Diss. ETH Nr. 20253

Ubiquitous Computing Technologies for Residential Energy Conservation

A dissertation submitted to
ETH Zurich

for the degree of
Doctor of Sciences

presented by

Markus Weiss
Diplom-Informatiker, Julius-Maximilians University Würzburg
born May 08, 1981, citizen of Germany

accepted on the recommendation of
Prof. Dr. Friedemann Mattern, examiner
Prof. Dr. Elgar Fleisch, co-examiner
Prof. Dr. Hans Gellersen, co-examiner

2012

Acknowledgements

This dissertation has been created during my time as researcher at the Bits-to-Energy Lab, a joint initiative of the Institute of Pervasive Computing and the Information Management Group at ETH Zurich as well as the Institute of Technology Management at the University of St. Gallen. Parts of the work were conducted during a stay as visiting scholar at the Engineering Systems Division, a part of the MIT Energy Initiative at the Massachusetts Institute of Technology in Cambridge, USA.

I wish to express my gratitude to all who supported me over the past years. This work could not have been completed without your help.

First, I want to thank my supervisor Prof. Dr. Friedemann Mattern for his great guidance, his inspiring thoughts, and the exceptional freedom that fostered creativity and made working a great pleasure. Further, I want to thank my co-supervisors Prof. Dr. Elgar Fleisch and Prof. Dr. Hans Gellersen for their valuable discussions and challenging feedback from different angles. I also want to thank my project manager Dr. Thorsten Staake for working close with me on each step along the way, for countless valuable and fun discussions, and for pushing at the right moments to proceed from initial ideas to tangible results.

Of course, all my colleges at the Bit-to-Energy Lab and the affiliated institutes at ETH Zurich and the University of St. Gallen as well as the ones I cooperated with at the Massachusetts Institute of Technology deserve a big thank you. It is impossible to mention each of you, but being a member of these groups and being in a team with you has been an incredible experience in an inspiring environment (be it during day or night times). I further want to thank my students for all their effort and enthusiasm they put into my project. Working with you always encouraged me in my work. Special thanks go to the industry partners that supported this work and helped realize an applicable system and enabled its evaluation in the field.

Last, I want to thank my family, especially my parents Renate and Siegfried, as well as my close friends for their support, encouragement, and for visiting me all over the world throughout my life – Danke, Ihr seid Hölle.

Zurich, March 2012 Markus Weiss

Abstract

Residential electricity consumption is continuously increasing and accounts now for about one third of the total electrical energy produced in Europe and the U.S. How much residential electricity is used depends primarily on the operated household appliances and the behavior of the residents. One major difficulty for individuals who are interested in saving energy in their household is the lack of information about their electricity consumption. Feedback on energy usage is typically only provided by a monthly (if not yearly) utility bill and thus remains rather vague and opaque to most residents. As a result, most individuals could reduce their electricity consumption, but few know how much they consume and even fewer know how much energy they consume for a particular purpose (e.g., lighting). And even those who do have a fair understanding of their consumption patterns rarely receive guidance about the changes that will have the biggest impact on their electricity bill.

Through recent technological advances in terms of cost, size, and computing power, Information and Communication Technology can help in many ways to address the challenge of making residential electricity consumption visible to individuals. Embedding computing and communication devices in everyday objects, as advocated by Ubiquitous Computing, can help to communicate the consumption, but also the most energy-efficient usage of a particular smart appliance. Smart meters that enable capturing fine-grained electricity consumption information at high frequency are currently replacing traditional electricity meters. Smartphones have become ubiquitous, powerful computing platforms that allow visualizing energy consumption on the spot without the need for external wall displays. By digitally enhancing physical devices that populate homes, Ubiquitous Computing is offering new possibilities to address the problem of residential energy conservation.

Applying Ubiquitous Computing technologies for residential energy conservation raises research questions about the most suitable overall system design of energy feedback solutions and the most appropriate modality of communicating the consumption and guidance information to the consumer. This thesis addresses these research questions by examining how Ubiquitous Computing can help provide effective feedback that goes beyond mere consumption values and is at the same time integrated into daily life. Following a user-centric approach that combines the use of smartphones and smart meters, we tackle some of the open challenges in residential energy conservation. The contributions of this thesis are threefold.

First, we design, develop, and evaluate an electricity sensing and feedback infrastructure that seamlessly integrates into the residential environment. It addresses the technical requirements that have been identified in previous research to enable users to better understand their energy consumption (i.e., integration into daily life, real-time information provisioning, low usage barrier, and fine-grained consumption information). At the same time, the infrastructure serves as an easily extendible framework that can be used by other researchers (e.g., to develop and test visualization concepts, to realize further automated energy savings, or to design behavioral science experiments). To demonstrate the feasibility of our approach, we implemented a prototype of the infrastructure and deployed and evaluated it in a laboratory setting as well as in four households in Switzerland. The architecture supports the interaction capabilities of mobile phones together with the integration of smart electricity meters and is used as the base for most other work done in the context of this thesis.

Second, we evaluate the potential of mobile phones to serve as portable electricity feedback monitors in two different experimental settings: a user study as well as a real-world deployment. In the user study, we analyze the perceived value of various feedback functionalities and identify which type of feedback is meaningful to users. Moreover, we evaluate the general usability, accuracy, and intention of use of such an electricity feedback application. The real-world deployment aims at characterizing different user types and providing qualitative results gathered through the use of the application. It shows that to foster long-term application of the system motivational concepts are required that engage users once their initial curiosity is satisfied. Overall, the results confirm the suitability of mobile phones as an energy feedback interface and provide insights for the design of future energy conservation applications. They outline that a clear and easy to explain use case scenario is key and that knowledge-increasing functionalities as well as those functionalities from which monetary savings can be directly implied are perceived as most important. To address technophobe users, action-guiding feedback that goes beyond displaying aggregated information in mere numbers is required.

Third, we develop, implement, and evaluate an algorithm that disaggregates the overall energy consumption to the consumption of individual devices. It enables users to link consumption with behavior and provides the base for automated energy recommendation systems. Compared to other load disaggregation approaches, our algorithm does not require additional hardware nor complex, time-intense calibration conducted by domain experts. Moreover, our approach is able to easily take new appliances into account where other systems require recalibration. With a simple yet powerful feature provided by the user interface on the mobile phone, users can incrementally integrate additional appliances into the disaggregation process. This is particularly important in a fast changing home environment. We evaluated the performance of our system in a laboratory test study with eight simultaneously running devices, achieving recognition rates of almost 90%.

Kurzfassung

Der Stromverbrauch privater Haushalte wächst kontinuierlich und macht heute bereits rund ein Drittel der gesamten elektrischen Energieproduktion in Europa und den USA aus. Die Strommenge, die ein Haushalt verbraucht, hängt in erster Linie von den betriebenen Geräten und von der Art und Weise, wie Bewohner diese Haushaltsgeräte verwenden, ab. Das Hauptproblem für Personen, die daran interessiert sind, in ihrem Haushalt Energie zu sparen, ist die fehlende Information über den individuellen Stromverbrauch. Die Rechnungsstellung mit zugehöriger Darstellung des Verbrauchs erfolgt heute höchstens monatlich (oft sogar nur jährlich) und lässt damit viele Haushalte über ihren Stromverbrauch lange im Unklaren. Grundsätzlich könnten viele Bewohner ihren Verbrauch reduzieren, allerdings wissen nur die wenigsten, wie viel und anteilmäßig für welchen Verwendungszweck sie Strom nutzen. Selbst diejenigen, die bereits ein gewisses Grundverständnis bezüglich ihres Verbrauchs besitzen, bekommen nur selten ausreichend Hilfe und Unterstützung, um konkrete Einsparpotentiale zu erkennen und auszunutzen.

Die jüngsten Fortschritte im Bereich der Informations- und Kommunikationstechnologie können diese Informationslücke im Haushalt schließen und die Bevölkerung bezüglich Energieverbrauch und Einsparmöglichkeiten sensibilisieren. Durch die Integration von Rechen- und Kommunikationsfähigkeiten in Alltagsgegenstände, wie es das Ubiquitous Computing propagiert, kann der Stromverbrauch erst erfasst und anschließend die Verbrauchsinformation, gemeinsam mit Hinweisen zur energieeffizienteren Verwendung eines Geräts, an den Benutzer kommuniziert werden. Intelligente Stromzähler (Smart Meter) können den Stromverbrauch in hoher zeitlicher Auflösung messen und ersetzen derzeit klassische Stromzähler flächendeckend. Mobiltelefone stellen heute allgegenwärtige, leistungsstarke Rechenplattformen dar, die den Energieverbrauch direkt darstellen können, so dass nicht auf zusätzliche Hardware, wie beispielsweise Wandbildschirme, zurück gegriffen werden muss. Durch die digitale Anreicherung von Alltagsgegenständen mit Rechenleistung und Kommunikationsmodulen bietet das Ubiquitous Computing eine neue Möglichkeit, das Problem des Energiesparens im Privathaushalt in Angriff zu nehmen.

Der Einsatz von Ubiquitous-Computing-Technologien mit dem Ziel, das Energiesparen in Privathaushalten zu vereinfachen, wirft aber auch einige Forschungsfragen auf, etwa, welche Architektur für Systeme, die Feedback zum Energieverbrauch liefern, am besten geeignet ist oder welche Art und Weise der Informationsdarstellung für den Benutzer am verständlichsten ist. Die vorliegende Dissertation befasst sich mit diesen Forschungsfragen, indem sie untersucht, wie Ubiquitous Computing helfen kann, effektives Feedback, das über die Darstellung der reinen Verbrauchsinformation hinausgeht, bereitzustellen und dieses dabei gleichzeitig im Alltag des Benutzers zu integrieren. Durch eine benutzerorientierte Herangehensweise, die den Gebrauch von intelligenten Stromzählern mit Mobiltelefonen verbindet, nimmt die Arbeit diese Herausforderungen an und liefert dazu die drei nachfolgend beschriebenen Hauptbeiträge.

In einem ersten Teil der Dissertation wird eine Infrastruktur, die Feedback über Stromverbrauch liefert und nahtlos in das Haushaltsumfeld integriert ist, auf Basis von Ubiquitous-Computing-Komponenten entworfen, entwickelt und evaluiert. Die Infrastruktur setzt dabei die aus der Literatur bekannten technischen Anforderungen im Bereich des Verbrauchsfeedbacks (wie Integration in den Tagesablauf, Bereitstellung möglichst feingranularer Verbrauchsinformation in Echtzeit und mit niedriger Nutzungsbarriere) um und ermöglicht es Benutzern, so ihren Stromverbrauch besser zu verstehen. Gleichzeitig stellt die Infrastruktur ein leicht erweiterbares Framework dar, das anderen Wissenschaftlern die Möglichkeit bietet, eigene Visualisierungskonzepte zu testen oder Verhaltensexperimente durchzuführen. Um die Umsetzbarkeit unseres Ansatzes zu demonstrieren, wurde ein Prototyp der Infrastruktur implementiert, in einer Laborstudie sowie in vier Haushalten installiert und anschließend evaluiert. Die Architektur setzt dabei auf die Interaktionsmöglichkeiten von Mobiltelefonen und die Integration von Smart Metern und dient damit als Basis für einen Großteil der weiteren Arbeit.

Im zweiten Teil der Arbeit wird das Potential des Mobiltelefons als mögliche Benutzerschnittstelle eines Systems, welches Feedback zum Stromverbrauch liefert, evaluiert. Zunächst wird in einer Benutzerstudie der wahrgenommene Wert unterschiedlicher Feedback-Funktionen untersucht und es wird analysiert, welche Art von Feedback für Nutzer relevant ist. Des Weiteren werden die allgemeine Benutzbarkeit und Genauigkeit sowie die Nutzungsabsicht der entwickelten Benutzerschnittstelle untersucht. Durch den Einsatz des Systems in einer Langzeitstudie in vier schweizer Haushalten charakterisieren wir unterschiedliche Benutzertypen und zeigen, dass weiterführende Konzepte aus dem Gebiet der Verhaltensforschung notwendig sind, um eine dauerhafte Verwendung des Feedback-Systems, die über die Phase der ersten Neugier hinausgeht, zu garantieren. Insgesamt bestätigen die Resultate die Eignung des Mobiltelefons als Benutzerinterface für Feedback zum Energieverbrauch und liefern wichtige Einblicke für den

Aufbau zukünftiger Energiesparapplikationen. Sie zeigen, dass ein klarer und einleuchtender Anwendungszweck entscheidend für deren Verwendung ist und dass Funktionen, welche das Wissen seitens der Benutzer erhöhen, gemeinsam mit Funktionen, von denen sich direkte monetäre Einsparungen ableiten lassen, als am wichtigsten eingeschätzt werden. Um auch weniger technisch versierte Anwender anzusprechen, ist es wichtig, Funktionen zu integrieren, welche über die Visualisierung der reinen Verbrauchsinformation hinausgehen und direkt handlungsleitende Maßnahmen bereitstellen.

Schließlich wird im dritten Teil der Arbeit eine Methode vorgestellt, welche es ermöglicht, den gemessenen Gesamtstromverbrauch eines Haushalts auf Geräteebene herunterzubrechen. Dies erlaubt es Benutzern, den Stromverbrauch einzelnen Geräten oder Handlungen direkt zuzuordnen und kann gleichzeitig als Basis für ein System dienen, welches automatisch Energiespartipps ableitet und Einsparmöglichkeiten aufgezeigt. Im Vergleich zu anderen Disaggregationsverfahren benötigt unser System keine zusätzlichen Systemkomponenten oder zeitintensive Kalibrierung durch Fachexperten. Die für die Disaggregation notwendige und sonst oft komplexe Aufzeichnung von Gerätesignaturen erfolgt im Hintergrund mit Hilfe von Ubiquitous-Computing-Technologien unter Einbezug des Benutzers. Dadurch ist es mit dem implementierten Verfahren möglich, auch Geräte zu erkennen, welche erst zu einem späteren Zeitpunkt nach der Systeminstallation im Haushalt verwendet werden. Dies ist besonders in einem sich ständig verändernden Umfeld wie dem Privathaushalt von großer Bedeutung. Eine Evaluation in einem Labortest mit bis zu acht gleichzeitig betriebenen Haushaltsgeräten bestätigt mit einer für viele Zwecke ausreichenden Erkennungsrate von fast 90% die Umsetzbarkeit des von uns entwickelten Ansatzes.

Contents

1 Introduction

1.1 Motivation

The worldwide yearly energy consumption has been steadily increasing and reached ~143000TWh in 2008, of which approximately 17% (~17000TWh) is final electricity use [8]. With a share of 39.3% in 2007, electricity generation is the leading source of carbon dioxide emissions in the U.S., and it is expected to grow stronger than any other final form of energy [9, 10]. In particular, the electricity consumption of the residential and the commercial sectors has been increasing over the last decades whereas the electricity use of the transportation and industry sectors has remained stable. In the U.S., residential and commercial electricity consumption now accounts for approximately 73% of total electricity use [11]. The residential sector alone has increased by 54% since 1991 [5], consuming 38% of final electricity use in the U.S. and 29% in the EU [11, 12].

To a large extent, this increase in residential electricity consumption can be traced back to the growing number of electrical appliances. Despite considerable efficiency gains with respect to the large and omnipresent household appliances (e.g., refrigerators, freezers, washing machines, and dishwashers), the total electricity use for household appliances in the IEA19[1] grew by 57% from 1990 to 2005 [5]. Improved living standards resulted in more households buying and using a growing pool of small appliances (especially consumer electronics such as large TVs, PCs, audio and communication devices, etc.) that are consuming significantly more electricity [4, 13-15] (see Figure 1.1). Heating, ventilation, and air conditioning (HVAC) systems are another substantial contributor to the residential energy bill. Their market penetration has almost tripled over the past 30 years. They account for 49% of the residential energy consumption in the U.S. and contribute even more significantly in Europe (e.g., 61% in the

[1] 19 out of 22 countries covered by the International Energy Agency (IEA) provided detailed insights into household energy use.

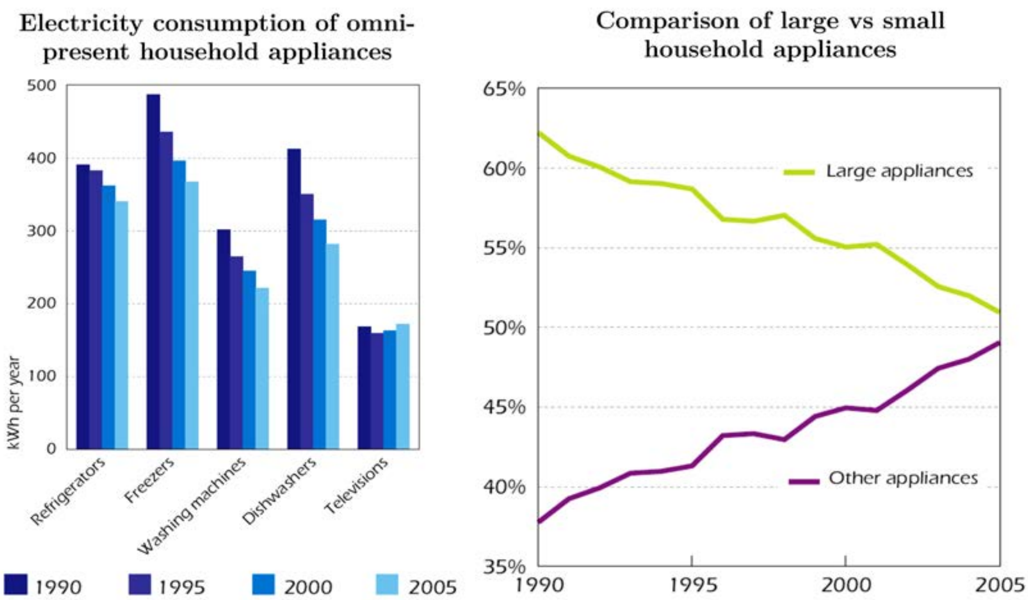

Figure 1.1 Trends in residential electricity consumption: The efficiency of large appliances has been increasing over the past decades. Except for TVs, for which the screen size is overcompensating technological advances in efficiency (left). At the same time, the increasing share in electricity consumption of small appliances starts dominating over the consumption of large appliances (right) [5].

U.K. and 70% in Switzerland) [4, 16, 17]. Despite the rising permeation, HVAC systems have matured and the share in residential energy consumption has been decreasing – in contrary to the consumption of electrical appliances whose share today is higher than ever (see Figure 1.2).

This increase in residential energy consumption along with the yet unused saving potential are only two reasons that demonstrate the need to address energy conservation in the residential domain. One major difficulty for residents who are willing to save energy at home is the lack of information about their energy consumption. Feedback on energy usage is typically only provided by a monthly (if not yearly) utility bill and thus remains rather vague and opaque to most residents. As a result, most individuals could reduce their energy consumption, but few know how much they consume and even fewer know how much energy they use for a particular purpose (e.g., lighting) [18]. And even those who do have a fair understanding of their consumption only rarely receive guidance or meaningful feedback for individual energy-saving actions that can be taken [19].

Ubiquitous and Pervasive Computing (Ubicomp), which advocates to digitally enhance physical objects that populate people's everyday life with computing, sensing, and communication capabilities, has the potential to help overcome the lack of meaningful residential electricity information and thus contribute to residential energy conservation. Continuous technical

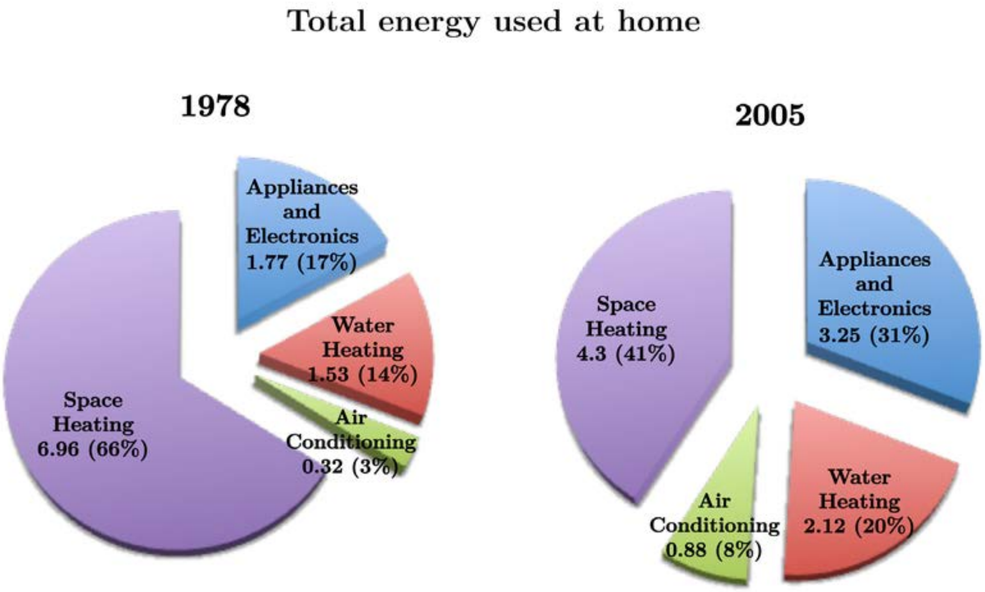

Figure 1.2 Development of the total energy used in households in quadrillion British thermal units (BTU) and percent in the U.S. While the share of electrical appliances has almost doubled, the share of HVAC systems has decreased. However, HVAC systems are still the most energy consuming entity in homes [4].

advances regarding computing power, storage, sensing capabilities, communication technologies, and size together with lower costs and energy usage have led to the rise of Ubicomp. In essence, Ubicomp enables merging the digital world of bits and bytes with the physical world. In contrast to the pre-internet era, where computers resided as standalone machines that filled rooms or people's desks, today very small and low-cost microprocessors and sensors integrate computing into every aspect of our environment [20]. Examples of these novel computing technologies range from RFID tags to tiny sensors that are integrated into appliances to embedded computing platforms (Figure 1.3). Even latest generation smartphones today incorporate a wide variety of sensing technology (e.g., a light sensor, an accelerometer, a gyroscope, a GPS module, etc.) that go far beyond what was originally necessary for communication.

Ubicomp technologies make more fine-grained information – contextual, temporal, and spatial – easily available and provide everyday objects with integrated services like Mark Weiser and others pioneers envisioned more than a decade ago [20, 21]. As such, Ubicomp has permeated various application areas across all domains [22]. For example, tagging objects with barcodes and sensors provides high-resolution data that increased transparency and has led to more efficient operation and new services [23, 24]. In retail and wholesale, RFID tags help optimize operation along supply-chains [25, 26], allow for assessing the ecological footprint [27, 28], and can be used to protect brands and trademarks [29, 30]. At home, Ubicomp technologies have created smart environments in which physical objects

Figure 1.3 Ubicomp technologies (from upper left to lower right): Tiny pressure sensors, an embedded device (Gumstix) that offers the computing power of a low-end netbook at the size of a stick of a gum, a sensor node (Java Sun spot) offering various sensing and communication means, an RFID tag that enables the tracing of retail products, Adruino open-source prototyping platform, Plogg smart power outlet, and smartphones that incorporate a wide variety of sensors (light, acoustic, location, acceleration, etc.).

automatically exchange information to collaborate and to offer a range of services [22, 31-33] that aim at making users' daily life more comfortable, easier, and safer [34-37].

With its typically small size and relatively low cost, Ubicomp technologies are quite suitable for applications in the residential environment. It unobtrusively integrates computing into physical objects used in everyday life thereby enhances the home environment with new computational and sensing capabilities. Smartphones have become ubiquitous and powerful computing platforms that allow visualizing data on the spot. Smart meters that enable capturing fine-grained electricity consumption information at high frequency are currently replacing traditional electricity meters, and recent home appliances feature built-in wireless communication modules. Through the integration of such communication and computing capabilities in everyday objects that populate homes, Ubicomp is offering new possibilities to address the problem of residential energy conservation.

Conserving energy in residential environments comprises two different approaches: The use of automated savings wherever possible as well as the provisioning of behavior-changing information where necessary [38]. Ubicomp technologies can help contribute to both. They enable automated energy savings as well as support users by supplying meaningful information that helps alter user behavior when purchasing and using the appliances (see Figure 1.4). Automated savings result from the coordinated control that helps to adapt to available resources and optimize consump-

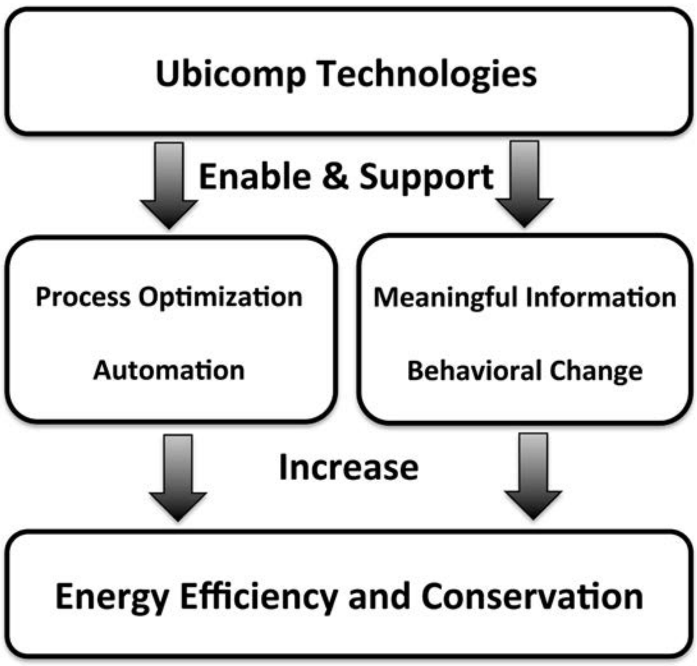

Figure 1.4 How can Ubicomp technologies potentially help increase energy efficiency and conservation?

tion [39, 40]. For example, information on location and velocity available from sensors integrated in mobile phones combined with the electricity data gathered by electricity meters can be used to automatically adjust heating to home occupancy and user preferences [41-43]. Automation can further enable the use of renewables whenever available, the immediate reaction to electricity pricing signals, and the integration of electric vehicles into the smart grid [38, 44-52].

When automation is not a possible solution, Ubicomp helps to achieve savings by bringing consumers "in the loop". Currently, most inhabitants lack meaningful information that would allow for taking effective measures to conserve electricity [53-55]. However, informed users, who understand the origin and the impact of their energy consumption, may act very differently [56]. User-induced saving effects mainly result from two factors: First, the energy usage of many appliances and systems is highly dependent on how we operate them. Without the required information, energy consumption in identical homes can easily differ by a factor of two or more, depending on the inhabitants' behavior [57]. Second, the decision to invest in efficient appliances and energy saving technologies is up to the user. Therefore, awareness and willingness to take action are crucial and can only be achieved with adequate information at hand. In addition, this can help to realize additional savings since users that often deal with their consumption are more likely to pay a premium for other energy-efficient products (e.g., an electric vehicle or solar cells) and services, and their willingness to spread the word can help to positively influence decisions of others[2] [58].

The recent move towards the smart grid is one important building block where Ubicomp can utilize its capabilities to contribute to residential energy efficiency. The requirement to conserve energy, the necessary modernization of the electrical grid infrastructure, and the integration of more sustainable energy generation (e.g., residential installations of solar and wind power) have started paradigm shifts in the energy domain that offer new opportunities for energy conservation [38]. As one consequence, smart electricity meters are being widely deployed (they are becoming mandatory for many households in Europe and the U.S.) and soon are going to generate enormous amounts of electricity consumption data. Together with the broad availability of technologies from the Ubicomp domain (in particular cheap sensors, low-power processors, wireless proximity communication, spontaneous networking, touch screen displays, mobile Internet connectivity, etc.), it becomes possible to leverage this information to provide real-time energy consumption feedback on the spot and further use the generated data to automatically optimize efficiency. Such energy consumption feedback systems, at least if designed appropriately, offer the potential to contribute considerably to a lower residential energy consumption (see Figure 1.5).

Despite all new opportunities, however, energy consumption feedback systems can easily become complex. To foster adoption, they have to seamlessly integrate into users' daily life and work reliably. At the same time, energy consumption feedback systems have to feature a low usage barrier to be utilized persistently, provide meaningful electricity feedback that goes beyond mere consumption values, and should be designed in a way that motivates users to engage themselves with their consumption over longer time periods. Furthermore, different users (e.g., technophiles compared to technophobes) essentially require different feedback [59]. Thus, the effectiveness of residential electricity feedback not only depends on the modality of the presented information, but also on the interaction capabilities and the functionality of a suited user interface. For that, different components (e.g., embedded systems, low-cost sensors, portable interfaces, and mobile phones) have to communicate amongst each other to gather and store data as well as present the information at a suitable user interface. Gathering and communication data in the residential environment also introduces a potential threat to privacy and security [38, 51, 52]. All of this raises inherent challenges on how to design energy consumption

[2] This indirect effect of one person's attitude on others is also known as spillover effect.

Figure 1.5 Traditional Ferraris electricity meter with limited consumption feedback capabilities (left) and Ubicomp-enabled consumption feedback on a mobile phone (right).

feedback systems without jeopardizing the ultimate goal of conserving residential energy.

In this thesis, we address some of these above-mentioned challenges. For that, we investigate how Ubicomp technologies such as embedded systems, low cost sensors, portable interfaces, and mobile phones can be applied as a means to foster residential energy conservation. To this end, we propose eMeter, a distributed electricity sensing and monitoring infrastructure based on Ubicomp components that seamlessly integrates into the residential environment and people's daily life. It is designed so that it can help to foster energy efficiency through both user-induced energy saving effects through meaningful information while at the same time opening new possibilities for automatic energy conservation. By testing the system in a user study and a real-world deployment, we confirm the suitability of mobile phones as energy feedback devices and the system's potential to disaggregate the overall electricity consumption to device-level information.

1.2 Objective and Approach

As outlined in the previous section, there exist various challenges that have to be addressed to contribute to the broader goal of long-term energy conservation at home. At first, only limited information was available and research within the domain focused on ways to provide this consumption information. Stand-alone consumption feedback is helpful, but only offers limited potential for energy conservation if users are not taken into account from the beginning of the design process of energy monitoring systems. To foster adoption of such systems, the feedback has to be easily available on a suitable interface with engaging interaction capabilities. And at least as important, energy monitoring systems have to provide meaningful information beyond mere usage data to support users' efforts to derive direct conservation measures.

In this work, we aim to address the need for meaningful electricity information through applying Ubicomp technologies to provide effective feedback beyond intangible consumption values. To do so, we pursue the following three steps: We first identify how Ubicomp technologies can help to create an electricity feedback system that drives adoption through a low usage barrier and helps users better understand where their energy use is occurring. Next, we seek to identify which information and functionality users require, and answer the question whether mobile phones are a suitable interface when it comes to electricity monitoring. Last, we look how to further automatically leverage the developed infrastructure and the data gathered by developing algorithms that disaggregate the overall electricity consumption to device level.

To achieve the goals of this thesis, our approach is twofold: We develop a prototypical implementation of the proposed system and together with that provide the infrastructure perspective. In addition, we pursue the human computer interaction perspective by comprising a user study and a real-world deployment using our system to obtain parts of the presented results. However, our prototypical implementation was not designed for energy-efficient and secure data communication nor trimmed to specifically use the least power-drawing hardware resources as possible to conserve energy (e.g., lowering CPU frequency or cutting Wi-Fi power when not in use). Nevertheless, all these are important aspects for larger trials that might follow up this work and as such remain future system optimization tasks. In addition, the integration of web services into the home, as provided by our prototype, raises the threat of privacy and security issues, which have to be taken seriously, but are beyond the context of this work.

This thesis was conducted in the Bits to Energy Lab of ETH Zurich and the University of Sankt Gallen and partially at the Massachusetts Institute of Technology. The Bits to Energy Lab conducts research on consumption feedback, customer engagement, and data analytics in strong collaboration with industry partners. As such, this thesis was supported by different companies form the telecommunication, metering, and utility sector.

Overall, this work is part of the more general theme of how to achieve sustainable long-term energy savings in residential environments. A future research goal would thus be the design of an experiment that aims at quantifying short-term and long-term energy savings and identifying what further engagement strategies are necessary to induce a pro-environmental behavioral change of users. However, this is outside the scope of this work, which focuses on identifying and providing meaningful electricity feedback based on an applicable Ubicomp infrastructure.

1.3 Contributions

In the previous sections, we outlined the importance of saving energy in the residential sector and highlighted some of the opportunities and inherent challenges that arise when applying Ubicomp technologies to the domain of residential energy management. In this section, we outline the three main contributions of this thesis, whose main goal is using Ubicomp technologies to leverage energy conservation in residential environments through meaningful information while providing a framework for further automated energy savings. In particular, the contributions are the following.

1.3.1 A Ubicomp Infrastructure for Providing Meaningful Electricity Consumption Feedback in Residential Environments

This thesis deals with the design, development, and evaluation of an energy sensing and feedback infrastructure. We focus on developing a system that features a low usage barrier and thus fosters the application of the resulting energy consumption feedback. For that, we investigate the potential of a diversity of Ubicomp technologies that can be applied to gather electricity consumption data and visualize feedback. As a result, we propose the eMeter infrastructure that fulfills the most important feedback criteria (e.g., in real-time, at hand when needed, allowing for a break-down of the overall energy consumption) and is unobtrusively integrated into users' daily life. At the same time, the system is easily extendible and serves as an enabling framework for researchers who investigate further aspects of residential energy conservation (e.g., to realize further automated energy savings or for the design of behavioral science experiments). The system consists of three components:

- A smart meter to acquire the electricity consumption data that compared to classical meters contains a communication interface for meter readings,

- a gateway implemented on an embedded device that takes care of data handling, storage, and processing, and
- multiple user interfaces on different devices that go beyond mere consumption visualization and additionally enable users to interact with the feedback system and the electricity usage of their appliances.

The eMeter system is complemented by a central community-based electricity feedback platform called PowerPedia. It enables users to compare the consumption of their appliances with that of others and collaboratively share energy saving tips. PowerPedia thus provides behavior-influencing feedback beyond dry numbers and intangible units that helps users to better understand the electricity consumption of individual appliances and take effective action to save electricity.

To validate the design, we present a prototypical implementation of the architecture and several different demonstrator interfaces. As an architectural style, we use the REST paradigm [60], which allows us to address the technical requirements the architecture should fulfill. In particular, it offers a lightweight access to data, which is important for resource-constrained devices (e.g., mobile phones, wireless sensor nodes, and embedded devices). In addition, relying on REST makes the system easily extendible by loosely coupling the individual components. When developing the system, this is important because it simplifies the use of different sensors for data acquisition, the development of different visualizations, and the connection of different user interfaces.

To demonstrate the feasibility of our approach, we deployed the complete infrastructure and evaluated it in a laboratory setting as well as in four private households in Switzerland. The architecture supports the interaction capabilities of mobile phones and the integration of smart electricity meters. The system thus serves as the base for the following two contributions of this thesis.

1.3.2 Evaluation of the Suitability of Mobile Phones as Electricity Feedback Devices

We evaluate the suitability of mobile phones as electricity consumption feedback devices in a user study and in a real-world deployment. Based on the implemented infrastructure, we have specifically developed a mobile phone application for residential energy monitoring. In an iterative approach, we used paper prototyping, a user survey, and a focus group to step-wise implement and refine the mobile phone user interface. It is designed to supply users with meaningful electricity consumption feedback that goes beyond the mere visualization of consumption values and aims at bringing users into the loop. The final design of the user interface was then used in a user study with 25 participants. In the study, we analyzed the perceived value of various feedback functionalities as well as the general usability, accuracy, intention of use, and word of mouth of such an electricity feedback application. The mobile phone application was also evaluated in a long-term real-world deployment in four households in Switzerland. The experiment aimed at characterizing different user types (e.g., number of sessions, average time spent per view) and qualitative insights gathered through the use of the application. In particular, this contribution confirms the suitability of mobile phones as energy consumption feedback devices.

Our evaluation shows that:

- Knowledge-increasing functionalities as well as those functionalities from which monetary savings can be directly implied are perceived as most important;
- to address technophobe users, action-guiding feedback that goes beyond displaying aggregated information is required;
- in comparison to different commercially available energy monitoring devices, the implemented measurement feature enables users to interactively determine the electricity consumption of an appliance and is easy to use, comfortable, and sufficiently accurate;
- the mere consumption value of an appliance is not meaningful enough to let users draw effective energy conservation actions (e.g., classify whether the consumption of a device is high or low for that particular class of appliance) and thus has to be enriched. This has led to the implementation of PowerPedia – a community-based electricity usage and saving tips platform for appliances that embeds residential electricity feedback in bigger, more tangible picture.
- the user interface implemented on a mobile phone increases users' electricity awareness and literacy, but different engagement strategies have to be implemented to permanently involve users once their initial curiosity has been satisfied.

1.3.3 An Algorithm for Automatically Disaggregating the Total Electrical Load to Device Level

To allow users to link the operation of their home appliances to the individual sources of electricity consumption, we developed and evaluated the potential of an algorithm that detects switching events of individual appliances. Based on the first contribution of this thesis, the eMeter electricity feedback system, the algorithm exploits the provided interaction capabilities and offers a user-friendly incremental way to facilitate load disaggregation. That is, a measurement feature on the user interface is used to acquire the required input knowledge, which enables the automatic recognition of load patterns of devices. The algorithm thus disaggregates the total electricity consumption and can serve as a basis for an automated recommendation system on how to save electricity (e.g., this allows deriving information about the device's energy efficiency and about conservation measures that can be applied).

Compared to existing single sensor approaches, our system addresses open challenges (e.g., the recognition of smaller loads and overlapping on/off events of multiple appliances) and offers additional advantages in terms of usability. Rather than discouraging users through a time-intensive calibration or a long training period after the initial system deployment, the proposed algorithm does not require a signature of every appliance in advance, but can grow its signature database over time. In addition, the signatures can be gathered in a simple, explorative way with the developed user interface, and the signature acquisition does not require specific knowledge of domain experts nor additional custom hardware. This incremental approach brings along another crucial benefit, which is particularly important in a fast changing home environment. It allows for easy integration of new appliances that are introduced at home. Where other systems need to completely recalibrate, our algorithm is able to incrementally integrate new signatures of newly bought devices.

In particular our contributions are as follows:

- The design and prototypical development of an algorithm that recognizes switching events of appliances and thus facilitates load disaggregation,
- the provision of a proof of concept implementation, and
- an evaluation of the algorithm through experiments in a laboratory study with eight simultaneously operating devices.

The results of our laboratory study are promising and confirm the suitability of the scheme. With a recognition rate of 87% in the laboratory environment, interesting applications, such as automatic recommendations for a more economic use of electricity in households, become possible. Furthermore, the information about the operation of appliances can be used to realize automated energy savings on top of the eMeter system (e.g., through inferring home occupancy state and driving a smart heating control strategy).

1.4 Outline

The remainder of the thesis is structured as follows: In Chapter 2, we propose the eMeter infrastructure for providing meaningful electricity feedback in residential environments. It is based on Ubicomp components and addresses some of the shortcomings of existing systems that are discussed and categorized first. We next focus on the functional requirements that are necessary to deliver the feedback features. We then present the system architecture and design in detail. That is, after providing a general overview, we elaborate on each of the three components that take care of data acquisition, data processing, and data visualization. We conclude this chapter with a section on the implementation details of the developed prototypes and their evaluation in a lab and real-world setting.

Chapter 3 deals with the evaluation of mobile phones as energy consumption feedback devices. We first describe the process that led to the final design of the user interface, before we explain the experimental setting. We conducted a user study in a laboratory environment with 25 participants and a real-world deployment with four households to evaluate our prototype and confirm its suitability as an electricity feedback device. We then present the quantitative and qualitative results of the evaluation followed by a general discussion on the perceived value of different feedback functionalities. As a direct outcome of the discussion, we complemented the eMeter system with PowerPedia – a community-based electricity usage and saving tips platform for appliances. Its functionality and integration to the eMeter system is described, followed by a summary of the chapter's key results.

In Chapter 4 we show how the developed architecture can be used to detect switching events of appliances and thus facilitate automated load disaggregation. After revisiting the architecture to highlight the relation of different physical quantities that are gathered by the system, we explain how load signatures can be used to classify appliances. We next introduce the key concept of the AppliSense load disaggregation algorithm and provide details on the algorithm design. We then report selected results of the algorithm evaluation that were obtained in a laboratory study. We conclude this chapter with a discussion of the results and existing limitations of the proposed disaggregation scheme.

Finally, in Chapter 5 we provide a summary of the contributions of this thesis and discuss future directions of research that are enabled by this work.

2 eMeter – A Pervasive Electricity Consumption Feedback System

In this chapter, we describe the first contribution of this thesis: eMeter – a pervasive electricity consumption feedback system. After an introduction on electricity consumption feedback systems in general, we outline limitations in terms of installation cost, usage barrier, and visualization capabilities of existing approaches. From that, we derive the functional requirements for our eMeter electricity sensing and feedback infrastructure. We then take an in-depth look on the single components of the architecture, their provided functionality, and the communication flow. That is, the data acquisition layer that is responsible for gathering the data, the data processing layer that analyzes and stores the electricity consumption data and handles incoming requests from the user interface, and the data presentation layer that takes care of interactively visualizing the consumption information. Last, we conclude this chapter with implementation details on the developed prototypes and results from our laboratory and real-world deployment that demonstrate the feasibility of our approach. Parts of this chapter were published in [1, 38, 61-64].

2.1 Electricity Consumption Feedback Systems

The existence of many unnecessary electricity loads can be attributed to a lack of transparency in energy consumption [54]. This leads, at least partly, to lost saving potentials, because users lack knowledge about their energy consumption in general as well as about the pool of devices used at home [18]. Electricity feedback systems aim at closing this gap.

Research on energy and electricity feedback dates back to the 1970s. After the energy crisis of 1973 and the oil price shock research focused on energy conservation. To address the concern of a possible exhaustion of fossil fuel, research identified two possible solutions: the development of new sources of energy and the reduction of energy use through feedback [1, 65]. More precisely, to maximize feedback researchers already envisaged an online meter that continuously reports the amount of energy used per day and month [65]. However, technical possibilities back then were rather limited. Kohlenberg et al. [66] were one of the first to design a very simple feedback device on electricity use. It consisted of a current-sensitive relay that used a dedicated light bulb to inform users when 90% of the previously recorded peak levels were reached. Their work provided first insights on electricity consumption feedback such as users started recognizing the high electricity consumption of a kettle when boiling water. Follow-up research provided first evidence that feedback can lead to energy reduction. Associating electricity consumption with cost (i.e., cents per hour) showed an electricity conservation effect and led to an average reduction of 12% in electricity use [67, 68].

Since then technological progress enabled more sophisticated forms of feedback. Its effectiveness was intensively studied and an excellent overview is provided by [53, 69-72]. The authors come to the conclusion that direct feedback can enable savings in the range of 5%-15%. However, these potential savings are hard to quantify and depend on various variables in the experimental setting (e.g., weather, time of year, observation duration, recruited participants, etc.), which in many cases remained to a certain extent uncontrolled or were in favor of the conducted experiments. For that reason, real savings through pure electricity consumption feedback reside more likely in a range between 1% and 3% [38, 73, 74].

Feedback has been shown to be one of the most effective strategies in reducing electricity usage in the home [75]. With the advent of low-cost sensing technologies, fast computation, and advances in machine learning, we now have the potential to provide electricity consumption feedback in real time for a variety of consumption activities [76]. In the following, we focus on the technological perspective of systems that enable electricity consumption feedback. Several electricity monitoring solutions already exist that can provide such feedback [77]. They aim to help users understand where energy wastage occurs and thus try to establish a basis for conscious energy usage. These electricity feedback solutions can broadly be classified into two categories according to the number (and type) of sensors used to acquire the electricity consumption information (see Figure 2.1)

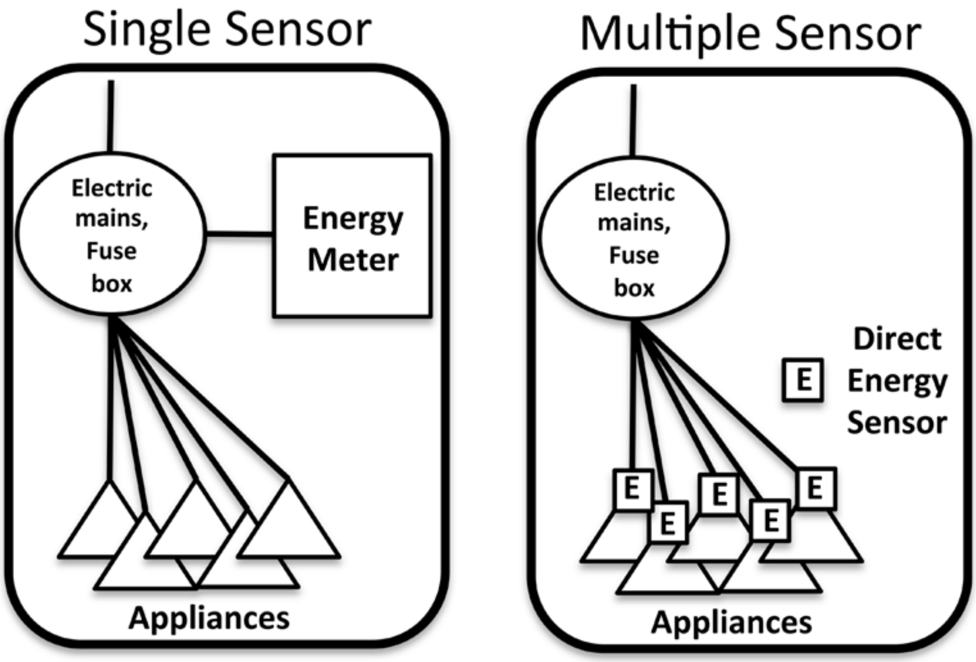

Figure 2.1 Overview of the two systematically different approaches to provide electricity feedback in residential environments.

