companies and other service providers.

4 Current status and outlook

We are currently investigating statistical classification methods to detect the occupancy state of an household from its electricity consumption profile. Further, we are conducting a user-study to quantify the ability of household occupants to predict their own schedules on a day-to-day basis. To measure the actual occupancy state of the household we rely on the information collected by an application running on modern smartphones. This application collects GPS and WiFi traces and also allows users to explicitly declare their status (e.g., at home, away, going to sleep) by using a phone-mounted NFC-reader to scan accordingly labeled tags.

As a next step towards the realization of our smart heating system we will test the robustness of our occupancy classification algorithm using the data collected in preliminary studies. Our goal is to make the algorithm incrementally integrate sensor information depending on the actual accuracy requirements. In particular, additional sensor data should be gathered only when the confidence in the current classification is not sufficient to reliably change the state of the thermostat. To quantify the achievable energy savings and test our smart thermostat in real-world settings, we are currently deploying it in test households equipped with digital electricity meters.

References

1. Mattern, F., Staake, T., Weiss, M.: ICT for Green – How Computers Can Help Us to Conserve Energy. In: Proceedings of the 1st international Conference on Energy-efficient Computing and Networking (e-Energy 2010). (2010)

2. Lu, J., Sookoor, T., Srinivasan, V., Gao, G., Holben, B., Stankovic, J., Field, E., Whitehouse, K.: The Smart Thermostat: Using Occupancy Sensors to Save Energy in Homes. In: Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems (SenSys 2010). (2010)
3. Kirchner, A., Kemmler, A., Hofer, P., Keller, M., Jakob, M., Catenazzi, G.: Analysis of Energy Consumption in Switzerland 2000–2009. Swiss Federal Office of Energy, Bern, Switzerland (2010) [Original publication in German].
4. Campbell, A.T., Eisenman, S.B., Lane, N.D., Miluzzo, E., Peterson, R.A., Lu, H., Zheng, X., Musolesi, M., Fodor, K., Ahn, G.S.: The Rise of People-Centric Sensing. Internet Computing, IEEE **12**(4) (2008) 12–21
5. Vasseur, J.P., Dunkels, A.: Interconnecting Smart Objects with IP - The Next Internet. Morgan Kaufmann (2010)
6. Guinard, D., Trifa, V., Wilde, E.: A Resource Oriented Architecture for the Web of Things. In: Proceedings of the 2nd International Conference on the Internet of Things (IoT 2010), IEEE (2010) 1–8
7. Fielding, R.T.: Architectural Styles and the Design of Network-based Software Architectures. PhD thesis, University of California, Irvine (2000)
8. Mayer, S., Guinard, D.: An Extensible Discovery Service for Smart Things. In: Proceedings of the 2nd International Workshop on the Web of Things (WoT 2011). (2011) 7–12
9. Gupta, M., Intille, S.S., Larson, K.: Adding GPS-Control to Traditional Thermostats: An Exploration of Potential Energy Savings and Design Challenges. In: Proceedings of the 7th International Conference on Pervasive Computing (Pervasive 2009). (2009)