2.1.1 Single Sensor Approaches

The first category consists of single sensor solutions, which are primarily limited to displaying the aggregated consumption of a circuit or even the entire electricity demand of a household. There are several products commercially available, such as Onzo[3], Current Cost[4], Power Cost Monitor[5], Wattson[6], and TED-1000[7] only to name a few. Once installed, they visualize the overall electricity consumption on central a display unit (see Figure 2.2). The number of kilowatt-hours and cost equivalents often enrich the provided feedback. However, installation at circuit or household level is complex. Typically these solutions as well as other scientific work (e.g., that focus on design aspects) [78-80] require a current clamp to be attached around the internal wiring of the electric mains (see Figure 2.2) and users are therefore often discouraged from installing such products. Furthermore, these solutions suffer from the fact that, mainly for safety rea-

[3] Onzo Ltd., www.onzo.co.uk
[4] Current Cost, www.currentcost.com
[5] Blue Line Innovations Inc., www.bluelineinnovations.com
[6] DIY Kyoto, www.diykyoto.com/uk/wattson/about
[7] Energy Inc., www.theenergydetective.com

Figure 2.2 Single sensor solutions for providing electricity feedback at home. The Onzo smart energy kit consists of an external display unit and a current clamp that has to be installed at the main electricity supply (left side). The Wattson electricity monitor (upper right) and the Blueline Power Cost Monitor (lower right) require a likewise setup, but differentiate in the design. Source: Manufacturer (see footnote previous page).

sons, the wiring around household meters is inaccessible in many countries and modifications require a technician. Another drawback is that they are unsuitable for providing users with feedback on the consumption of individual devices, which, from a feedback perspective, would be necessary to draw conclusions on how consumption and behavior relate to each other [53]. As a consequence of the latter, trials show that 50% of the energy monitors are no longer used once the battery is depleted, indicating that meanwhile users lost interest and that these overall electricity feedback solutions are not capable motivating users for a longer time period [81].

Some experimental systems attempt to disaggregate the total consumption measured by a single sensor to provide more specific information about electricity consumption of individual devices [82]. The aim of these non-intrusive load monitoring systems is to keep equipment costs and installation effort to a minimum, but still provide detailed energy usage data. To determine which appliances are currently running, some of these systems simply measure the overall power difference from one point in time to the next; a principle that has been investigated by several researchers in the past [83-86]. To infer device-level consumption information from a single sensor, other more sophisticated approaches use statistical signature analysis and pattern detection algorithms to identify what appliances are currently operated from the current and voltage wave forms [87] or listen for unique noise changes on the power line that are caused by the abrupt switching of devices [88, 89]. To achieve disaggregation, these systems require either a priori knowledge about the household devices and their electrical characteristics, or entail a complex calibration and training phase involving the user, in which the system learns about specific device characteristics. However, a priori knowledge is difficult to obtain in a world of fast-changing small appliances, and manual training is a significant barrier to usage. Furthermore, appliances with varying power consumption that overlaps with the one of other devices pose a particular challenge for disaggregation algorithms [90].

2.1.2 Multiple Sensor Approaches

Multiple sensor approaches can be subdivided into direct and indirect systems. Direct systems require an in-line sensor to be installed for every device or circuit. Indirect sensing systems use a central electricity meter together with additional context sensors to monitor energy consumption.

Direct sensing systems mostly come in the form of smart power outlets (see Figure 2.3). They are relatively easy to deploy and several products exist[8]. Once installed, they measure the attached load and display the measurement data on the unit itself or transmit it wirelessly to a remote display. However, these systems are not able to aggregate consumption from multiple sensors and combine the different data to form a comprehensive picture.

To overcome this limitation, other work has focused on developing systems that integrate multiple individual power sensors. One way is to combine the central electricity meter with several other in-line sensors that submeter individual major appliances [91]. The other way is to create an infrastructure that displaces the central electricity meter, but focuses on developing expandable systems that can integrate numerous individual sensors. Pioneer work from Lifton et al. incorporated various sensing and communication capabilities into a power outlet [92]. Similarly,

[8] For example “Kill a Watt”, www.p3international.com/products/special/P4400/P4400-CE.htm

Figure 2.3 Smart power outlets (Plogg (left) and AcMe (right) [2]) are the most common multiple sensor approaches. Devices are attached directly to the outlet, which incorporates a sensor and a communication module to transmit the measurements to a display.

Paradiso [93] developed a wirelessly-networked electricity sensor was followed up by Jiang et al. [2, 94]. They developed a system based on the Epic mote platform [95] where sensors measure power consumption directly at the outlets and communicate their readings over a wireless IPv6 network to a server that populates a central database. An extension to the system was presented in [96]. The authors propose a method for decoupling of voltage and current measurements, but still being able to obtain real, reactive, and apparent power. Other work [97] takes this approach further and integrates sensor measurements directly into a 3D visualization in Second Life. Then, in [98] a system is presented that integrates commercially available smart power sockets ("Ploggs"[9]) which communicate their measurements via Bluetooth or Zigbee. A gateway is responsible for identifying smart sockets that are within range. It also makes their functionality available as resources on the Web and provides local aggregation of device-level services (e.g., the accumulated consumption of all sockets). By providing this Web-API on top of an otherwise closed proprietary system the authors aim to open up the space to a broader community. A similar concept, but more in the context of automation and smart homes, comes from Jahn et al. [99]. The authors built a system on top of a middleware framework with an interactive user interface. It facilitates intelligent communication with heterogeneous embedded devices through an overlay P2P network. One of the application scenarios that is being enabled is energy monitoring. Reinhardt et al. have focused on building a low-cost wireless sensor for distributed power metering that can be integrated into wireless

[9] Plogg, www.plogginternational.com

communication networks of smart homes. It features a reprogrammable (but also low-level compared to the just mentioned approaches) microcontroller, which allows developers to easily deploy new algorithms [100].

Indirect sensing systems try to remedy the drawbacks of direct sensing systems by keeping intrusion into the electrical system at a minimum. While direct sensing systems all suffer from the fact that deploying a large number of electricity sensors (i.e., meters) throughout a house quickly becomes expensive, indirect systems combine one power sensor with multiple, low-cost context sensors. In [101], Kim et al. describe a system that uses a single electrical sensor to measure the entire electricity consumption of a household together with additional context sensors (such as light, sound, and electromagnetic sensors) that help to infer which appliance is currently operating from the measurable signals it emits. Within a defined set of appliances, the authors show that the system can estimate device-level power consumption within a 10% error range. However, the system's performance depends greatly on the correct calibration and placement of the distributed context sensors, which is not an easy task for the average user.

2.2 Limitations of Existing Approaches

In the last section, we had a look at different types of electricity feedback systems. We next characterize the systems in terms of installation complexity, cost, usage barrier to apply the system, calibration effort, and the capability to provide overall and device-level usage information. Table 2.1 summarizes the main advantages and disadvantages of the various electricity consumption feedback systems.

Since access to the electric mains and technical knowledge about wiring is required, single sensor systems are typically hard to deploy. However, to provide feedback they only rely on a single sensor, which makes them reasonably priced. Once setup, these systems have a low usage barrier and since the single sensor is typically installed close to the household meter or in the fuse box, overall electricity consumption is easy to monitor. However, to obtain information about the individual consumption at device level calls for more sophisticated approaches that require algorithms to be calibrated. This often involves a longer-term training period of the system or the submetering of the most important appliances for a limited time during the system initialization period. In addition, due to the wide variety of electrical devices involved, the accuracy of these systems is somewhat limited and there currently does not exist a general solution [102-104].

In contrast, direct in-line electricity monitoring systems are very accurate at device level since the electricity is measured at the device itself. However, this advantage comes at a high cost, as in principle every appliance has to be equipped with a electricity sensor. At the same time this increases the usage barrier, since most users are not willing to install a large number of sensors or smart power outlets throughout the house. Some large household appliances that account for a significant part of the residential energy consumption can simply not be equipped with a sensor since their wiring is hard to reach or behind walls (e.g., lighting, oven, etc.). Therefore such systems will typically only cover a subset of all electricity consuming devices in a household. To capture the consumption of multiple outlets together, some of these devices can communicate with a gateway that offers a data aggregation layer that is responsible for processing the consumption of multiple outlets. Another crucial issue is connectivity among the individual sensors and their integration into home networks [100]. For non-technical users, setting up a multi-hop communication network throughout their home often poses problems.

Table 2.1 Properties of different energy monitoring solutions

Characteristics	Single sensor	Multiple sensors	
		Direct in-line	Indirect
Installation	Hard	Medium	Hard
Cost	Low	High	High
Usage barrier	Low	High	High
Calibration	Hard	Easy	Hard
Device-level accuracy	Low	High	Medium
Household-level accuracy	High	Low	High

Finally, indirect systems are theoretically able to provide feedback both on overall electricity consumption and, to a certain extent, on device-level electricity usage. However, they require users to deploy various context sensors in the right places and necessitate complex calibration often through domain experts. This leads to both high costs and a high usage barrier. Moreover, such systems often require a recalibration when new devices are bought and connected to the electrical grid at home.

The way forward for electricity monitoring systems involves a scenario in which household appliances, which today have only limited capabilities, become more powerful and smart. Through the integration of small, inexpensive embedded ICT components, they would sense and transmit their current energy usage together with other status information. Within the house, appliances could communicate with each other via an established protocol (e.g., powerline, Zigbee, WLAN), although dedicated new technologies such as digitalSTROM[10], rivaling traditional domestic network technologies (BACnet, EIB, KNX, etc.), might also feature. An more general overview and detailed analysis of the capabilities of the individual protocols can be found in [34].

Moreover, the cost of integrating embedded Web servers (based on REST and IPv6 / 6LowPAN) into household appliances should in future be low. This would lead to a wide variety of application scenarios in which the smart domestic electricity meter (or a similar device) could serve as a central component for data aggregation and analysis. At the same time, embedding a Web interface into appliances would enable them to be fully integrated into the Internet. As well as the allocation of a device-specific Web page for status information, this would allow the device to be controlled and its data to be processed using the full power of Web 2.0 tools [31]. It is obvious, however, that with such possibilities we would need to pay serious attention to privacy and security issues, which is beyond the context of this work.

2.3 Feedback Requirements

We previously outlined the technical capabilities and limitations of energy monitoring systems. However, feedback has only proven to be effective if the right modality of electricity information is provided [53, 55, 105, 106]. Therefore, we discuss in the following what is necessary from a feedback perspective to get “users in the loop”.

User-induced saving effects mainly result from two factors: First, the energy demand of many loads (including heating, air conditioning, ventilation, warm water systems, driving habits, etc.) is highly dependent on how we operate them. Virtual identical households (same buildings, same number of inhabitants, identical age groups, same location etc.) can vary by factor 2.6 in energy usage [57]. Second, the decision to invest in efficient devices and energy saving technologies (including heat pumps, thermal insulation, etc.) is, within the limits set by the regulator, up to the user. Therefore, awareness and the willingness to take action are crucial. Feedback ideally has to address both short-term savings enabled through meaningful information that helps adapting to more energy-efficient behavior as well as long-term conservation measures from sensitized users who know about their impact on the environment.

---

[10] www.digitalstrom.org

Changing user behavior in the domain of residential energy conservation, however, represents a major challenge since energy saving is often regarded as inconvenient and seen as more of a necessary constraint than a key objective [107]. Nevertheless, providing consumers with feedback on the energy consumption can increase awareness, strengthen energy literacy, and motivate – at least some addressee – to change their habits and ultimately save energy. It is generally expected that through detailed and immediate feedback, some of the residential electricity consumption can be conserved [53, 69]. However, to maximize saving potentials, technology itself is not sufficient, nor is the mere visualization of consumption values in some "obscure" electrical measurement unit [106]. Feedback design should be centered on user preferences [108]. However, instead feedback on energy consumption is often presented in a rather technical and non-interactive way on somber devices that lack the ability to motivate users. According to literature, effective feedback has to feature the following [53-55, 102, 109]:

Low usage barrier. To be applicable by a large user base, systems ideally should rely on the least special purpose hardware as possible and the amount of necessary installation to setup and maintain the system should be limited. However, many commercially available energy monitoring systems require either complex installation around the central fuse box or the use of many electricity sensors. These systems typically induce a high usage barrier because the electric wiring around the circuit breaker or the fuse box is – at least across Europe – only accessible to technicians, and because equipping most appliances throughout the house with a dedicated sensor is costly and rather burdensome. However, since energy monitoring is a low involvement topic for many people, systems should be designed to allow for easy interaction.

Integrated in daily life. Integration of feedback into users' daily life is important for long-term energy conservation. If not achieved, trials have shown a decline in involvement regarding the energy conservation [38]. When using an additive battery-dependent display for electricity feedback, in 50% of all cases users do not replace the battery once it is depleted [81]. This indicates a loss of interest after the users' initial curiosity has been satisfied. Thus, since not being integrated into users' daily life, these additive displays seem not capable to motivate users for longer time periods.

Frequently, in real time, and at hand when needed. Feedback should be provided frequently and in real-time allowing for users to relate feedback to a certain behavior or device usage [56]. Continuous feedback has been proven to be most effective. The authors of [110] investigated the effects of continuous versus monthly feedback on gas usage. The results show that people confronted with continuous feedback save more (12.3%) than those who had received monthly feedback (7.7%). In addition, only feedback that is at hand when needed (e.g., directly after a particular control action) is able to satisfy users' spontaneous curiosity [111].

Break down of the entire energy consumption. Besides visualizing the overall electricity consumption in real-time, it is important to provide the possibility for an apportionment of the overall consumption [102]. A breakdown, e.g., for specific rooms, appliances, or times of the day, is a powerful way of establishing a direct link between action and effect / result [112]. This considerably improves the intensity of reflection and interpretation of a measure or omission [53] and should be ideally combined with tips for direct energy saving measures [108].

2.4 eMeter System Requirements

The before-mentioned characteristics – low usage barriers, high degree of integration in every day's life, timeliness of informational support, and allocation to specific loads – are believed to be important for energy feedback systems to be effective. Due to their importance for adoption and application of the resulting electricity feedback system, we consider these user-centric feedback requirements as essential design criteria for the development of our eMeter electricity feedback system. In the following, we explain how we took these into account in the design process of the eMeter system.

We try to achieve a low usage barrier by utilizing components that are ubiquitous in the residential environment. We first considered using smart power outlets as electricity sensors and actually started developing our early feedback prototype utilizing commercially-available power outlets as data source [98]. However, the disadvantages regarding a high usage barrier and cost led to an alternative design at the end. By using a smart electricity meter as single data source, which is going to be installed in households throughout Europe by law anyhow and a mobile phone as user interface, users' effort to setup the system is limited to a minimum. Compared to other systems that require experts or domain knowledge for setup, all that users have to do, is install a mobile phone application that can easily be downloaded from the Internet. The eMeter system does not rely on further special purpose hardware and requires no modification by users – neither to the electrical wiring, nor by deploying additional hardware at device level. The only additional component our system utilizes is based on a commercially available embedded device. It is integrated to the smart electricity meter as a clip-in module at the time of installation and thus does not bother users.

By using a mobile phone as user interface, the system features not only feedback on a device that is already part of users' life, but also the ability to provide instantaneous feedback that is at hand when needed. Other potentially suited technologies we considered as user interface were a picture frame, a TV, a web portal on a computer, and a dedicated stand-alone device. On some interface technologies, we partially started developing user interfaces to gather hands on experience that we took into account in our user interface assessment presented in Table 2.2. Picture frames are becoming more and more popular in homes and are able to address a broad user base of all age groups. However, they are typically not used on a daily basis and offer hardly any interaction capabilities. Most of the currently available devices have a fixed power supply and portability as such is limited. TVs are even less portable, but feature a high integration into daily habits and probably address the largest user base among the considered technologies (on average there exist 2.5 TVs per household [4]). Some TVs started providing interactive features, such as the integration of RSS feeds or the possibility to play simple games, which makes them more engaging than the pure display of consumption on a picture frame. Computers in contrast offer typically high interaction capabilities and integration into daily routines, but lack the possibility to provide immediate feedback to satisfy initial curiosity. This may be better achieved with mobile phones. They are highly portable and used in daily life, but suffer from the fact that up to now smartphones are typically used by mostly tech-savvy people. However, the penetration of smartphones is currently strongly rising, which should most likely result in a larger user base across different target groups as time progresses. Presenting the consumption on a dedicated in-home display can address a wide user base. However, compared to smartphones, in-home displays lack the interaction capabilities, and since not yet being integrated into users' daily life, these additional displays do not seem capable of motivating users for long periods of time [81]. Taking all the facts in to account, we decided to go with a mobile phone as user interface, which at the moment offers the highest benefits despite the smaller user base.

Table 2.2 Comparison of potential electricity feedback technologies

Characteristics	User base	Daily life integration	Interaction capabilities	Portability
Picture frame	High	Medium	Low	Medium
TV	High	High	Medium	Low
Computer	Medium	High	High	Low
Mobile phone	Medium	High	High	High
Dedicated device	High	Low	Medium	Medium

Lastly, useful feedback has to link specific actions to their effects by providing the ability to disaggregate overall electricity consumption. In order to take effective measures, it is vital to understand how much power individual devices consume in standby mode or while operating [80]. In fact, users have rather limited possibilities to investigate their household's efficiency with simple measures [110]. Although relying on a single sensor, the system as such has to provide a way to achieve a device-specific consumption breakdown with simple means. We accomplish this by enriching the user interface with an interactive measurement feature that allows for users to determine the consumption of a switchable appliance just by turning the appliance on or off.

Besides the described user-centric requirements, our system design has to meet several system-centric requirements that ensure the technical ability to deliver the feedback as well as the system's applicability in the real world. In particular, this includes modularity, expandability, and future-proofness, which we considered to a certain extent in our system design.

Modularity. When developing an electricity feedback system, it is important to be able to exchange individual components. It allows the easy integration of different user interfaces and electricity sensors, and is typically enabled through standardized interfaces. This not only improves maintainability, but also makes the system far more reusable than a traditional monolithic (closed) design.

Expandability refers to the possibility to extent the system by adding components, peripherals, or capabilities to it. For example, this allows incorporating new services or components that were not envisaged at the time of the system design. In the context of the smart grid, this ensures that residential energy feedback systems are ready to support features that are currently developed but not yet implemented.

Future-Proofness. Compared to advances in information technology, buildings face a much longer life cycle. Where technological development, for instance, replaces consumer electronics within years, buildings typically remain more or less untouched for decades. Taking this into account, it becomes evident that systems integrated in the residential environment should ideally feature future-proofness that guarantees system's operability over the lifetime of the building.

The above-listed requirements do not present a complete list. Indeed, there exist further technical, behavioral, economic, and social factors that should be ideally met to foster adoption of energy feedback solutions [113]. For example, system-centric technical requirements also include privacy, security, communication efficiency, and energy-efficiency in terms of the system's power usage. However, fulfilling all these requirements goes beyond the scope of this work and is thus not an objective of our prototypical implementation.

2.5 eMeter System Architecture

In this section, we explain the architecture of the designed, developed, and implemented residential electricity feedback system in-depth. We first provide an overview on the system design and the communication flow between the individual components, before we explain the functionality of each of the three layers, i.e., data acquisition, data processing, and data visualization in detail.

2.5.1 Overview

The eMeter system consists of three loosely coupled components that take care of data acquisition, data handling, and data visualization (see Figure 2.4). The first component is a smart electricity meter that measures the total electrical load of all attached devices in a household. It logs the total electricity consumption at a frequency of one sample per second. The utilized smart meter has an integrated communication interface that is connected to a gateway, which is responsible for continuous data acquisition processing, and storage from the electricity meter, and also for the handling of the incoming requests of the user interface. The gateway itself contains an Ethernet and Wi-Fi module for communication purposes. The

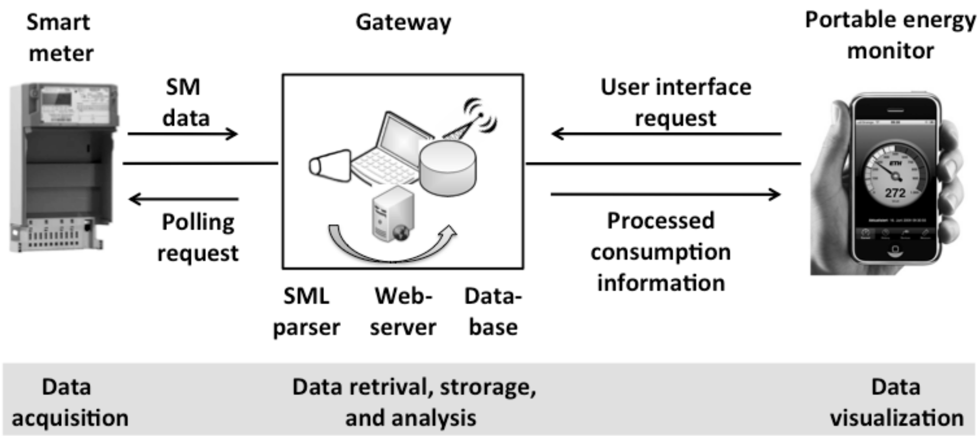

Figure 2.4 The eMeter system architecture consist of three loosely coupled components: a smart meter for data gathering, a gateway on an embedded device for date processing, and a portable energy monitor on a smartphone that visualizes the electricity consumption.

third component of the system is the user interface that is implemented as a mobile phone application for iPhone, Android, and Windows Phone 7 handsets. It provides users with portable real-time feedback on their electricity consumption and allows them to interactively explore from where their electricity demand is occurring.

The system architecture is based on the REST (Representational State Transfer) paradigm [60]. REST is a resource-oriented approach that enables physical resources to be easily and seamlessly integrated into the Web [31]. For this purpose, REST proposes two basic principles. First, transferring the conventional operation-centric model view into a data-centric view, which essentially means that services now become resources that can be identified and manipulated (i.e., transferred, indexed, put on Web pages etc.) by using URLs. Second, the only available operations to access, update, delete, and create resources are the four main operations provided by HTTP (GET, POST, DELETE, PUT).

Consequentely, the communication between the three components is realized using http over TCP/IP. The gateway provides a RESTful Web API to access its functionality and the meter's resources. Compared to proprietary protocols and to other existing communication standards in the building domain that come and go, we believe HTTP to be an established future-proof protocol that might gain greater attention in future home energy management applications [34]. Following the concept of REST also decouples the three components, which benefits the overall system design because it makes the individual entities location-independent and easily exchangeable. Thus, the system does not rely on the meter of a particular manufacturer and can support different user interfaces and communication scenarios.

In general, the proposed system design enables three different communication scenarios (e.g., all three components are operated on the same local area network), which are illustrated in Figure 2.5. First, the system can be setup to communicate within its own local area network. If done so, the gateway opens up an own wireless access point (WAP) and allows users to directly connect to its WIFI. The gateway is connected to the smart meter over the Ethernet interface and locally hosts all required data and services. Hence, all information stays local under the full control of the user. Second, the gateway can be connected to the user's home area network. This approach is beneficial in case users have already established a wireless network (e.g., operate a WIFI router). Instead of directly talking to the gateway, the communication is realized via the user's home area network. The consumption data still remains local. However, users can enable the system for remote monitoring. Then the consumption data becomes accessible from locations outside the home and thus enables users to inspect their residential electricity consumption on the user interface from any location

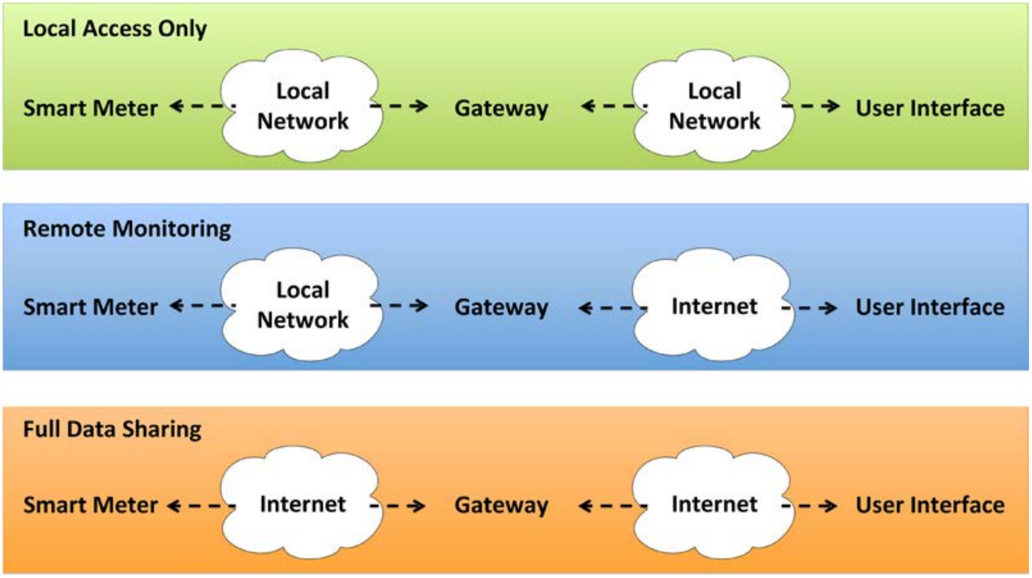

Figure 2.5 The design of the eMeter system supports three different operation modes: Local access, remote monitoring, and full data sharing. This is achieved through the loose coupling of the individual components that makes them location–independent.

that provides mobile Internet connectivity. This option essentially requires some consumption information to leave the residential network, which can potentially lead to the introduction of privacy threats. The third setup features full data sharing over the Internet. The meter itself is connected to the web, sending the data to the gateway that is located on a dedicated sever on the Internet and can be accessed from anywhere. In this scenario, all consumption data is first communicated over the Internet to the server that further processes the information and provides universal access. This guarantees full access to all information at all times, which is important when first deploying and refining the prototypical implementation and when further developing algorithms on top of the electricity consumption data.

2.5.2 Data Acquisition

The first component of the eMeter infrastructure is the smart meter (model ZMK420/E750 by Landis + Gyr). The eMeter system uses it as a single sensor to gather the electricity consumption information of private households. Compared to a classical meter, it contains an Ethernet interface for remote meter readings and advanced capabilities in terms of sampling frequency and physical quantities that are measured. Once installed it logs the electricity consumption of all residential appliances that are attached to the electric circuit of the household on a second by second basis.

For communicating, the meter uses the Smart Message guage[11] (SML) protocol, which is a non-proprietary communication protocol developed and applied by leading European utilities (e.g., RWE, E.ON, etc.). It is implemented in meters of various manufactures (e.g., Landis + Gyr, EDF Froeschl, Hager, etc.) and has been specifically developed to meet the requirements of smart grids and smart metering. Most importantly this includes the suitability for remote readings via IP-based connections.

The smart meter records the electrical information as overall consumption as well as individual values per electrical phase (i.e., two separate phases in the US and three in Europe). Table 2.3 provides an overview of the physical quantities that are directly available from the meter. Their name extensions L1, L2, or L3 indicate the phase to which the respective value corresponds (e.g., powerL1 represents the recorded power at the first electrical phase).

Table 2.3 Measurement data directly acquired from the meter.

<table>
<tr><th>Value</th><th>Quantity</th><th>Type</th><th>Unit</th></tr>
<tr><td>smartMeterId</td><td>identifier</td><td>int</td><td>ID</td></tr>
<tr><td>createdOn</td><td>timestamp</td><td>long</td><td>Date</td></tr>
<tr><td>powerL1, powerL2, powerL3</td><td>P1, P2, P3</td><td>double</td><td>Watt</td></tr>
<tr><td>currentNeutral</td><td>Ineu</td><td>double</td><td rowspan="2">Ampere</td></tr>
<tr><td>currentL1, currentL2, currentL3</td><td>Ieff1, Ieff2, Ieff3</td><td>double</td></tr>
<tr><td>voltageL1, voltageL2, voltageL3</td><td>Ueff1, Ueff2, Ueff3</td><td>double</td><td>Volt</td></tr>
<tr><td>phaseShiftVL2L1, phaseShiftVL3L1</td><td>φ21, φ31</td><td>double</td><td rowspan="2">Degree</td></tr>
<tr><td>phaseShiftCVL1, phaseShiftCVL2, phaseShiftCVL3</td><td>φ1, φ2, φ3</td><td>double</td></tr>
</table>

The attribute smartMeterId is a numerical number that identifies the smart meter, and createdOn holds the instant in time when the measurement was recorded.

PowerLx (P_x). The real power P_x on L_x is expressed in the derived SI unit Watt ($[P_x]_{SI} = W$).

CurrentNeutral (I_{neu}) relates to the current on the wire which is connected to the neutral point ($[I_{neu}]_{SI} = A$).

[11] www.t-l-z.org/docs/SML_080711_102_eng.pdf

CurrentLx (I_{effx}). The quantity I_{effx} denotes the current on circuit L_x, and is expressed in the SI unit Ampere ($[I_{effx}]_{SI} = A$). The following applies:

$$I_{effx} = \sqrt{\sum_{i=0} I_i^2}.$$

where I_0 is the DC (direct current) component, I_1 the current component of the fundamental frequency, and $I_j (j > 1)$ the current component of the j-th harmonic.

VoltageLx (U_{effx}) corresponds to the voltage over L_x, and is expressed in the SI unit Volt ($[U_{effx}]_{SI} = V$).

PhaseShiftVLyLx (φ_{yx}) denotes the phase shift between the voltage on L_x and the voltage on L_y. The unit is expressed in degree ($[\varphi_{yx}]_{DIN1301} = °$).

PhaseShiftCVLx (φ_x) relates to the phase shift between the current component of the fundamental frequency (I_1) and the voltage U_{effx} on L_x, and is expressed in degree ($[\varphi_x]_{DIN1301} = °$).

2.5.3 Data Processing

The second component of the eMeter architecture is the gateway. It manages and handles the captured data and is implemented on an embedded device. It is connected to the smart electricity meter to collect the consumption information and handles the incoming request from the user interface. For that, the gateway consists of an SML parser, a lightweight web server, a database, and auxiliary services, which are necessary for the gateway's operation and administration (see Figure 2.6). The first three are described in detail in the following.

2.5.3.1 SML Parser

The SML parser provides four basic functionalities. It automatically collects all available measurement data, preprocesses the received data into a JSON[12] object, and thereafter stores the measurement in the database or on a backup server. Last, the parser is responsible for post processing (e.g., updating tables in the database that use the inserted measurement). JSON is a lightweight data exchange format – comparable to XML, but with a smaller message body – suited for simple data structures. It can be easily

[12]JavaScript Object Notation, http://json.org

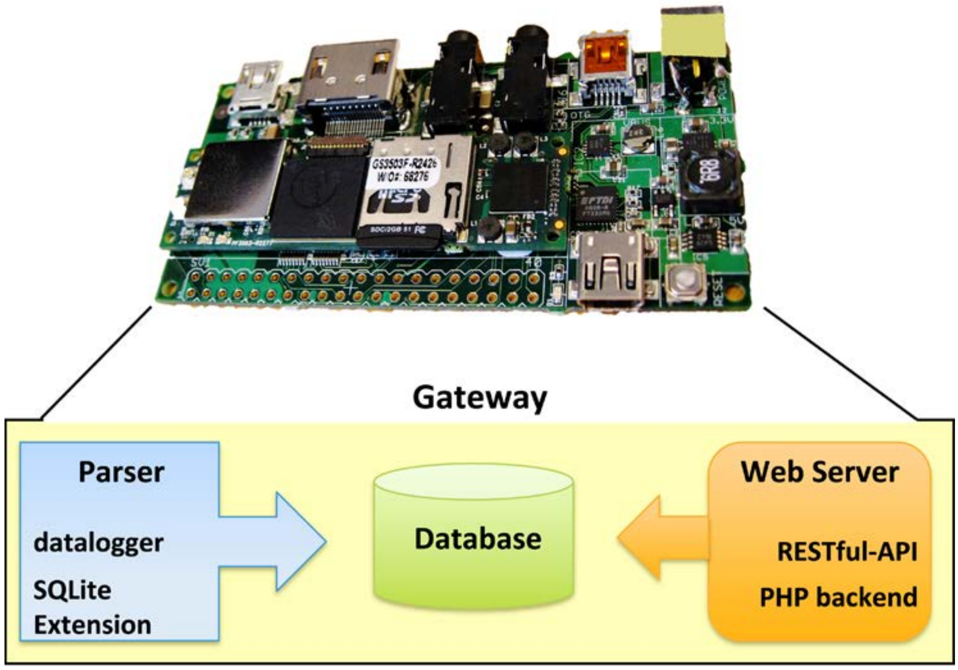

Figure 2.6 Overview of the main modules of the eMeter gateway that is programmed on the depicted development board for embedded devices. The SML parser preprocesses and automatically stores the consumption data in the database that is accessed by the web server upon requests from the user interface.

parsed by machines, but at the same time remains human-readable for developers.

The SML parser consists of two modules to acquire and handle the meter readings: the *datalogger daemon* and the *eMeter SQLite Extension.* The *datalogger daemon* first connects to the meter's Ethernet interface and establishes a TCP/IP connection. To acquire all recent measurement information, it then sends a special data packet consisting of three SML messages (PublicOpenRequest, GetProcParameterRequest, Public CloseRequest). In return, the *datalogger* receives a 580 bytes large binary SML encoded data packet that contains the requested electrical information. Using C, the response is translated into a JSON representation of the received measurement data. An example of such a binary encoded SML data packet and the respective JSON representation is illustrated in Figure 2.7. The JSON message with a size of approximately 450 bytes is then inserted in the database or transmitted to an external backup server. The code snippet for the overall power is highlighted in blue in both representations to enable a better comparison of both formats. Since the smart meter measures the electricity consumption of the household in fixed intervals, the performance of the SML parser is critical. Figure 2.8 illustrates an overview of the *datalogger* run loop that is executed to collect the measurement data. To guarantee real-time electricity consumption feed-

Figure 2.7 Comparison of a SML encoded measurement directly available form the utilized smart meter and the resulting JSON output of the SML parser [1].

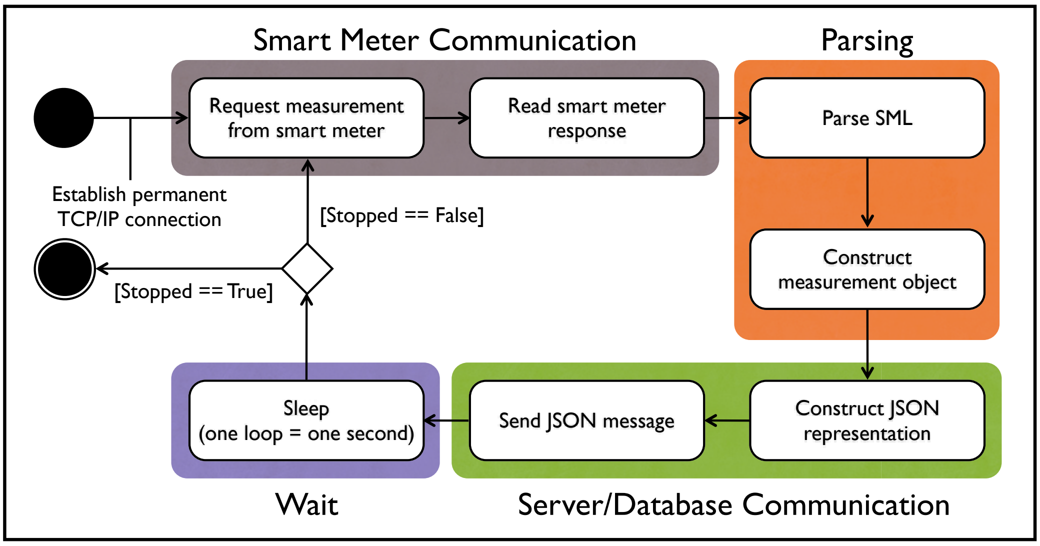

Figure 2.8 UML activity diagram of the *datalogger* run loop [1].

back the parser has to ensure that the polling request is issued in synchronized intervals with the meter logging and that the response can be processed in the time a new measurement is recorded. We implemented a timer-based approach that determines the time the parser has to fall asleep at the end of the run loop before issuing a new request. In case a measurement cannot be processed in the meantime the current loop is aborted.

The *SQLite eMeter Extension* updates the database upon arrival of a new measurement. It is used to calculate statistics of the stored historical electricity consumption data and to update cache tables that are required by the user interface. The benefit of moving this functionality into an external module – separate of the database and the *datalogger* process – is that the performed calculations and updates can be easily modified without requiring a recompilation of the complete SML parser module.

2.5.3.2 Database

The second component of the gateway is the SQLite3 database that is used to store and administrate the data. It features the Write-Ahead-Logging (WAL) journaling mode control mechanism that enables having multiple database content readers and one writer at the same time. This most importantly enables the web server to asynchronously read the database and still serve incoming requests from the user interface although a write process might be currently ongoing. The SML parser inserts the acquired measurement directly into the database whereas almost all other database related operations are moved to the dynamically loadable *eMeter SQLite3 Extension*.

2.5.3.3 Web Server

The third gateway component is the lighttp web server. In order to enable interoperability with other applications, such as the developed user interface, we use a lightweight solution based on PHP. The web server provides access to the gateway's functionalities and the smart meter's sensor values using URLs. The benefit of making the meter application accessible through a simple web API is the direct integration of the smart electricity meter to the web, which eases the development of applications and prototypes on top of the smart electricity meter.

Such an approach that makes an application's functionalities accessible through an interface respecting the core principles of the web is often referred to as RESTful. Traditionally, this type of approach is used to integrate several websites together. However, in recent research [114-116] REST is used to seamlessly connect real-world objects, embedded devices, and sensor nodes to the web. Systems using the REST paradigm and HTTP as a basic architecture for communicating with smart objects are subsumed under the term Web of Things [31].

In our case, the most interesting benefit of this approach is the seamless integration of the gateway functionality, the meter and its values to the web. In fact, it allows us to monitor the sensor values and control the gateway components simply by calling the corresponding URL in a web browser. In response to the call, the RECESS-framework[13] wraps the results in form of a JSON message. This is in contrast with the proprietary closed protocols used by most commercially available solutions that do not allow for easy integration. It is worth noting that for low-level solutions that require ultra lightweight data access, customized approaches might be better suited than REST. Using HTTP introduces an overhead when compared to ultra-optimized proprietary solutions. Thus, systems relying on a very small latency or requiring efficient communication in terms of the data that is transmitted might require a different solution than our RESTful approach. However, for us the benefit of providing lightweight data access, which especially important for constraint devices and the greater ease of interaction outweigh this drawback [117].

While the gateway can support multiple formats, we decided to use JSON (as a lightweight alternative to XML) for interaction with other applications, and HTML for providing a human readable representation in a Web browser. Figure 2.9 shows the HTML representation of the last five measurements of the currently available monitoring data that can be received in response to simply calling the following URL:

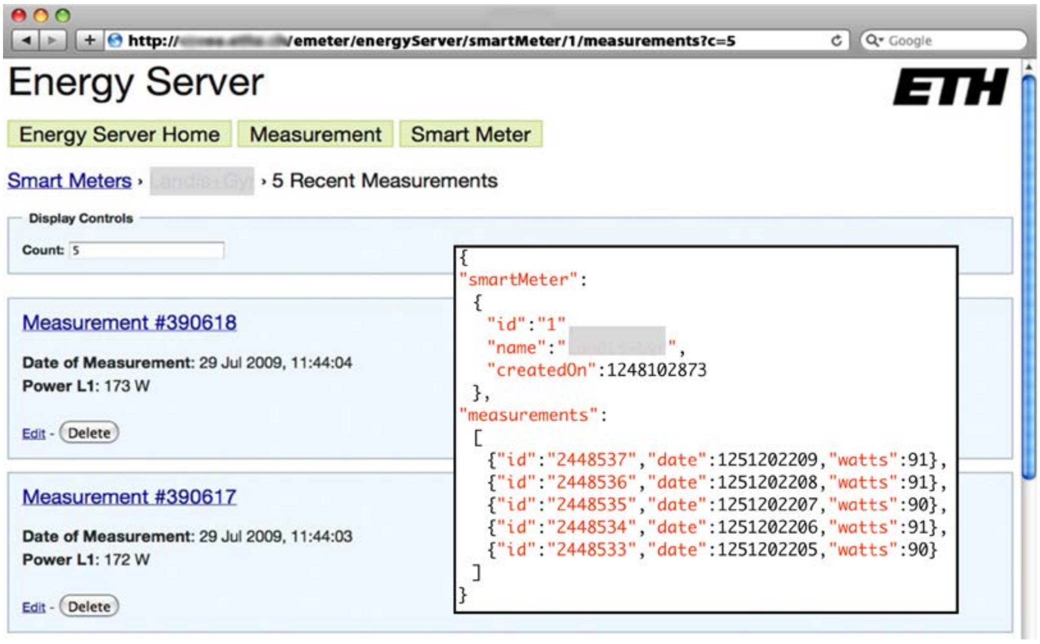

Figure 2.9 JSON and HTML representation of the GET request that obtains the last five measurements of smart meter number one from the database.

[13]www.recessframework.org

http://[serverAddress]/emeter/energyServer/smartMeter/1/measurements?c=5

The JSON-data, which is processed by the user interface, can be obtained by extending the URL with ".json" (see Figure 2.9). Thereby, the structure of the naming scheme we used to access the server and its components can be specified as follows:

http://[serverAddress]/emeter/energyServer/...
...[gatewayControl]
...smartMeter/[resourceControl]
...smartMeter/[UID]/[measurementControl]

Figure 2.10 provides an overview of the API EnergyServer. All resources can be represented in any of the following three formats: HTML, XML, or JSON. An exemplary selection of the functionality of the individual resources is provided in Table 2.4. UID refers to the unique identifier of the particular smart meter. Note that the table does not provide a complete overview of all available commands; its goal is more to provide a general idea of the actual functionality.

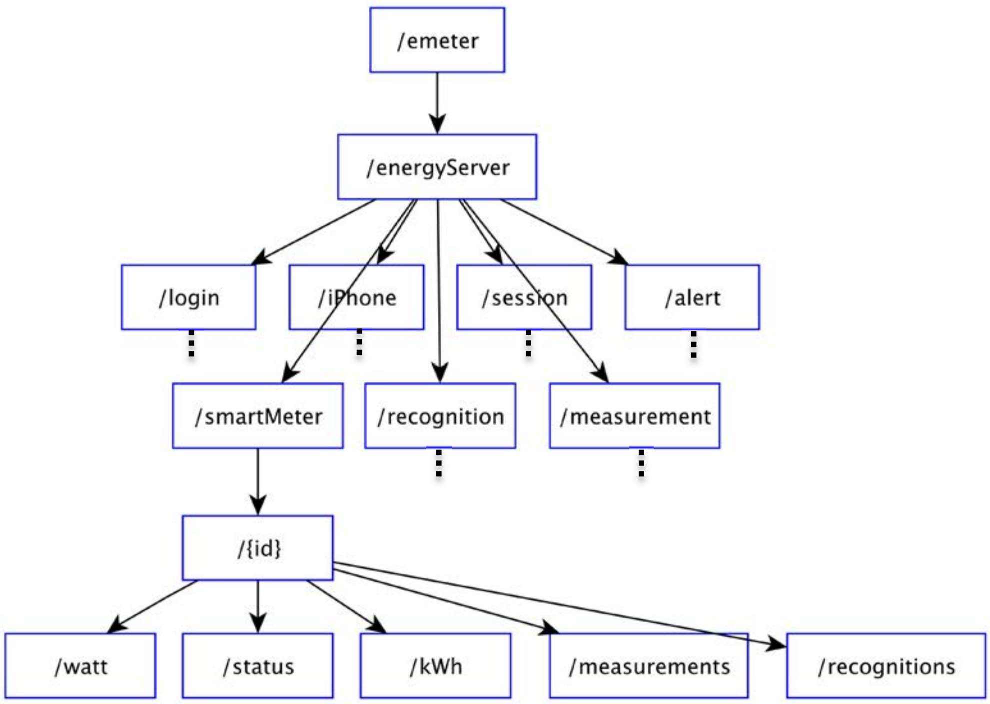

Figure 2.10 RESTful Web API of the energyServer. Table 2.4 provides the corresponding command overview.

Table 2.4 Control command overview provided by the RESTful gateway Web API.

Component	Keyword	Action
gatewayControl	status	shows the status of the gateway components
	restart	restarts the gateway
resourceControl	*	lists all smart meters
	new	creates a new resource
measurement Control	kWh?timespan	kWh per timespan
	measurements?c/*	c last/all measurements
	measurements?d/m/y/avg	measurements of the day /month/year/average

2.5.4 Data Visualization

The third component of the eMeter system is the user interface. It uses the functionality provided by the gateway to access the database and to dynamically present real-time information on the electricity consumption of the household. Besides the possibility to follow the consumption and control the gateway using any web browser, we developed a content-rich user interface as a smartphone application for different mobile platforms, i.e., iPhone OS, Android, and Windows Phone 7.

To provide the important feedback features mentioned in Section 2.3, the eMeter user interface consists of the following five main views (Figure 2.11): Live visualization of current electricity consumption (a), a historical view of electricity consumption (b, c), a device inventory view that displays energy usage and costs per measured device (d), a measurement view (e) which enables the user to interactively measure the consumption of almost any switchable electrical appliance in the house, and an engagement view (f) which aims at involving users to learn more about their electricity consumption.

The *current consumption view* (Figure 2.11a) shows current consumption in real time. The color-coded, self-learning scale – it automatically adopts colors and range – allows users to assess how their current consumption compares to their historical consumption (green to red). The blue part of the scale depicts the level of electricity base load in the home, i.e., the usage that is still active when the user is not at home. It is determined by a weighted moving average of the consumption values measured between 2am and 4am. The weight for each new value taken into account thereby corresponds inversely proportional to the difference between the new value and the previous average consumption value. This allows us to assure that atypically high consumption values during night time – which in most cases are not caused by the standby consumption – have little influence, while small changes – e.g., due to a new device attached to electric circuit – cause fast adoption of the scale.

Figure 2.11 eMeter iPhone user interface (from upper left to lower right): current consumption view, history view (aggregated consumption), history view (budgeting snapshot), device inventory view, measure view, and engagement view.

The *history view* (Figure 2.11b, 2.11c) shows a line chart of historical consumption. Users can choose between different time periods, e.g., previous hour, previous day, etc. Together with the chart, this view displays equivalents such as kWh and cost for the accumulated consumption over the last five selected periods (Figure 2.11b, lower part). The color-coded bars allow users to compare their historical consumption to that of a typical average household of the same size in the same location. To enable this comparison users have to detail characteristics of their household on the first start of the application (e.g., how many persons are living in the household, is cooking done with gas or electricity, etc.). The historical consumption view also provides a snapshot on budget calculations and projections (Figure 2.11c lower part). When tapped, a new side view is opened that helps users better understand the impact of their energy consumption (Figure 2.12 left). Users can adjust their budget for different time spans by setting a saving target using the blue slider. The view directly visualizes the impact of the savings in terms of conserved electricity, money, and CO_2. Users can also set a budget or saving target they want to achieve. If activated, an alarm reminds them once their defined budget is close to being empty or their saving target is projected to overshoot (Figure 2.12 right).

The *device inventory view* (Figure 2.11d) lists all previously measured devices. In addition, it allows users to view device details and assign a location (e.g., a room) as well as a particular utilization scheme (upon which the device's cost calculations are based) to the device. It also enables users to sort the readings by location or the amount of power used, so that the biggest energy guzzler appears at the top.

The *measurement view* (Figure 2.11e) enables users to interactively measure the electricity consumption of most switchable appliances in the household. Figure 2.13 illustrates the measurement process from a user's point of view. To perform a measurement, the user simply activates the process by pressing the start button and then turns the device being measured on or off. The corresponding result is shown on the display within seconds (Figure 2.14). The necessary calculations for this are performed on the mobile phone – as soon as the user initiates the measurement, the current consumption value determined by the smart meter is stored, and the measurement algorithm on the phone then waits for a significant change in this value. It then calculates the difference between the two values. After the measurement, the user interface provides additional options for personalization and the possibility to store the measured appliance in the device inventory.

is

Figure 2.12 Budget functionality (left) and alarm functionality (right) on the eMeter Windows Phone 7 user interface. Users can set a saving target using the blue slider to adjust the targeted limit. Conserved energy, cost, and CO_2 are automatically reflected on the right. The alarm view provides an overview on activated alerts and allows activating new and editing old ones.

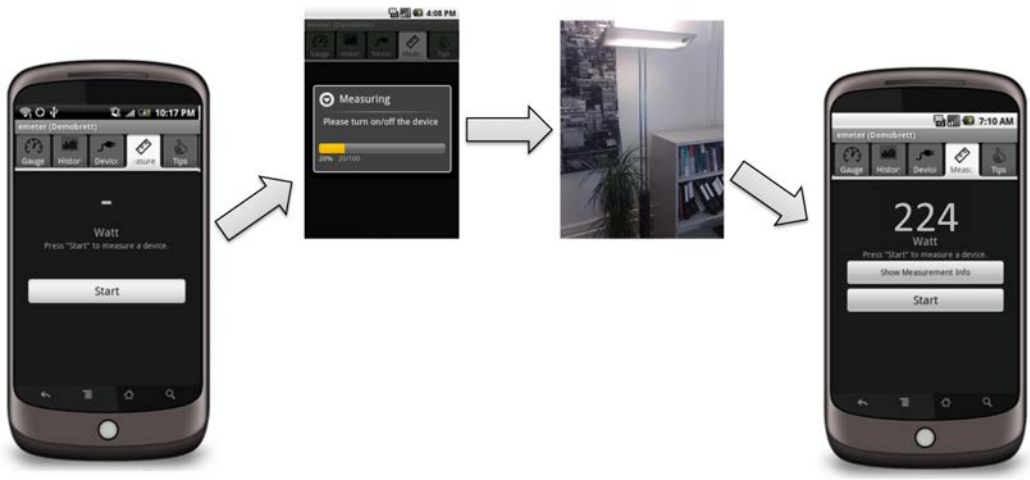

Figure 2.13 Measurement functionality on the eMeter Android user interface. Users first have to initialize the measurement by pressing the start button and thereafter turning the device of interest on/off. Within seconds the result gets displayed on the user interface.

Figure 2.14 Users measuring the power consumption of different appliances with the measurement functionality of the eMeter iPhone user interface.

larger than a predefined threshold value (dP > Threshold). The algorithm then assumes that the user has switched a device on/off, and in the following, waits for the power to settle within a particular range (dP < Settle). Once this has happened, the algorithm stores the last measurement and checks the validity of the whole measurement. In case another device is incidentally switched on or off during the measurement interval the result may be incorrect. To detect this second simultaneous event, the algorithm takes besides the increase or drop in real power the different electric circuits and additional physical variables, such as apparent power and power factor, into account. This not only allows to determine on which line the switching event has occurred, but also enables detecting in case two appliances that are attached to different lines are switched on at the same time. In that case, the algorithm detects an edge on two lines and users are prompted to repeat the measurement process. By comparing the value of the second measurement after normalization to the one of the previous measurement (compare (J1,J2)), the phase on which the switching event has occurred can be identified and the consumption of the device can be identified (compare dP(M1,M3)).

The *engagement view* (Figure 2.11f) automatically displays notifications and tips on how to conserve electricity. It aims at involving users with their electricity usage once their initial curiosity has been satisfied. Push notifications with quiz questions and competitions (e.g., find a device in your home that consumes 200W) are used to remind users of the application and their energy consumption. Further meaningful feedback is provided through general, appliances-specific, and household-specific energy conservation tips.

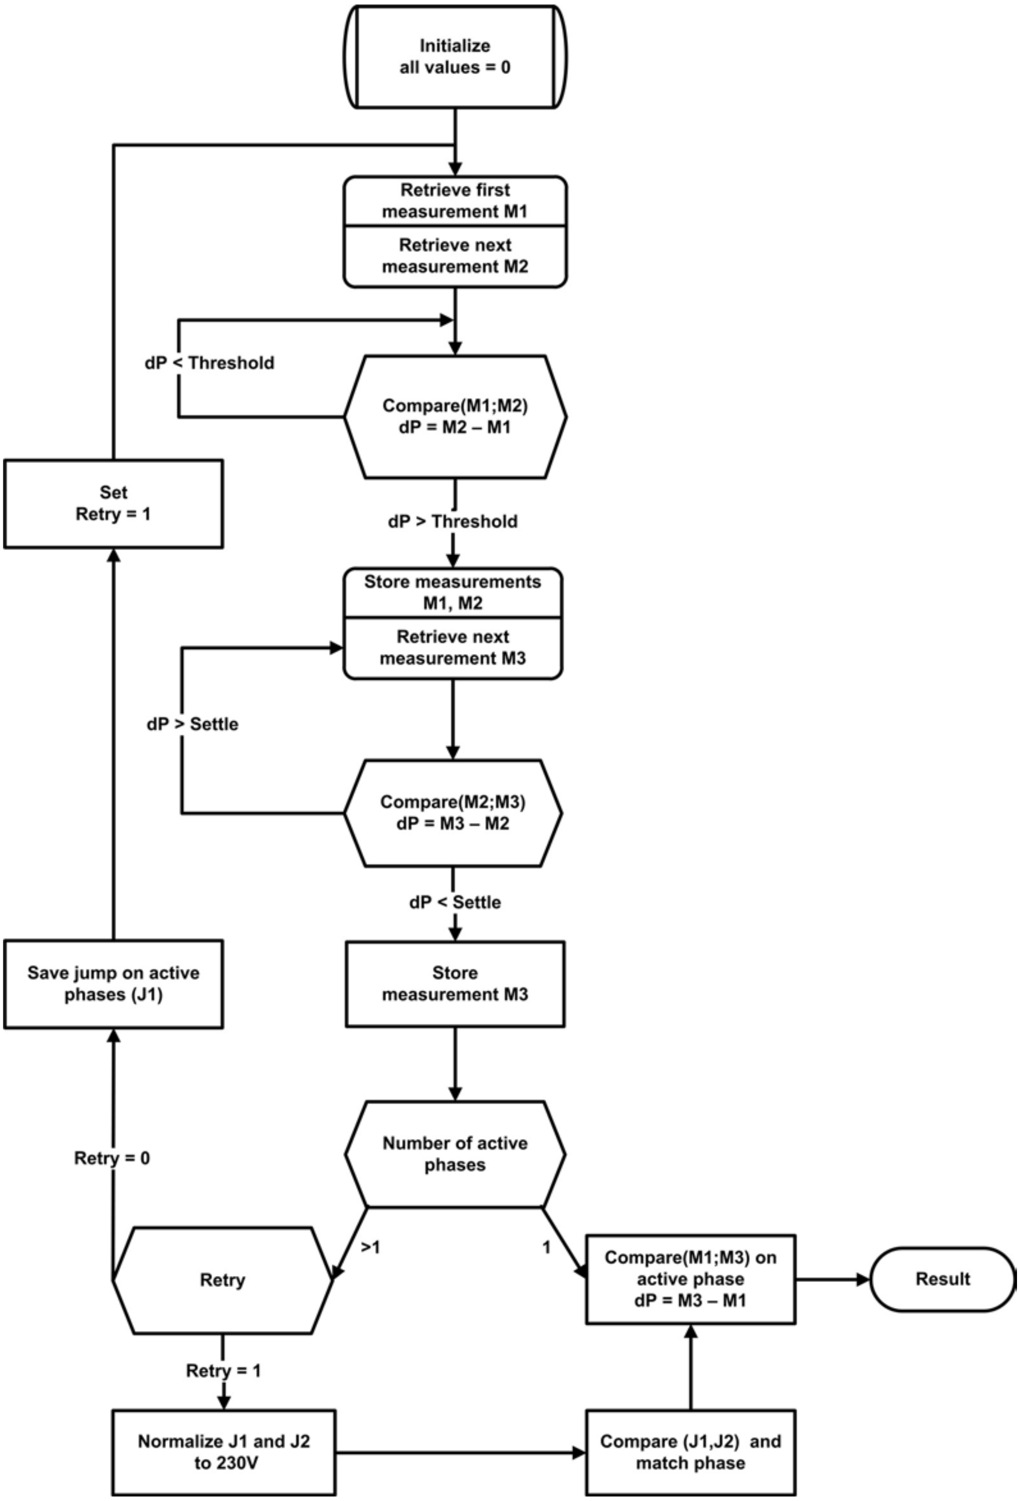

Figure 2.15 Flow diagram of the measurement algorithm that is used to determine the power consumption of individual switchable appliances.

2.6 eMeter Prototypes: Implementation Details and Evaluation

In this section, we describe implementation details of our prototypical development of the eMeter system and the laboratory setting as well as the real-world deployment we used to evaluate the full functionality of our prototype.

2.6.1 Implementation Details

A detailed view on the composition of the eMeter infrastructure is given in Figure 2.16. We use the ZMK420 smart meter developed by Landis + Gyr (Figure 2.17 right), which is the successor of the ZMK400[14] to measure the electricity consumption of the entire household. It is a SyM^2 (Synchronous Modular Meter[15]) compatible device, a standard that was introduced by a consortium of mostly German electricity suppliers. It contains an Ethernet interface for remote meter readings based on the SML protocol. SyM^2 devices use a second counter to reliably maintain the correct time reference of the current measurement even in case of multiple power outages. As such, the ZMK420 logs the electricity consumption in one-second intervals. The communication capabilities of the meter can be extended with different modules (e.g., a GSM/GPRS module for communication on mobile networks).

The second component of the infrastructure, the embedded gateway (Figure 2.17 right), is implemented on a Gumstix Overo[16], an embedded device that offers netbook-like computing power at the size of a chewing gum (Figure 2.17 left). It is based on a Texas Instruments System-on-a-Chip (SoC) architecture with a 720MHz CPU and 256MB storage of RAM and flash memory. An Ethernet and a WIFI module enable communication and interconnectivity with the other eMeter system components. Besides the local flash storage, it also contains a microSD card slot that can be used to further extent local drive space (e.g., for hosting the database). The SML parser that is implemented in C, handles the communication to the smart meter via the Ethernet interface and to external backup servers

[14] Landis+Gyr. ZMK400. http://www.landisgyr.eu/apps/products/data/pdf1/D000028192_en_E750.pdf
[15] SyM^2. http://www.vde.com/de/fnn/extras/Sym2/Seiten/default.aspx
[16] Gumstix Overo. www.gumstix.com

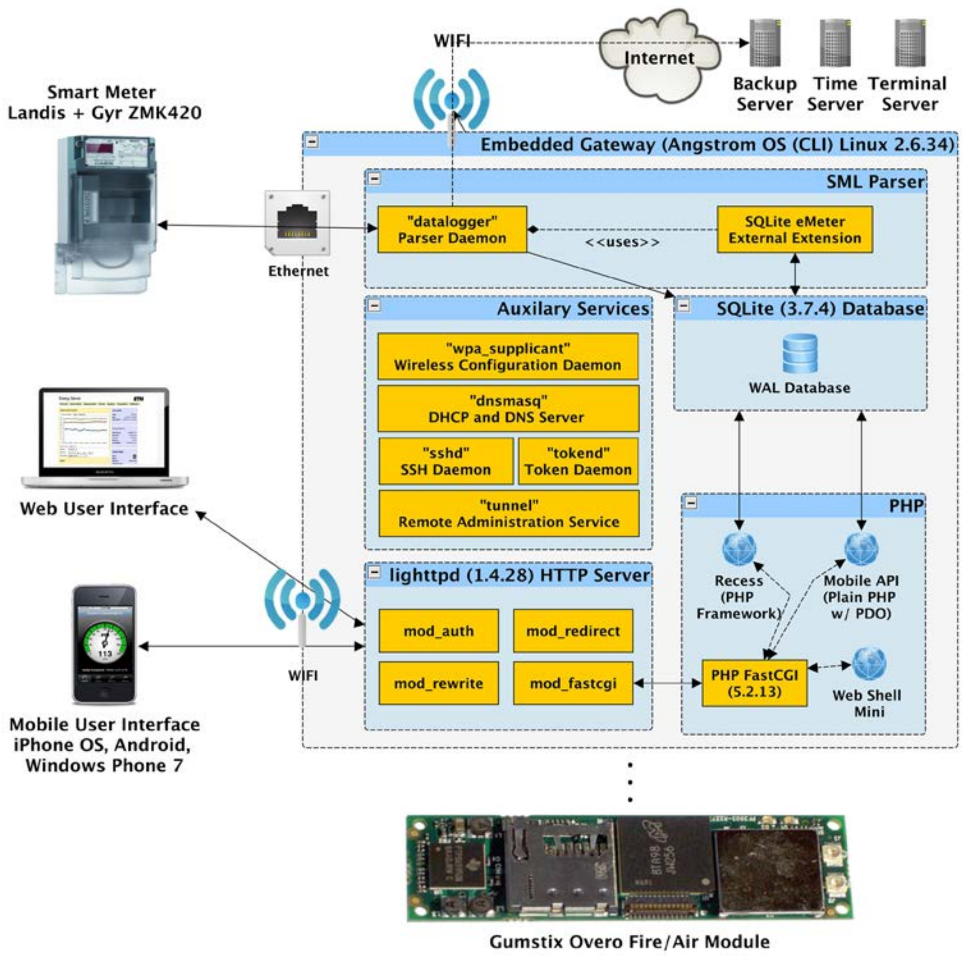

Figure 2.16 Detailed view of the eMeter gateway architecture. The SML parser handles all consumption data related processing and has write access to the database. The lighttp web server and its PHP backend serve the incoming user interface requests. Auxiliary services handle services important for the gateway operation and remote administration.

Figure 2.17 Gateway hardware: Embedded device (Gumstix) that is used to host the different gateway software modules (left). It is incorporated in the developed smart meter clip-in module of the final eMeter prototype (right).

via its WIFI module. For that, the cURL[17] library issues the respective HTTP requests. To handle the incoming user interface requests, we impelmented a full HTTP software stack using a lighttp (1.4.28) web server in combination with PHP (5.2.13) as backend. The server was built with the following four modules enabled:

- mod_auth: limits the access to the available URLs of the API;
- mod_fastcgi: enables running of PHP scripts via the FastCGI interface;
- mod_rewrite: rewrites the URLs of the RESTful-APIs;
- mod_redirect: for request-handling between the Recess-API and the mobile-API.

We use the Recess 0.12 PHP framework for the web user interface implementation, which requires 5.2.x versions of PHP. Hence, PHP 5.2.13 was built with FastCGI enabled, and a single child process and a main process for request management. The child process is restarted every 10000 requests to ensure that all leaked resources are freed back to the OS. The web server communicates with the PHP process via a UNIX socket. However, using the implemented Recess-API led to performance problems with the mobile user interface, which requires the delivery of data in real time. For example, the measurement functionality requires the current electricity consumption value in less than a second, a fact that could not be met using the Recess framework. We solved this problem by implementing a dedicated mobile-API for the mobile user interfaces (Figure 2.16) that uses plain PHP in combination with PDO (PHP Data Objects). Figure 2.18 shows a comparison of the performance of a RESTful-API call on the embedded gateway using a pure C implementation, PHP with PDO, and the Recess PHP framework. The results were obtained using the API call that is particularly critical since it is used by the measurement functionality: /emeter/energyServer/currentPhases.php?sm=<SM_TOKEN>. We conducted 3000 consecutive calls per implementation with first a single and then ten concurrent requests at a time. We find that even in case only one client is present the median to process the call of the Recess-API is significantly higher than using plain C or PHP with PDO (755ms vs. 18ms vs. 23ms). The longest response time (after removing very few outliers) for Recess is 24.6 times longer than PHP and 37.5 longer than C respectively. Moreover, an internal processing time of above 700ms is already critical to deliver the feedback in time. Using ten concurrent connections (i.e., more users utilizing the system simultaneously), the response time strongly in-

[17] cURL. http://curl.haxx.se/

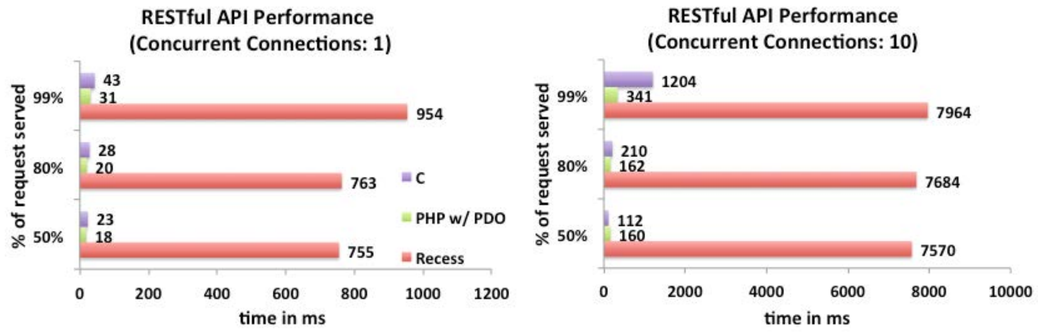

Figure 2.18 Comparison of the API performance using different software implementations: Response time in ms for one connection (left) and ten concurrent connections (right).

creases over all technologies. While the Recess framework cannot process the API call in less than the required one second, C and PHP both show better performance. However, the C implementation has a worse performance for an increasing number of served requests and a worst-case response time of over one second compared to significantly lower response time for PHP with PDO. As result, we implemented the mobile-API using PHP with PDO technology.

Besides the RESTful-APIs, we developed a mini web shell that enables administration of the embedded gateway over the web browser. It allows executing simple Unix commands without the need to access the Gumstix via SSH. This is especially handy for small setting changes, bug fixes, or status retrieval with the smartphone, which usually does not feature SSH.

The auxiliary services (see Figure 2.16) are important for the operation (e.g., WFI configuration, DHCP, etc.) and administration (e.g., SSH, Tunnel) of the embedded gateway. For example, the *tokend* process periodically sends a token (i.e., the MD5 encrypted MAC address of the connected smart meter) to the IPv4 multicast group. The user interfaces on the smartphones that are currently connected to the same network (i.e., the same IPv4 multicast group) listen for this token, and if detected automatically pair the application with the gateway. This simplifies the setup process for users because the setup effort on the first startup is limited to starting the application. Once the token is detected, access to the electricity consumption data of the meter is automatically granted.

The user interface was developed as a native application for three different mobile platforms: iPhone OS, Android, and Windows Phone 7. Since smart meters will be installed through the energy utility, user installation effort is limited to simply install a downloadable smartphone application. Figure 2.19 provides a top-level view of the implementation details of the user interface. On the first application start, the user is requested to speci-

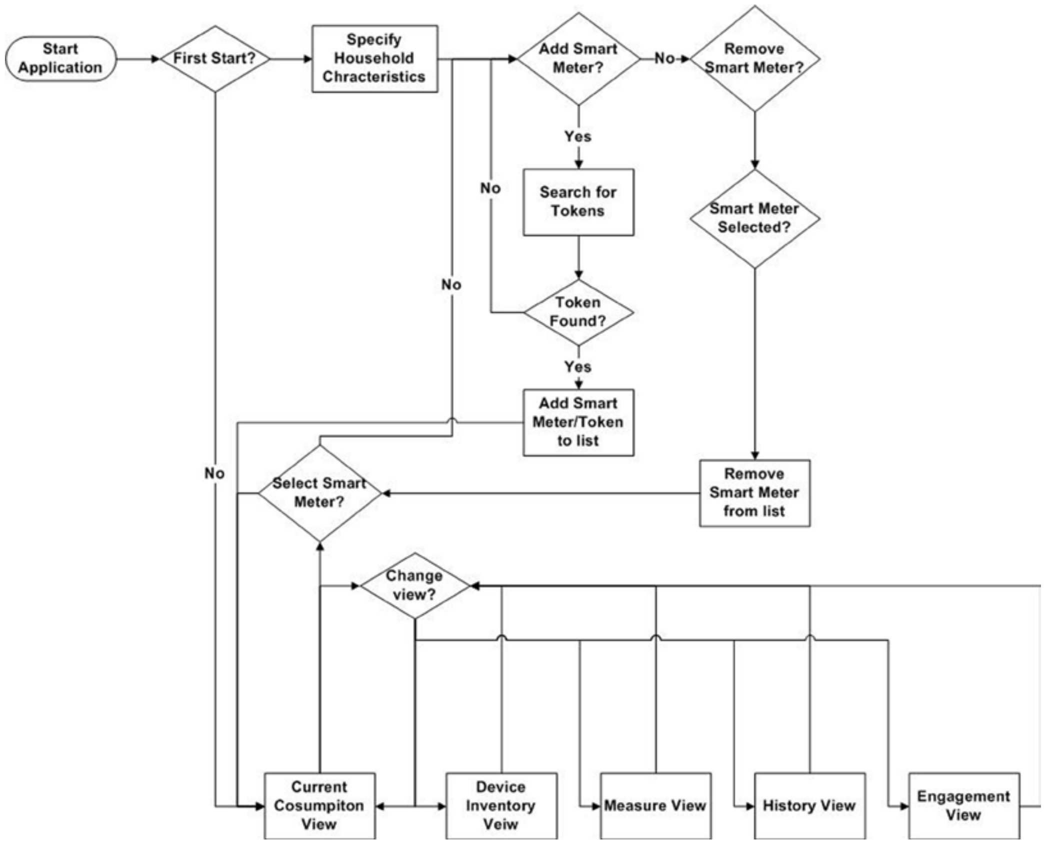

Figure 2.19 SDL action diagram of the top layer of the mobile user interface.

fy certain characteristics of the household that are later used for comparing the electricity consumption to the one of other households for budget calculations etc. The user interface then automatically tries to identify a smart meter by listing for a token on the IPV4 multicast address. If such a token is discovered, the meter is added to the list of accessible data sources and the application continues in the current consumption view. From that view, users can switch to any of the other main views or add a new smart meter to its list of data sources. On the next application start, the application automatically switches to the current consumption view with the last meter selected as data input. Each of the views contains several sub views that implement the functionality described in Section 2.5.4 and use the resources provided by the RESTful gateway API. Table 2.5 presents an overview of gateway URIs used by the smartphone applications. In order to get the required information in real-time (e.g., the current overall consumption) the user interface issues a GET request to the corresponding gateway URI and converts the server-generated JSON response that contains the requested values from the database (Figure 2.20). Using the same principle (i.e., PUT and POST verbs), it is also possible to store user-generated content in the database. This enables the system to store data (e.g., household characteristics) in the database on the gateway.

Table 2.5 Resources of the mobile API of the gateway utilized by the mobile user interfaces.

URI	HTTP Method	Description
/emeter/energyServer/current Phases.json?$sm_token	GET	Information used by the measurement functionality
/emeter/energyServer/current .json?$sm_token	GET	Current consumption details used in the corresponding view
/emeter/energyServer/smart Meter/$id/kWh.json	GET	Total kWh used in the history view
/emeter/energyServer/smart Meter/$id/projection.json	GET	Projections for the budget view
/emeter/energyServer/smart Meter/$id/watt.json?$time	GET	Returns electricity consumption for the line chart of the history view
/emeter/energyServer/iPhone /$phoneID/notifications.json	GET	Current notifications for the engagement view
/emeter/energyServer/recogni tion	POST	Send a measurement to the database on the gateway
/emeter/energyServer/smart Meter/$id.json	GET	Smart meter status information

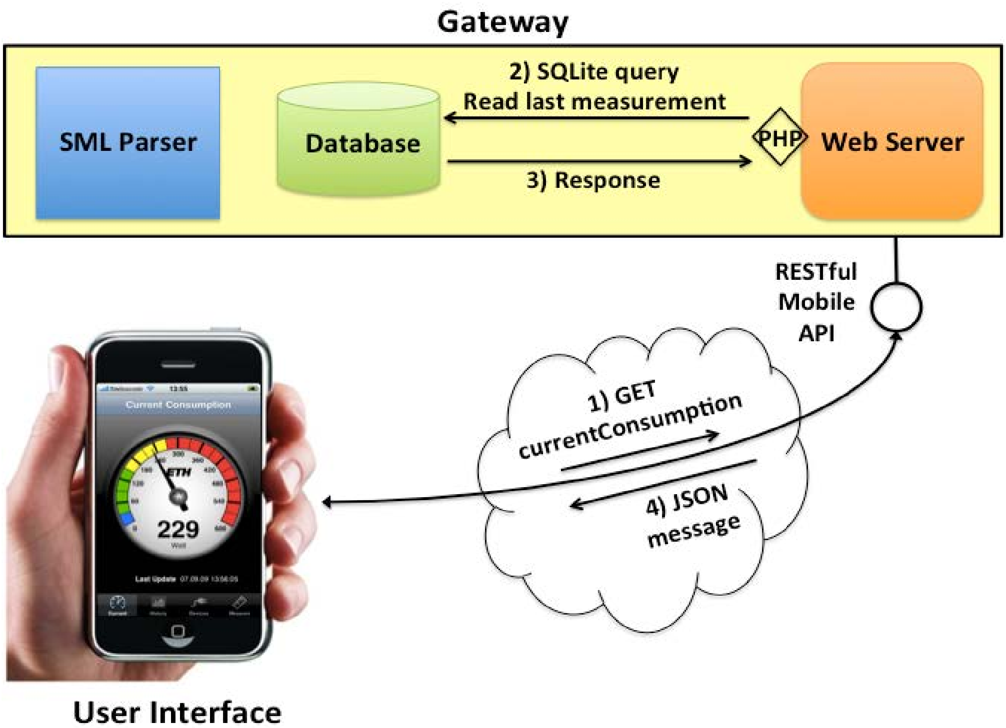

Figure 2.20 To obtain the current consumption that is visualized in the respective view, the user interfaces issues a GET request to the mobile API (1), the web server processes the call and uses a SQLite query issued by PHP to get the requested data in response (2-3). Last, a JSON message containing the information is formed and transmitted back to the user interface (4).

2.6.2 Architecture Prototypes and Evaluation

Based on the proposed design of the architecture, we used the discussed components and realized a prototype of the system in three steps. We first experimented with the implementation of the user interface on the iPhone. This enabled us to gather early feedback on the user interface design and the required electricity feedback functionalities. In order to quickly realize a working prototype that features feedback of real consumption data, we decided to connect the user interface to an existing RESTful API of a commercially available smart power outlet [118] (see Figure 2.3 left).

We refined the user interface based on the first feedback from a focus group and a paper-based survey (see Section 3.2) and built the second prototype that utilized a smart meter to log the electricity consumption. The meter and further hardware components required for operation and safety purposes were mounted to a portable wooden board (Figure 2.21 right). Besides better portability and demonstrability, this offers the benefit that individual components of the backend could easily be replaced if necessary. Figure 2.21 (left) depicts the front end of the demonstration setup of the eMeter system used at public demonstrations. It consists of a large LCD TV screen that shows the web user interface and the mobile phone user interface. The mobile phone user interface is running in the iPhone emulator. In addition, multiple phones on stands were used to provide visitors a real look-and-feel experience of the eMeter system.

Since developing on embedded devices is time-intense, the gateway logic that handles the communication to the meter and to the user interface was split into two locations at this point. The Gumstix hosted the SML parser, but transmitted the data over its WIFI interface to a dedicated machine

Figure 2.21 eMeter demo web user interface and mobile user interface used at larger public demonstrations and as permanent setup (left). First prototypical implementation of the eMeter backend on a wooden board (right).

on the Internet (instead of directly storing and processing it internally). On that server, the remaining parts of the gateway (i.e., web server database, etc.) were realized.

To demonstrate the feasibility of our approach as well as to evaluate the correct functionality and stability of the system, this version of the prototype was deployed for several months in a laboratory on our group's floor. Colleagues, visitors, and industry partners used the prototype daily and helped to further refine and improve the usability, stability, and functionality of the eMeter feedback system.

At the end, the system was running reliably for several months and led to the final prototypical implementation of the eMeter system, which is illustrated in Figure 2.22. The "demo case" consists of the smart meter that can be plugged into any electrical outlet to measure the electricity consumption of the attached devices. It communicates via its Ethernet interface with the Gumstix that hosts the full gateway functionality (i.e., web server, SML parser, and database). Its WIFI interface can be configured to operate in different communication modes (see Section 3.5.1) allowing direct or infrastructure-mediated access to the data through the RESTful API.

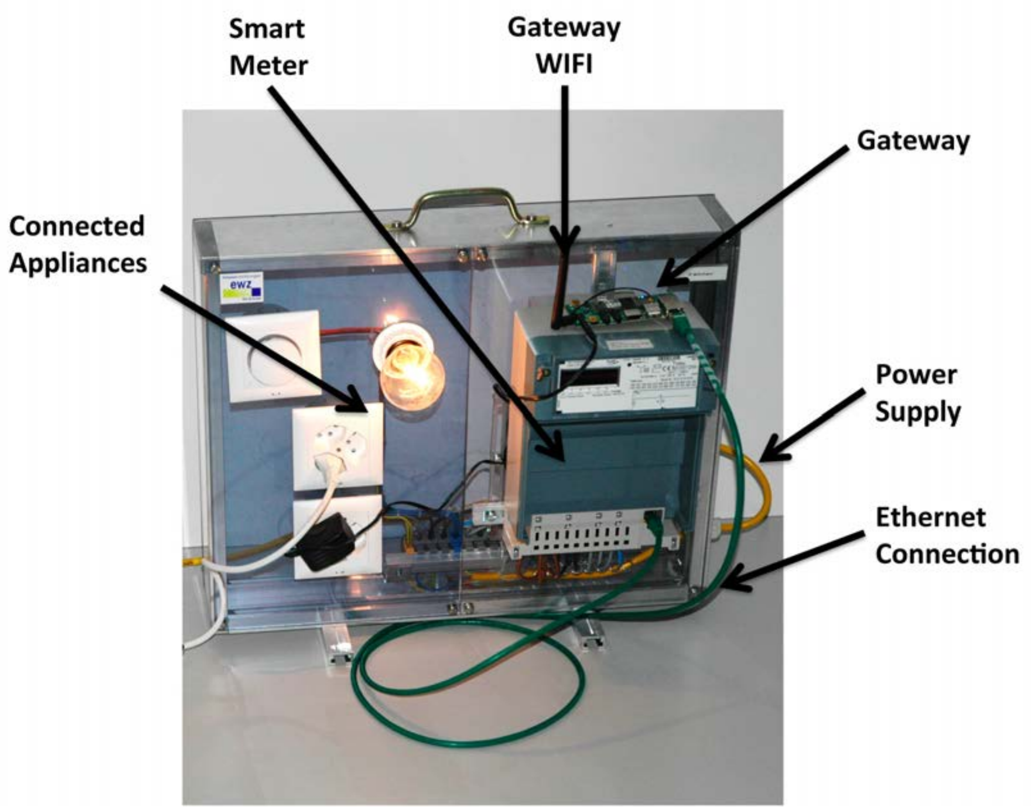

Figure 2.22 Final prototypical implementation of the eMeter architecture (backend only): It consists of the smart meter that records the electricity consumption of all connected appliances. The gateway connects to the meter via Ethernet and serves the incoming user interface requests over the WIFI interface.

We evaluated the gateway performance based on the final prototype. The system was set up in the laboratory environment and the user interface was constantly requesting data from the gateway for the observation period of one day. Over the whole period, we observed that the mobile user interface was constantly connected to the gateway's WIFI interface and no loss in connectivity occurred. During operation all major processes, such as the SML parser, the web server, and PHP, all show stable process memory footprint, which is important to rule out memory leaks. The load averages in terms of CPU utilization of the system remain low for all observation periods (e.g., 15 min load averages: 0.05 minimum, 0.18 maximum, and 0.066 average). Further in-depth evaluation shows that the SML parser that is continuously handling the data acquisition process causes most of the CPU load:

- PHP CPU usage: min – 0%; max – 3.8%; avg – 0.96%.
- SML parser CPU usage: min – 0%, max – 11.4%; avg – 3.13%.
- web server CPU usage: min – 0%, max – 1.9%, avg – 0.46%.
- overall CPU usage: min – 0%; max – 19%; avg – 4.96%.

This behavior seems natural, since the SML parser contains the eMeter SQLite Extension module that is responsible for the management of the database. With a user interface constantly pulling, it has to perform a lot of modifications and updates to the database, which explains the relatively high load averages. Under constant strain from the user interface, we can confirm that the web server, PHP, and the database only cause a minimal load on the system. In summary, we can conclude that the CPU resources are more than sufficient for our gateway implementation.

Another important performance aspect of the gateway is the disk usage of system. The Gumstix Overo is equipped with a 256MB flash memory that can be extended with up 8GB of storage using a micro SD-card. After the initial system setup 240MB are available for general system use. With full debug information of the most important system services enabled, drive space decreases at a rate of 29.8MB/h. However, this effect can be limited utilizing the "logrotate" package that limits the amount of the most recent debug information to fixed threshold (e.g., 24 to 48 hours). Then disk space highly depends on the size of the database that increases at rate of 2.4MB/h. This means that without further optimization on the database structure the gateway is capable of storing 7 days of real-time electricity consumption information on its internal file space. In case this is not sufficient, the corresponding database file could be moved to the external SD-card that provides additional space.

Since the utilized smart meter records the electricity consumption information on a second by second basis, one of the key criteria for the gateway is to retrieve the measurement, process it, and subsequently store it in the database in less than one second (i.e., the datalogger run loop (compare Figure 2.8)). We evaluated the performance of the measurement retrieval time over a 24h period. For multiple runs, we analyzed the syslog and database log messages. Figure 2.23 depicts the measurement retrieval times for one typical run. Out of 86400 possible measurements we recorded 85960 in the database. That corresponds to a loss of 439 measurements or 0.51% respectively. 67.5% of all measurements were stored within 200ms, 87.3% within 300ms. No successful retrieval process took longer than 500ms; in fact, the number of measurements with a time to store of over 400ms is negligible. This result is sufficient for providing users with real time feedback.

We repeated that test a number of times. Taking all test runs into account, the loss rate varies between 0,5% and 0,75%. This loss is largely due to the small time difference between the measurement interval of the meter and the parser interval of the gateway (i.e. there is yet no synchronization implemented). In a few cases, the run loop of the datalogger has a slightly larger interval than one second, which causes the small loss of measurements. Lower loss rates can be achieved through the implementation of the following two measures. The time overhead needed for initiating the TCP/IP connection to the meter is currently unknown. Measuring this overhead and incorporating it into the time calculation of the datalogger run loop could help to better calibrate the timing. Instead of measuring the overhead, TCP/IP connection keep alive could be used to achieve a persistent connection between the meter and the gateway. Alternatively, a message buffer that queues several measurements together with a shorter run loop interval could be implemented.

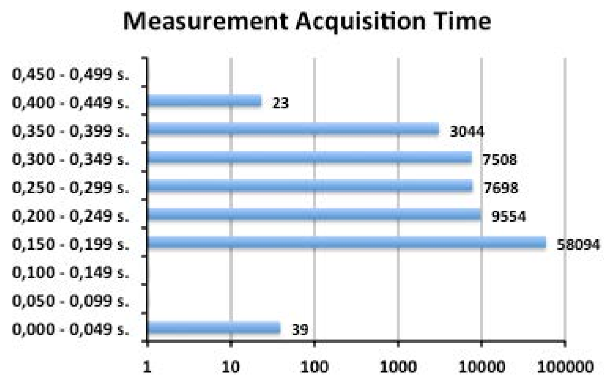

Figure 2.23 Measurement acquisition time corresponds to the time the datalogger module needs for requesting the data from the meter, parsing it into a JSON message, and subsequently storing it in the database.

The evaluation shows that the performance of the designed infrastructure meets the technical requirements of an electricity consumption feedback system. The utilized hardware components together with the implemented software modules are well suited to provide sufficient space for data hosting and are powerful enough for request handling and real-time information delivery. These results encouraged us to deploy the eMeter system in four households in Zurich, Switzerland to demonstrate the real-world feasibility of our approach. The system was rolled out together with industry partners, which took care of correct installation and ensured functionality of the meter. In order to not interfere with the currently installed meter that had to remain active for billing purposes, the smart meter was installed parallel to the existing meter (Figure 2.24 left). The gateway was operated in full data sharing mode. That is, its internal WIFI was hooked up to the WIFI of the participants forwarding all data to our central server on the Internet.

After the initial setup with the partners, we ensured the system was fully functional and showed the same behavior as in the laboratory environment. To do so, we spent several hours in the households operating (Figure 2.24 middle) and measuring all kinds of different devices (Figure 2.24 right). In order to be able to access the system from outside the local network (e.g., in case of system failure or maintenance), we implemented a reverse SSH tunnel that auto-establishes a connection to our server. Since being initiated from inside the home area network, it was able to bypass existing firewalls.

The system was running and permanently used from June 2010 until August 2011. During this period, we collected over 100 million measurements and almost 400 user sessions of the smartphone application. While the deployment confirms the feasibility of our approach also in a real-world scenario, it also contributes to other findings discussed in course of this thesis: the evaluation of effective forms of meaningful electricity consumption feedback and the design of an algorithm that disaggregates the overall energy consumption to the consumption of individual devices.

Figure 2.24 Real-world deployment of the eMeter system. The smart meter was installed next to the existing meter (left); demonstration of the current consumption view with the stove running (middle); utilization of the measurement functionality to determine the consumption of a LCD TV (right).

2.7 Summary and Discussion

In this chapter, we presented one of the three main contributions of this thesis: the design and implementation of a pervasive electricity sensing and feedback system. It consists of three components – a smart meter, an embedded gateway, and a user interface on a mobile phone – and provides ground for the other contributions of this thesis. The developed architecture shows a possibility how future electricity monitoring systems can be composed to provide real-time and fine granular electricity feedback while respecting user-centric and system-centric design criteria. By developing a portable user interface on a mobile phone, we provide the electricity feedback on a device that is already integrated in many users' daily life. The feedback itself comes in a way literature suggests it is ideally desired by users to better understand the origin of their consumption. That is, in real-time, on device level, and at hand when needed.

To overcome the limitations of electricity existing monitoring solutions (e.g., battery dependency, complex installation, etc.), our architecture uses components that will be integrated in many households in the course of smart meter implementation. Moreover, through the development of a RESTful API, the gateway integrates the readings of the smart meter into the Web and makes them easily accessible for humans through a web browser or a mobile phone application. Following the REST paradigm, this also decouples the individual components of the architecture, while relying on HTTP as a future-proof communication protocol. Furthermore, through the use of JSON, other applications can be easily developed on top of the system [31, 98]. We have demonstrated this by the development of user interfaces on three different mobile platforms as well as a web user interface. The system can thus serve as an enabling framework for other researchers that can easily develop and test their own visualization concepts (e.g., for investigating behavioral change effects through different visualizations or forms of engagement) or to investigate automated energy saving applications on top of it [42, 61]. Instead of developing ever more proprietary solutions and wireless protocols, the developed eMeter system shows a way of how future energy monitoring system can be implemented more efficiently.

Using a smart meter as single data source, the system not only visualizes mere consumption data in real time, but also helps users to put their consumption in a bigger picture (e.g., comparison to others, budgeting). However, the system's device-level accuracy, suffers from the single sensor approach. We try to surpass this by integrating a power measurement functionality for appliances into the user interface. This simple yet powerful tool allows users to interactively explore their electricity consumption. It aims at providing users with an initial idea how much different devices consume. Provided that the small measurement interval is representative, it enables users to measure the electricity consumption of any switchable or pluggable device. However, there also exist some devices in the household that consume a non-negligible amount of energy, such as the washing machine or the freezer, which usually cannot just be turned on or off. For such devices that cannot easily be measured by users, automatic device identification methods could be envisaged that automatically detect and match the electricity patterns in the overall electricity consumption to the device that caused the load [90, 119]. Alternatively, we suggest combining the eMeter system with the use of a smart power outlet [98, 118]. This would enable users to gather device-level electricity consumption information for both switchable and non-switchable appliance without tremendously increasing the complexity of the eMeter system.

To validate the design, we presented a prototypical implementation of the whole infrastructure that was deployed and evaluated in in a laboratory setting. Our results show that the system is well dimensioned for its purpose and functions reliably as originally intended for several months. This encouraged us to deploy the system in practice in four households to gather data and further make use of the infrastructure in our other contributions. In the following two chapters, the proposed eMeter system is applied to evaluate the designed user interface as well as to investigate the potential of using this infrastructure to facilitate the automated detection of appliance switching events.

3 Evaluating Mobile Phones as Energy Feedback Devices

In this chapter, we evaluate the suitability of mobile phones as energy feedback devices. The work of this chapter was partly published in [63, 98, 120, 121]. In Section 3.1, we start by reviewing relevant work in terms of information provisioning on mobile phones and its evaluation. We then describe the goal-driven development of a mobile phone application for residential energy monitoring based on electricity consumption data acquired by the eMeter infrastructure in Section 3.2.

Next, Section 3.3 explains the user study we conducted with 25 participants to evaluate the designed user interface and the perceived value of various feedback functionalities. After elaborating on the experimental setting, we present selected study results. Section 3.4 reports on the experimental setting and findings obtained form the long-term real-world deployment of the eMeter infrastructure in four households in Switzerland. Both sections show how mobile phones can help users monitor and control their energy consumption and confirm the suitability of mobile phones as energy feedback devices.

Section 3.5 follows with a discussion on the results of the conducted user study and the real-world deployment. After that, Section 3.6 describes PowerPedia that was implemented as direct measure of the results of the preceding evaluation. It aims at putting the electricity consumption in a bigger, more tangible picture beyond mere numbers by providing meaningful feedback on the consumption of electrical appliances. Last, this chapter concludes with a summary in Section 3.7.

3.1 On the Information Provisioning on Mobile Phones and its Evaluation

With the rise of Ubicomp, data about real-world events is being captured at an increasingly detailed level. Together with the rapid growth of the mobile phone market and mobile Internet access, this has led to a large number of mobile applications, which aim to support users' daily life in a wide range of areas. To name a few, this ranges from insurance claims assistance [122] over shopping assistance [123] to emergency response [124].

A relatively new direction focuses on providing information about the personal environmental impact of travel, shopping, and residential resource consumption [125]. Ecorio[18] and Carbon Diem[19] for example allow for tracking the personalized carbon footprint with the help of the smartphone´s GPS sensor. The GreenMeter[20] aims at reducing the fuel consumption and resulting cost by using the mobile phone's internal accelerometer to measure forward acceleration and calculate fuel economy as well as carbon footprints. The Carbon Tracker[21] application serves a similar purpose, but bases the calculation mainly on self-reporting. Another mobile application semi-automatically senses and reveals information about personal transportation behavior and tries to motivate users to choose green forms of transportation [6].

Information provisioning with respect to residential resource consumption has received considerable attention lately. There exist numerous mobile phone applications that allow users to monitor the electricity consumption of individual household devices. These solutions are often based on smart power outlets, like Tendril[22] or the Energy UFO[23]. Once installed, they record the attached load and are capable of transmitting the measurement data wirelessly to a remote user interface. However, these products typically lack the possibility to aggregate the consumption of multiple sensors and to fuse the different data into a comprehensive picture. To surpass this limitation, different projects worked on integrating data from multiple sensors into one mobile application [7, 98, 99, 126]. The developed

[18] Ecorio. http://www.ecorio.org
[19] Carbon Diem. http://www.carbondiem.com
[20] Green Meter. http://hunter.pairsite.com
[21] Carbon Tracker. http://www.clearstandards.com/carbontracker.html
[22] Tendril Inc. http://www.tendril.com
[23] Energy UFO. http://www.visiblenergy.com

mobile prototypes are targeted to establish easier access to energy consumption data for users and allow developers to easily build their own application on top. While the concept is interesting and helps to provide important findings for future work, one has to keep in mind that deploying a large number of sensors in a residential environment is cumbersome and expensive. This implies a high usage barrier that typically hinders adoption.

Besides the work that focus on the technical progress to deliver feedback on portable devices, other research investigates the modality how energy consumption data should be presented [53, 55, 69, 108, 127-130]. For example, the choice of measurement unit strongly influences the comprehension of feedback, as some units (e.g., kWh or carbon dioxide emissions) are more difficult to understand than others (e.g., money). Although kWh is often not thoroughly understood [129], this unit has established a sense of trust due to its scientific basis. Money as a unit is easier to understand, but is only suitable if potential savings are large. Otherwise, people might consider their energy expenditure as a minor and less important part of their total expenditure. Data granularity with regard to for example a specific source or source category enables users to prioritize their energy saving effort. To direct the consumers' focus towards saving energy, push messages or prompts can be a helpful tool, but one has to assure that users are not flooded with messages. Even lightweight push mechanisms have proven to be effective [128]. To be able to judge whether one's consumption is either high or low, current consumption has to be contrasted with past performance. For that purpose, bar charts or line charts have been proven to be most effective [106]. To display development there have also been attempts to use artistic visual designs instead of pragmatic (numeric) designs [128]. Research on the modality of energy feedback has helped to better understand what works in terms of data representation, however, it is important to note that there does not exist a "one-size-fits-all" solution. That is, different feedback has to be presented to differently motivated people in order to be effective and target a wide user base [59].

In order to evaluate the usability of mobile phone applications as feedback technology, different methodologies are applied in the mobile Human-Computer-Interaction domain (i.e., case studies, field studies, laboratory experiments, applied research, action research, surveys, basic research, and normative writings). Applied research builds on trial and error on the basis of researchers' capabilities of reasoning through intuition, experience, deduction, and induction is used most often (42%) to evaluate mobile phone applications. It is followed by laboratory experiments (24%) and 15% field studies. Only 6% of the work report from surveys, 5% are normative writings (concept development writings), and 5% report from case studies. Hardly any work based on action or basic research could be identified [131].

The evaluation of mobile phone applications is usually done with regard objective criteria as well as subjective performance criteria. For example, [132] compares 11 state-of-the-art barcode scanners in a laboratory study with 20 participants focusing on objective performance criteria such as time per scan and reliability. The authors found that the reliability ranges between 50 and 100% and the average time per scan lies between 3.5 and 10.4 seconds with the most reliable scanners having the shortest time per scan. The authors of [122] conducted a survey study with a commercial sample of 2000 people to find out how people rate a mobile claim assistance application with regard to subjective criteria only. Participants were asked to judge the applications' usefulness, ease of use, credibility, and their intention to use. UbiGreen [6] is a mobile application that targets environmental aspects and has been evaluated in a user study. The developed prototype semi-automatically senses and reveals information about the user's transportation behavior. In a field study, the authors tested how two kinds of eco-feedback (representing carbon dioxide emissions with a tree gaining or loosing leaves and an ice bear sitting on melting or growing ice floes) were perceived by 14 participants (Figure 3.1 left).

In terms of portable energy feedback, there exist few practical evaluations of user interfaces. One notable is the evaluation of Energy Life (Figure 3.1 right), a mobile phone application that can give feedback on the electrical consumption of individual devices. To do so, the mobile phone user interface is being connected to a server, which in turn is wirelessly linked to a variety of plug sensors. The mobile application aims at helping users to monitor their consumption and quality of their implemented conservation practices by providing feedback and so called awareness tips. The latter are meant to increase the users' knowledge on the consequences of their electricity consumption in general and of that of specific devices in particular. In a laboratory study, objective and subjective criteria were evaluated with 20 users. Participants were asked to fulfill several different tasks (e.g., identify the highest consumption in the past). On average 8.6 out of 12 tasks were performed successfully, and the average number of errors for each task was 18.8. After the practical part, a 41-item questionnaire was used to evaluate Energy Life with regard to the following qualitative subjective criteria: navigation, comprehensibility, structural clarity, pleasantness, satisfaction, learnability, feedback, consistency, control, and usefulness [108].

The applied methodology of the last-mentioned work inspired us for the design of our own study that is discussed later in this chapter. In contrast to the wide amount of research that has been conducted on data representation, we aim at contributing on the usability and the required functionality of portable energy feedback systems. Compared to other work within the usability and functionality domain, our applied methodology for evaluating the suitability of mobile phones as energy feedback devices uses pa-

Figure 3.1 Examples of information provisioning on mobile phones with respect to green applications: UbiGreen (left) is an application that provides green travel information on its mobile phone user interface. A filed test has been used to evaluate the prototypical implementation. The Energy Life application (right) provides electricity feedback on individual appliances and has been evaluated through a user study. Source: UbiGreen [6], BeAware [7].

per based prototypes and surveys only to achieve better results during the design process of the application. The final evaluation then is based on a user study and a long-term field test where users get in touch with the look and feel of a real prototype application (including a fully functional backend) and compare it to commercially available energy monitoring products.

3.2 From Paper to Practice

Before designing and developing the user interface of our electricity consumption feedback prototype, we conducted a survey to provide us with an idea what functionalities users would expect. The survey design was developed in three steps. First, we initiated and led a discussion with industry experts in order to identify two applications where the participants were confident that these were easy to explain in a paper-based survey and offered varying degrees in terms of their feedback characteristics to obtain two diametric applications. Applications, which were found to be suitable, were:

1) A mobile phone application that utilizes a smart metering infrastructure that allows users to get feedback on the consumption of individual appliances (high timeliness of informational support, suitable for investigation of specific loads, high degree of interaction, and portability).

2) A washing machine with a simple display, which provides feedback on energy consumption of specific programs and information on the energy that is saved by choosing eco-programs (low timeliness of information, low degree of interaction, and no portability).

In a second step, we extended the question catalogue by a set of constructs to evaluate general functionalities of consumption feedback. We used established constructs taken from the Technology Acceptance Model (TAM) [133], constructs on the word of mouth, and questions concerning the willingness to pay. To evaluate the validity of the reported willingness to pay, we used a technique called framing. Framing means that information can be presented in different contexts, which affect the perception of the information [134]. In a third step, we evaluated the comprehensibility by reviewing the constructs of the questionnaire with non-experts.

The survey was conducted at lively points throughout the city center of Zurich, Switzerland. 185 persons participated in the survey (50.3% male) with all age groups evenly represented. The sample was slightly biased as respondents with a higher education degree and an above–average income were overrepresented. However, we do not expect this to considerably reduce the validity of the findings. In the following, selected results of the survey that lead to the design of the prototype are presented. Regarding the general attitude towards conserving energy, roughly 50% of the participants think it is currently rather cumbersome to save energy and not fun. In doing so, 89% like to be supported by innovative technology.

The left side of Figure 3.2 depicts the comparison of the perceived usefulness and the intention to use of the mobile phone application with the washing machine display. The figure shows the mean as well as the standard deviation (as black error bars). The ratings for both applications are above average, with the mobile phone application performing worse over all the ratings. It was surprising to us that the washer application performed better than the mobile device with regard to the likelihood to recommend the application to a friend, the likelihood to use the application, and the expectation that the information would lead to energy savings. A potential explanation for this unexpected outcome is that the washer application had an inherent use case (the application domain “washing” and saving by selecting an “eco-program” was clear), while the innovative phone application had a more general and not very obvious use case (“you

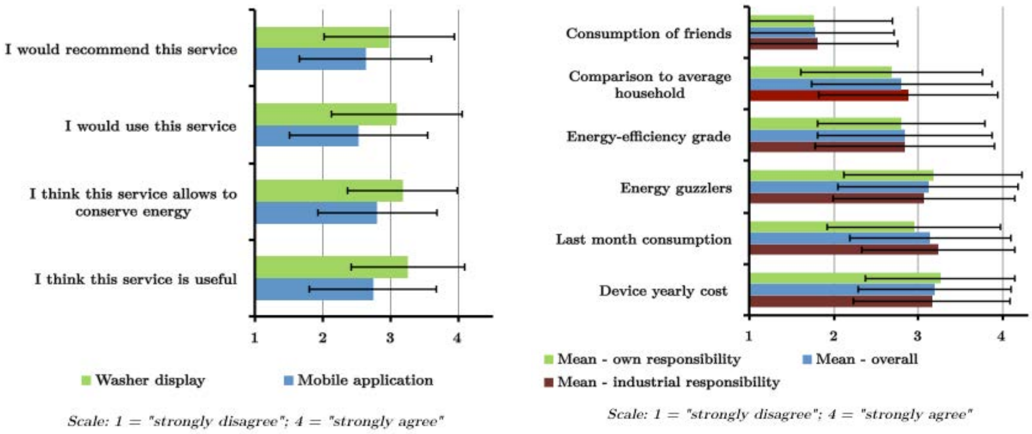

Figure 3.2 User assessment of the usefulness and intention to use of two future energy saving technologies (left). Evaluation of the desired functionalities of such an application (right).

can follow your energy consumption and measure how much energy your devices consume, which can help you save").

In order to conserve energy at home, it is important to be aware of how much particular appliances consume as well as of effective measures that increase the energy-efficiency. Thus, we asked users with what information they would like to be provided on a mobile phone application. The right part of Figure 3.2 shows the participants' assessment of the different functionalities. Besides depicting the overall mean of the whole sample and the standard deviation in error bars, the figure provides a more in-depth view on the mean of participants that regard energy conservation as their own responsibility (N=65) and those that believe it is industry's responsibility (N=120). Those who strongly or slightly disagreed on the statement "The industry is primarily responsible for saving energy" (1 or 2 points on the four point Likert scale) belong to the first group. The second group consists of people who slightly or strongly agreed with the statement mentioned above (3 or 4 points). Overall, we find that participants prefer to be provided with the yearly cost of single appliances followed by the last month consumption and the biggest energy guzzlers of the household. However, users do not want to compare their consumption to the one of their friends or family. A closer look on the attitude of the participants reveals that participants who regard saving electricity as their own responsibility rate the identification of the biggest energy guzzlers significantly higher than others. We assume this goes along with their higher involvement and interest for both, their personal energy efficiency and defined measures that allow for conserving energy.

In addition, we investigated the self-reported willingness to pay for a mobile application that allows for measuring the electricity consumption only. Knowing of the difficulty to obtain a reasonable indicator for a sales price with this method, we used the findings to get an indication for the relative price range, and for developing a better understanding on how stable the perceived value is. For that, we used two different, slightly modified versions of the questionnaire and distributed them half-half amongst the participants. The first version asked people for their willingness to pay, not indicating the possible saving potential of such a technology measure. The second version indicated the monetary amount of potential savings (85$), thereafter asking for the amount people would spend. Figure 3.3 illustrates the results. While the mean values do not vary significantly between both versions, the median increases from 16$ to 30$ when presented the version with frame. Thus, by providing an annual saving value, participants were willing to pay a higher price for the application. This effect is referred to as framing effect and highlights that the expectations towards the pricing of such an innovative energy feedback application are only vague. From these results, we can conclude that the price regarding such a mobile phone application is not yet set. Participants have a general understanding that mobile phone applications are low-priced. On the other hand, the results also show that it is hard to determine the price for an innovative product that is not yet touchable and the price can be varied according to the context the application is presented.

The survey results served as an indicative basis for our prototypical development. It encouraged us to offer a simple use case that makes the value of an innovative energy consumption feedback application clear even to individuals who are not familiar with the system. As an easy to explain use case, we identified a measurement function that can be explained as “learn how much a device consumes by just switching it on or off”. On that basis yearly costs can be calculated and the biggest energy guzzlers can be identified.

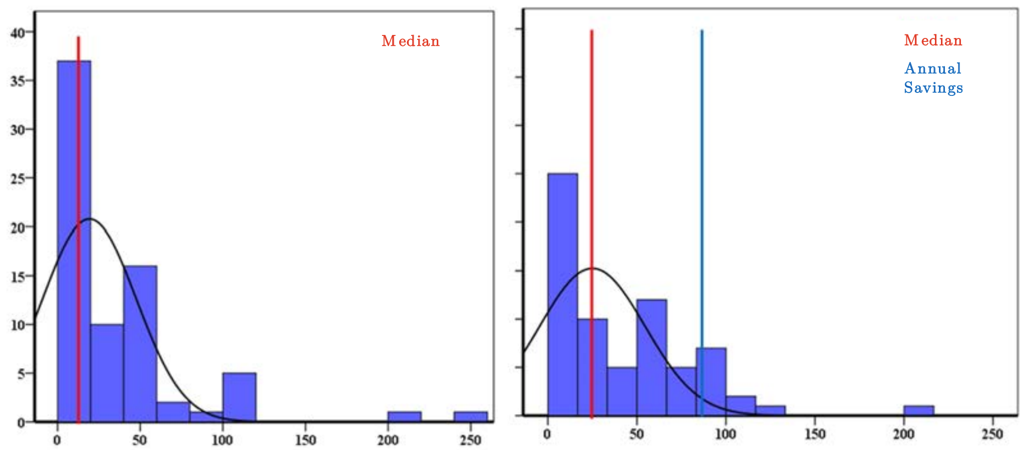

a) without frame

b) with frame

Figure 3.3 User Assessment: Willingness to pay for an electricity monitoring application on a mobile phone.

However, the results on willingness to pay, perceived usefulness, and the intention to use the system also showed us that innovative services seem to go beyond the imagination of users. Thus, it is necessary to investigate possibilities of electricity consumption feedback on mobile phones in a controlled environment, ideally in a user study where users can get in touch with the application and where the understanding of the application`s functionalities can be guaranteed.

For that, we started prototypically developing our electricity feedback user interface as a native application on a smartphone. Paper-prototypes were used to present first ideas on the visualization and possible functionalities to industry and research experts. Paper-prototypes are a known concept in user interface design and usability engineering that is typically used at an early stage of the design process of applications [135]. It helps to quickly visualize design ideas that result in a prototype drawing which is simple and whose development does not need a lot of time [136]. The benefit of being able to quickly incorporate changes into the prototype design is supposed to achieve better results when later designing the full application in the interface builder of the mobile phone development framework.

We draw paper prototypes of all views to determine the functionality and the amount of information that could be visualized as well as the necessary controls that are adequate to quickly navigate in the application. The upper part of Figure 3.4 illustrates the initial design ideas of the four main views of the user interface that were influenced by the survey results and thereafter iteratively reshaped through expert discussions. The current consumption view is designed in the style of a speedometer similar to the one known from vehicles. The evolvement of the view over different stages of the design process is documented in Figure 3.5. The discussion with industry experts resulted after several iterations in a more simplistic view that is easily readable. It focuses on the overall electricity consumption and on how it compares to the household's typical usage. While the first implementation was hard to read and understand, the second design was considered too dark and still too complex. The third implementation had all buttons removed and offered a much clearer design (i.e., more contrast and larger numbers). It automatically updates the current consumption and auto-adapts the color-coding of the scale. The final version sticks to the design and additionally offers a zoom-in feature (see magnifying glass in the rightmost picture of Figure 3.5) that enables users to have a more precise look on their consumption. Similar refinements were conducted for the other three main application views. The history view is based on the standard stocks application of the iPhone, while the design of the device inventory uses a customized iPhone list view. As a direct result of the user survey, the measurement view was incorporated as one of the main features of the smartphone application.

After several rounds of feedback and follow-up refinements of the design and the functionality of the user interface, this process led to the final implementation of the prototype (see lower part of Figure 3.4). It was used in the following in a user study and a real-world deployment in four households. Both experiments aim at confirming the suitability of electricity feedback on mobile phones. We next report on the conducted user study and selected results, before presenting findings from the real-world deployment.

Figure 3.4 From paper prototypes to practice implementation: Development of the eMeter user interface as mobile phone application. Current consumption view, history view, device inventory view, and measurement view (left to right) [1].

Figure 3.5 Refinement of the current consumption view: Different stages of the user interface when designing the application in iterative feedback rounds with industry and research experts [1].

3.3 User Study

In the following, we focus on the user study we conducted to evaluate the suitability of mobile phones as energy feedback devices. We first report on the sample that participated and the setting that was used to evaluate the eMeter electricity consumption feedback. We then briefly discuss the energy monitoring products used to provide users a general idea of what available solutions look and feel and to compare their potential to determine the electricity consumption of individual appliances. Last, we explain the procedure of the experiment and the tasks users had to accomplish during the study, before we report on selected study results with respect to the application and the perceived value of different feedback functionalities.

3.3.1 Sample and Setting

To evaluate the general usability of the mobile phone user interface and which feedback functionalities are perceived as most valuable on a mobile phone, we conducted a user study with 25 participants. Their different background ranged from students over marketing and sales persons to industry experts. Twelve of the participants were male (48%). We covered all age groups: 32% were between 18 and 25 years old, 28% between 26 and 35 years, 36% between 36 and 49 years, and 4% between 50 and 70 years.

The study took place in a neutral environment meaning that there were no disturbing factors like colleagues watching the experiment or noise. The appliances used during the study were chosen with regard to two criteria: Firstly, we were looking for appliances that are well known in a residential environment and thus integrated into the daily life and secondly cover the most important categories (kitchen equipment, office use, and consumer electronics). Thus, we selected a kettle, an office lamp (standard light bulb), a computer screen, a game console (Nintendo Wii), and an energy saving lamp (Figure 3.6). For the study, we set up the prototypical implementation of the eMeter architecture and placed the appliances on a table in a row with each one being plugged in separately below the table.

The iPhone eMeter user interface and two other electricity measurement technologies that were used to compare the eMeter prototype with commercially available products were placed on a second table opposite to the appliances. In order to guarantee a reproducible procedure with each participant, we fixed the distance between the measurement technology and the appliances at three meters. Hence, we were able to rule out that differences in measurement time are actually due to distance effects instead of technology effects. Apart from the participants two other people were in

Figure 3.6 User study setting: The eMeter architecture was set up together with different household appliances in a controlled, neutral environment.

the room: the experimenter provides instructions regarding the study and a second person measures and notes the objective performance criteria.

3.3.2 Utilized Energy Monitors

In order to familiarize participants with the features of electricity monitoring solutions, we selected two popular energy monitoring solutions that at the same time can be used to evaluate the measurement function (eMeasure) of the eMeter user interface.

The first technology, the Wattson (Figure 3.7 right), consists of a display unit and a single sensor that communicates wirelessly and is used to derive information on the entire energy consumption of a household in near real time. The sensor is rather difficult to install. It has to be clipped around a single phase of the electric mains or around a circuit in the fuse box. This modification around the electric wiring is in many countries inaccessible mainly for safety reasons. The battery powered display unit is portable, but it can also be connected to the residential power circuit. The overall consumption load is relatively easy to determine directly from the display unit. To derive the actual consumption of individual devices further manual calculations by the user are required. Besides indicating the level of consumption as ambient light, no further functionality is present on the device itself.

The second technology, the Click (Figure 3.7 left), is a commercially available smart power outlet. It has to be installed in-line with the application which electricity consumption should be measured, and it aims at

Figure 3.7 Commercially available energy monitoring solutions used in the user study to determine the consumption of individual appliances. Display unit and sensor of Click (left) and Wattson (right).

monitoring, controlling, and automating the attached device. Once a device is connected, the Click allows for direct feedback on the attached load within seconds. The portable display unit is solar powered and enables users to remotely switch the device on or off within close distance (e.g., the same room). However, to measure the electricity consumption of multiple devices, the user has to attach these separately or attach a multi-outlet power strip, which can be rather cumbersome in a residential environment with power supply cables of different consumer electronics twisting around each other.

3.3.3 Experimental Design: Procedure and Tasks

In this section, we describe the experimental design we used to evaluate the suitability of electricity consumption feedback on mobile phones. The user study was divided up into two parts. The first part aimed at interactively evaluating eMeasure using the iPhone eMeter user interface and the two above-introduced electricity monitoring systems. The second part of the study was designed to provide insights on the general performance of the user interface and on the perceived value of different feedback functionalities of portable electricity monitors.

3.3.3.1 Part 1: Evaluation of eMeasure

During the first part of the study, each participant first had to estimate the energy consumption of different appliances before confirming the estimate utilizing different measurement technologies (i.e., Wattson, Click, and eMeasure[24]). This enables us to check whether the participants already have a vague impression of the electricity consumption of each appliance. Hence, the participants estimated the consumption of the office lamp, the kettle, and the computer screen as well as the accumulated consumption of the game console and the energy saving lamp. Thereafter, users were asked to verify their estimate by determining the consumption utilizing the three above-mentioned technologies (see Figure 3.8). To prevent an order effect, we alternated the sequence participants were using the three measurement technologies. That is, a third of the participants started to measure devices with the iPhone, followed by the Wattson and finally the Click, the second group started with the Wattson, then Click, and then iPhone, and the third group started with the Click, then iPhone, and last Wattson. The order in which the appliances had to be measured was kept constant.

Before starting the measurement process, we explained the procedure to the participants and told them that both speed and accuracy of the measurement are of special interest. We then gave participants an instruction for each measurement technology to ensure a common understanding and comparable actions. Before starting the study, we asked the participants to read the manual carefully and pose questions, if they do not understand the procedure. To ensure that the measurement procedure was really understood, participants had to briefly summarize their task. That is, participants had to perform the different measurement tasks one by one with each technology. To do so, they had to walk over to the table with the five devices, measure electricity consumption of the appliance, and finally loudly indicate their determined result.

Figure 3.8 Participants measuring the electricity consumption of different residential appliances by utilizing different technologies.

[24] Note that during part one of the user study the eMeter user interface was limited to the measurement view to avoid confusion among the participants.

We started to count the measurement time exactly when “start” was pronounced and stopped the time as soon as the participants indicated the measured result. Apart from measuring the time needed to complete the single tasks of determining the consumption of the individual appliances, we noted measurement mistakes and other relevant observations like the participants’ spontaneous reactions.

After this interactive start, we handed the participants a questionnaire, which was designed to reveal how the measurement technologies are perceived regarding subjective evaluation criteria such as fun of use, attractiveness, comfort, and comprehensibility[25]. The participants had to rank the three technologies regarding each criterion and we used the borda count ranking method to analyze the questionnaires [137, 138]. To validate and better understand the results, we additionally asked to rate the importance of the subjective evaluation criteria and additional criteria like price, time delay in feedback, accuracy, and availability. This allowed us to obtain a list with general requirements regarding energy measuring devices. Besides, people were asked which measurement technology they would actually buy and what their willingness to pay is like.

3.3.3.2 Part 2: Evaluation of the eMeter user interface and functionality

The second part of the study consisted of a guided interview to explore and discover the different functionalities of the eMeter user interface and their meanings. For that, each user had to accomplish different tasks, which involved various implemented features (e.g., determine highest historical consumption or current consumption, how current consumption compares to historical, standby consumption, etc.) and aimed at gaining a solid understanding of the application.

Then, we handed participants a questionnaire that aimed at a general evaluation of the application and at assessing the functionalities that are perceived most valuable from a user perspective. Again, the questionnaire was anonymously completed in an unobserved environment. We asked the participants to rate the importance of implemented as well as possible future functionalities. Moreover, users had to rate the complexity, usefulness, ease of use, ease of learning, and satisfaction of the mobile phone application. The latter four factors are taken from the USE, an established questionnaire for measuring the usefulness, satisfaction, and ease of learning of a user interface by Lund [139]. Additionally, we asked the participants to

[25] The criteria were previously obtained in discussion of a focus group that consisted of 3 experts from academia, 4 industry experts, and 1 employee of consumer organizations.

indicate their intention to use (once a week, once a month, or never), and the willingness to tell their friends (word of mouth). All items were rated on a five-point Likert scale.

The questionnaire closed with items regarding age, sex, and technological affinity. Technological affinity was assessed based on four items [140]: 1) my friends and colleagues often ask me what I think about new telecommunication technologies, 2) my friends and colleagues are better informed about new technologies than me (inverse), 3) I am always up on the latest technologies in my area, and 4) I think it is fun to test new technologies.

3.3.4 Selected Study Results

The conducted user study aimed at evaluating mobile phones as energy feedback devices. In this section, we focus on selected results of the evaluation and discuss them with regard to general requirements people have concerning energy feedback systems. We first present results directly related to the eMeter user interface and its measurement functionality eMeasure. Thereafter, we report on results concerning the general functionality and the design of energy monitors.

3.3.4.1 Evaluation of eMeasure

The first task of the study consisted of estimating the consumption of different appliances typically used in a residential environment. The results show that users have only limited understanding regarding the consumption of household appliances. These findings are in line with [141]. The consumption and the difference of standard light bulbs and energy saving lamps seems well known (error of ~10% and ~25% respectively), the consumption of other electronics is estimated well of the mark (error of ~55%, ~1100%, and ~1400% for the kettle, the aggregated consumption, and the game console respectively). Due to the limited knowledge about energy consumption, people are not able to tell whether the gathered measurement result utilizing one of the three different energy monitoring technologies was right or wrong. The spontaneous reactions of the participants indicated that they were often surprised by the small energy consumption of the game console and the kettle.

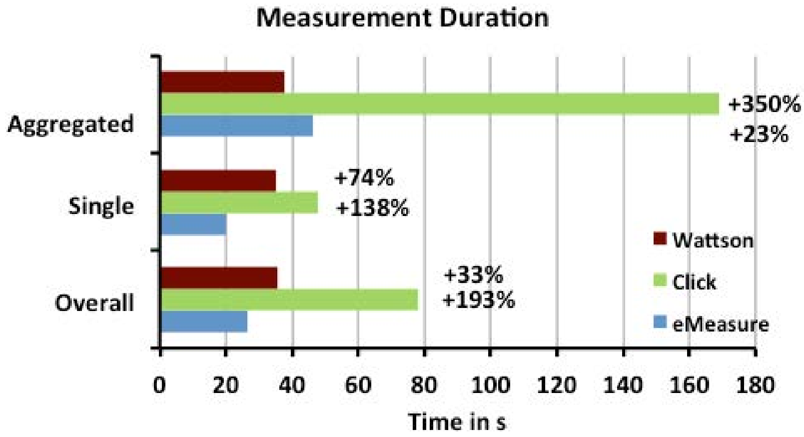

Figure 3.9 Measurement duration participants required for determining the electricity consumption for the three different energy monitoring solutions utilized in the user study.

During the study, we measured different objective performance criteria for the three energy monitoring tools and differentiated between the required time to perform a measurement, the number of wrong measurements, and the measurement accuracy user achieved utilizing the different technologies. Figure 3.9 shows the mean time in seconds participants needed to determine the electricity consumption for each technology as well as the percentage difference compared to the fastest technology. The iPhone application eMeasure performs best over all measurements and for single measurements. Participants needed on average 27 seconds to perform a measurement whereas it took 33% (36 seconds) longer utilizing Wattson and over twice as long (78 seconds) utilizing Click. As expected, Click also performs worst regardless of the task (i.e., measuring the power consumption of an individual or consecutive appliances). A more detailed look on the tasks performed reveals that participants were faster using Wattson when determining the consumption of two devices in a row. This is due to the problems users experienced with the mobile phone application when measuring the second appliance. Then, participants often forgot the procedure or did not know how to continue with the second measurement. This is also reflected in the high number of measurement failures. In roughly one third of the cases (8/25) participants could not complete the task of measuring the consumption of two appliances with eMeasure. Over all 125 measurements per technology, users were not able to determine the consumption utilizing Click twice, made ten errors using Wattson, and conducted 19 errors when measuring with eMeasure. Additionally, we asked participants whether they had used an energy measurement technology before and could confirm that this item has no influence on the measurement duration.

Figure 3.10 shows the measurement error, which is an admeasurement for the accuracy of the technology. For each technology, the depicted bars indicate the difference between the consumption determined by the participants and the real consumption verified with an electricity meter before the study was conducted. For all tasks, participants achieved the best results using Click slightly ahead of eMeasure and Wattson. For the first two technologies the measurement error resides within 5% for individual measurements. Utilizing Wattson the error ranges from 5% to 16%. This is mainly due to the continuously fluctuating consumption on the display unit. The high error (~15%) of the iPhone application when measuring the aggregated consumption of two devices originates from the difficulties users experienced with built-in support for subsequent measurements and the measurement procedure. The users' consumption estimate shows that they were not able to justify whether their determined result is incorrect and thus requires repeating the measurement.

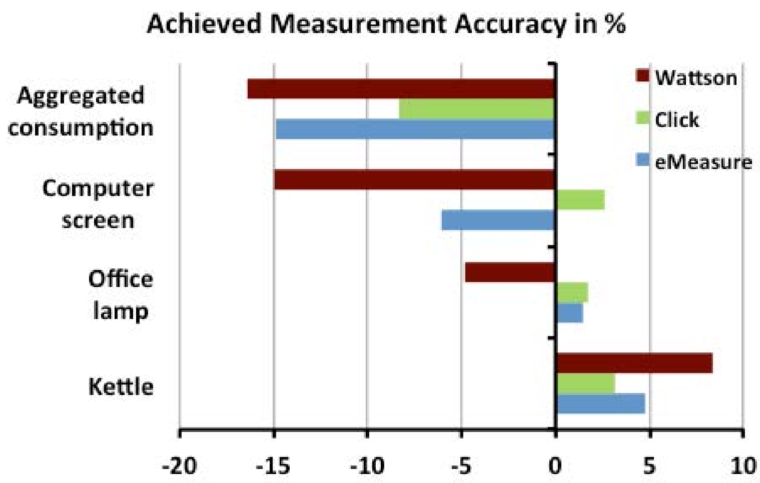

Figure 3.10 Measurement accuracy achieved by participants utilizing Wattson, Click, and eMeasure relative to the actual consumption of the appliance verified before the experiment.

3.3.4.2 Evaluation of the eMeter user interface

The general evaluation of the mobile application is depicted in Figure 3.11. It shows that participants had understood the user interface and the underlying functionalities. On a scale from one (lowest) to five (highest), participants rated the ease of use (4.04), ease of learning (4.04), and satisfaction with the application (4.16) all significantly above average (means in brackets). Taking into account that we covered a wide age range (18 to 51 years) and only five participants were iPhone users, we regard this as a

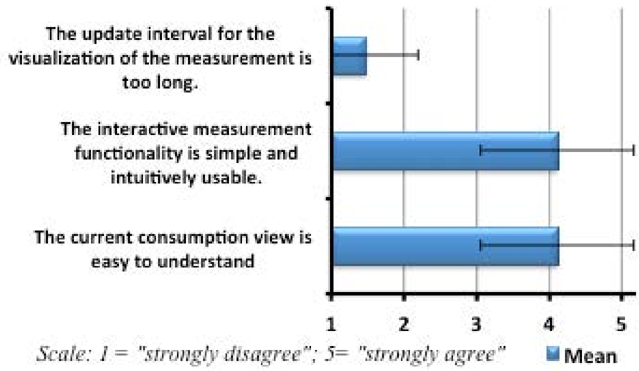

Figure 3.11 General evaluation results of the eMeter user interface.

very positive response for a prototypical application. General results also indicate that the feedback latency was perceived as more than satisfactory, the measurement functionality as easy to handle, and the individual views as easy to understand.

The left of Figure 3.12 illustrates the results for the mobile application assessment in terms of usefulness, intention to use, and word of mouth. Besides depicting the overall mean of the whole sample and the standard deviation in error bars, the figure provides a more in-depth breakdown in terms of technological affinity. It was assessed based on four items mentioned in Subsection 3.3.3.2. People with a scale mean of 3.64 or less were grouped to “not technological affine” (N=10), whereas people with a higher mean were assigned to the group “technological affine” (N=15). The evaluation on the perceived usefulness, the intention to use, and the word of mouth was only marginally affected by the technology affinity (except for the frequency of use where the technophile users indicated higher scores). The application reached high scores especially for positive expectations towards saving effects and knowledge gains. The word of mouth effect is significantly high. It thus offers potential for utilities or smart meter manufactures to positively influence their image providing such an application. A large part of the participants agrees that the application is useful. The prime use of the application – especially with technophobe participants – is seen in increasing the knowledge about the electricity consumption of individual devices. In consistence with the survey results, the claimed external social motivation (“demonstrate good behavior to others”) to use such a mobile application remains low. Users in general do not feel the need to express their pro-environmental behavior to others (or are not willing to admit it), but technophile users would rather do than technophobe.

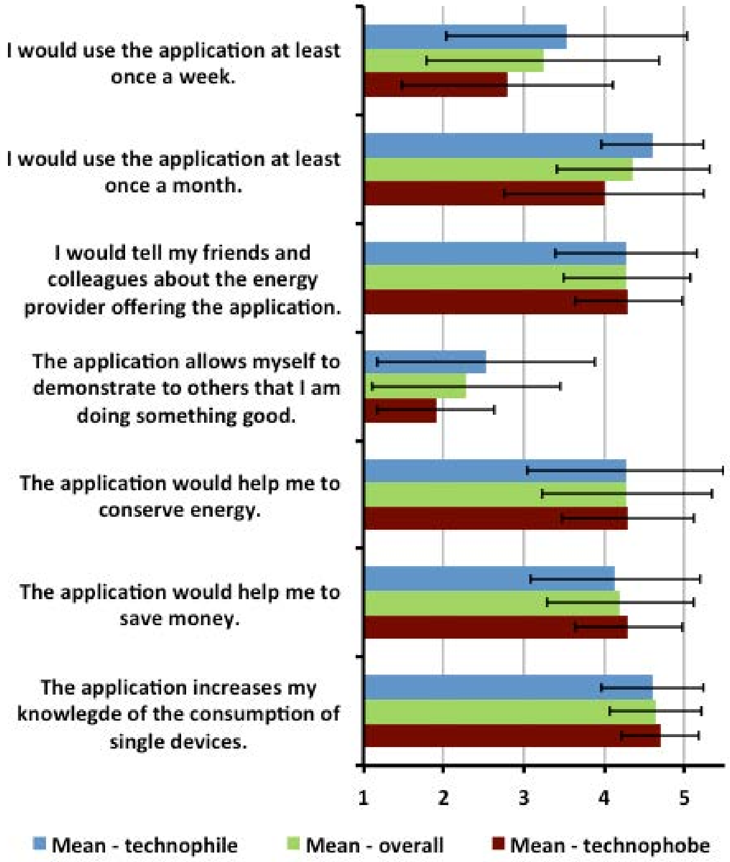

Figure 3.12 Mobile application assessment: Usefulness, intention to use, and word of mouth.

3.3.4.3 Evaluation of Objective Performance Criteria and Functionalities of Portable Energy Monitors

Figure 3.13 illustrates the results of the importance of the different performance criteria on a scale from 1 – most important to 9 – least important. Besides depicting the overall mean of the whole sample and the standard deviation in error bars, the figure provides a more in-depth split up on the sex. Overall participants perceive comprehensibility, comfort, and ease of learning to be the most important characteristics. Less important are design and fun of use. Women show a slightly different perception than men. The top two items remain, but then price, availability, and feedback time become slightly more important factors.

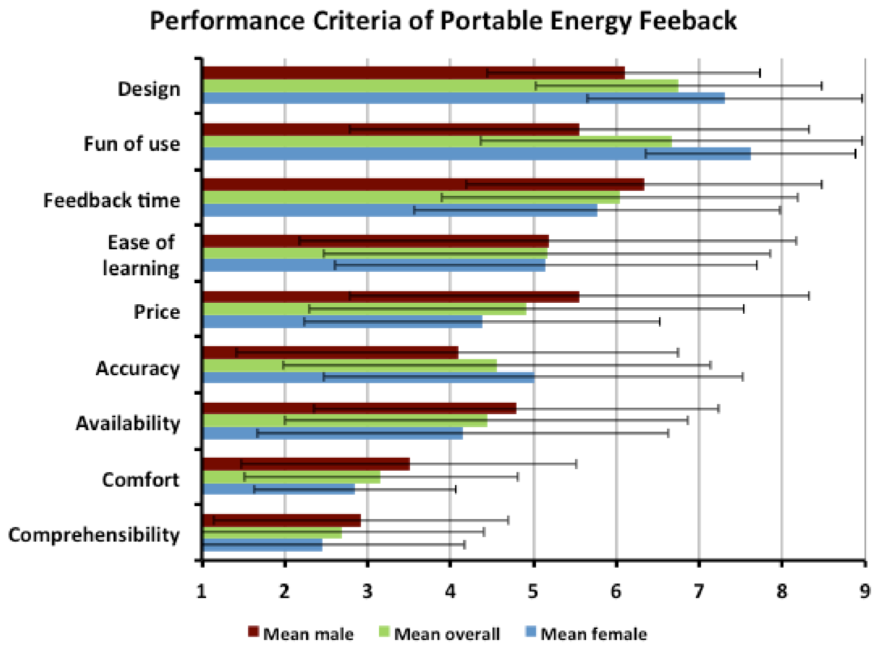

Figure 3.13 Performance criteria of portable energy feedback systems.

In order to evaluate which functions with respect to energy consumption feedback are perceived as most valuable on a mobile phone, we asked the participants to indicate their impressions on the following functionalities: Real-time visualization of the total consumption; visualization of the household's standby consumption; comparison of the current consumption with the historic consumption; costs of recent months; consumption of individual devices; consumption of recent months; projections of yearly cost on device level; efficiency grade of appliances; overview of biggest energy guzzlers; comparison of the consumption with the one of friends; possibility to show others my appliance pool; possibility to set a saving target.

Figure 3.14 provides an overview of the assessment of mobile electricity consumption feedback features. It depicts an in-depth view on the rating per functionality sorted in an ascending order according to the overall mean value (shown on the right). Overall, we find that participants value at-a-glance-feedback on their most prominent energy guzzlers most (mean of 4.72), followed by those functionalities that increase the knowledge about consumption or cost. The real-time view of the entire consumption achieves similar ratings with a mean of 4.6. For both, 96% of the participants indicated the importance of these functionalities. Down to a mean value of 4.16, still 80% of the participants perceive functionalities such as standby consumption and consumption of an individual device important. All these functionalities have in common that they provide an action-guiding feedback from which users can directly draw effective measures to

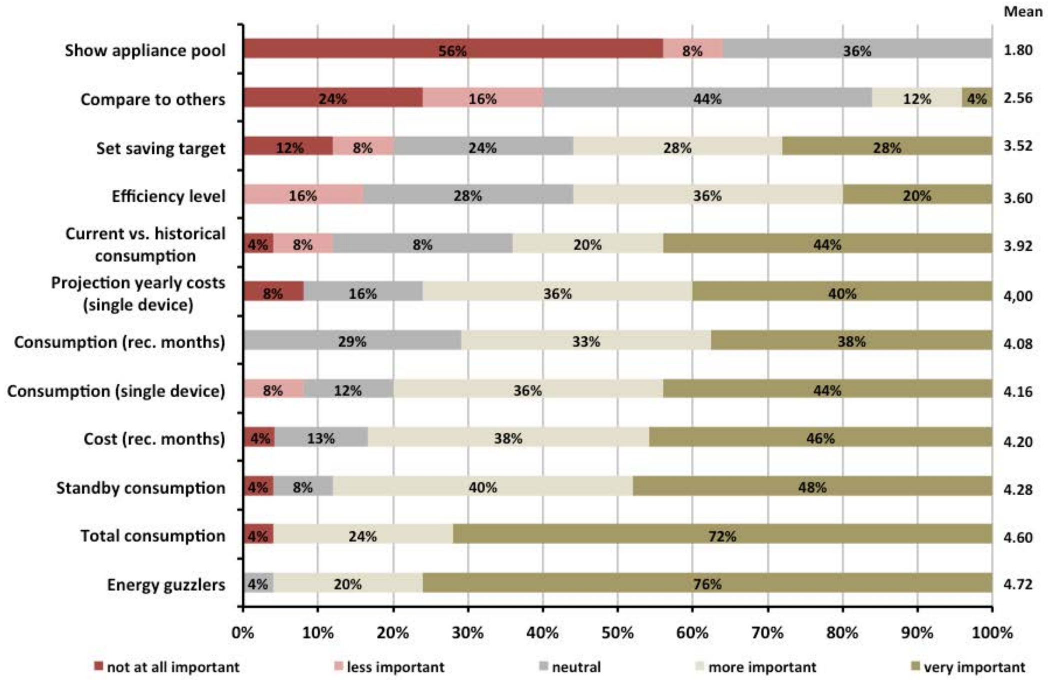

Figure 3.14 Evaluation of the functionalities provided by portable energy feedback systems.

lower their electricity usage. Surprisingly, cost of the recent months receives a high importance ranking of 84%, although the feature itself is not action-guiding. Below a mean value of 4.16 the picture changes. Those functionalities that present aggregated information from which people cannot imply a direct action (e.g., consumption of recent months, comparison of the current vs. historical consumption) receive significantly lower ratings and reside in the bottom half of the ranking. Functions aiding users through motivational support (e.g., set a saving target) are not perceived as important, nor are those that deal with social aspects (e.g., compare to others). They reside in the bottom third and especially the latter two receive a low importance rating of 16% and 0% respectively. A closer look on the technological affinity reveals that in general technophiles rate the functionalities higher. However, technophobes value features that present action-guiding (e.g., energy guzzlers) and device-level information (e.g., consumption of an individual device) over aggregated information (Figure 3.15).

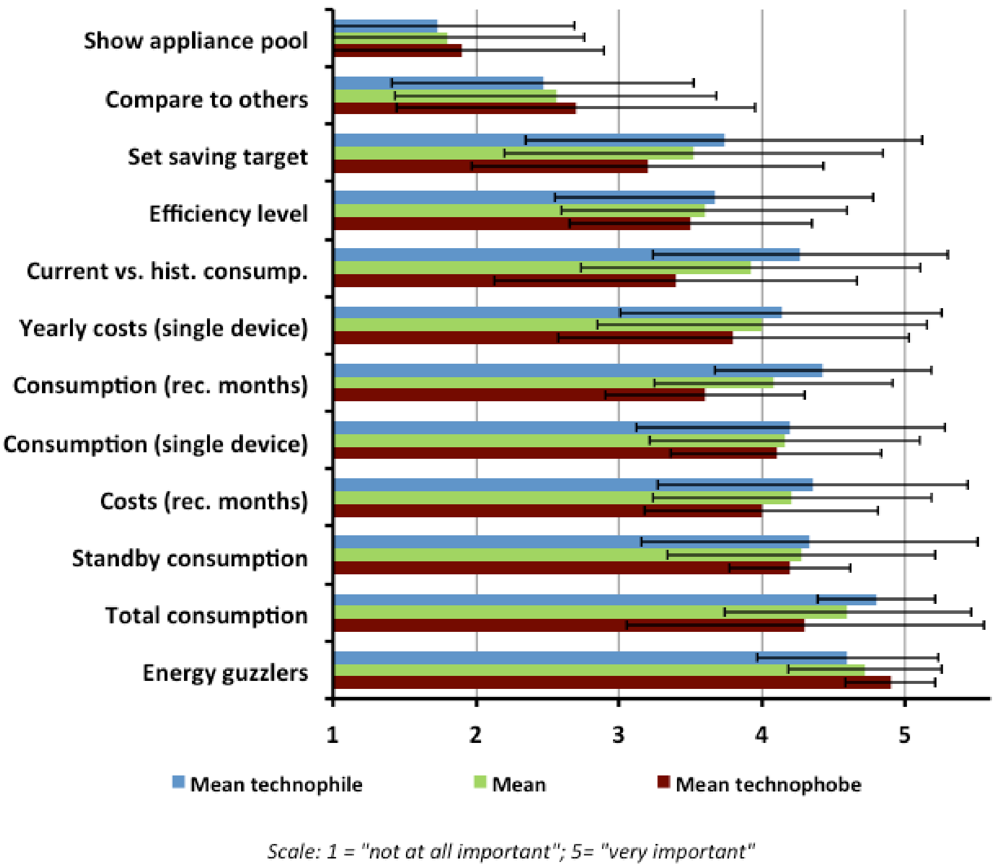

Figure 3.15 Evaluation of the functionalities provided by portable energy feedback systems.

3.4 Real-world Deployment

In this section, we explain the experimental setting of the real-world deployment we used to gather further insights on the long-term use and usability of the eMeter smartphone application. Thereafter, we report on selected quantitative and qualitative deployment results.

3.4.1 Experimental Setting

We installed the eMeter system in different households throughout Zurich, Switzerland (see Section 2.6.2). The participants for the real-world deployment were employees (and their families) of a local utility. Four households with different characteristics were selected. Table 3.1 provides an overview of the type of participating household as well as the major electrical appliances that characterize the electricity consumption of the household.

Table 3.1 Overview of the participating households of the real-world deployment.

Number	**Type**	**Inhabitants**	**Washer**	**Water heater**	**Dishwasher**	**Stove**
1	House	2	Yes	Yes	Yes	Yes
2	Apartment	1	Yes	Yes	Yes	Yes
3	Apartment	1	Yes	No	Yes	Yes
4	Apartment	4	Yes	No	Yes	Yes

The eMeter system setup period dated from April to June 2010. After the initial installation, the system kept running in every household for at least one month before the experiment started. During this time, we ensured the backend of the system and the user interface were operating correctly. The experiment started in the middle of July 2010 and lasted for 105 days until November 2010. On the first day, we held an introductory workshop, in which participants installed the eMeter iPhone application and were briefed on the upcoming experiment. In addition, we asked users to complete a questionnaire to reveal their expectations of the application and the experiment, their intention to use as well as to check their current knowledge about the electrical consumption of their household. Note, at the time of the real-world deployment the engagement view of the user interface was not yet fully developed and thus had been removed to avoid confusion.

During the whole time of the experiment, users were free to use the application to get aware of their electricity consumption. In addition, we asked the participants to solve one given task per month of which they were informed one week ahead of time via email. We thus ensured participants had sufficient time to understand the task or ask questions if necessary. The following three tasks had to be conducted during the experiment:

- Device Hunt: Participants had to explore their house in order to discover how many electrical devices they own.

- Meter Sleep: Participants had to try to switch off as many electrical appliances as possible or even disconnect them form the mains in order to determine the lowest possible power consumption of their household.

- Daily Routines: On one typical weekday morning and one evening, participants had to log in a notebook all operation events of every appliance that was used during that time in the household.

In November 2010, we concluded the experiment with a result workshop. First, we conducted a guided interview for one member of each household. For roughly 30 min participants, for instance, reported on technical difficulties and their typical uses cases of the eMeter application. We further asked about experienced any changes in their energy literacy. As final task of the interview, participants had to try to identify their household out of a set of different electricity load profiles. Last, the interviewees had to complete a questionnaire, which aimed at identifying how far their original expectations were met, and further whether their consumption awareness had increased over the time of the experiment. The whole interview and questionnaire was conducted in a neutral environment. Apart from the interviewee only the experimenter who was leading the interview and a second person who was taking notes were present. The workshop concluded with a general presentation of the results in front of a bigger audience.

In order to be able to measure participants' use of the eMeter user interface, we slightly modified the mobile phone application. That is, we incorporated timers and counters in the background of the application and invisible for users that measured the use of the application and its different views (e.g., time spend in a view, number of application starts, number of conducted device measurements, etc.). At the end of the experiment, this data was collected and used for the quantitative analysis of the application use.

3.4.2 Selected Deployment Results

In the following, we report on selected quantitative and qualitative results of the real-world deployment. The numeration of the households corresponds to Table 3.1. Over the course of the experiment (105 days), participants accounted for 318 application sessions that resulted in a total application use of roughly 16 hours (955 minutes). One of the tasks (namely the device hunt) provided insights on the amount of electrical appliances operated in the households. The number varied between 37 (household four) to 64 (household one). Table 3.2 provides an overview on the electricity use of the participating households. Besides the consumed electricity over the duration of the experiment, it shows the projected yearly consumption and the comparison to the consumption of an average household. It is interesting to see that all households are most likely to consume more electricity than the average Swiss household of corresponding characteristic (i.e., same size range, persons, etc.). The base load that has been determined by the eMeter system varies between 73W and 299W. The latter translates to costs up to 524CHF just for the electricity use of appliances when occupants are not at home or sleeping. In particular, the base load of household three accounts for 98% of the total consumption of the household.

The second task of the experiment aimed at lowering the electricity consumption in the household as much as possible. We find that the consumption of all households can be significantly reduced when switching all major appliances off. For example, household three lowered its base load to 25W. The lower boundary for the other households resided in the same range between 20W and 30W, which represents a large electricity conservation potential.

Table 3.2 Comparison of the electricity use of the participating households.

Number	Observ. Period	Projection Year	Average Use[26]	Base Load	Meter Sleep	Cost CHF
1	1429 kWh	4954 kWh	2900 kWh	295W/49%	-	517
2	892 kWh	3091 kWh	2550 kWh	73W/19%	32W	128
3	714 kWh	2476 kWh	1550 kWh	299W/98%	25W	524
4	1116 kWh	3868 kWh	3300 kWh	178W/37%	20W	312

Having a closer look on the daily consumption, Figure 3.16 illustrates a comparison for weekdays and weekends of the average consumption of the participating households. The black bar indicates the base load of the corresponding household. The family house (H1) uses most electricity with an average weekday consumption of just under 14kWh and 16.6kWh on weekends respectively. Each household shows an increase in consumption on weekends between 9.6% and 43.1%, which is a typical behavior due to higher occupancy compared to weekdays. It is worth noting the high difference regarding the contribution of the base load to the overall electricity use. Household two has the lowest base load and hence seems to have energy conservation measures (e.g., switchable power plugs) in place. Its higher average consumption compared to household three can be explained through the electrical water heater that is in charge for the warm water supply.

[26] Source: http://www.ckw.ch/internet/ckw/de/privatkunden/service/tipps/stromverbrauch.html

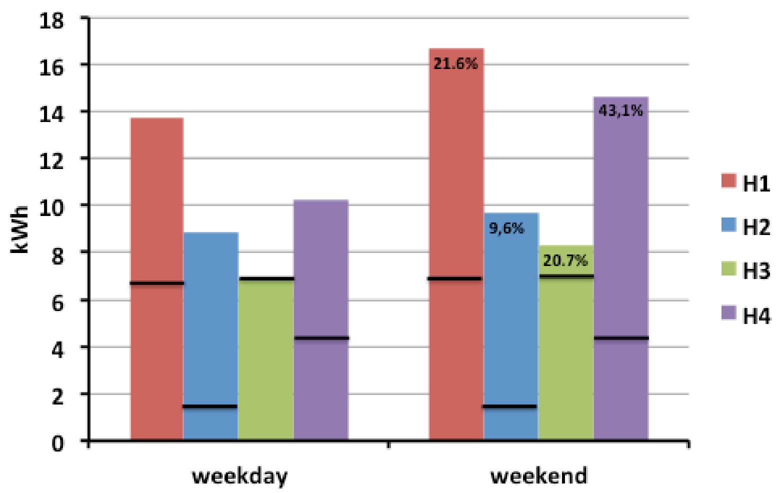

Figure 3.16 Electricity use details of the individual households (H1 to H4 from left to right) participating in the real-world deployment. Comparison of weekday and weekend use shows a higher consumption for all households (increase outlined in per cent). The black bar indicates the base load of the household.

Figure 3.17 illustrates the evolution of the utilization of the eMeter user interface. The figure shows the time in seconds individually for each view and calendar week aggregated over all participants. We see that participants were using the smartphone application a lot at the beginning of the experiment. However, usage already decreased rapidly after the first week and after the first 1.5 months the application was only used occasionally to monitor the electricity consumption. We further find that the current consumption view and the history view were used most frequently and continuously. They also account for the largest absolute share of time.

Figure 3.18 shows a more in-depth view on the utilization of the application. The left part of the figure depicts the average time per session in seconds over the total number of sessions per individual household. The radius of the circle provides a measure for the number of appliances that have been measured with the user interface. The average usage time of the application varies between 100 and 525 seconds. Household three measured most devices out of all participants. While roughly accounting for the same amount of appliance sessions than household three, household two spent significantly more time using the application (524s compared to 221s per session). Household four was most actively using the application and accounted for more sessions the rest of the participants together. However, the energy monitor is used more for quicker checks on the electricity consumption (e.g., to determine the power consumption of a device), while the other households took more time per session to explore what is going on in the household. Note, that household one only accounted for 19 sessions and stopped using the application due to personal reasons.

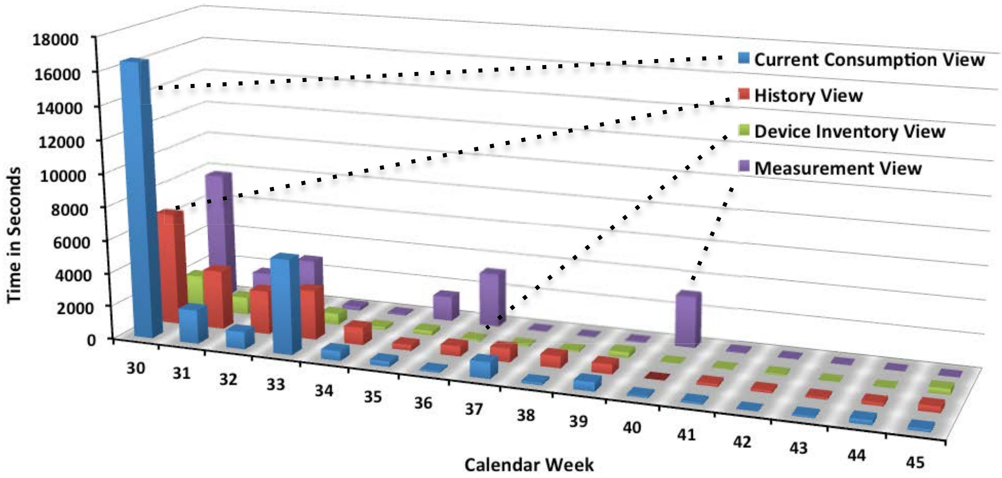

Figure 3.17 Evolution of the utilization of the different application views (aggregated sum over all participants per calendar week).

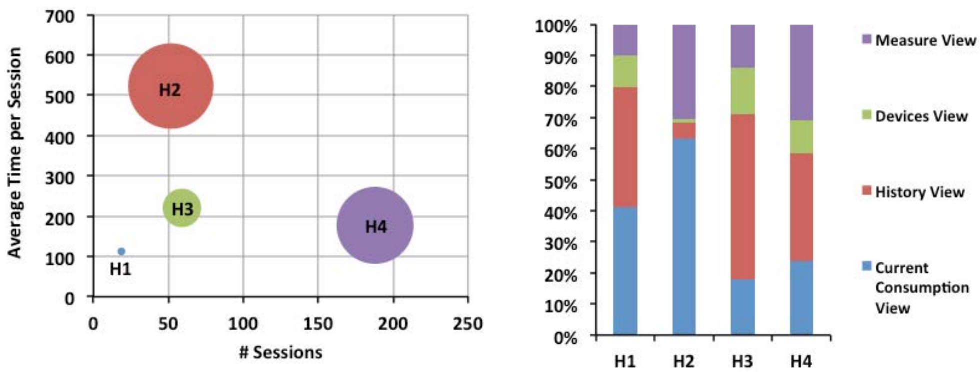

Figure 3.18 eMeter user interface utilization details: Frequency of use (left) and average time per view (right).

On the right of Figure 3.18, an analysis of the average time spent per view and individual household is visualized. Out of all participants, household four spent most time in the measurement view, which confirms the previous finding. We further find that there is no single view dominating over all households, but each household is using the application differently. Household two and four spend roughly the similar time using the application, but show diverging usage profiles. Household two mainly uses the current consumption view (~70% of the time) whereas household uses the historic electricity consumption feedback most. For household three, the use of the historic feedback also dominates over the use of the other views.

The general evaluation of the complexity of the mobile application was evaluated in a questionnaire at the result workshop. It confirms the previous results of the user study. On a scale from 1 (lowest) to 5 (highest), participants rated the ease of use (4.5) and the simplicity to learn how to use the application (4.5) well above average (means in brackets). The current consumption view was perceived as simple to understand (4.25), but participants indicated they had problems relating the historic consumption to the incurred cost (2.5). In addition, participants thought it took too long to measure the consumption of individual appliances (4.25). However, compared to the user study, participants had no direct comparison for the time it typically takes to get feedback on the consumption of an individual appliance.

The concluding interview was designed to provide further insights on the purpose of application use, eye openers participants experienced over the course of the experiment, and general improvements of the application. Besides the general use case of monitoring their residential electricity consumption and increasing consumption awareness, participants reported they used the application to identify the power consumption of individual appliances, often demonstrated the application to colleagues and friends, and used the application to observe what was going on at home at time the participant was not present.

One participant explicitly stated: *"I also used the application to check my electricity consumption when I was not at home. I like that I can see whether my kids are home yet or whether they have already left home and even if the lights are still burning. When they forgot to switch them off, I immediately called them to enforce that they switch lights off. [...] In this context, an alert functionality would be helpful that notifies me in case my current consumption is above my base load at a specified point in time"*. Another participant acknowledged, *"the application increases the awareness tremendously at first, but after two weeks, it needs further mechanisms that motivate me to engage with my electricity consumption"*.

The survey and the interview show that participants improved their energy literacy through the application, although they already had a profound knowledge compared to average users because of their profession (i.e., employees of a local utility). In the interview, we presented all participants four typical load profiles. Out of these load profiles one belonged to their individual household. All participants were able to identify their household based on different characteristics. Some were looking for the typical time they get up or leave the house, other recognized specific loads of appliances in the load curve that allowed them to identify the correct profile (Figure 3.19).

Participants also reported on specific eye openers that occurred during the experiment. Household four identified its fridge to be inefficient. In the history view, the participant observed the duty cycle of the appliance, which appeared to be atypically long. Household two identified an auxiliary water heater under the kitchen sink to contribute considerably to the overall energy consumption. It kicks in every hour consuming a significant amount of electricity that was unknown before (Figure 3.19). Household three realized its particularly high base load that resides above the one typical for a one-person household. Various IT equipment is known to be the cause, but the participant claims the PC, router, etc. are required to operate 24/7. Thus, although the potential to save energy was realized, no direct conservation measure was applied.

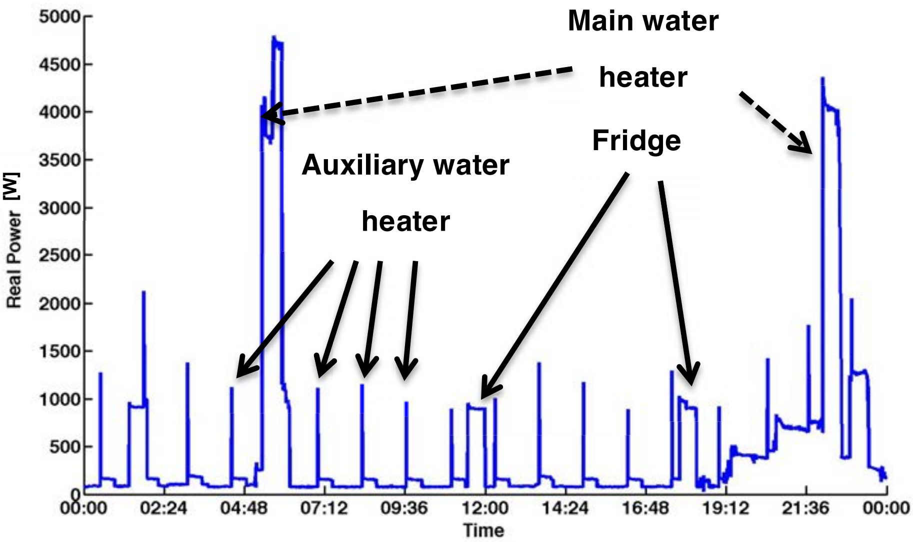

Figure 3.19 Typical load profile of one of the test households on a weekday. The participant is getting up at around 5am, leaving the house at 6am, and is returning from work at 5.30pm. The small peaks were identified to originate from the auxiliary water heater, which was unknown before the experiment and is kicking in almost every hour.

3.5 Discussion

The evaluation of the user interface that uses real-time data from the eMeter system in a user study and in the real-world deployment in four households confirms the suitability of mobile phones as energy feedback devices. Exposing people to a functioning prototype was crucial for us to gather experience with the application, while at the same time participants better understood the usefulness of the application. Our results thus extend the body of work on energy feedback systems and can serve as a starting point for further application development in this field. In particular, our findings that directly relate to the developed smartphone application are the following:

- In a pretest, before the development start, we found that the existence of a clear and simple to explain use case behind energy consumption feedback systems is a key success factor, whereas when left to the imagination of potential users, energy consumption feedback applications receive only medium ratings.
- Overall, the eMeter application receives positive ratings regarding general subjective performance criteria (e.g., complexity, usability, and understandability). Participants perceived the application as more than satisfactory for an energy feedback monitor and the individual views as easy to understand. These results are very affirmative for a prototypical application and confirm the suitability of the eMeter user interface as electricity feedback monitor.
- We could further confirm the functionality of the measurement feature eMeasure and suitability of the measurement process by comparing eMeasure to two other commercially available energy measurement solutions in terms of different objective performance criteria. Our results show that to determine the consumption of individual appliances, participants were fastest using the smartphone application, while the error margin (~5%) resided within the same range of applications that have been specifically developed for that use case. Moreover, eMeasure is perceived as simple and intuitively useable. We assume this is sufficient to increase the transparency when it comes to the energy usage of individual devices and provide users with a general understanding on the consumption of appliances.
- The user study revealed that users experienced problems using eMeasure to consecutively determine the consumption of multiple appliances. This was partly due to unclear instructions provided in the user study manual before the experiment, but also because the designed measurement flow turned out to be misleading. To pre-

vent these errors, we redesigned the eMeasure and simplified the corresponding view of the user interface to just provide one measurement option at a time.

- The application helps people to achieve a better understanding of their electricity consumption and the consumption of individual devices. Participants reported that their energy literacy has increased through using the application in the real-world deployment. After the experiment, participants were able to specify their base load, the consumption of individual devices (e.g., hair drier, fan, and TV) and could identify the load profile of their household.

- The real-world deployment showed the lower use of the application as time of the experiment progressed. At the start of the experiment, participants were using the application on a daily basis. However, usage decreased fast once the initial curiosity of the participants had been satisfied. This fact is a known problem of mobile applications and highlights the need for concepts that aim at engaging the user over extended periods of time [142, 143].

We implemented the most promising feedback features and evaluated the perceived usefulness of the different functionalities with our application in a user study. Moreover, to address different target groups appropriately, we focused on the individual difference between technologically and non-technologically oriented people. In the real-world deployment we monitored the use of the application and the different application views. Findings related to the individual functionalities of energy monitoring applications are:

- The real-world deployment confirmed that participants use the application for different use cases. This is reflected by the time participants were utilizing the different views of the application. We conclude that different user types have to be addressed with different functionalities. This goes along with the results of [59], which clearly state there exists no one-size-fits-all solution for energy feedback systems.

- In the user study, we found that the knowledge-increasing functionalities as well as those from which monetary savings can be directly implied are perceived as most useful. In contrast, functionalities that present aggregated information receive lower scores.

- The survey results of the pretest as well as the user study indicate that social motivation is so far not an important factor – at least not consciously – in terms of energy consumption feedback.

- To target technophobe users, a closer look revealed that it is important to implement simple, easy to understand, and action-guiding feedback that goes beyond aggregated information, such as a list of energy guzzlers.
- The measurement functionality that enables users to interactively determine the consumption of switchable appliances received good ratings. However, when compared to the other functionalities, its perceived value was not ranked amongst the top features. In the interview, participants of both the user study and the real-world deployment revealed that they think the interactivity introduced through the measurement feature makes it one of the key components of the eMeter user interface. However, being only provided with the mere consumption value at the end of the measurement process is neither meaningful nor action-guiding, which led to lower scores for the feature.

3.6 PowerPedia – A Collaborative Platform for Providing Meaningful Appliance Feedback

When it comes to conserving electricity, it is crucial for users to know how much energy is consumed by individual appliances [53]. For this reason, we equipped the user interface of the eMeter system with a measurement functionality. The simple feature enables users to determine the power consumption of individual appliances. However, the user study showed and the real-world deployment confirmed that the technical feedback (i.e., the mere consumption value that is displayed at the end of the measurement process) provided by the application is too dry and intangible for most users. Participants lacked the ability to position the consumption of an appliance in a bigger picture that would allow them to draw conclusions and take effective measures.

The comment of one participant clearly describes the situation: *"The measurement functionality is a nice interactive feature that helped me to realize there is quite a difference in the consumption of electrical appliances. However, it would be very helpful to have a measure for the efficiency of a particular device compared to similar devices of the same type."*

To address this shortcoming, we extended the eMeter system and developed PowerPedia. It aims at providing behavior-influencing feedback over and above mere consumption values. By integrating a community platform

– a Wikipedia for electrical appliances – PowerPedia enables users to identify and compare the consumption of their residential appliances with that of others. It thus helps users to better understand their electricity consumption and take effective action to save electricity.

In the following, we first describe the general concept of PowerPedia before presenting its architecture and finally the extensions to the user interface. The key idea is that PowerPedia continues to support users after the measurement process is completed. Then, users can upload the measured device to PowerPedia and compare the consumption against the consumption of other devices in the same device category that have been previously published by other users. PowerPedia provides an efficiency ranking based on the device category and specific energy-saving possibilities of devices. The content is thereby generated through the collaboration of users. For example, device and category-specific energy-saving measures as well as product ratings can be uploaded and thus shared with others. As add-on functionality, PowerPedia offers users direct integration into social networking platforms such as Facebook and Twitter. In order to keep track of the most energy-efficient devices in each category, PowerPedia embeds a harvester module that automatically updates the appliance efficiency ranking by incorporating the best-performing devices gathered from different consumer organizations[27]. The harvester also initializes PowerPedia with a first set of energy-efficiency measures.

The integration of PowerPedia in the eMeter architecture is illustrated in Figure 3.20. The platform is realized as an additional component to the eMeter system, which we described in full detail in Chapter 2. PowerPedia runs as a central entity on a dedicated server on the Internet. It consists of the SignatureServer that stores information about residential appliances, a lightweight harvester module that is used to automatically update PowerPedia, and an user interface to access the provided functionality.

The harvester is used to initialize PowerPedia with a first set of devices and energy-saving measures as well as to update the database on a monthly basis with the most energy-efficient appliances in each category. To do so, the harvester scans dedicated external consumer organization websites to acquire and extract the data before translating it into JSON. The result of the scan is then passed to the SignatureServer, which updates its list of devices in the database.

[27] e.g., Top10, www.top10.ch; EnergyStar, www.energystar.gov.

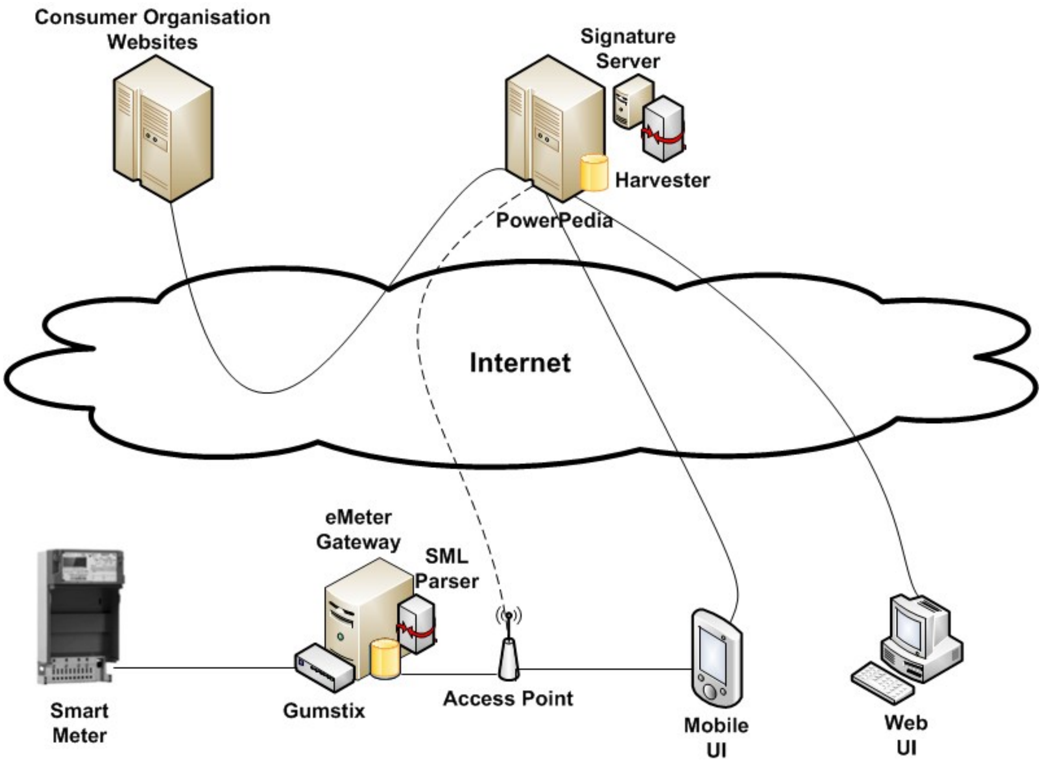

Figure 3.20 System overview of the eMeter system after the extension through PowerPedia. The mobile user interface and the web user interfaces access the PowerPedia directly. The PowerPedia consists of the SignatureServer that hosts all functionality and the harvester that automatically updates the SignatureServer with the most energy efficient devices and conservation tips gathered from consumer platforms in the Internet.

The SignatureServer incorporates the main functionality of PowerPedia. It follows the RESTful system design of the eMeter architecture, and is written in PHP using the same Recess framework as the eMeter system. Table 3.3 provides an overview of the most important functionality that is provided by the RESTful PowerPedia API. It details the URI that can be called, together with the corresponding HTTP verb to perform the action indicated. As an example, Figure 3.21 shows the JSON representation of device number 96 that is stored on the SignatureServer. It is the response to the following GET request: http://[IP]/powerPedia/device/96.json.

The architectural structure of the SignatureServer is shown in the UML diagram in Figure 3.22. The following object models are implemented:

- User: The user model is used to store user authentication information. This includes user name and password as well as first and last name. Each user can have multiple smartphones that are linked to the user's id.

- Category: The categories are used to group appliances of the same category (e.g., lights). Categories are structured hierarchically, meaning that a category can have multiple sub-categories.
- Device: The device model contains fields for the device name and description, a picture, the manufacturer, the type, the consumption value, time information, and an efficiency rating (see Figure 3). Each device is linked to exactly one category.
- Recognition: The recognition model is used to store the data that is collected when users measure electricity consumption and subsequently upload it to PowerPedia. A recognition is linked to a particular device, to the meter providing the information, and to the model of the user who uploaded the recognition.

Table 3.3 Overview of the RESTful API of PowerPedia.

URI	HTTP Method	Description
/powerpedia/category/	GET, POST	Get index, insert new category
/powerpedia/category/$id	GET, PUT, DELETE	Modify category
/powerpedia/category/$id /allDevices(Details)	GET	Get information on all devices (details) in the selected category
/powerpedia/category/$id/allSub	GET	List all sub categories of category id
/powerpedia/device/	GET, POST	Get index, create new device
/powerpedia/category/$id	GET, PUT, DELETE	Modify category
/powerpedia/device/$id/compare/ $category	GET	Compare whether device is already existing
/powerpedia/device/$id/rate/$watt	GET	Get the efficiency of a device
/powerpedia/recognition/addToDevice/$deviceId	POST	Add a recognition to an existing device
/powerpedia/recognition/$id	GET, PUT, DELETE	Modify existing category
/powerpedia/smartmeter	GET, POST	Get index, add smart meter
/powerpedia/smartmeter/$id	GET, PUT, DELETE	Modify smart meter information
/powerpedia/tip/	GET, POST	Get index, insert new energy conservation tip
/powerpedia/tip/$id	GET, PUT, DELETE	Get tip or modify tip
/powerpedia/user	GET	Index of all users
/powerpedia/user/$id	GET, PUT, DELETE	User management

```
{
    "device":{
        "id":"96",
        "name":"Osram - Mini Globe",
        "description":"Power (Watt)",
        "picture":"www.topten.ch/[...]/mglobe-e27.jpg",
        "manufacturer":"Osram",
        "type":"Mini Globe",
        "categoryId":"19",
        "createdOn":1272312189,
        "avgWatts":"7.0000",
        "efficiencyrating":{
            "count":4,
            "position":1,
            "best":7W,
            "worst":9W,
        },
        "userId":"5"
    },
    "user":{
        "id":"5",
        "username":"m",
        "password":"x",
        "createdOn":1273790766
    }
}
```


Figure 3.21 JSON representation of device number 96 stored on the SignatureServer.

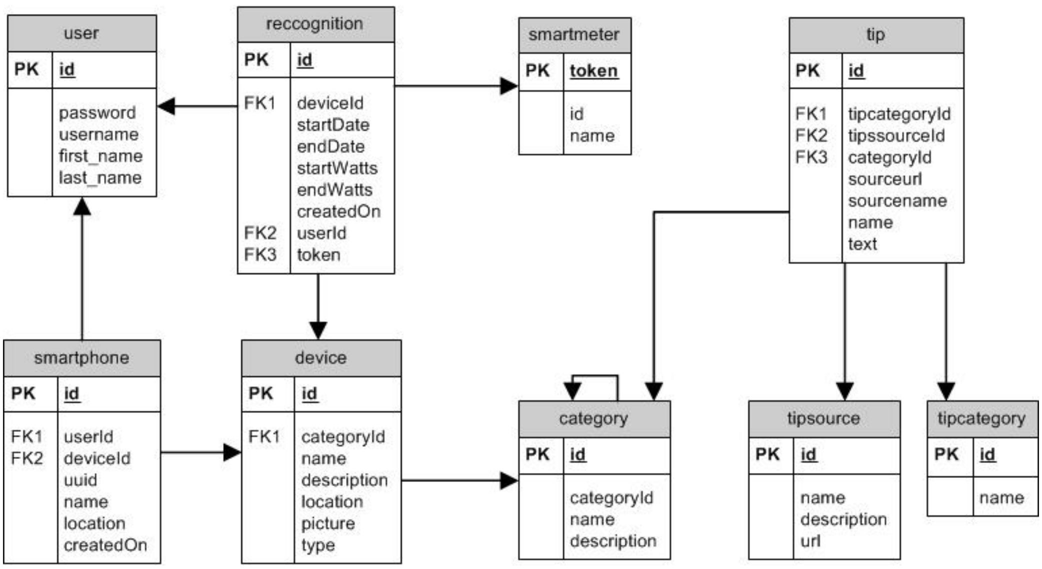

Figure 3.22 UML object model of the SignatureServer.

- Tip: The tip model contains energy conservation tips that consist of a tip name and source information. It is linked to different device categories and tips categories. Tips categories are used to group different energy saving tips (e.g., all tips from a particular user).

PowerPedia offers a stand-alone web user interface that enables users to browse through the different appliance categories and check for the most energy-efficient appliances. It requires creating an account before users can contribute to the platform. To make full use of the functionality provided by PowerPedia, we extended the mobile user interface and incorporated features of PowerPedia into the eMeter system. Let us briefly recall the process for measuring the power consumption of an appliance, to see how PowerPedia is supporting users with meaningful information within this process (Figure 3.23). Users initialize the measurement by pressing the start button in the measurement view. After switching the appliance under measurement on or off, its power consumption is displayed on the user interface within two to ten seconds. Where users before were confronted with the mere power value of the appliance, we now integrated the functionality of PowerPedia. When storing the measured device in the device inventory, users are now offered the possibility to upload the measured appliance to PowerPedia in exchange for an efficiency ranking and specific energy con-

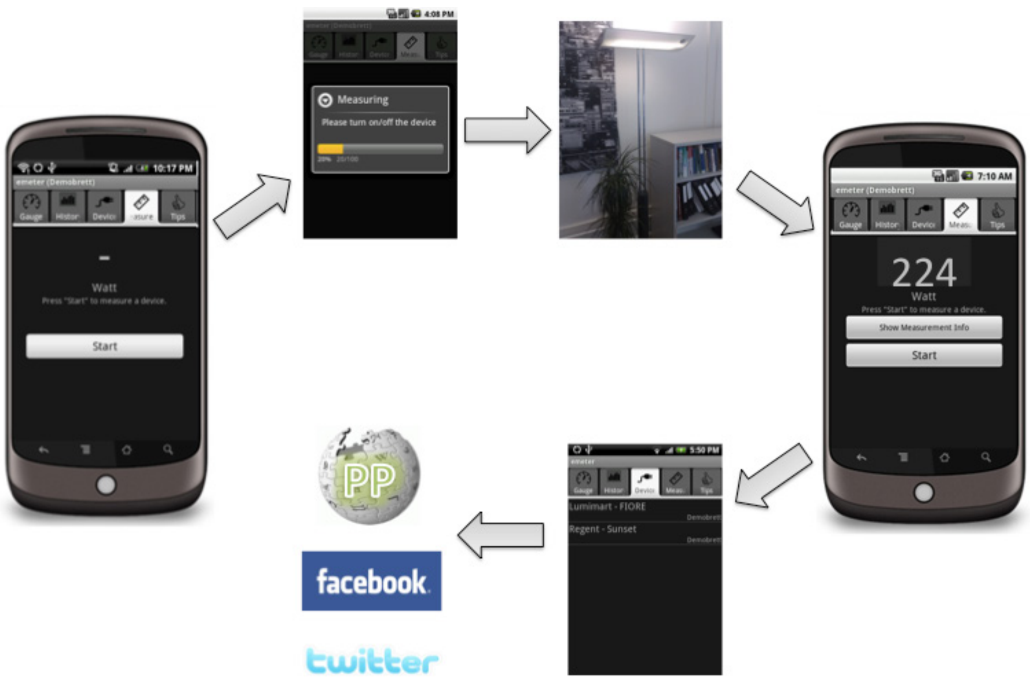

Figure 3.23 Measurement process for determining the power consumption of an appliance and its subsequent publication on PowerPedia using the eMeter Android user interface. The integration of the functionality offered by PowerPedia supports users with meaningful information beyond the mere consumption value after the measurement is completed.

servation tips. Moreover, users can publish their measurement including the device specifcs on Facebook or Twitter, given their credentials were provided when setting up the application.

After initiating the publishing process by pressing the corresponding green button (Figure 3.24 left), users have to specify the device category and manufacturer. In order to simplify the process for users, location, categories, and device names are downloaded from PowerPedia in the background and made available through an auto-completion feature (Figure 3.24 middle). For example, the feature leads users through the hierarchically structured device categories (e.g., consumer electronics, TV, LCD, 34-inch, Samsung, model type). By uploading the device to PowerPedia (Figure 3.24 right), users can compare their consumption against the consumption uploaded by other users as well as the consumption of the most energy-efficient appliances in the category harvested from consumer organization websites. The more precisely the information is provided when uploading the device, the more specific is the reflected information from PowerPedia. For example, if users detail the model of the measured appliance, PowerPedia is able to show an efficiency ranking of the selected device category together with category and appliance specific energy saving tips as well as all measurements of the device conducted by other users. The efficiency ranking based on all PowerPedia's community entries aims at placing the consumption within a more tangible context (Figure 3.25 left). It shows users the efficiency of their uploaded device together with information on the best and worst performing devices and the total number of uploaded devices in the selected category. In addition, users can share their information on popular social networks such as Facebook and Twitter.

The tips view (Figure 3.25 right) displays energy-saving measures downloaded from PowerPedia. It consists of tips that can be applied in general as well as tips relating to specific device categories. In addition, users can publish energy-saving tips on PowerPedia and so share their experiences with other users. Tips can be flagged to indicate that they have been applied. This allows PowerPedia to indicate the percentage of users who have already applied a particular measure. In a similar manner, users can also use the publish feature to share their experiences directly with their friends on their preferred social network.

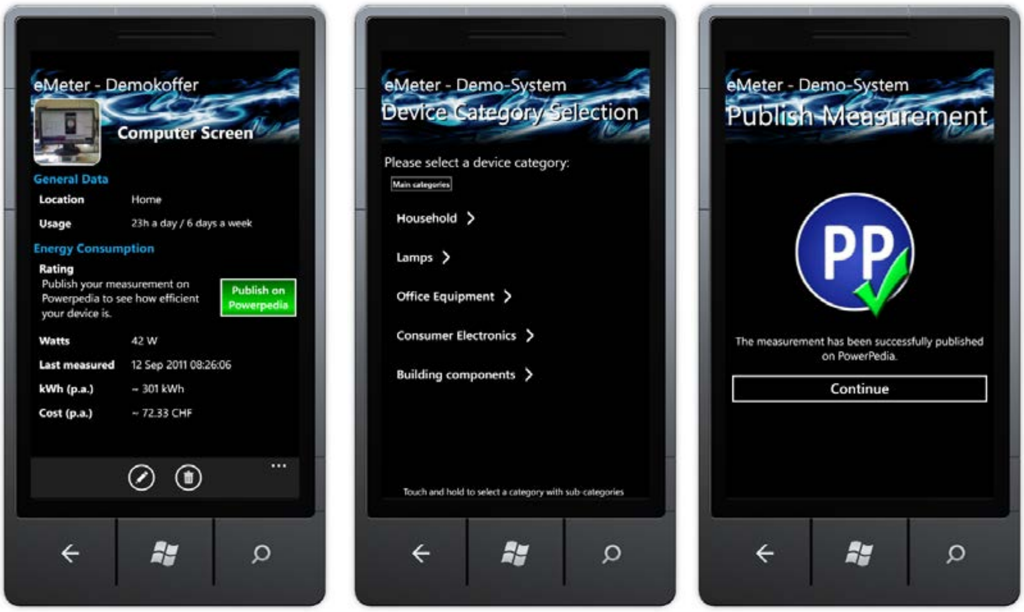

Figure 3.24 PowerPedia integration in the Windows Phone 7 eMeter user interface. After pressing the green publish button, users are provided with an auto-completion feature that supports them in finding the correct device category and model. Once selected, the conducted measurement is published on PowerPedia and others can see the measurement.

Figure 3.25 PowerPedia integration in the eMeter Android client. After publishing the measurement the device inventory reflects an efficiency ranking (left) and an additional view with further energy conservation tips (right). Both help users to put the consumption in a bigger picture.

3.7 Summary

In this chapter, we confirmed the suitability of mobile phones as energy feedback devices. We consider the applied methodology to develop a prototype application based on preliminary interviews and a survey with diametric application scenarios to be well suited for application development at an early stage. It helped us to critically assess the user requirements and to extend the application's functionalities. Exposing people to a functioning prototype was crucial for us to gather experience with the application, while at the same time participants better understood the usefulness of the application. In order to get users in the loop, we implemented the most promising feedback features and evaluated the different functionalities with our application.

We first presented the goal–driven design of a user interface as portable energy monitor on a mobile phone. In the conducted user survey, we found that the existence of a clear and simple to explain use case behind energy consumption feedback systems is a key success factor. However, when left to the imagination of the potential users, energy consumption feedback applications receive only medium ratings. Based on that finding, we developed a user interface that is supported by live data from the eMeter infrastructure. It implements a variety of other promising feedback features including a measurement functionality (eMeasure), which enables users determining the consumption of individual switchable appliances. The application thus offers the clear use case: "Learn how much a device consumes by just switching it on or off".

In a user study with 25 users, we evaluated the user interface with respect to its general usability (e.g., ease of use, understanding, etc.) and the perceived value of different feedback functionalities. In particular, we evaluated the performance of eMeasure in comparison with two other commercially available energy monitoring solutions. The results confirm that the interactive feature of the eMeter user interface allows users to quickly determine the consumption of electrical appliances with in an acceptably small error margin. We further tested the benefits and capabilities, confirmed the suitability of such a mobile phone application to serve as an energy feedback system, and identified the functionalities that are perceived most valuable by users, in general and with regard to individual differences concerning technological affinity. Tailoring energy feedback functionalities to different user groups is important to achieve effective and encouraging energy feedback applications [103, 107, 144]. It allows addressing a wide user base (beyond typical green users), which is doubtless key for the large-scale success of energy feedback systems.

The long-term deployment in four households showed that different users use the application for different use cases. It confirmed the application's ability to raise consumption awareness and energy literacy, but also clearly showed that the usage decreases once the user's initial curiosity is satisfied. Most energy feedback applications acutely suffer from the fact that their rate of use drops significantly soon after their initial installation and user behavior relapses [145]. Thus, it will be crucial to embed the eMeter system in a broader context as well as to develop concepts that aim at keeping users in the loop [72, 146, 147].

This confirms our expectation that not only differently motivated users expect different functionality, but also that research on engagement concepts (e.g., rewards, games, coupons, competitions, etc.) for energy saving applications is required [148, 149]. Such engagement strategies counter user fatigue and motivate long-term use of an energy feedback application. They thus foster behavior change that ultimately leads to energy conservation [150]. These concepts are without doubt important, however they mainly reside outside the computer science domain (e.g., behavioral science) and thus go beyond the scope of this work.

Another important aspect is the interactivity of energy consumption feedback that is introduced through the mobile phone application. We believe that this is key to get users involved into energy conservation. eMeasure is a good example how interactivity can be used in this context. It easily enables users to familiarize with their energy consumption. However, both experiments, the user study as well as the real-world deployment, revealed that mere power consumption values are not meaningful enough for most users.

To surpass this drawback, we implemented PowerPedia and integrated its functionality into the eMeter user interface. In a collaborative manner, PowerPedia aims at putting the electricity consumption in a bigger, more tangible picture that enables users to derive direct measures to reduce their consumption. Designing a game around eMeasure (e.g., "find an appliance in your house that consumes 100W") is another potential application where interactivity could help increase residential energy awareness and if devised properly (e.g., in a fun and competitive way), motivate users to conserve electricity [151].

The presented results show that Ubicomp can help to foster energy efficiency in residential environments. The developed mobile phone application is based on loosely coupled components that are already integrated in the daily life or will become ubiquitous in the course of smart metering. We confirm the suitability of mobile phones as energy feedback devices and as a technology to measure the electricity consumption of residential appliances. This should help to drive adoption of such energy monitoring systems. However, the results also reveal that different functionalities are required to target different user groups and long-term application of energy monitoring systems requires concepts that motivate people after their initial curiosity has been satisfied. These aspects should be confirmed and further investigated from a behavioral science perspective in a larger experiment, which allows for controlled study conditions (e.g., has a control group).

4 Leveraging Smart Meter Data to Recognize Home Appliances

The requirement to conserve energy, the modernization of the electrical grid infrastructure, and the growing share of electricity from intermittent sources (e.g., wind and photovoltaics) gave rise to paradigm shifts in the energy domain [38]. One building block of the move towards smart grids is the worldwide adoption of smart meters that measure and communicate residential electricity consumption. Originally intended to simplify the meter reading processes for energy utilities, smart meters are nowadays seen as an enabler for new energy efficiency services, flexible tariffs, and demand response programs.

In this chapter, we present a set of algorithms that make use of smart meters and together with the interaction capabilities of smartphones to leverage residential electricity data to recognize home appliances. Based on the eMeter infrastructure, we show the potential of applying Ubicomp technologies for residential load disaggregation. The disaggregation of individual appliances within a particular household in terms of their energy demand enables several particularly promising applications in the residential domain, be it for delivering itemized electricity bills or for providing targeted energy saving advice.

In Section 4.1 we first review related work with respect to residential load disaggregation that help set the context for our proposed disaggregation scheme AppliSense. Next, Section 4.2 revisits the architecture to recall the physical quantities that are important for the disaggregation process. Thereafter, Section 4.3 presents the fundamentals that can be used to classify electrical home appliances, before Section 4.4 explains the AppliSense disaggregation scheme in detail. Section 4.5 reports on the evaluation of AppliSense that was tested in a laboratory study with eight simultaneous running devices, achieving recognition rates almost 90%. In Section 4.6 we discuss the results and the limitations of the proposed disaggregation scheme, before we last conclude this chapter with a summary and a brief outlook in Section 4.7. Parts of the work of this chapter have been recently published in [152].

4.1 On Residential Load Disaggregation

Device-level electricity consumption information is essential for users to establish the link between consumption and device utilization, to enable sophisticated energy efficiency services (e.g., targeted, automated recommendations), and to reduce residential electricity consumption by providing users with direct conservation measures. Depending on the applied modeling approach, different data input is required to provide users with disaggregated electricity use in the residential domain [153]. We can broadly differentiate between offline and online techniques to disaggregate overall electricity use to device level (see Figure 4.1).

A traditional way derives household energy models based on offline data from field surveys using Conditional Demand Analysis (CDA). The method was invented by [154] on a detailed data set of 5000 households and later similarly applied by [155, 156]. By comparing the load profiles of households with known appliances gathered from a survey to those without a statistical analysis can be computed. The strength of this method is the ease to obtain the data input required for the analysis [157]. However, the accuracy of CDA is limited because of the diversity of devices and appliances in homes and because many end-uses share temporal load profiles, making aggregate load profiling relatively inaccurate. Furthermore, CDA has traditionally relied on self-report surveying for load disaggregation, which provides a relatively sparse dataset containing various self-reporting biases [119]. Different modifications of the approach have been proposed to surpass the above-mentioned limitations (e.g., subjective appliance usage estimation) [158-160] and shaped the technique in a way that it is used more for the extrapolation of household energy use to regional or country-wide energy use [161, 162].

In contrast to offline-based methods, online approaches disaggregate the electricity consumption based on sensor data. They can be classified into two different domains: Distributed direct sensing and single-point sensing. While a detailed overview of the distributed sensing systems is provided in Section 2.1.2, we only mention these approaches here briefly for the sake of completeness, before we focus on single sensor approaches that are important in the context of our developed disaggregation scheme. A large variety of distributed sensing approaches provide device-level consumption information by deploying sensors at each appliance or power outlet. Having

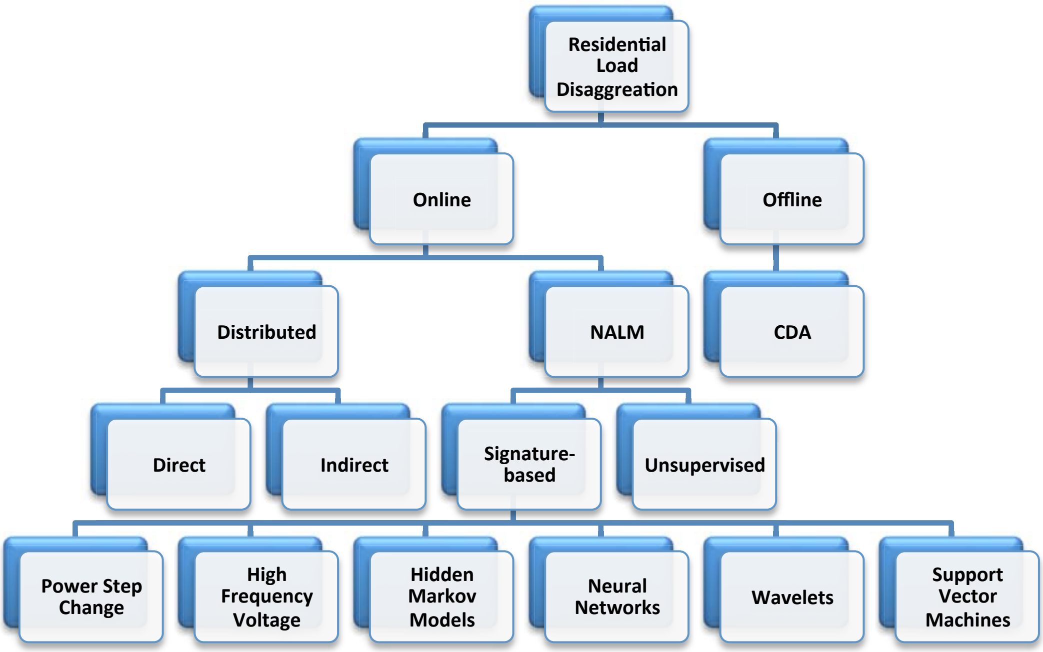

Figure 4.1 Overview and classification of residential load disaggregation approaches.

a sensor installed directly at the appliance enables the accurate measurement and feedback of the electricity consumption and offers the benefit of being able to directly control the appliance (e.g., a relay that allows switching the appliance on or off). While the concept seems straight forward, it is costly and the installation of a large number of sensors imposes a high usage barrier. The sensor deployment at each appliance is not only difficult and discouraging for users, but also often not feasible for large household appliances in the residential environment that consume a significant amount of residential energy because they are hard-wired (e.g., lighting or hot water heaters) or difficult to reach (e.g., refrigerator, dryer, or washer) [119]. Another burdensome drawback of distributed sensing methods is the communication network that has to be established between the individual sensors located throughout the house.

Single sensor approaches are typically subsumed under the concept of Nonintrusive Appliance Load Monitoring (NALM). These techniques disaggregate overall electricity use of a household by extracting characteristic, observable features from the recorded signal of a single sensor. The input data signal and the applied methodology differ by approach. Typically overall power consumption, current measurements, or voltage fluctuations are used to facilitate the disaggregation process with the help of hidden Markov models, wavelets, neural networks, and support vector machines [104]. Only recently a number of papers have been published that provide a great overview on the research that has been conducted in the field of load disaggregation [90, 104, 119, 163].

The initial work dates back to the 1980ies, where Hart [164] tried first to match a-priori known appliance signatures to the step change in the overall power signal by using real and reactive power measurements at a rate of 1Hz. The concept proved to be effective in various field tests – at that time especially for the disaggregation and fault detection of larger loads [165-169] – and paved ground for various other works in that field. Hart's method can detect on and off events of steady-state appliances (e.g., lights) that consume a relatively constant load, but faces difficulties detecting multi-state appliances (e.g., laptops) that have a varying power consumption. Note that appliances also change their resistance after startup, which can affect the power consumption by as high as 10% [170]. Another downside is that two appliances that are operated at the same time might be classified as one appliance, which is especially crucial when detecting smaller loads. Moreover, when two identical devices are operated in the home (e.g., lights) in different locations, this approach is not able to identify which of the two is appliances is currently operated.

A lot of work has been done to overcome some of the outlined limitations. Norford and Leeb introduced transient event detection at high sampling rates to disaggregate devices with similar power consumption [171, 172], and follow-up work by Laughman et al. [173] explained how to use current harmonics for detecting continuous variable loads. However, to disaggregate the overall load the power consumption of the detected device has to be estimated, which is a non-trivial, recently addressed problem [174-176]. A variant of Hart's scheme focused on the separation between simultaneous on/off events of appliances. The same authors also proposed an extension to the original algorithm that focuses on disaggregating large power drawing appliances such as heat pumps. Instead of only regarding the step change to identify an appliance, they suggest using an appliance signature that consists of an edge as the appliance is turned on and a slope as the appliance operates [177-179]. Marchiori et al. propose decreasing the complexity by using one sensor per circuit. Their developed scheme combines on a heuristic and a Bayesian approach that disaggregate electricity consumption using historical steady-state energy use pattern for each device [180, 181]. While all proposed concepts so far seem promising to solve some of the problems, they all suffer from several drawbacks. All approaches require excessive training before the disaggregation start, their performance has not been tested in practical scenarios and hardly anything is known about their accuracy. Last, the robustness of these methods is unknown (e.g., how is the presence of new appliances affecting the previously recorded signatures, etc.) [90].

Other work utilizes rule-based algorithms and neural networks to disaggregate overall residential energy consumption data. Early approaches were typically bound to low-resolution data. Powers [182] iterative algorithm is based on real power only. It tries to analyze the energy consumption top-down at a low sampling interval of 15 minutes. However, his approach is based on a large a-priori known reference database that requires monitoring of each appliance in the home for several days. Prudenzi disaggregated consumption data for large loads at the same sampling rate by using a neural network approach [86]. Ruzzelli et al. used a special purpose sensor that has to be installed at the circuit breaker. The consumption information is post-processed in an artificial neural network that requires a lengthy training process to achieve the appliance break down [183]. Other rule-based work focuses on the possibility to differentiate between appliances with similar power consumption by taking into account their frequency of use [184], on disaggregation in the industrial domain [185], and on using pattern recognition methods to disaggregate the overall electricity consumption into major energy end-uses [85]. The latter were one of the few to report explicit results. Namely, large loads (washer, dryer, etc.) could be detected with an accuracy of up to 90%. However, the necessity to develop appliance-specific decision rules and the extended training period (i.e., all appliances had to be continuously monitored for one week), which is characteristic for most approaches, makes the technique hardly applicable in practice.

Later, more sophisticated approaches dealt with the analysis of data sampled at higher frequency using wavelets. Wavelets allow simultaneous time and frequency location, which is a significant advantage over approaches that use Fast Fourier Transformation [171, 186] where time localization is not possible. The authors of [187] started experimenting with fuzzy numbers that are used for harmonics signature recognition. They found that each type of current waveform polluted with power harmonics could be represented by a five level wavelet decomposition. Based on that finding, the authors developed an algorithm that uses discrete wavelet transforms to identify the harmonics signature of non-linear loads. However, harmonics pollution in the electric system through the diversity of home appliances makes this approach difficult. Instead of using discrete wavelets, Norman's disaggregation approach [188] relies on continuous wavelet transforms were claimed to perform better compared to the previous approach. Other authors explored statistical signature analysis to infer the devices operating from the current and voltage waveforms [87, 189]. This is beneficial because the resulting waveform is not a function of time. This may allow identifying two devices that have similar current/voltage waveforms because their current-voltage curve might be different. Although appliance signatures based on time independent current-voltage curves of appliances seem promising, this approach like the ones for wavelets have never resulted a working disaggregation algorithm. Srinivasan combined harmonic signature analysis with neuronal networks. In his work, several different classification models for signature extraction and device identification were developed and tested. For a set of appliances, a detection accuracy of around 80% was reported [190]. The authors of [191] propose a combination of three optimization algorithms and a neural network. In a second work, the same authors propose a household simulator that generates data that can be used to test disaggregation algorithms. The simulator allowed the authors to assess many different situations and confirmed that the disaggregation performance can benefit from the fusion of different techniques [192]. However, to detect appliances based on any kind of signature the strongest drawback remains: significant exploratory work is required to gather the signatures in the first place.

In contrast to these high frequency sampling approaches that usually rely on special purpose sensors, Kolter et al. [193] recently investigated the possibility of disaggregation using discriminative sparse coding based on hourly consumption data of over 10000 appliances from different homes. While the approach is performing better than other sparse coding techniques, its accuracy suffers from the low frequency despite the extremely rich training set. The authors try to predict the share of ten appliance categories over one week given only the signal that contains the aggregated consumption of all devices of one home and achieve an accuracy of around 50%. In addition, the authors released a rich data set containing detailed power usage information from several homes. It contains the overall home electricity consumption recorded at a high frequency (15kHz) and sub metered data from individual circuits in the home, each labeled with its category of appliance or appliances, recorded at 0.5Hz as well as up to 20 plug-level monitors that directly monitor the connected appliances at sampling frequency of 1Hz, with a focus on logging electronics devices where multiple devices are grouped to a single circuit. The data set is freely available and should boost further research in the domain of energy disaggregation [194]. First applications of the data set have recently been published. Kolter used Hidden Markov Models to break the electricity consumption down into different use categories. At this aggregated level, the approach achieved an accuracy of 82%. Parson [195] uses the same technique for disaggregation. General models of appliance types are tuned to specific appliance instances using only signatures extracted from the aggregate load. Their evaluation is limited on to the three highest energy-consuming appliances of the above-mentioned data set showing that it can disaggregate 35% of the total energy consumption to an accuracy of 83%.

Alternatively, some recent work has advocated unsupervised approaches that do not rely on a long and complex training period. Clustering methods are applied to automatically classify appliances and their corresponding power consumption. Goncalves et al. published first results that indicate the potential of their linear blind source separation strategy. The small number of existing large household appliances with a power consumption of above 400W could easily be classified, but small appliances were typically clustered together and were difficult to separate [196]. Kim et al. [197] propose three new unsupervised Markov models that are tested on data from seven small households with four to ten appliances. Their findings indicate that the proposed methods can potentially outperform standard Factorial Hidden Markov Model approaches for appliances with easy load profiles, but disaggregation becomes challenging as the number of devices increases and the signatures of appliances become more complex. More research is required to fully evaluate the potential of the proposed techniques as well as to develop a method that estimates the number of operated appliances at home making unsupervised clustering algorithms more accurate.

Instead of unsupervised learning, Berges et al. propose a user-centered high-frequency disaggregation approach. Their system builds upon a commercially available oscilloscope that samples current and voltage at frequency of 15kHz. Their signature-based decomposition method is based on real and reactive power as well as on transients that can be gathered at these high sampling frequencies. The recognition strategy for predetermined signatures follows the classical decomposition chain: Event detection, feature extraction, and matching based on nearest neighbor and Bayes. The system was tested in a lab setting with nine devices and achieved an accuracy of 85%. The authors further envision that unmatched, but correctly detected events are forwarded to the user, who can label the event as a new signature [198-200].

Another high frequency idea (i.e., up to 500kHz sampling rate), but based on voltage sensing instead of current, has been explored by the authors of [88] and [89]. They developed and combined two complementary approaches in a system that relies on a single sensor that can be plugged-in anywhere to the electric circuit of a household. It listens on the residential power line to detect unique noise changes [201] and electromagnetic interference that occur through the abrupt switching of devices and switch mode power supplies respectively. Mechanical switches produce electrical noise [202], which varies significantly by appliance [91], and the electromagnetic interference is caused by switch mode power supplies incorporated in many modern consumer electronics. The system can be used to infer about appliance operation which in combination with the measurements of an electricity meter can reveal the consumption of particular devices. An evaluation of the concept in sample households shows an accuracy of to 94%. Although the method seems superior compared to previous approaches, the same drawbacks such as complex and time-intense training remain and the overlapping of noise signals. In addition, new open questions are introduced: It is unclear how the correlation to the power signal exactly achieved; not all devices are causing a noise or interference; and the recorded noise signatures depend on the location of the appliance relatively to the location of sensor. A relocation of either as well as the presence of new devices inevitably requires a complete recalibration of the system.

Baranski and Voss use an inexpensive optical sensor that is attached to a traditional utility meter to disaggregate the consumption [203]. Their scheme combines different concepts such as clustering, finite state machines, and heuristics [83, 84]. The method of detecting on/off events is to some extent similar to the approach by Hart. But instead of trying to match every individual event, the authors incorporate an optimization algorithm that tries matching a largest possible set of events. Unlike most other approaches so far, their method does not rely on complex training. However, traditional meters nowadays become more and more obsolete and are currently replaced by smart meters.

Summarizing the related work, existing approaches can be broadly divided in online and offline-based techniques. Offline approaches use surveys to discover what appliances are operated at home. Originally intended to statistically compare load profiles of households with known appliances to those without, offline methods have become more and more obsolete due to the rise metering and sensing technology that made real consumption data available. Such online methods for consumption disaggregation do not rely on self-reporting and can potentially provide extremely rich end-use datasets [119]. They consist of distributed and single-sensing techniques: Distributed approaches can rather easily achieve a consumption breakdown, but these systems typically suffer from complex calibration procedures and deploying a large number of sensors in the residential environment quickly leads not only to high cost but also to a discouraging high usage barrier [183]. In contrast, single sensor systems represent a more cost-effective solution, but often rely on custom hardware (e.g., for high sampling rates) and presuppose either a-priori knowledge about the household devices and their electrical characteristics, or require a complex training phase involving the user where the system learns about the specific device characteristics.

However, a global signature database is difficult to obtain in a world of fast changing small appliances, and training procedures at the initial deployment are discouraging users and hinder adoption [204]. To address this problem, recent research has focused on intelligent training methods and building large signature databases that can ideally be shared across homes [205-207]. But appliance signatures are often influenced by the local environment and are thus to a certain extent bound to the system they were recorded (e.g., sampling frequency, physical quantities, etc.). These limitations increase the complexity and might make a central appliance signature database impossible [82]. In addition, current disaggregation approaches cannot take devices into account that are introduced into the residential environment and were not present at the initial calibration of the system. Overall, we find that no disaggregation method currently exists that is easily deployable, highly accurate for all household appliances, and cost-effective, and most existing approaches fail to meet usability requirements that are essential for fast adoption. Lastly, but most importantly, there is a great need for improved user interfaces, which effectively improve the algorithms by utilizing user-feedback. This feedback mechanism should be non-intrusive and must be designed to easily blend with the resident's normal daily activities [82, 206].

The approach presented in this thesis addresses this need for improved user interfaces and demonstrates the potential of Ubicomp to facilitate load disaggregation in a simpler and user-friendlier way. The proposed disaggregation scheme AppliSense is an integrated solution to disaggregate the electricity consumption of households to device level. It uses a single sensor and builds on the early principles described by Hart [164] to not only address some of the above-mentioned shortcomings in terms of usability, but also remaining technical challenges (e.g., the recognition of smaller loads and overlapping on/off events of multiple appliances). For this, we use the eMeter system that does not rely on custom hardware and on top designed AppliSense, a Ubicomp-enabled disaggregation scheme that utilizes user feedback to avoid complex training. More concretely, we make use of smart electricity meters, which are going to be installed in large numbers in the U.S. and Europe over the next years, together with a user interface on a mobile phone, which much simplifies the appliance signature acquisition process because this is done as a side effect, invisible to users. Our approach thus shows how Ubicomp can enable load disaggregation in a more applicable, simpler, and user-friendlier way.

4.2 Revisiting the Architecture

In this section, we briefly revisit the data acquisition component of the eMeter infrastructure because the physical quantities directly available from the meter as well as those that can be derived thereof are important for the later on presented disaggregation approach.

We use a smart meter as a single sensor to log the total electricity consumption of the household. The utilized meter measures the following physical quantities at phase level, i.e., each quantity has one representation per connected phase L_x:

- Effective voltage U_{eff} in Volt and effective current I_{eff} in Ampere;

- Phase shift φ between current and voltage in degrees;
- Real power P in Watt;
- Neutral conductor current I_{neu} in Ampere for grounding the meter.

From the measured quantities, further information such as apparent, reactive, and distortion power can be derived. These characterize the electrical load behavior of appliances and can be used to differentiate between different device categories detailed in the next section. The following explains how these quantities are derived.

- PowerAllPhases (P_{all}) represents the aggregated sum of real power of all three electrical phases $L_1 - L_3$, and is expressed in Watt ($[P_{all}]_{SI} = W$).
- Apparent power on line S_x, ($[S_x]_{SI}$ = VA), is the product of the effective values of current (I_{effx}) and voltage (U_{effx}):
  $S_x = U_{effx} \times I_{effx}$.
  For the absolute value of apparent power S, the real power P, and the total of reactive power Q_{tot} the following holds:
  $|S| = \sqrt{P^2 + Q_{tot}^2}$.

  Q_{tot} consists of two components, the reactive power Q_{trans} and the distortion power D:
  $|Q_{tot}| = \sqrt{Q_{trans}^2 + D^2}$.
- Total reactive power. The absolute value of the total reactive power Q_{totx} on L_x, $[Q_{totx}]_{SI}$ = Var, can be computed with the measured value P_x as follows:

  $|Q_{totx}| = \sqrt{S_x^2 - P_x^2}$.
- Translative reactive power. Through nonlinear consumers, such as inverters, power supplies, or inductivities, non-sinusoidal currents I_{eff} can occur at sinusoidal voltages U_{eff}. These non-sinusoidal currents are composed of sinusoidal parts of different frequencies (I). If I_{1eff} is the sinusoidal current part of the fundamental frequency and φ the phase shift between current and voltage, then for real power P and reactive power Q the following equations hold:

$$\left.\begin{aligned} S' &= U_{eff} \times I_{1eff} \\ P &= S' \times \cos(\varphi) \\ Q_{trans} &= S' \times \sin(\varphi) \end{aligned}\right\} \qquad \frac{Q_{trans}}{P} = \frac{S' \times \sin(\varphi)}{S' \times \cos(\varphi)} = \tan(\varphi).$$

This allows computing the translative reactive power Q_{transx} on L_x ($[Q_{transx}]_{SI}$ = Var) from P_x and φ_x even if the current curve is non-sinusoidal:

$Q_{transx} = P_x \times \tan(\varphi_x)$.

- Distortive reactive power. The value of distortion power on L_x (D_x, $[D_x]_{SI}$ = Var) caused through nonlinear consumers (such as inverters, power supplies, or inductivities), can be computed as:

  $|D_x| = \sqrt{Q_{totx}^2 - Q_{transx}^2}$.

4.3 Classification of Residential Appliances

In the following, we explain how domestic appliances can be classified according to their characteristic load signatures based on the above-mentioned physical quantities measured by the smart meter. Depending on its characteristic electric and electronic, an appliance can be of resistive, inductive, or capacitive nature. For example, a standard light bulb is purely resistive whereas a vacuum cleaner is predominantly inductive. In general, incandescent appliances (e.g., a kettle or a light bulb) are mostly resistive (ohmic), motors (e.g., a fan or a heater) predominantly inductive, and devices containing a power supply or electronic frequency converters (e.g., laptops) mainly capacitive. Figure 4.2 illustrates the resulting signature space for residential electric appliances. Resistive appliances typically reside on the x-axis whereas predominantly inductive appliances have a positive, predominantly capacitive a negative ordinate. Since the power consumption of appliances crucially depends on the utilized internal components, it can vary quite strongly between different models of the same appliance category (e.g., TVs). For that reason, the depicted boundaries for the individual appliances are fuzzy and only provide an indicative overview on the situation.

Figure 4.3 illustrates exemplary appliance power signatures at the standard sampling frequency for different appliance categories over different operation lengths. The signatures were recorded with our system in a controlled environment at a rate of f_s = 1 sample/sec, i.e., f_s = 1Hz. If the load is purely resistive, then the voltage and current are in phase (e.g., the iron (Figure 4.3 (lower left)). The reactive component Q of the apparent power is zero, meaning all power is transferred to the load. While appliances work through the real (active) power, the reactive (passive) power caused by inductors and capacitors does not drive the load, but heats

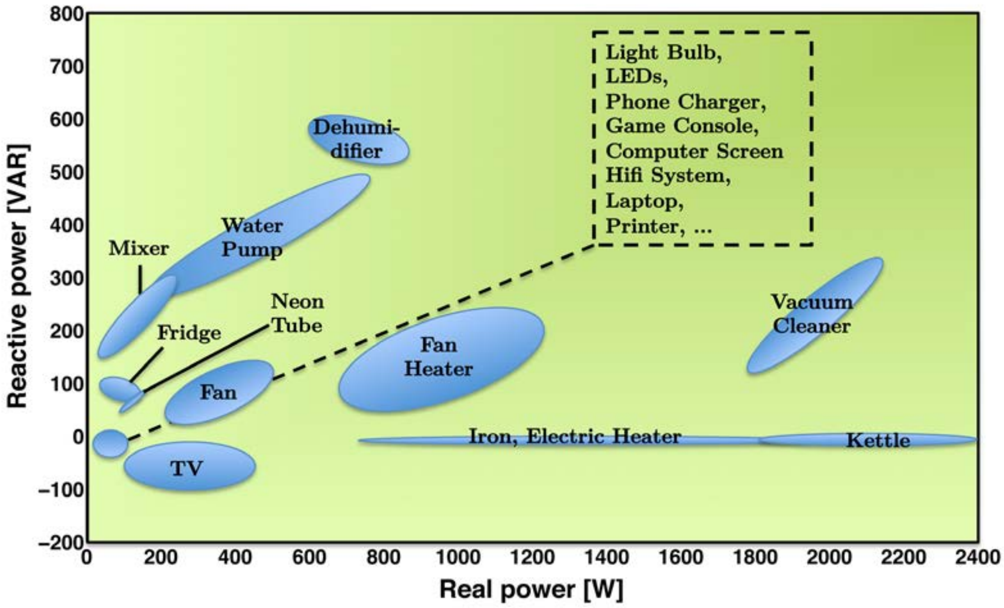

Figure 4.2 Signature space of residential appliances: Resistive appliances typically reside on the real power axis because their reactive power is typically zero. Appliances with a dominant inductive component are characterized by a positive reactive power whereas predominantly capacitive devices show a negative reactive power.

wires and thus wastes energy. However, capacitors and inductors affect the power consumption of a device by shifting the current with respect to the voltage. In case the electric load is completely reactive, voltage and current are 90 degrees out of phase. This leads to the fact that per cycle the product of voltage and current is positive for one half, but negative for the other half. This results in a net energy flow of zero as on average exactly the same amount of energy flows towards the appliance as flows back. A consumer with reactive components is either of type ohmic-inductive (e.g., Figure 4.3 (middle)) with a typical phase shift of $0 < \varphi < \pi$ between current and voltage or ohmic-capacitive (e.g., Figure 4.3 (lower right)) characterized by a negative phase shift ($0 > \varphi > -\pi$). In addition, in electrical networks there may exist non-sinusoidal currents and voltages (e.g., caused by inverters in switching events) that result in harmonics. These harmonics cause an additional reactive component, the so-called distortion power (Figure 4.4).

In mathematical terms this can be expressed as:

$$|S| = \sqrt{P^2 + Q_{trans}^2 + D^2},$$

where S is the apparent power, P is the real power, Q the translative component, and D the distortive component of the total reactive power.

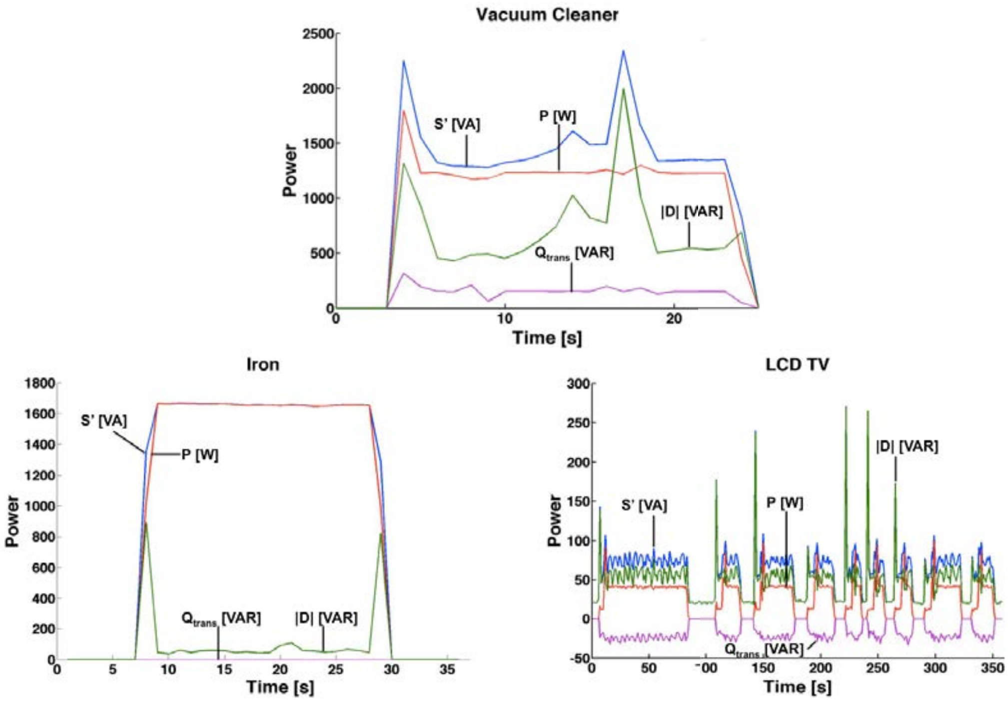

Figure 4.3 Power signatures of three different residential devices from different appliance categories and for different operation periods. The iron (left) is a ohmic device, which is characterized by a translative reactive power Q_{trans} equal to zero. The vacuum cleaner (middle) features a positive translative reactive power Q_{trans}, which is typical for predominantly inductive appliances. The power consumption of a TV (right) shows a negative translative power component characteristic for predominantly capacitive appliances.

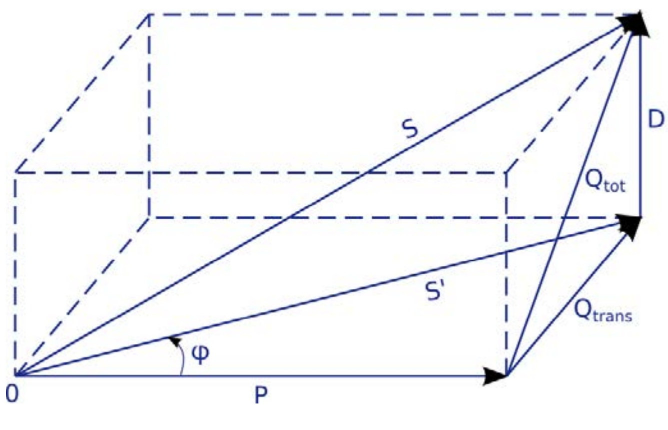

Figure 4.4 Relation between different power quantities that can be derived using the eMeter system.

Based on its internal composition and its possible modes of operation (e.g. static, multi-level, or variable), an appliance imposes a characteristic load profile on the electric circuit. This signature depends on the relation of the different power components and can be used to discriminate between appliances when disaggregating the total consumption. Our prototype system measures these parameters either directly or indirectly. In addition to these physical quantities, the signature length, peak voltage, and current are also important in terms of the appliance signature.

4.4 The AppliSense Disaggregation Algorithm

The AppliSense algorithm uses electricity consumption data recorded by the smart electricity meter to detect switching events of appliances to automatically break down the total electricity consumption to device level. In the following, we first outline the basic idea of our system that is based on the early principles of Hart and pays particular respect to usability. We then explain how we use Ubicomp to much simplify the process of recording appliance signatures on which the algorithm crucially depends is acquired. Last, we discuss details of the algorithm design, such as the applied filtering and the technique used for matching detected switching events to existing appliance signatures.

4.4.1 Key Concept

The electricity consumption of a household fluctuates over time based on the operation of individual devices used by the residents (see Figure 4.5 (left)). For example, switching on a light induces the depicted change in the load curve. Having a more detailed look on the consumption data, the figure shows that there exist intervals where the load remains more or less constant on a stable level. A black bar marks two of these levels. The difference in real power (dP) between these levels indicates the change in electricity consumption due to the operation of the light. Our system not only measures the total load of the household, but the load characteristics (i.e., apparent power, real power, etc.) of each of the three phases separately. This phase-level data allows us to split up the overall electricity to get an even more detailed view.

These considerations lead to the following key concept of AppliSense, which is a variant of the Hart scheme to recognize device switching events in the load curve based on an appliance signature database: First, we identify points in time where significant changes between two levels of power consumption in the load curve occur. Second, once such an edge is detected, we compute the differences of the different physical quantities between these two consecutive levels and classify the change as a potential appliance switching event. And third, we compare each of these differences with a known set of differences from an appliance signature database and map the edge to an individual device according to its load characteristics.

The right of Figure 4.5 illustrates these steps. It shows the electricity consumption (red) at a certain time interval in which five load levels (black bars) were identified. For simplicity, only the real power is visualized in this example. From this we can compute four deltas: dP1, dP2, dP3, and dP4. Each of these deltas corresponds to a potential on/off event of a device. The algorithm tries to match these with a known device signature from the signature database. For that, each entry dP_i in a column of the matrix on the left symbolizes a delta which was extracted from the load curve at time i. The operator represents a detector logic that compares the rows of the matrix to the signature vector with the known deltas. The resulting vector holds the best matching entry, in case a matching appliance could be identified. In the example, two matching signatures of a known device (a turning on and a turning off event) are detected for dP2 and dP3. However, no signature is matching the switching events at time instants one and four.

After having explained the basic concept of the AppliSense algortihm, we now focus on how Ubicomp technologies can be used to simplify the appliance signature acquistion process in the next subsection.

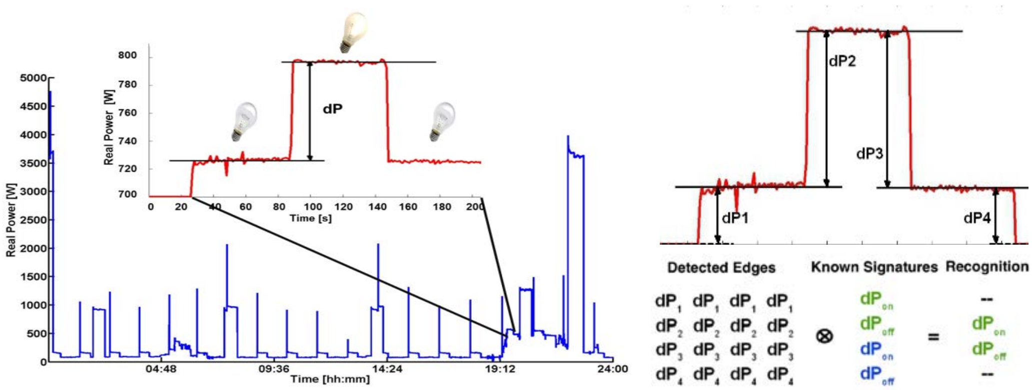

Figure 4.5 Key idea of the AppliSense algorithm. A switching event of an appliances cause a change in power consumption (left). The signal is divided into an alternating sequence of levels and edges. Then features are extracted to detect switching events and match them to a known appliance signature (right).

4.4.2 Signature Database

In contrast to other load disaggregation systems, which often discourage users by requiring a long training period or complex calibration at the time of system installation, we wanted to develop a system that is easy to use. This is particularly important for the generation of the signature database that is used to identify a detected appliance switching event. Hence, our system makes use of Ubicomp technologies to offer a user-friendly way for signature recording. For that, we use the measurement functionality embedded in the eMeter user interface. It allows users to identify the consumption of an individual appliance in a simple, explorative way while at the same time logging the signature in the background, invisible to the user. This also facilitates the easy integration of new appliances that are introduced at home (i.e., appliances that were not present at the time of the initial system setup). Whereas other systems need to completely recalibrate, our approach is able to incrementally acquire signatures and thus integrate new devices. In particular this means that it is not necessary to take signatures of every appliance in advance, but the signature database is established with simple means over time, which is crucial in a fast changing home environment. Moreover the signatures can be recorded with the user interface of an existing energy monitoring system and load disaggregation becomes an add-on feature.

The measurement process from a user's perspective is illustrated in Figure 4.6. To measure the consumption of a device, users initialize the measurement by pressing the start button on the user interface. After operating the selected appliance (either on or off), the system then computes the power consumption of the appliance based on the measurement algorithm within a few seconds (see Subsection 2.5.4). During the measurement, the signature acquisition process (see Figure 4.7) runs in the background (only real power depicted for clarity reasons). Every appliance switching event causes a change in power consumption that is measured by the smart meter. It logs the whole appliance signature (i.e., change in apparent, reactive, and distortion power, power factor, voltage, etc.) that is stored in the appliance signature database. In addition, we can classify whether an on (dP>0) or off (dP<0) appliance switching event has occurred. AppliSense uses this information later as input knowledge to match detected edges with operated appliances.

The idea of our approach is to systematically improve the detection algorithm. Through user feedback the number of signatures in the database is increased while the system is being used. This leads to higher precision in recognizable operation events over time, and at the same time avoids a time intense training period at the beginning.

Figure 4.6 User-friendly signature acquisition process: From a user point of view, appliance signatures are recorded invisible in the background as part of the measurement functionality embedded in the eMeter user interface.

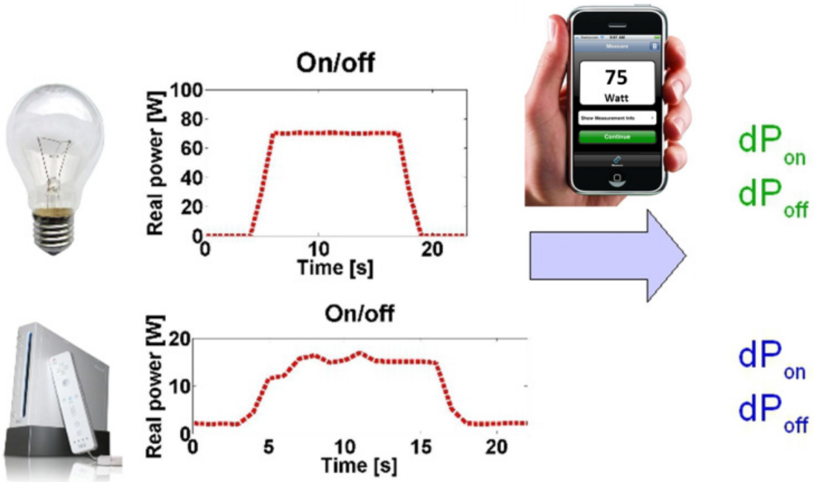

Figure 4.7 While users are utilizing the eMeter smartphone application to learn more about the power consumption of their individual home appliances, the eMeter backend records the appliance signature in the background and stores the in the appliance signature database [3].

4.4.3 Algorithm Design

In this subsection, we explain the algorithm design of the AppliSense load disaggregation algorithm. The algorithm is implemented in Matlab and takes csv-flies as input that contain electricity consumption data extracted from the MySQL database with a Python script. Next, the AppliSense core algorithm is running, before the results are visualized by ePlot. Figure 4.8 depicts the just-described AppliSense tool chain, which individual components we describe next.

4.4.3.1 Data Extraction

AppliSense takes a time slot (i.e., begin and end time) and a smart meter id as input. Originally, the AppliSense core algorithm was accessing the MySQL database directly via Microsoft's Open Database Connectivity (ODBC[28]), but over time the archive table that holds all the measurements of all connected meters in data sharing mode (see Section 2.5.1) grew to over 13 million entries. This makes data retrieval at run time through the corresponding SQL queries slow and locks the database while the query is processed. Hence, we separated the data retrieval process. A Python script now is responsible for extracting the measurement data from the database offline. The output is stored in CSV formatted files that serve as input for the AppliSense core algorithm. The data is split into different files according to a user-defined time period (e.g., calendar days) and smart meter ID. The files hold all corresponding measurements in ascending order. The script can be called from the command line with the following parameters:

Figure 4.8 AppliSense tool chain. A Python script is responsible for data extraction from the MySQL database. The AppliSense algorithm takes csv-files as input and produces an output stream that consists of electricity consumption data and event labels that correspond to identified device switching events. The ePlot visualization framework is responsible for depicting the results.

[28] Microsoft Open Database Connectivity: http://msdn.microsoft.com/en-us/library/windows/desktop/ms710252%28v=vs.85%29.aspx.

- --user USER: MySQL server user
- --host HOST: MySQL server hostname
- --port PORT: MySQL server port (optional)
- --passwd PASSWD: MySQL server password
- --db DB: MySQL database
- --sm SM: Smart Meter number (ID)
- --fromdate FROMDATE: Starting date (optional)
- --todate TODATE: Finishing date (optional)
- --period <day|week|month|year> (default: day)

4.4.3.2 AppliSense Algorithm

The AppliSense load disaggregation algorithm itself consists of six steps that are subsequently discussed in the remainder of this subsection. The algorithm follows the early principles discovered by Hart, but much simplifies the signature acquisition process for users. Figure 4.9 provides an overview on the individual steps. After the retrieval and normalization of the electricity consumption data, the algorithm requires three steps for feature extraction (depicted in orange) and then tries to match detected events to known signatures from the appliance signature database.

1) Normalization: In power circuits, load-dependent voltage drops can occur (e.g., in reaction to a switching event of an appliance). From

$$I = \frac{U}{R} \text{ and } S = U \times I$$

for apparent power S and effective values of voltage U and current I, a quadratic relation arises:

$$S = \frac{U^2}{R}.$$

Hence, voltage drops can lead to large differences in power consumption, which we have to account for by normalizing the all measured and calculated power values to a constant voltage (of 230 V):

$$S'_n = \left(\frac{230}{U}\right)^2 \times S.$$

The eMeter system measures the electricity consumption data for each electrical phase separately. In the following, each of the remaining steps of

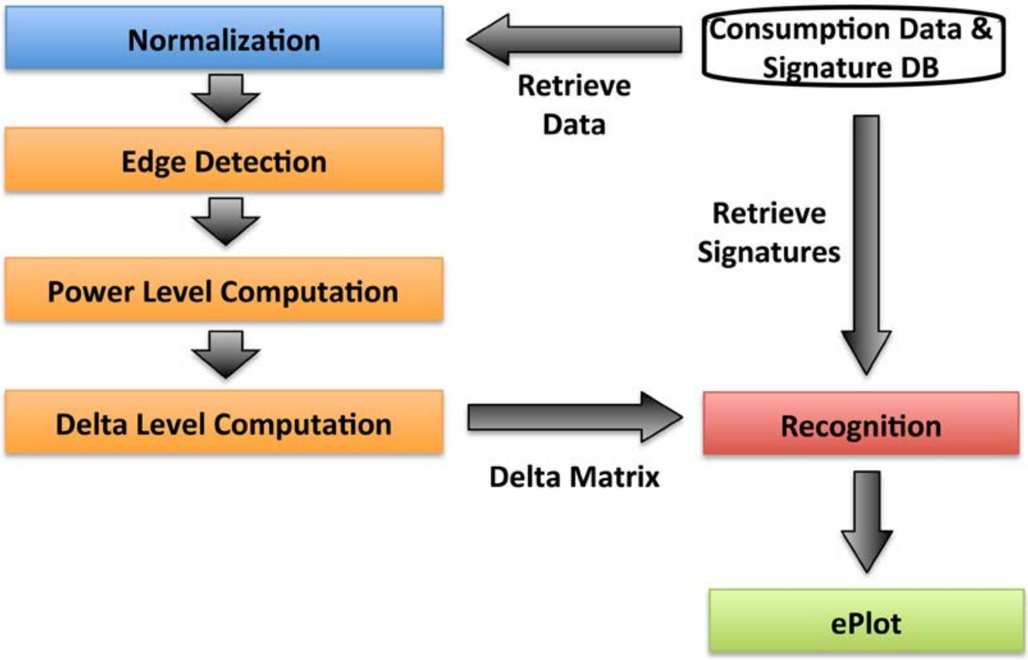

Figure 4.9 Overview of the individual steps of the AppliSense load disaggregation algorithm.

the AppliSense algorithm is performed for each electric phase individually. Instead of analyzing the high level overall power consumption, this enables AppliSense to use more fine-grained data to recognize appliance switching events.

2) Edge Detection: In order to identify edges in the recorded electricity consumption data that correspond to switching events of appliances, we use the normalized apparent power S'_n as input vector. The algorithm computes the absolute values of the differences between two consecutive values of normalized apparent power S'_n in the data series. If the absolute value of such a difference is larger than a predefined threshold f_th, then the value potentially belongs to an edge. However, depending on the applied strategy for detecting appliance switching events, there typically exist more potential edges than real appliance switching events. The threshold f_th is to be chosen carefully. It has to be sufficiently large to guarantee robustness against small changes in apparent power that occur due to noise on the electric power line. However, if the value is chosen too large, real appliance switching events might be missed.

Figure 4.10 depicts the apparent power of a Nintendo Wii usage cycle over a time span of 180 seconds. The two distinctive edges are related to operating the game console. It was turned on after 22 seconds and off after 106 seconds. The figure also shows the relatively strong fluctuations in apparent power during the device start phase compared to the relative

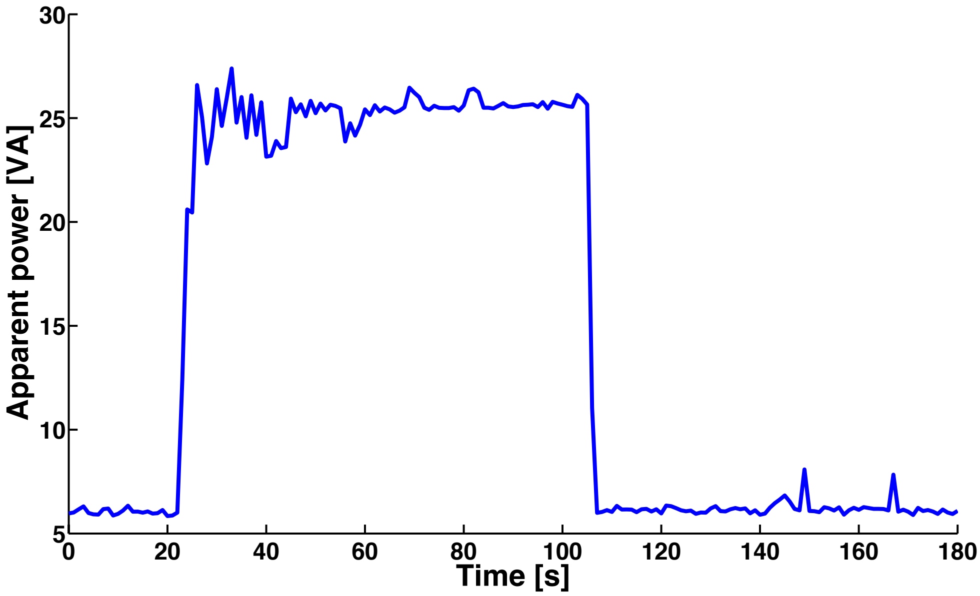

Figure 4.10 Apparent power for a 180 second duty cycle of a game console.

constant load during standby (from 0s to 22s). These fluctuations should ideally be filtered because they do not belong to a device switching event. The left of Figure 4.11 shows a histogram that depicts the difference of two subsequent apparent power values over the same duty cycle and time frame of the game console. We find larger changes in apparent power when turning the application on/off compared to times of operation or standby. We experimented with different thresholds and generally achieved best results applying a filter with a threshold f_th of 2VA. This removes a large number of time steps that do not correspond to a switching event (Figure 4.11 right). However, due to the transient behavior of the particular appliance, there persist some peaks (e.g., between 24 to 45 seconds) in the graph although no switching event occurred. In general, such oscillations during operation can be even stronger and more frequent in real-world scenarios, which would result in a high number of spurious events. Applying a smoothing filtering mechanism can help remove these false positives. However, the filter application also bears the risk of cancelling out edges (typically small ones) that correspond to real device switching events. Consequently, these switching events would not be identified and the operation of the corresponding appliance would be missed.

In order to decrease the number of spurious events, we investigated different smoothing filters. We tested a median filter, a mean filter, a kernel-weighted average filter (Nadaraya-Watson filter with Gaussian kernel), and different combinations of these on the apparent power signal. An advantage of a median filter is the ability to remove outliers. However, periodic curves (e.g., sine, triangle, saw tooth, square, etc.) could be removed.

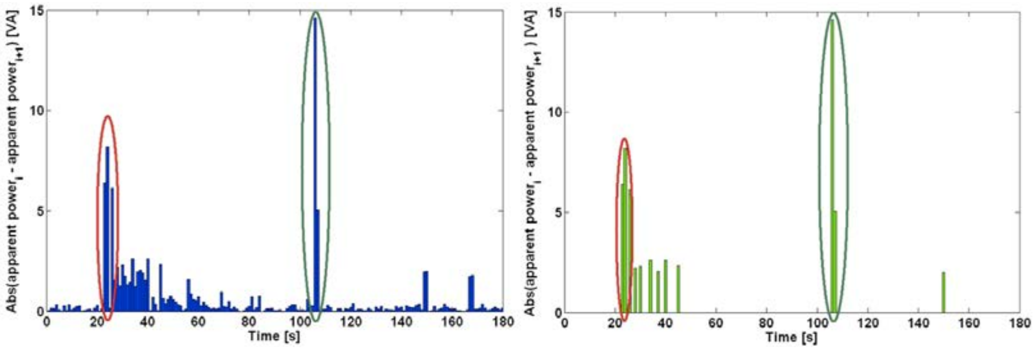

Figure 4.11 Resulting absolute differences in apparent power for the same duty cycle as in Figure 4.10 (left). An application of a filter of 2VA helps narrowing down the number of potential switching events (right). The red circle marks the turn on event whereas the green circle indicates the turn off event [3].

On the other hand, a mean filter which computes equally weighted averages of a sliding window of values has the ability to smoothen periodic oscillations, but may not always remove large outliers. Even worse, it might erase edges which correspond to an actual on/off switching of a device. A kernel-weighted average filter adds more complexity compared to the previously mentioned filters. Different kernel functions can be applied for signal filtering. We experimented with different ones and observed best results when applying a Gaussian kernel [208]. It allows preserving edges while attenuating oscillations of the original signal. The extent to which the filter smoothens the signal is determined by the kernel bandwidth, which relates to the window size.

In order to evaluate the influence of the filters on the edge detection and to find the most appropriate combination of filtering, we simulated a typical household usage scenario over 30 minutes in a controlled lab environment. During that period appliances of different characteristics were used and 12 appliance switching events occurred. Table 4.1 shows the results when applying the above-mentioned filters to the signal. The number in brackets corresponds to the window size/kernel bandwidth of the respective smoothing filter. The table displays the number of changes of apparent power values larger than 2VA, the achieved percentage in reduction compared to the original, and the number of missed appliance on/off events. Overall, the original signal contained 709 changes in apparent power S with a delta larger than 2VA.

Using a median filter or a mean filter alone reduces the number of potential edges by 74% and 70% respectively without missing a device switching event. The performance of kernel smoothening strongly depends on the bandwidth of the kernel. The potential reduction varies between 35% and 94% depending on the kernel bandwidth. A combination of mean and median filter achieves slightly better results (76%) than the two filters separately at no extra cost in terms of computation complexity. The reason for this relatively small improvement is due to the fact that the possibility to remove outliers is constrained by the small window size used during the evaluation. We experimented with different windows sizes and ultimate choose a window size of 5. Adding a kernel filter to the smoothening strategy leads to higher reduction in potential edges (between 3% and 17%). From a bandwidth of 60 on, we observe that the smoothening starts canceling out true switching events. Independent of the bandwidth parameter, however, using a kernel filter increases the computational complexity significantly.

Table 4.1 Performance comparison of different smoothening filters.

Filtering Method	**$\Delta S > 2VA$**	**Reduction**	**Missed**
Median(5)	185	73.9%	0
Mean(5)	218	69.2%	0
Kernel(3)	459	35.3%	0
Kernel(100)	46	93.5%	0
Median(5), Mean(5)	174	75.5%	0
Median(5), Mean(5), Kernel(3)	151	78,7%	0
Median(5), Mean(5), Kernel(60)	78	89%	1
Median(5), Mean(5), Kernel(70)	52	92.7%	4

Overall, we achieved best results using a kernel filter. However, this comes at high computational cost due to the quadratic complexity of the filter. Hence, we decided to go for a more efficient solution that performs close to optimum. It combines a median filter that removes outliers with a mean filter that further smoothens the signal (see line 5 of Table 4.1). The result of this smoothening strategy is illustrated in Figure 4.12. In our evaluation scenario, we used a notebook, several different lights, and a kettle to obtain the original power signal (blue). The blue markers correspond to the 709 points in time where the absolute difference of two subsequent apparent power values are greater than 2VA. Applying a median filter of five followed by a mean filter of the same size results in the green markers. The reduction gain (75.5%) of the filter can be seen by comparing the red with the green markers. The edge detection interprets the remaining 174 green markers as a binary vector which indicates at position i that the smoothed estimate of the apparent power at time step i differs by more than 2VA from the value at position i-1. Hence, the corresponding measurement at time i belongs to a potential device switching event.

3) Power Level Computation: Having identified the relevant edges, the next step of the algorithm extracts power levels that connect two edges in the smoothened signal. An overview of the relation between the key classes is provided in Figure 4.13. The sequencer uses the binary output of the edge detection to compute the individual levels including their particular characteristics (i.e., standard deviation, mean values, etc.).

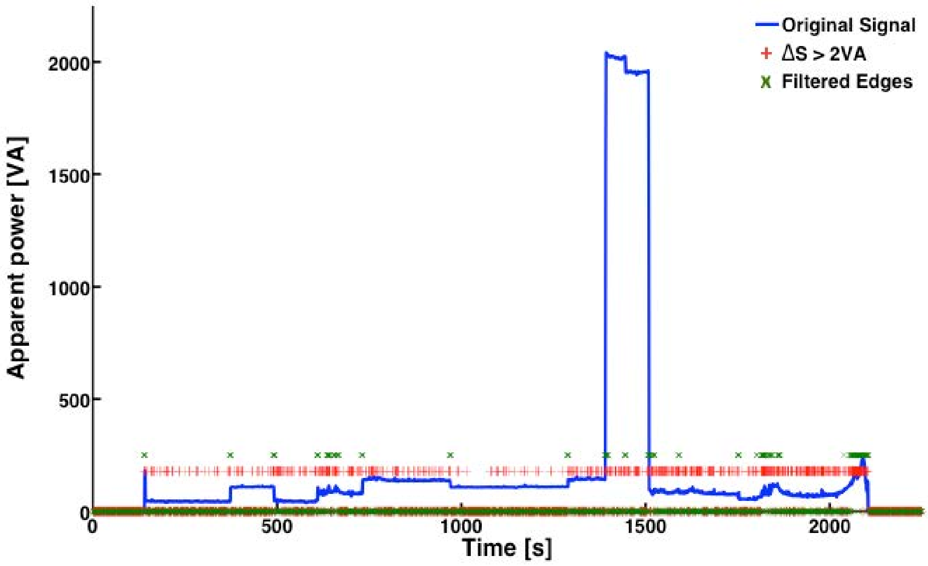

Figure 4.12 Evaluation of different filter strategies. Depicted is the application of a combination of a median/mean filter with a window size of 5 (see line 5 of Table 4.1).

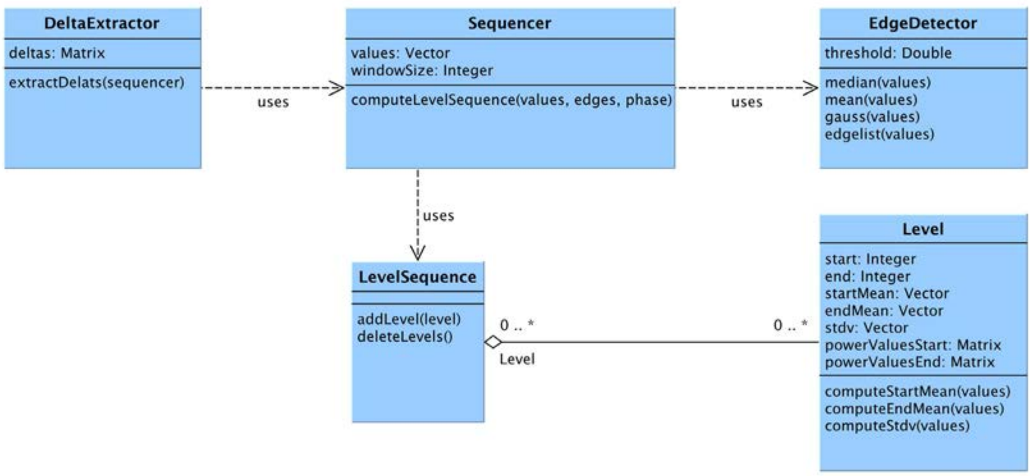

Figure 4.13 UML class diagram illustrating the key components of the AppliSense algorithm.

Each power level consists of a start and an end time, a vector with component-wise means of real, reactive, and distortion power for the first five measurements at the start and the last five measurements at the end of the interval (start mean (sm) vector and end mean (em) vector, respectively), and a three-by-five matrix which holds the real, reactive, and distortion power values for the start and the end of the interval. The component-wise standard deviation of all power values is also calculated.

4) Delta Level Computation: From two consecutive power levels, the algorithm computes the difference vector for real, reactive, and distortion power. These vectors are used for matching the edge to a particular device in the recognition step that follows thereafter. To take oscillations during start up and shut down of an appliance (e.g., due to heating up at the start of a kettle) into account, we not only calculate one difference vector for level i to i+1 (e.g., end of level i (em_i) – start of level i+1 (sm_{i+1})), but four difference vectors $d_{i,j}$ that include the start and the end values of both levels (see Figure 4.14):

$\overrightarrow{d_{s\_\iota,s\_\iota+1}} = \overrightarrow{sm_\iota} - \overrightarrow{sm_{\iota+1}}$,

$\overrightarrow{d_{s\_\iota,e\_\iota+1}} = \overrightarrow{sm_\iota} - \overrightarrow{em_{\iota+1}}$,

$\overrightarrow{d_{e\_\iota,s\_\iota+1}} = \overrightarrow{em_\iota} - \overrightarrow{sm_{\iota+1}}$, and

$\overrightarrow{d_{e\_\iota,e\_\iota+1}} = \overrightarrow{em_\iota} - \overrightarrow{em_{\iota+1}}$.

For each edge, we add these four vectors to a result matrix that is used as input for matching the device signatures in the next step.

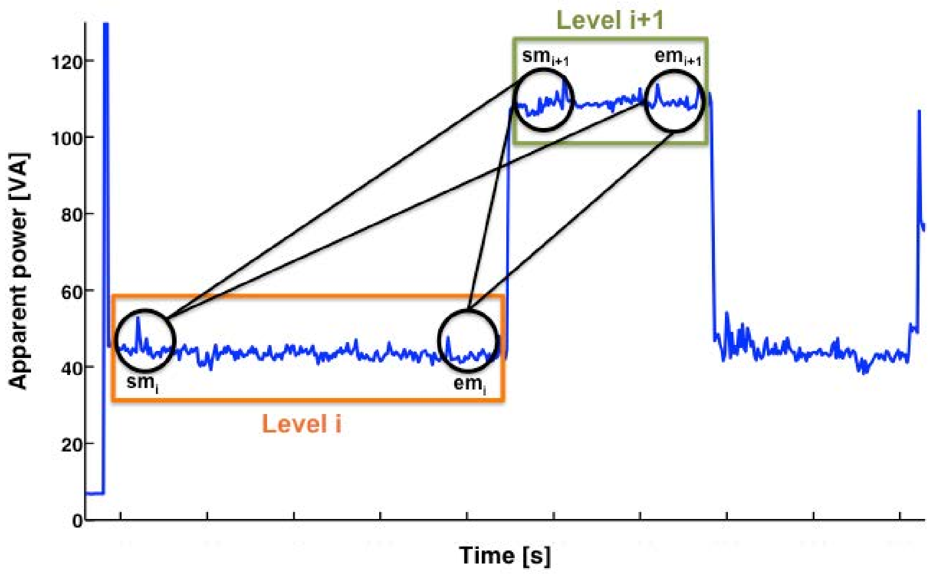

Figure 4.14 Computation of the difference vectors of two consecutive levels. Four vectors are extracted for each edge and compared to existing signatures.

5) Recognition: The recognition step of the algorithm tries to match known appliance signatures $\overrightarrow{k_j}$ from the signature database with extracted delta vectors $\overrightarrow{d_i}$ of the delta matrix D obtained as a result in the previous step. In order to identify an appliance on/off event, we perform a nearest neighbor search in the two-dimensional dQ/dP space (see Figure 4.15). First, the algorithm computes for every $\overrightarrow{d_i}$ its Euclidean distance to every $\overrightarrow{k_j}$ in the two-dimensional vector space. If this is smaller than a predefined value (r) of the length of $\overrightarrow{k_j}$ plus an oscillation value (osc), a potential matching is identified:

$$\|\overrightarrow{d_i} - \overrightarrow{k_j}\| < r \cdot \|\overrightarrow{k_j}\| + osc \begin{cases} if\ true, \overrightarrow{k_j}\ is\ a\ potential\ match\ for\ \overrightarrow{d_i} \\ if\ false, \overrightarrow{k_j}\ is\ not\ a\ match\ for\ \overrightarrow{d_i} \end{cases}$$

The oscillation term (osc) is the length of a vector which consists of the maximum of the standard deviation in the real power at level i or i+1 as first component, and of the maximum of the standard deviation in reactive power at level i or i + 1 as second component:

$$osc = \begin{pmatrix} \max\left(std(P\ at\ level\ i), std(P(at\ level\ i+1)\right) \\ \max\left(std(Q\ at\ level\ i), std(Q(at\ level\ i+1)\right) \end{pmatrix}$$

After this, every $\overrightarrow{d_i}$ contains a set of associated possible recognition candidates $\overrightarrow{k_j}$ from the signature database. Note that this set of possible associated recognitions could also be empty. In such a case, the corresponding $\overrightarrow{d_i}$ could not be related to a known signature. This could be caused for example by a detected edge which does not correspond to an appliance switching event, or by the non-existence of a corresponding signature in the database that matches $\overrightarrow{d_i}$. Second, for each $\overrightarrow{d_j}$, a nearest neighbor match

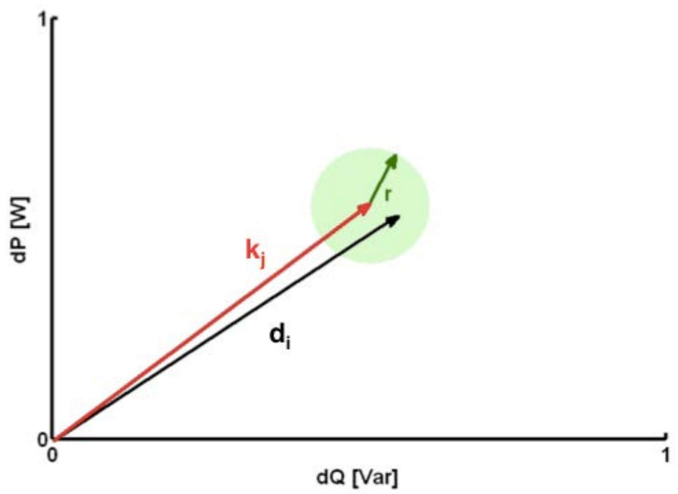

Figure 4.15 In the recognition step, known appliance signatures k_j are compared to extracted deltas d_j in the dQ/dP space.

is performed over all potentially matching signatures $\overrightarrow{k_\iota}$ that have been associated with $\overrightarrow{d_\jmath}$. In case the Euclidian distance of two candidates resides within an uncertainty range, it can eventually help to take the distortion power into account when conducting the nearest neighbor match. Finally, the algorithm writes the disaggregation results into a text file (see Figure 4.16) and launches the visualization process through the ePlot framework. The text file contains a human-readable output of the detected appliances in the load profile. It specifies the time, the event (i.e., either on (1) or off (0)), the phase the appliance was detected and the appliance name.

4.4.3.3 ePlot Visualization Framework

The goal of the ePlot framework is to provide a flexible way for visualizing electricity consumption data recorded by the eMeter system. The framework takes the extracted and processed electricity consumption data that is passed along the tool chain as input vector and supports two different functionalities for visualization (see Figure 4.17).

Date time	Event	Circuit	Device
2009-11-12 09:32:24	1	1	Office lamp (fluorescent)
2009-11-12 09:34:01	0	1	Office lamp (fluorescent)
2009-11-12 09:34:27	1	1	Light bulb 70W
2009-11-12 09:42:01	1	1	Kettle
2009-11-12 09:42:34	0	1	Kettle
2009-11-12 09:54:25	1	1	Light bulb 70W
2009-11-12 09:54:57	0	1	Light bulb 70W
2009-11-12 09:55:53	1	1	Kettle
2009-11-12 09:56:20	0	1	Kettle
2009-11-12 09:58:42	1	1	Office lamp (fluorescent)
2009-11-12 10:01:18	1	1	Light bulb 80W
2009-11-12 10:01:32	0	1	Light bulb 70W
2009-11-12 10:02:26	1	1	Kettle
2009-11-12 10:02:45	0	1	Kettle
2009-11-12 10:06:20	1	1	Light bulb 80W
2009-11-12 10:07:15	0	1	Kettle
2009-11-12 10:09:00	1	1	Light bulb 80W
2009-11-12 10:31:01	1	1	Light bulb 70W
2009-11-12 10:31:53	0	1	Light bulb 70W
2009-11-12 10:32:09	1	1	Light bulb 70W
2009-11-12 10:32:39	0	1	Light bulb 70W

Figure 4.16 Sample of the human-readable output of the AppliSense recognition step.

First, direct data plotting allows directly plotting the electricity consumption information in a flexible way. Users can arbitrarily choose from all physical quantities that are available in the eMeter system and limit the visualization to individual phases, to the aggregated consumption of all phases, or to any combination of the just-mentioned. Additional parameters easily enable to change the style of the presentation. For example, the figure can be automatically saved, the color of individual data plots can be changed, or the x-axis can be transformed to rather display the date than continuous seconds, which is useful when analyzing behavior of inhabitants. The second functionality – labeling – works after the same principle, but additionally assigns the results of the AppliSense algorithm as labels to the plotted load curve. For each detected switching event the corresponding edge in graph is associated with the corresponding device that caused the switching event. An example of this visualization is depicted in the next section as part of the evaluation of the AppliSense algorithm. The functionality of the ePlot framework is implemented in Matlab and can be called within the workspace of the program.

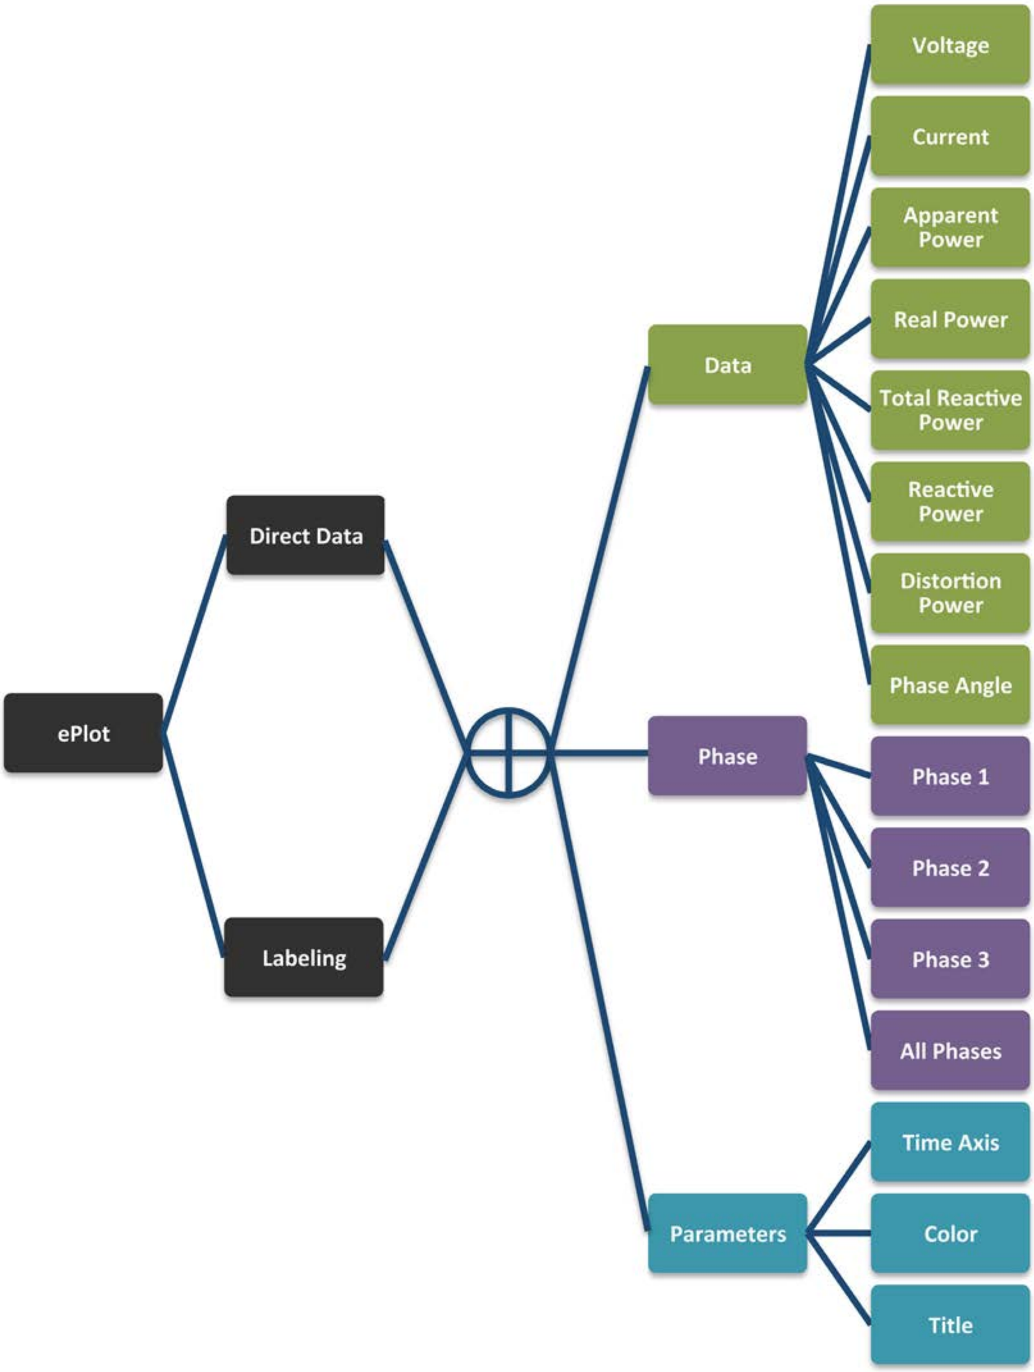

Figure 4.17 Overview of the ePlot functionalities. For the two different modes – Direct Data and Labeling – the options Data, Phase, and Parameters can be used to flexibly specify how the corresponding data should be plotted.

4.5 Evaluation

In the following we describe the evaluation we conducted to assess the performance of the AppliSense disaggregation scheme. Thereafter, we conclude this section with a brief discussion on the steps we took to prepare AppliSense for the real world.

4.5.1 Laboratory Study

In order to analyze the performance of the AppliSense algorithm, we installed the eMeter system in a laboratory environment (see Figure 4.18). For the evaluation, we used a controlled set of appliances, which typically occur in a student's household. Table 4.2 provides an overview of the appliances, their real power consumption stated on the manufacturer label, their verified real power range in operation (measured by a separate power monitor), the appliance category (O for ohmic, I for ohmic-inductive, and C for ohmic-capacitive), and the real power that is obtained as part of the power signature using the eMeter smartphone application. All devices were connected to the same phase during the whole evaluation. Some of the appliances have power consumptions within the same range. However, if belonging to different appliance categories, we should still be able to differentiate the corresponding events.

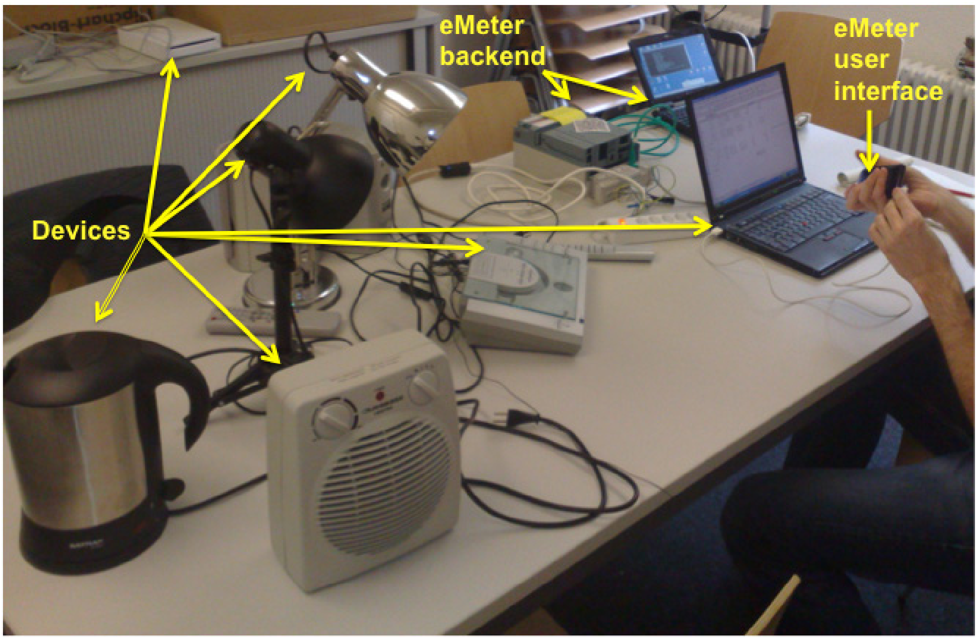

Figure 4.18 Laboratory setting for the evaluation of the AppliSense algorithm. We used different typical household appliances along with the eMeter system measure the performance of our proof of concept implementation.

During times when only a single appliance was active, the algorithm identified the on/off events of all devices except the CD player correctly. Every device was turned on and off at least three times. The edges caused by the CD player were not recognized neither when being turned on nor when being turned off. This can be explained through the limitations introduced by the filtering. Using a window size of 5 in our test scenario leads to a lower boundary of 10VA for edges that can be recognized. The CD player has a relatively high standby consumption of 6W compared to its 3 – 7W in operation. While the median filter does not influence the signal, the constant 3W during operation result in a step-wise increase of 0.6VA after application of the mean filter. This increase is too small (<<2VA) to be detected as an event by the algorithm using the chosen combination of the median/mean filter.

Table 4.2 Appliances used for the evaluation of the AppliSense algorithm.

Appliance	Labeled Power	Power Range	Category	Consumption
Light bulb	75W	70W	O	70W
Kettle	2200W	1855 – 1933W	O	1900W
Heater	2000W	1619 – 1667W	O	1635W
CD player	13W	9 – 13W	I	3W
Fan	50W	45W	I	45W
Notebook	72W	30 – 35W	C	35W
Fluorescent lamp	35W	21 – 28W	C	25W
Wii	52W	10 – 45W	C	15W

Next, we combined the use of multiple devices in a random order. Although the CD player cannot be recognized, we operated it and other devices with unknown signatures from time to time to vary the base line consumption and to have more appliances concurrently running. Over a time span of several hours, we documented 80 switching events. 77 times the algorithm detected a switching event correctly, which means that for the time stamp t an on/off switching was conducted on circuit i of appliance X, the algorithm output corresponds (t, i, X, on/off). Figure 4.19 shows a sample labeling output of the algorithm for a simulated office environment. After the notebook has been turned on, different devices were concurrently used and a kettle was operated. However, the red circle highlights a moment at which the office lamp is turned on but the event is not detected. This is due to the oscillations caused by a device that was operating at the same time. A second (not depicted) problem occurred when switching on and off the notebook. Due to the different battery levels, the power consumption had varied compared to the one registered in the appliance signature database. This led to the correct detection of the corresponding edge, but no appliance signature was matching.

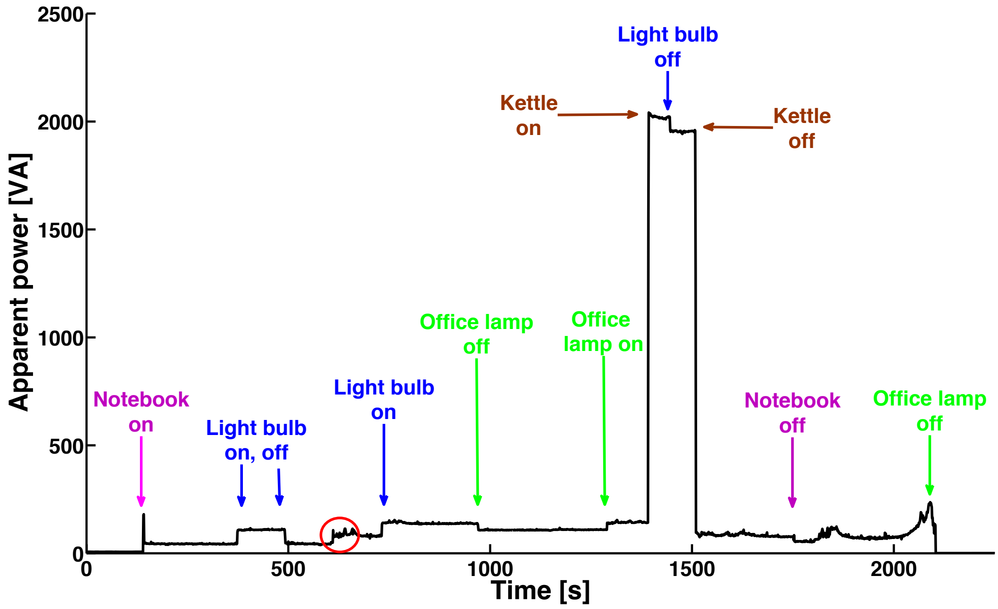

Figure 4.19 Labeled load curve as output of the AppliSense algorithm.

Overall, the evaluation shows promising results. We generated 144 device switching events in our test scenario. 16 of these came from devices with a small consumption so that the corresponding edges were cancelled out. When subtracting theses events, the algorithm identified 125 out of the remaining 128 events correctly, which results in a recognition rate of about 90%. In practice this enables interesting applications, such as automated recommendations for a more economic use of electricity in households.

4.5.2 Preparing AppliSense for the Real World

The results of the laboratory study encouraged us to investigate the real-world applicability of our signature acquisition process and prepare AppliSense for the real world. We used two households (household number two and three (see Section 3.4.1)) of our running real-world deployment to test the appliance signature recording in a more dynamic residential environment. Since people were permanently living in these homes, extended access was rather limited. We spent several hours in each of the two households taking signatures of various appliances.

Overall, we the signature acquisition process was working reliably for most appliances even while other appliances were operated in different rooms. However, appliances with varying power consumption were tricky to capture. For example, a PC produced different signatures when waking up from standby. We verified that in case two switching events occurred at the same time, but on two different electrical phases the measurement functionality on the mobile phone worked reliably. That is, it indicated that an error during the measurement had occurred and the user was asked to repeat the appliance switching event. On the second try, the signature acquisition process was then able two identify the correct switching event (for algorithm details see Section 2.5.4). We also discovered some anomalies that were not visible in the laboratory deployment. For example, the stove is using more than one electrical phase at all times. However, evidence from more households would be needed to draw conclusions on how to adapt the signature acquisition algorithm to record the signature of this device as well as on how appliance signatures fluctuate as new appliances are introduced.

Evaluating the performance of AppliSense in a real-world deployment also requires ground truth information on the operation times on which devices are running. Ideally, such an experiment should be conducted in a fully controlled environment (e.g., a living lab). Then every appliance could be equipped with a sensor (e.g., a smart power outlet) that individually monitors the consumption of the device. This information could then be compared to the recognition output of the AppliSense algorithm to analyze the energy identification ratio of the proposed disaggregation scheme.

We thought on less invasive alternatives that enable recoding ground truth information even in settings where equipping every appliance with an individual sensor is impracticable and not an option. One alternative that is often used but less precise is the use of logbooks in which residents manually keep track of the devices they operate [209]. Main drawbacks result form the lapse of time that occurs between the operation of an appliance and the manual entry in the logbook as well as from the time difference between the user and the system clock. This is a significant drawback in a dynamic environment in which time accuracy plays an important role.

In order to make this process more convenient for users and at the same time more accurate, we developed an electronic logbook application on smartphones. The interface provides a list of the most common residential appliances together with on and off buttons on the left and right respectively (see Figure 4.20). In case the utilized device is not yet listed, the application offers the possibility to add new appliances in a single step at the bottom of the application. A total count of appliances is displayed on the upper left. It is intended to inform users on the number of devices they operate in their household. Instead of manually recording an appliance switching event with pen and paper, users only have to press the corresponding button on the mobile phone and an appliance switching record is generated in the database of the underlying eMeter system. This should

Figure 4.20 User interface of the logbook application that can be used to track appliance switching events.

help mitigate the above-mentioned drawbacks. First, it improves the time accuracy, since the application uses the eMeter system clock. Second, it helps decrease the time between the event and the record entry since users can provide feedback on the spot.

The residents of three different households tested the functionality of the logbook application. Figure 4.21 illustrates the results of an individual household that was tracking its appliance switching events with the application. The figure depicts the overall real power consumption of the household over a duration of roughly 3 hours. The labels were extracted from the eMeter system database and correspond to appliance switching events. They were manually generated by the user utilizing the logbook application when operating different appliances at home. Most of the time, the user was keeping track the individual switching events and the labels could easily be assigned to the corresponding edges. However, the figure also shows that there exist edges with no corresponding label (see yellow circles in Figure 4.20). These either originate from appliances that operate automatically or were accidentally omitted by the user, and shows one of the drawbacks of this approach (another one arises from relying the user and the potential of user mistakes when handling the application). Overall, our insights demonstrate the correct operation of the developed logbook application and outline such a Ubicomp-enabled logging approach is a well suit-

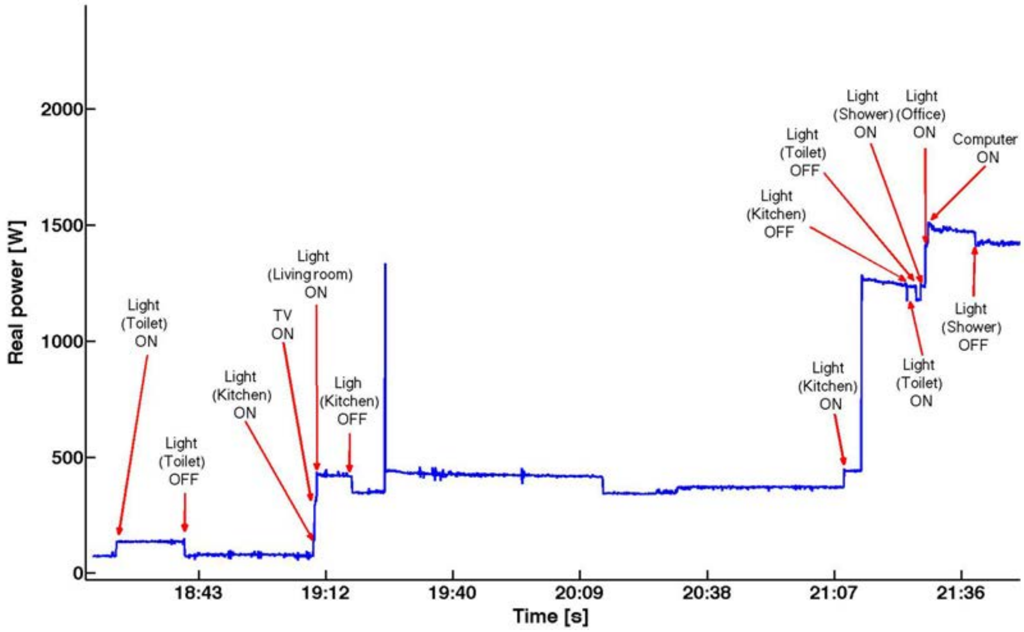

Figure 4.21 Appliance switching events labeled by the user utilizing the logbook application.

ed alternative for gathering ground truth information in real-world environments in which equipping each appliance with an individual power sensor is not an option.

The system is thus prepared for a future evaluation in a real-world environment. In order to measure the system accuracy and the energy identification ratio our qualitative insights show that a more controlled environment is needed. To achieve this, we envision an experiment in which smart power outlets individually meter every appliance of the household or if this is not possible, the use of the logbook application to generate ground truth information (e.g., residents can be assigned designated days or time frames in which they use the logbook application to label each appliance switching event). The first option is without doubt preferable over relying on users to provide the ground truth information, but might be hard to realize and most likely is not possible outside living lab environments.

4.6 Discussion and Limitations

The evaluation of AppliSense through the conducted laboratory study confirmed the operability of our proof of concept implementation. It shows a way a how Ubicomp can help simplify the signature acquisition process and as such contributes to the better applicability of load disaggregation systems. However, the evaluation also shows that there exist different limitations in the current implementation of AppliSense. Some are bound to the assumptions made during the design, some are due to the nature of the approach (i.e., recognizing device switching events), and others originate from the details of the current implementation of AppliSense itself. In the following, we are going to discuss the limitations and room for future improvements of AppliSense.

The proposed AppliSense disaggregation scheme is based on the early principles of Hart. We use steady-state power draw levels to identify edges in the overall power signal and thereafter match these edges to previously recorded appliance signatures. For that, we designed a Ubicomp-enabled, incremental signature acquisition method, which is based on the assumption of steady-state devices (i.e., devices that have discrete power states that can be measured with the measurement feature of the mobile phone application). As a result, users can train AppliSense with a portable mobile application without the need for special-purpose hardware or complex calibration by domain experts. This simplifies the signature acquisition process and fosters the applicability of such disaggregation systems, but some limitations of the underlying principle remain:

- Devices that have multiple power states (e.g., a hair drier with different fan speeds and heat options) require one recorded signature per state. Moreover, our evaluation showed that oscillations caused by operating devices can mask the switching event, particularly of low power drawing appliances. This could especially be a problem in larger households (e.g., family houses) with lots of appliances and activity.

- Continuous variable loads (e.g., a personal computer or a washer) pose another significant problem. Like most other disaggregation approaches, AppliSense cannot detect devices that do not have well-defined operation states, but have a continuously changing consumption. This is due to the initial assumptions regarding the algorithm design and the tradeoff for relying on a single sensor system with a 1Hz sampling frequency.

- Last, our disaggregation scheme is bound to switchable appliances, which at the moment excludes non-switchable devices such as fridges or freezers from the disaggregation process.

Different measures can be applied to overcome some of the abovementioned disadvantages and thus help enhance the disaggregation performance of AppliSense. In particular, the improvements could consist of the following:

- AppliSense could be extended with a module for the recognition of non-switchable, hard-wired heating and cooling appliances, which account for a large portion of the residential electricity bill. The signatures of such appliances typically become particularly evident during nighttime. The module could filter these switching events and automatically (i.e., without relying on the user) attribute them to respective appliances, since there only exist a limited number of such appliances in typical residential environments (i.e., one or two fridges, a freezer and eventually an electric water heater) and these devices usually contain different power-states and operating cycles. Since being based on the same edge detection principle, the extension could be integrated in the current disaggregation scheme and allow AppliSense to attribute another significant portion of residential electricity use.

- We envision the just-described module as a first step towards a more sophisticated, unsupervised learning mechanism that aims at classifying reoccurring edges in the overall power signal. Reoccurring edges raise the probability level that they are related to a real appliance switching event. Once a certain probability threshold for one particular event is reached, user could be prompted to confirm the operation and details of the utilized device on the portable user interface. Even more sophisticated, this prompt could suggest what device just has been used based on the appliance signature´s characteristics right after the switching event occurred. All this would help streamline the interaction and further minimize the burden for users.

- With respect to continuous variable loads, we foresee extending the eMeter system a limited number (one to three) smart power outlets that could be used to individually monitor variable loads that have a significant impact on the residential electricity bill (e.g., washer, drier, etc.). This would potentially allow improve the disaggregation performance of AppliSense without much increasing the complexity for users.

Other drawbacks result directly from the implementation of the applied filter design and the clustering method for detecting edges in the overall power signal. Addressing these shortcomings would require changing the current implementation of the algorithm. The current thresholds for the utilized quantities have proven to be effective in the lab environment, but may require refinements in other settings (e.g., in real-world deployments), and as such have not been tested for their optimal effectiveness. Originally, we envisioned the distortion power to be a potential additional classifier for appliance switching events, but we observed that it is not a reliable feature that often depends on other appliances that are currently operated.

Thus, we developed the signature matching based on a nearest neighbor search in the dP/dQ-pane without further taking the distortion power into account. Limitations directly related to the current implementation are the following:

- Devices that contain a relatively small power consumption are not detectable by AppliSense. The lower boundary of 2VA in the current filter design for two consecutive power values together with a mean filter that is calculated of the last five values leads to a lower detection boundary of 10VA. This is sufficient for most home electronics, but not for low power-drawing equipment such as phone chargers, small radios, or LEDs. These devices are per se not detectable with the current design.

- Short operating times of devices can lead to missed appliance switching events. Due to the current implementation of the power level computation, the load has to remain within a relatively stable interval for at least five seconds or more to be considered as a level. If not, the corresponding time span in the power signal might still be classified as potential edge and AppliSense does not detect the operation of the device. In practice, this should have a relatively low impact on the recognition performance, since most devices should be running for more than five seconds after their initial start.

- Several devices that are switched on or off concurrently or within five seconds and that reside on the same electrical phase cannot be detected because AppliSense would interpret the two contiguous switching events as one. This would result in an edge that cannot be matched to an existing appliance signature.

- In the conducted laboratory study we observed that the appliance signatures recorded with the mobile phone application were very reliable. That is, the delta vectors obtained with the measurement function when turning an appliance on/off are stable and reproducible over time. However, the measurements conducted during the limited time in the household provided first evidence that this may be different in a more dynamic home environment. In addition, the presence of new devices can have a disturbing effect on previously recorded signatures - there the AppliSense algorithm may need several (slightly different) signatures per device to reliably recognize appliances.

4.7 Summary and Outlook

In this chapter, we gave a detailed description and evaluation of the third contribution of this thesis: AppliSense, a Ubicomp-enabled disaggregation scheme that facilitates automatic recognition of switching events of electric appliances to disaggregate the total electrical load to device level. In contrast to existing approaches, our objective was to develop an applicable system that achieves this without requiring a complex training period and with minimal user involvement. AppliSense does not rely on special purpose hardware, instead it makes use of components that are becoming ubiquitous in home environments: a smart meter and mobile phones. It shows how Ubicomp-enabled feedback systems can be used to much simplify the signature acquisition process, which is mandatory for many load disaggregation systems while at the same addressing the great need for user interfaces, which effectively help improve disaggregation algorithms by utilizing user-feedback [82, 206].

A measurement feature on the mobile phone is used to acquire the appliance signatures. Its non-intrusive design easily blends with residents' daily life. By utilizing the user input, AppliSense can incrementally establish its appliance signature database over time and avoid long and complex training at the beginning after the system setup. Moreover, where other systems require a complete recalibration, AppliSense enables to take new appliances into account that were not present at the initial setup of the system. This is particularly important in fast changing residential environments where old devices are frequently replaced.

Applying data analytics to the gathered metering data allows the system to raise energy awareness by providing better-tailored energy feedback. With a recognition rate of about 90% in the laboratory environment, the results of our evaluation study confirm the suitability of the general scheme of AppliSense and provide ground for many interesting applications on top of the system. An automated recommendation service could derive household-specific energy saving measures and provide residents with tips on the use of their appliances. In combination with actuation capabilities, we can foresee the information provided by AppliSense to be used to automatically optimize energy consumption and hence increase residential energy efficiency. For example, opportunistic sensing methods could be used to derive occupancy state from electricity and appliance use data, to use this information in a smart heating control strategy [41, 42, 210]. Not least, appliance level consumption information can give rise to new business models (e.g., providing cross-selling offers for non-energy-efficient devices).

The conducted test in the real-world deployment provided first qualitative insights on the stability of the signature acquisition process in practice and helped prepare the system for the real world. However, the dynamic residential environment also showed that a more controlled setting (e.g., a living lab or houses with individually submetered appliances) is ideally required for the further evaluation of AppliSense. Such a deployment would also allows analyzing the algorithm's dependency on the number of manually recorded signatures and how the presence of new appliances is affecting these previously recorded signatures. This also includes the possibility to open the scheme for the combination with other disaggregation techniques, which is most likely to achieve better disaggregation results. As a first step to overcome existing limitations, we envision a module for auto-identification of hard-wired heating and cooling devices. Further improvements can result from using unsupervised learning methods to attribute detected edges that do not yet correspond to an existing signature in the database to appliances [196, 197]. The application of clustering concepts that automatically classify these events together with the possibility to prompt users for their confirmation on the portable user interface once a certain probability level is reached can help increase the electricity identification ratio (i.e., the amount of total electricity that can be classified by AppliSense).

On a larger scale beyond household-level, combining the signatures identified by AppliSense in different households with PowerPedia can help building a central appliance signature database [205]. This database then can serve as a central repository of appliance signatures, and feedback the information as input knowledge for AppliSense itself or other disaggregation approaches.

5 Conclusion

Residential electricity consumption has been continuously increasing over the past decades mainly for two reasons. First, the increasing number of electrical appliances in homes (especially due to the rising amount of small consumer electronics) and second, the behavior of residents and the way they operate their appliances. The lack of transparency regarding the residential electricity consumption means that even people that are willing to save electricity are not aware of possible electricity conservation measures. This motivates the need for meaningful residential electricity feedback and guidance on how to save electricity and money. To address this situation, we investigated how Ubicomp technologies, which enable to digitally enhance physical objects with computing, sensing, and communication capabilities, can help provide meaningful electricity feedback that goes beyond the mere visualization of consumption values and at the same time is unobtrusively integrated into daily life. In the remainder of this chapter, we will first summarize the contributions of this thesis, and then we will discuss limitations and open challenges for future work.

5.1 Contributions

In this thesis, we presented a user-centric approach that combines the use of smartphones and smart meters to demonstrate how Ubicomp can help foster residential energy conservation. Applying Ubicomp in the residential domain raises inherent challenges on how to design applicable electricity feedback systems that provide meaningful information. This comprises the components that are used to integrate the electricity feedback into people's daily lives as well as the modality of information and the functionality that support users in their conservation efforts. In particular, the individual contributions of this thesis can be summarized as follows:

- We designed, developed, and evaluated a pervasive electricity sensing and feedback infrastructure that makes use of Ubicomp technologies. The resulting eMeter system combines the use of a smart meter and mobile phones to provide meaningful electricity feedback on a portable device in a way literature suggests. At the same time, the eMeter system serves as an easily extendible framework that can be used by other researchers for their experiments. We demonstrated its usefulness through the implementation of three different user interfaces and the integration of PowerPedia – a central collaborative platform for appliance-specific electricity consumption feedback. Our evaluation of the eMeter system in a laboratory study and in four Swiss households confirms the real-world applicability and feasibility of our approach. Overall, this contribution shows how Ubicomp can in the future help to realize electricity feedback systems that feature a low usage barrier and enable users to better understand effective measures for conserving electricity.

- We confirmed the suitability of mobiles phones as energy feedback devices in a user study with 25 participants and in a long-term real-world deployment. For that, we specifically developed a smartphone application which uses the data provided by the eMeter system. This Ubicomp-enabled electricity feedback features real-time information provisioning on the spot and offers interaction possibilities that allow users to engage with their consumption. Using this feedback interface, we identified which electricity feedback functionalities are perceived most valuable by users and what is necessary to address different user types. In the long-term study, we confirmed that there does not exist a one-size-fits-all feedback solution, motivational concepts are required to address user fatigue once the initial curiosity with the application has been satisfied, and meaningful feedback that enables users to derive direct measures is required to foster energy conservation.

- We developed and evaluated AppliSense, a disaggregation scheme that leverages the total residential electricity load measured by a smart meter to automatically recognize home appliances. By making use of the interaction capabilities of the mobile phone user interface, we show how Ubicomp can help simplify the otherwise cumbersome (and often also costly) appliance signature acquisition process. The specifically designed measurement functionality that is used for the signature acquisition method is integrated in the resident's normal daily life. We thus addressed the need for improved user interfaces that allow training of recognition algorithms based on user feedback. The proposed disaggregation scheme also allows addresses technical challenges. For example, it enables taking devices into account that are introduced at home after the initial system setup without requiring a recalibration of the disaggregation system. A laboratory study with up to eight consecutively running devices was used to confirm our proof-of-concept implementation achieving recognition rates of almost 90%, which is sufficient for many interesting applications (e.g., an automated recommendation service on how to conserve energy).

5.2 Limitations and Future Work

In this thesis, we pursued an applied approach that focused on identifying and providing meaningful electricity information based on an applicable electricity sensing and feedback infrastructure for residential environments. To do so, we focused on the infrastructure perspective during the first part and used human computer interaction methods in other parts of this work. This led to a number of open issues that remain unaddressed in the context of this work, but nevertheless may be interesting for future work. In the following, we first explain how the eMeter system can help enable automatic energy conservation, before we discuss limitations and future work directly related to the three contributions of this thesis.

Leveraging energy conservation through automation. Ubicomp can contribute to energy conservation through supplying users with meaningful information and supporting behavioral change, but also by enabling automated energy savings and resource optimization. Technical measures like automation are often preferred by users and offer significant energy saving opportunities [211] that have not been investigated in this thesis. Indeed, automated energy savings are a desirable form of energy conservation since it can be realized in the background so that it is mostly invisible for users. At the same time, such savings can be achieved even at times when users are not present, and neither rely on users' knowledge on how to conserve energy nor require behavioral change, which is hard to induce [56, 212]. We envision leveraging residential energy efficiency through combining the meaningful feedback of the eMeter system with automated energy conservation [61]. Ubicomp technologies incorporated into devices can form networks and make use of energy data (e.g., provided by the eMeter system) that potentially help adapt to available resources and optimize consumption [39, 40]. This becomes particularly relevant in the context of Heating, Ventilation, and Air Conditioning (HVAC) systems. Their penetration has almost tripled over the past 30 years and they account for 49% of the residential energy consumption in the U.S. and contribute even more significantly in Europe (e.g., 61% in the U.K. and 70% in Switzerland) [4, 16, 17]. Over the past decades, the share of energy consumption from HVAC systems is decreasing as the technology matures and efficiency increases, but a lot of energy is still wasted because HVAC systems are typically not turned off or down when occupants are sleeping or away, which leaves a lot of room for automated optimization [213]. For example, information on location and velocity available from sensors integrated in the eMeter mobile phone user interface combined with the electricity and appliance use data gathered by the eMeter backend system and AppliSense could be used to automatically adjust heating to home occupancy and user preferences [41-43].

Moreover, combining the eMeter system with networked sensors integrated into appliances that enable communication and automated control can be used to coordinate the decision making of when and how to operate these appliances (e.g., time to start and mode of operation). This can enable smart grid features such as the use of renewables whenever available, the immediate reaction to pricing signals, and the integration of electric vehicles [44-50]. However, savings through automation are difficult to achieve and not applicable for every household [34]. Barriers, among others, include the large variety of appliances and communication standards that hinders adoption [34], the different personal preferences of inhabitants, the physical limitations of the environment, and high upfront costs. Moreover, even in fully automated systems the human factor can lead to so-called rebound effects[29] and thus should be taken into account when designing such systems. More globally, there additionally exists some risk that ICT has only a low overall effect because positive and negative environmental impacts partially cancel each other out when aggregated [214].

Advancing the eMeter system. Designing an electricity feedback infrastructure for residential environments ideally should take the lifetime of buildings into account that is typically much longer (i.e., 40 – 50 years) compared to the innovation cycle of today's information technology. Updating components or exchanging hardware is rather difficult in typical homes with mostly non-techy residents. Energy consumption feedback systems, like the eMeter system developed in this thesis, thus have to be future-proof with respect to the installed components and communication protocols. We only partly addressed this issue by relying on TCP/IP and HTTP over WIFI for communication, which are established protocols at the time of the system implementation. In addition, the ongoing miniaturization of information technology might lead to web servers being incorporated in many home appliances and thus these protocols might be used more widely in future smart home environment. However, this can be subject to change, as hardware components become obsolete and new technologies and communication standards arise (e.g., ones that are specifically designed for residential areas or even for the information flow along the electrical grid) [34]. Moreover, the eMeter system currently is not designed for energy efficient communication nor trimmed to use the least power-

[29] E.g., a person with a fuel economic car, for example, might partly compensate the savings of the technology by simply driving more, because it is now cheaper.

drawing hardware resources to conserve energy whenever possible. While the energy use of the electricity feedback system itself is without doubt important for commercial products, it seemed impractical to pursue this issue (e.g., by designing an embedded device with the amount of resources specifically needed or by optimizing the systems internal energy usage by cutting back CPU frequency or WIFI power) in the context of this work.

Future-proofness also has to be kept in mind with respect to security. In particular, when looking at data communication across the electrical infrastructure, security becomes necessary through the recently started transformation towards the smart grid. Like buildings, the electricity grid is intended to be operated for a long time, but network security concepts and means (e.g., key lengths or encryption algorithms) – and the possibilities of attackers – change at a much faster pace [38]. Through the use of smart meters and the potential extension of the information flow into the home, security issues have to be incorporated into the design process. In the future, security leaks also might rise from devices that primarily had nothing to do with the stability of the electrical grid itself. At home, interconnected smart spaces are evolving in which devices such as televisions, digital picture frames, Wi-Fi-enabled radios, media centers, entertainment systems, and alarm clocks make use of integrated communication interfaces and computing power to offer a wide range of new services [22, 33, 34, 37]. A virus provoking malfunction of these interconnected devices or a denial of service attack could lead to serious damage not only at home, but also at the electrical grid level (e.g., cause an unexpected grid overload) [38].

Gathering data in the residential environment also raises a potential threat to privacy and as such is a tradeoff. Knowing much about the electricity consumption and the appliances used within the home may reveal much about occupancy state and the standard of living, but also enables interesting opportunities for remote services, global optimization, and stability of the electrical grid [51, 52]. We did not specifically focus on privacy within this thesis, but the developed eMeter system offers different communication modes that are useful in this context. For example, one leaves most of the recorded consumption information inside the home. Communicating only data that is absolutely essential with respect to the application scenario in mind (e.g., for billing purposes only some of the collected data has to be exchanged with the energy utility) might be part of the solution. Establishing people's trust in the individual entities that are most probably distributed across different industries might be another difficult challenge [38].

Utilizing eMeter for behavioral science research. Our evaluation of the user interface of the eMeter system confirmed the suitability of mobile phones as energy feedback devices and shows what functionality is valued and required by different user types. The results from the conducted user study and the long-term deployment thus provide insights for the design of future residential energy saving applications. It would further be interesting to confirm the results in a larger trial from a behavioral science perspective, which offers the possibility to control the accompanying study conditions (e.g., has a control group, can take weather effects into account, etc.). Such an experiment would also allow a more detailed identification of the feedback preference of different user types, which is crucial for the large-scale success of residential energy feedback systems. Since there exists a high diversity among the targeted user base, energy saving applications have to be able to address all users adequately, regardless of their doubtless different personal preferences in terms of provided functionality and design. Thus, to effectively encourage household energy conservation electricity, feedback has to be tailored to specific user groups and use case scenarios [144].

In addition, the eMeter system is built as an easily extendible framework. It enables exposing people to a fully functioning user interface with a rich set of functionalities instead of relying on paper prototypes or surveys. As such, it would be interesting to make further use of the system to quantify achievable short-term and long-term energy savings, experiment with different forms of visualization and functionalities, and identify what further motivational concepts are necessary to not only counter user fatigue and establish a permanent use of the system, but also induce a pro-environmental behavioral change of users [72, 146, 147]. Moreover, the use of the system can be combined with flexible energy tariffs [215-217]. So far, research has lacked the ability to test how users react in response to different price structures when being supported through a portable user interface that is able to provide on the spot information and interaction capabilities independent of the current location of users.

Enhancing the AppliSense disaggregation scheme. The disaggregation of the overall electricity consumption to specific end uses is an open challenge that has been around for more than 30 years. We contributed to this challenge with our own disaggregation concept AppliSense that shows how Ubicomp can help much simplify the appliance signature acquisition process required for most disaggregation approaches. We evaluate our concept in a laboratory study and provide first insights on the real-world applicability. Thus, it would be interesting to follow up on the results by testing AppliSense's real world energy identification ratio [200] in a controlled experiment and thereafter implementing the necessary refinements. Such an experiment would also allow for further research on the appliance signatures and their characteristics in dynamic real-world scenarios. For example, recording multiple signatures of the same appliance could help increase the accuracy of AppliSense. Most likely, a successful load disaggregation scheme will combine a number of different approaches in the end [82, 90]. Indeed, we envision the extension of AppliSense by a module that is responsible for identifying non-switchable heating and cooling appliances (e.g., fridge, freezer, etc.) and the combination with smart power outlets that measure the consumption of individual appliances that contribute significantly to the residential energy bill (e.g., washer or dryer). From a user perspective, this should not dramatically increase the complexity of the system, but help to cover large, currently unattributed parts of the electricity disaggregation.

In this thesis we showed a way of how utilizing Ubicomp technologies enables integrating the user and a portable user interface into the disaggregation process. This offers new possibilities not only during the signature acquisition process, but also during the run time of the disaggregation system. Unsupervised learning could be used to classify appliance switching events and prompt users for their confirmation once a certain probability level is reached. This could substantially enhance the disaggregation performance. We also believe that the proposed method for signature acquisition is a significant improvement compared to traditional training methods and can help in building a central appliance signature database. For example, PowerPedia that has been developed as part of this thesis could serve as such a central database and signatures could be transmitted automatically when users upload their devices. At the same time, AppliSense and PowerPedia provide a framework for an automated recommendation system that derives energy conservation measures upon the specific load profile of the household and the operated appliances.

6 Appendices

6.1 Offline User Survey

Universität St.Gallen

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Umfrage zu Energiespar-Technologien
und Energiespar-Dienstleistungen

Version A

Bits to Energy Lab
Lehrstuhl für Informationsmanagement, ETH Zürich
Institut für Technologiemanagement, Universität St. Gallen
Institut für Pervasive Computing, ETH Zürich

www.bitstoenergy.ch

Liebe Teilnehmer

Der verantwortungsvolle Umgang mit Energie wird immer wichtiger. Für viele gestaltet sich das Energiesparen jedoch schwierig, da oft nicht bekannt ist, welche Massnahmen wirkungsvoll sind und welche Einsparungen dadurch erzielt werden können. Moderne Technologien und intelligente Verbrauchsanzeigen können hierbei unterstützen.

Mit dieser Umfrage wird das Interesse an Produkten und Dienstleistungen zum Energiesparen untersucht. Es gibt keine richtigen oder falschen Antworten – Ihre Meinung zählt!

Vielen Dank für Ihre Unterstützung.

Einstellung zum Energiesparen

A1. Bitte bewerten Sie folgende Aussagen:

Energiesparen …	Stimme gar nicht zu	Stimme eher nicht zu	Stimme eher zu	Stimme voll zu
…ist nötig, um die Klimaerwärmung einzudämmen	☐	☐	☐	☐
…ist ein gutes Mittel um Geld zu sparen	☐	☐	☐	☐
…ist selbstverständlich	☐	☐	☐	☐
… ist wichtig, um Kindern ein gutes Vorbild zu sein	☐	☐	☐	☐
… kann Spaß machen	☐	☐	☐	☐
… sollte durch technische Innovation geschehen	☐	☐	☐	☐
… wird von Familie oder Freunden als sehr wichtig angesehen	☐	☐	☐	☐
… sollte primär durch die Industrie erfolgen	☐	☐	☐	☐

A2. Welche Massnahmen haben Sie in den letzten fünf Jahren vorgenommen? (Mehrfachnennungen möglich)

		Weiss nicht
Stand-by Verbrauch vermieden	☐	
Energiesparlampen gekauft	☐	
Energieeffiziente Geräte gekauft	☐	
Energieberatung in Anspruch genommen	☐	
Solaranlage installiert	☐	
Heizungsanlage optimiert	☐	
Haus / Wohnung isoliert	☐	
Ökostrom bezogen	☐	☐
Energiekostenzähler gekauft	☐	☐

A3. Welche Gründe halten Sie davon ab, zu Hause noch mehr Energie zu sparen?

	Stimme gar nicht zu	Stimme eher nicht zu	Stimme eher zu	Stimme voll zu
Ich habe nicht genügend Zeit	☐	☐	☐	☐
Ich weiss nicht, wie ich sparen kann	☐	☐	☐	☐
Ich möchte das Geld nicht ausgeben	☐	☐	☐	☐
Ist mir zu viel Aufwand	☐	☐	☐	☐
Mich interessiert das Thema nicht	☐	☐	☐	☐

Technologien zum Energiesparen

Im Folgenden werden verschiedene Produkte und Dienstleistungen beschrieben, die Sie beim Energiesparen zu Hause unterstützen können.

Visualisierung von Energieverbräuchen

B1. Der Stromverbrauch Ihres Haushalts kann mit Hilfe eines neuen „intelligenten“ Stromzählers sehr genau gemessen und gespeichert werden. Die Messwerte können auf verschiedene Art und Weise in Ihrem Wohnbereich angezeigt werden. So bekommen Sie ein Bild, wie viel Strom aktuell verbraucht wird und welche Kosten daraus resultieren.

Ich finde diese Dienstleistung sinnvoll	☐	☐	☐	☐
Ich denke, dass man mit dieser Dienstleistung Energie sparen kann	☐	☐	☐	☐
Ich würde diese Dienstleistung nutzen	☐	☐	☐	☐
Ich würde diese Dienstleistung weiterempfehlen	☐	☐	☐	☐

B2. Im Durchschnitt sparen Haushalte mit intelligenten Stromzählern 52 CHF pro Jahr.
Ich würde für diese Dienstleistung pro Jahr max. ______ CHF bezahlen.

B3. Wo würden Sie sich die Anzeige des aktuellen Stromverbrauchs wünschen? (Nur eine Nennung)

Separates Display (ähnlich Wetterstation)	Fernseher	Im Internet	Mobiltelefon
☐	☐	☐	☐

Interaktives Messen

C1. Mit Hilfe einer Mobiltelefon Anwendung können Sie den Stromverbrauch einzelner Geräte messen und unmittelbar auf dem Display Ihres Mobiltelefons ablesen.

Ich finde diese Anwendung sinnvoll	☐	☐	☐	☐
Ich denke, dass man mit dieser Anwendung Energie sparen kann	☐	☐	☐	☐
Ich würde diese Anwendung nutzen	☐	☐	☐	☐
Ich würde diese Anwendung weiterempfehlen	☐	☐	☐	☐

C2. Im Durchschnitt sparen Haushalte mit einer solchen Anwendung 85 CHF pro Jahr.
Ich würde für diese Anwendung pro Jahr max. ______ CHF bezahlen.

Unterstützung beim Energiesparen

D1. Um effektiv Energie zu sparen, ist es hilfreich zu wissen, welche Geräte besonders viel Strom verbrauchen und welche Massnahmen (z.B. Austausch von Geräten) dazu beitragen, Energie einsparen.

Welche Informationen möchten Sie zusätzlich auf Ihrem Mobiltelefon angezeigt bekommen?

Effizienzklasse (A++, A, B, C …) eines Gerätes	☐	☐	☐	☐
Gesamtverbrauch mehrerer zeitgleich genutzter Geräte	☐	☐	☐	☐
Jährlichen Kosten, die durch ein Gerät entstehen	☐	☐	☐	☐
Top fünf Geräte, die am meisten Energie verbrauchen	☐	☐	☐	☐
Gesamtverbrauch des letzten Monats	☐	☐	☐	☐
Vergleich mit einem durchschnittlichen Haushalt	☐	☐	☐	☐
Aktueller Energieverbrauch Ihrer Freunde / Bekannte	☐	☐	☐	☐
… andere: ____________________	☐	☐	☐	☐

Motivation zum Energiesparen

E1. Geräte können so gestaltet werden, dass Sie zum Energiesparen motivieren. So könnte z.B. das Display Ihrer Waschmaschine anzeigen, wie viel Strom und Wasser Sie mit dem Öko-Waschprogramm bisher schon gespart haben.

Ich finde diese Zusatzfunktion sinnvoll	☐	☐	☐	☐
Ich denke, dass man mit dieser Zusatzfunktion Energie sparen kann	☐	☐	☐	☐
Ich würde diese Zusatzfunktion nutzen	☐	☐	☐	☐
Ich würde diese Zusatzfunktion weiterempfehlen	☐	☐	☐	☐

E2. Im Durchschnitt sparen Haushalte mit solch einer Funktion 20 CHF pro Jahr.

Ich würde für eine Waschmaschine (ca. 1200 CHF) mit dieser Zusatzfunktion

max. ☐ CHF mehr bezahlen.

Unterstützung beim Energiesparen als Dienstleistung

F1. Wie wollen Sie beim Energiesparen unterstützt werden?

Durch individuelle Energiespartipps	☐	☐	☐	☐
Durch unabhängige Experten	☐	☐	☐	☐
Durch Freunde und Bekannte	☐	☐	☐	☐
Durch innovative Messgeräte und Technologien	☐	☐	☐	☐
Durch effizientere Geräte	☐	☐	☐	☐

F2. Was denken Sie, wäre die Motivation eines Energieversorgers, Ihnen die vorher beschriebenen Produkte / Dienstleistungen anzubieten?

Er will sein Image verbessern	☐	☐	☐	☐
Er will Geld verdienen	☐	☐	☐	☐
Er hat echtes Interesse, dass seine Kunden Strom sparen	☐	☐	☐	☐

F3. Angaben zur Person
(Sie helfen uns, Ihre Antworten besser zu verstehen)

Geschlecht:	weiblich	männlich
	☐	☐

Ich bin Mieter:	ja	nein
	☐	☐

Bildungsabschluss:	Grundschule	Lehre	Matura	Hochschule
	☐	☐	☐	☐

Alter:	unter 18	18-25	26-35	36-49	50-70	über 70	keine Angabe
	☐	☐	☐	☐	☐	☐	☐

Möchten Sie über die Ergebnisse der Umfrage informiert werden?

Name: ________________________ Email oder Tel. Nr.: ________________________

Kommentare:

Universität St.Gallen

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Umfrage zu Energiespar-Technologien und Energiespar-Dienstleistungen

Version B

Bits to Energy Lab
Lehrstuhl für Informationsmanagement, ETH Zürich
Institut für Technologiemanagement, Universität St. Gallen
Institut für Pervasive Computing, ETH Zürich

www.bitstoenergy.ch

Liebe Teilnehmer

Der verantwortungsvolle Umgang mit Energie wird immer wichtiger. Für viele gestaltet sich das Energiesparen jedoch schwierig, da oft nicht bekannt ist, welche Massnahmen wirkungsvoll sind und welche Einsparungen dadurch erzielt werden können. Moderne Technologien und intelligente Verbrauchsanzeigen können hierbei unterstützen.

Mit dieser Umfrage wird das Interesse an Produkten und Dienstleistungen zum Energiesparen untersucht. Es gibt keine richtigen oder falschen Antworten – Ihre Meinung zählt!

Vielen Dank für Ihre Unterstützung.

Einstellung zum Energiesparen

A1. Bitte bewerten Sie folgende Aussagen:

Energiesparen …	Stimme gar nicht zu	Stimme eher nicht zu	Stimme eher zu	Stimme voll zu
…ist nötig, um die Klimaerwärmung einzu-dämmen	☐	☐	☐	☐
…ist ein gutes Mittel um Geld zu sparen	☐	☐	☐	☐
…ist selbstverständlich	☐	☐	☐	☐
… ist wichtig, um Kindern ein gutes Vorbild zu sein	☐	☐	☐	☐
… kann Spaß machen	☐	☐	☐	☐
… sollte durch technische Innovation ge-schehen	☐	☐	☐	☐
… wird von Familie oder Freunden als sehr wichtig angesehen	☐	☐	☐	☐
… sollte primär durch die Industrie erfolgen	☐	☐	☐	☐

A2. Welche Massnahmen haben Sie in den letzten fünf Jahren vorgenommen? (Mehrfachnennungen möglich)

		Weiss nicht
Stand-by Verbrauch vermieden	☐	
Energiesparlampen gekauft	☐	
Energieeffiziente Geräte gekauft	☐	
Energieberatung in Anspruch genommen	☐	
Solaranlage installiert	☐	
Heizungsanlage optimiert	☐	
Haus / Wohnung isoliert	☐	
Ökostrom bezogen	☐	☐
Energiekostenzähler gekauft	☐	☐

A3. Welche Gründe halten Sie davon ab, zu Hause noch mehr Energie zu sparen?

	Stimme gar nicht zu	Stimme eher nicht zu	Stimme eher zu	Stimme voll zu
Ich habe nicht genügend Zeit	☐	☐	☐	☐
Ich weiss nicht, wie ich sparen kann	☐	☐	☐	☐
Ich möchte das Geld nicht ausgeben	☐	☐	☐	☐
Ist mir zu viel Aufwand	☐	☐	☐	☐
Mich interessiert das Thema nicht	☐	☐	☐	☐

Technologien zum Energiesparen

Im Folgenden werden verschiedene Produkte und Dienstleistungen beschrieben, die Sie beim Energiesparen zu Hause unterstützen können.

Visualisierung von Energieverbräuchen

B1. Der Stromverbrauch Ihres Haushalts kann mit Hilfe eines neuen „intelligenten" Stromzählers sehr genau gemessen und gespeichert werden. Die Messwerte können auf verschiedene Art und Weise in Ihrem Wohnbereich angezeigt werden. So bekommen Sie ein Bild, wie viel Strom aktuell verbraucht wird und welche Kosten daraus resultieren.

	Stimme gar nicht zu	Stimme eher nicht zu	Stimme eher zu	Stimme voll zu
Ich finde diese Dienstleistung sinnvoll	☐	☐	☐	☐
Ich denke, dass man mit dieser Dienstleistung Energie sparen kann	☐	☐	☐	☐
Ich würde diese Dienstleistung nutzen	☐	☐	☐	☐
Ich würde diese Dienstleistung weiterempfehlen	☐	☐	☐	☐

B2. Ich würde für diese Dienstleistung pro Jahr max. ☐ CHF bezahlen.

B3. Wo würden Sie sich die Anzeige des aktuellen Stromverbrauchs wünschen? (Nur eine Nennung)

Separates Display (ähnlich Wetterstation)	Fernseher	Im Internet	Mobiltelefon
☐	☐	☐	☐

Interaktives Messen

C1. Mit Hilfe einer Mobiltelefon Anwendung können Sie den Stromverbrauch einzelner Geräte messen und unmittelbar auf dem Display Ihres Mobiltelefons ablesen.

	Stimme gar nicht zu	Stimme eher nicht zu	Stimme eher zu	Stimme voll zu
Ich finde diese Anwendung sinnvoll	☐	☐	☐	☐
Ich denke, dass man mit dieser Anwendung Energie sparen kann	☐	☐	☐	☐
Ich würde diese Anwendung nutzen	☐	☐	☐	☐
Ich würde diese Anwendung weiterempfehlen	☐	☐	☐	☐

C2. Ich würde für diese Anwendung pro Jahr max. ☐ CHF bezahlen.

Unterstützung beim Energiesparen

D1. Um effektiv Energie zu sparen, ist es hilfreich zu wissen, welche Geräte besonders viel Strom verbrauchen und welche Massnahmen (z.B. Austausch von Geräten) dazu beitragen, Energie einsparen.

Welche Informationen möchten Sie zusätzlich auf Ihrem Mobiltelefon angezeigt bekommen?

	Stimme gar nicht zu	Stimme eher nicht zu	Stimme eher zu	Stimme voll zu
Effizienzklasse (A++, A, B, C …) eines Gerätes	☐	☐	☐	☐
Gesamtverbrauch mehrerer zeitgleich genutzter Geräte	☐	☐	☐	☐
Jährlichen Kosten, die durch ein Gerät entstehen	☐	☐	☐	☐
Top fünf Geräte, die am meisten Energie verbrauchen	☐	☐	☐	☐
Gesamtverbrauch des letzten Monats	☐	☐	☐	☐
Vergleich mit einem durchschnittlichen Haushalt	☐	☐	☐	☐
Aktueller Energieverbrauch Ihrer Freunde / Bekannte	☐	☐	☐	☐
… andere: ____________________	☐	☐	☐	☐

Motivation zum Energiesparen

E1. Geräte können so gestaltet werden, dass Sie zum Energiesparen motivieren. So könnte z.B. das Display Ihrer Waschmaschine anzeigen, wie viel Strom und Wasser Sie mit dem Öko-Waschprogramm bisher schon gespart haben.

	Stimme gar nicht zu	Stimme eher nicht zu	Stimme eher zu	Stimme voll zu
Ich finde diese Zusatzfunktion sinnvoll	☐	☐	☐	☐
Ich denke, dass man mit dieser Zusatzfunktion Energie sparen kann	☐	☐	☐	☐
Ich würde diese Zusatzfunktion nutzen	☐	☐	☐	☐
Ich würde diese Zusatzfunktion weiterempfehlen	☐	☐	☐	☐

E2. Ich würde für eine Waschmaschine (1200 CHF) mit dieser Zusatzfunktion max. ☐ CHF mehr bezahlen.

Unterstützung beim Energiesparen als Dienstleistung

F1. Wie wollen Sie beim Energiesparen unterstützt werden?

	Stimme gar nicht zu	Stimme eher nicht zu	Stimme eher zu	Stimme voll zu
Durch individuelle Energiespartipps	☐	☐	☐	☐
Durch unabhängige Experten	☐	☐	☐	☐
Durch Freunde und Bekannte	☐	☐	☐	☐
Durch innovative Messgeräte und Technologien	☐	☐	☐	☐
Durch effizientere Geräte	☐	☐	☐	☐

F2. Was denken Sie, wäre die Motivation eines Energieversorgers, Ihnen die vorher beschriebenen Produkte / Dienstleistungen anzubieten?

	Stimme gar nicht zu	Stimme eher nicht zu	Stimme eher zu	Stimme voll zu
Er will sein Image verbessern	☐	☐	☐	☐
Er will Geld verdienen	☐	☐	☐	☐
Er hat echtes Interesse, dass seine Kunden Strom sparen	☐	☐	☐	☐

F3. Angaben zur Person
(Sie helfen uns, Ihre Antworten besser zu verstehen)

Geschlecht:	weiblich	männlich
	☐	☐

Ich bin Mieter:	ja	nein
	☐	☐

Bildungsabschluss:	Grundschule	Lehre	Matura	Hochschule
	☐	☐	☐	☐

Alter:	unter 18	18-25	26-35	36-49	50-70	über 70	keine Angabe
	☐	☐	☐	☐	☐	☐	☐

Möchten Sie über die Ergebnisse der Umfrage informiert werden?

Name: ________________________ Email oder Tel. Nr.: __________________________

Kommentare:

6.2 Questionnaire of the User Study

Universität St.Gallen

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

User Study zu innovativen Strommesstechnologien und portablem Feedback zum Stromverbrauch

Bits to Energy Lab
Lehrstuhl für Informationsmanagement, ETH Zürich
Institut für Technologiemanagement, Universität St. Gallen
Institut für Pervasive Computing, ETH Zürich

www.bitstoenergy.ch

Liebe Teilnehmer

In dieser Studie werden drei verschiedene Technologien zur Messung des Energieverbrauchs von Haushaltsgeräten getestet.

Der Ablauf ist wie folgt:

1) Zunächst werden Sie gebeten, den Verbrauch von verschiedenen Geräten mit Hilfe von drei unterschiedlichen Technologien zu messen. Danach bitten wir Sie, einen kurzen Fragebogen zur Bewertung der drei Technologien zu beantworten.
2) Als nächstes werden Sie, anhand verschiedener Aufgaben, gebeten, einen mobilen Energiemonitor auf dem iPhone zu bewerten.
3) Zum Schluss bitten wir Sie noch um einige persönliche Angaben

Das genaue Vorgehen erklärt Ihnen der Versuchsleiter.

Teil 1: Vergleich von Technologien zur Messung des Energieverbrauchs

A0 Bitte tragen Sie die Messergebnisse in folgende Tabelle ein:

Technologie 1			
Verbrauch Gerät 1	Verbrauch Gerät 2	Standby – Verbrauch	Kombinierter Verbrauch

Technologie 2			
Verbrauch Gerät 1	Verbrauch Gerät 2	Standby – Verbrauch	Kombinierter Verbrauch

Technologie 2			
Verbrauch Gerät 1	Verbrauch Gerät 2	Standby – Verbrauch	Kombinierter Verbrauch

A1. Bitte ordnen Sie die drei Technologien gemäss Ihres Gesamteindrucks, indem Sie Ihre 1. Wahl mit 1, Ihre zweite mit 2 und Ihre dritte mit 3 bewerten.

	iPhone Applikation
	Click
	Wattson

A2. Bitte ordnen Sie die drei Technologien nach ihrer Verständlichkeit (= Sie wussten zu jedem Zeitpunkt, was zu tun ist), indem Sie Ihre 1. Wahl mit 1, Ihre zweite mit 2 und Ihre dritte mit 3 bewerten.

	Click
	Wattson
	iPhone Applikation

A3. Bitte ordnen Sie die drei Technologien nach dem Komfort der Messung (= Der Aufwand für die Messung war gering), indem Sie Ihre 1. Wahl mit 1, Ihre zweite mit 2 und Ihre dritte mit 3 bewerten.

	Wattson
	iPhone Applikation
	Click

A4. Bitte ordnen Sie die drei Technologien nach ihrer optischen Attraktivität, indem Sie Ihre 1. Wahl mit 1, Ihre zweite mit 2 und Ihre dritte mit 3 bewerten.

	iPhone Applikation
	Wattson
	Click

A5. Bitte ordnen Sie die drei Technologien nach dem Spass bei der Nutzung, indem Sie Ihre 1. Wahl mit 1, Ihre zweite mit 2 und Ihre dritte mit 3 bewerten.

	Click
	iPhone Applikation
	Wattson

A6. Haben Sie schon einmal einen Energiekostenzähler benutzt?

☐ Ja, welchen ____________________

☐ Nein

A7 Stellen Sie sich vor, Sie messen nach dem Kauf eines Messgeräts die wichtigsten Verbraucher ihres Haushalts, um "Stromfresser" zu finden und verwenden das Messgerät danach regelmässig (z.B. monatlich, um neu gekaufte Geräte zu inventarisieren).

Bitte bewerten sie die folgenden Kriterien nach Ihrer persönlich wahrgenommenen Wichtigkeit!

	gar nicht wichtig				sehr wichtig
Design	☐	☐	☐	☐	☐
Genauigkeit	☐	☐	☐	☐	☐
Verständlichkeit	☐	☐	☐	☐	☐
Verfügbarkeit	☐	☐	☐	☐	☐
Komfort der Messung	☐	☐	☐	☐	☐
Spass bei der Nutzung	☐	☐	☐	☐	☐
Erlernbarkeit	☐	☐	☐	☐	☐
Preis	☐	☐	☐	☐	☐
Zeit, bis Feedback angezeigt wird	☐	☐	☐	☐	☐

A8 Bitte ordnen Sie die oben genannten Kriterien noch nach ihrer Wichtigkeit (d.h. das wichtigste Kriterium bekommt eine 1, das zweite eine 2 etc.)

☐	Design
☐	Genauigkeit
☐	Verständlichkeit
☐	Verfügbarkeit
☐	Komfort der Messung
☐	Spass bei der Nutzung
☐	Erlernbarkeit
☐	Preis
☐	Zeit, bis Feedback angezeigt wird

A9. Welche der drei Messtechnologien würden Sie kaufen? (nur eine Auswahlmöglichkeit)

☐ Click
☐ Wattson
☐ iPhone Applikation
☐ gar keine!

Teil 2:
Leitfaden zur iPhone Applikation

Nun möchten wir Sie bitten, eine Reihe von Aufgaben mit Hilfe der iPhone Applikation zu lösen um folgende Fragen zu beantworten.
Die Applikation besteht aus 4 Hauptansichten. Die erste zeigt eine Übersicht über den aktuellen Verbrauch in Echtzeit, die zweite Ansicht stellt Historische Verbräuche dar und vergleicht diese mit den verursachten Kosten. Die dritte Ansicht gibt einen Überblick über die bereits gespeicherten Messerwert und die vierte Ansicht ermöglicht das Messen und abspeichern einzelner Geräte.

1) **Bestimmen Sie den aktuellen Gesamtverbrauch.**

   __________________________________ Watt

2) **Wofür stehen Ihrer Meinung nach die Farben des Tachos?**

   ____

   ____

   ____

   ____

3) **Wann war der Verbrauch innerhalb der letzten 5 Wochen am größten?**

   ____

4) **Welche Kosten waren damit verbunden?**

   ____

5) **Vergleichen Sie den Energieverbrauch der letzten Woche mit dem eines typischen Haushaltes ihrer Größe!**

   ____

   ____

6) **Messen Sie einen beliebigen Verbraucher und speichern Sie ihn als neues Gerät mit Bild, Namen und geschätzter Nutzungsdauer ab. Schauen Sie in der Geräteübersicht nach ob Sie das Gerät abgespeicherte Gerät wieder finden.**

Fragen zur Mobilfunk-Applikation

B1. Komplexität
Bitte kreuzen Sie an.

	stimme gar nicht zu				stimme voll zu
Die Applikation ist mir insgesamt technisch zu anspruchsvoll.	☐	☐	☐	☐	☐
Die Zeit, die bis zur Anzeige des Messwertes verstreicht, ist zu lang	☐	☐	☐	☐	☐
Die Anzeige des Gesamtverbrauch-tachos ist verständlich.	☐	☐	☐	☐	☐
Es viel mir leicht die historischen Verbräuche den Kosten zuzuordnen.	☐	☐	☐	☐	☐
Die Messfunktion ist einfach und intuitiv bedienbar.	☐	☐	☐	☐	☐
Das abspeichern eines Geräts fällt mir leicht.	☐	☐	☐	☐	☐

B2. Funktionen
Bitte bewerten Sie, wie wichtig für Sie die folgenden Funktionen sind.

	gar nicht wichtig				sehr wichtig
Anzeige des Gesamtverbrauchs	☐	☐	☐	☐	☐
Anzeige des Standby-Verbrauchs des Haushalts	☐	☐	☐	☐	☐
Vergleich des aktuellen Verbrauchs mit dem historischen Verbrauch	☐	☐	☐	☐	☐
Übersicht über den Verbrauch der letzten Monate	☐	☐	☐	☐	☐
Kosten der letzten Monate	☐	☐	☐	☐	☐
Stromverbrauch einzelner Geräte	☐	☐	☐	☐	☐
Kosten einzelner Geräte	☐	☐	☐	☐	☐
Hochrechnung auf jährliche Kosten einzelner Geräte	☐	☐	☐	☐	☐
Effizienzklasse einzelner Geräte	☐	☐	☐	☐	☐
Übersicht über die größten Stromfresser	☐	☐	☐	☐	☐
Vergleich mit anderen(z.B. Freunden)	☐	☐	☐	☐	☐
Möglichkeit, anderen zu zeigen, welche Geräte ich verwende	☐	☐	☐	☐	☐
Möglichkeit ein Sparziel zu setzen	☐	☐	☐	☐	☐

B3. **Angenommen, Sie dürften sich nur DREI Funktionen auswählen, welche wären dies?**
Bitte kreuzen Sie an.

☐	Anzeige des Gesamtverbrauchs
☐	Anzeige des Standby-Verbrauchs des Haushalts
☐	Vergleich des aktuellen Verbrauchs (sekundengenau) mit dem historischen Verbrauchs
☐	Übersicht über den Verbrauch der letzten Monate
☐	Kosten der letzten Monate
☐	Stromverbrauch einzelner Geräte
☐	Kosten einzelner Geräte
☐	Hochrechnung auf jährliche Kosten einzelner Geräte
☐	Effizienz einzelner Geräte
☐	Übersicht über die größten Stromfresser
☐	Vergleich mit anderen (z.B. Freunden)
☐	Möglichkeit, anderen zu zeigen, welche Geräte ich verwende
☐	Möglichkeit ein Sparziel zu setzen

B4. **Welche Funktionen wünschen Sie sich noch zusätzlich?**

____

B5. **Einstellung zur Applikation**
Bitte bewerten Sie, wie wichtig für Sie die folgenden Funktionen sind.

	Stimme gar nicht zu				Stimme voll zu
Die Applikation wird mir dabei helfen, Energie zu sparen.	☐	☐	☐	☐	☐
Die Applikation wird mir dabei helfen, Geld zu sparen.	☐	☐	☐	☐	☐
Die Applikation erhöht mein Wissen, über den Verbrauch einzelner Geräte	☐	☐	☐	☐	☐
Es macht Spaß, die Applikation zu nutzen.	☐	☐	☐	☐	☐
Mit der Applikation kann ich anderen zei-gen, dass ich Gutes tue.	☐	☐	☐	☐	☐
Die Applikation ist nutzerfreundlich.	☐	☐	☐	☐	☐
Die Bedienung ist einfach zu lernen.	☐	☐	☐	☐	☐
Ich bin insgesamt sehr zufrieden mit der Applikation.	☐	☐	☐	☐	☐
Ich würde die Applikation gerne meinen Freunden zeigen.	☐	☐	☐	☐	☐

B6. Nutzungsabsicht
Bitte kreuzen Sie an.

Stimme gar nicht zu — Stimme voll zu

Ich würde die Applikation mindestens einmal im Monat nutzen	☐	☐	☐	☐	☐
Ich würde die Applikation mindestens einmal in der Woche nutzen	☐	☐	☐	☐	☐
Ich würde die Applikation gar nicht nutzen	☐	☐	☐	☐	☐
Ich würde die Applikation eher als ein Webportal mit gleichen Funktionen nutzen	☐	☐	☐	☐	☐

B7. Stellen Sie sich vor, Ihr Energieversorger würde eine solche iPhone Applikation kostenfrei für ihre Kunden anbieten. Wie beurteilen Sie dazu die nachfolgenden Aussagen?

Stimme gar nicht zu — Stimme voll zu

Ich würde meinen Freunden und Kollegen von diesem Energieversorger erzählen.	☐	☐	☐	☐	☐
Ich würde bei dem Energieversorger bleiben wollen, der mir derartige mobile Services anbietet, selbst wenn er geringfügig teurer ist als die Konkurrenz.	☐	☐	☐	☐	☐
Ich hätte Bedenken, dass diese Anwendung meine Stromkosten unnötig erhöhen würde.	☐	☐	☐	☐	☐
Ich würde die Applikation gerne im Rahmen eines Servicepakets nutzen und wäre bereit einen kleinen monatlichen Aufpreis zu zahlen.	☐	☐	☐	☐	☐

B8. Wie viel würden Sie für die iPhone Applikation bezahlen?

☐ gar nichts
☐ bis 5 CHF
☐ bis 10 CHF
☐ bis 20 CHF
☐ über 20 CHF

Teil 3:
Angaben zur Person
(Sie helfen uns, Ihre Antworten besser zu verstehen)

Geschlecht:	weiblich	männlich
	☐	☐

Alter:	unter 18	18-25	26-35	36-49	50-70	über 70	keine Angabe
	☐	☐	☐	☐	☐	☐	☐

C1. Was ist Ihre Position?

☐ selbstständig
☐ Angestellte/r
☐ leitende/r Angestellte/r
☐ Rentner/in
☐ Student/in

C2. Besitzen Sie ein iPhone?

☐ Ja, seit __________ Monaten
☐ Nein

C3. Besitzen Sie ein anderes Smart Phone? (HTC, Android, etc.)

☐ Ja, seit ___________ Monaten
☐ Nein

C4. Wie oft nutzen Sie Ihr Mobiltelefon?

☐ Mehrmals pro Stunde
☐ 1 x pro Stunde
☐ Mehrmals am Tag
☐ 1x pro Tag
☐ Mehrmals pro Woche
☐ 1 x pro Woche
☐ seltener

C5. Wie hoch sind im Schnitt Ihre monatlichen Handygebühren?

☐ Weniger als 20 CHF
☐ 20 bis 50 CHF
☐ 50 bis 100 CHF
☐ Mehr als 100 CHF

C6. Wie gross ist Ihr Interesse an Ihrem Energieverbrauch auf einer Skala von 1 (wenig) - 10 (viel)?

____

C7. Technologieeinstellung
Bitte kreuzen Sie an!

	stimme gar nicht zu				stimme voll zu
Freunde und Kollegen fragen mich oft nach meiner Meinung zu neuen Telekommunikationstechnologien.	☐	☐	☐	☐	☐
Meine Freunde und Kollegen wissen meist besser über neue Technologien Bescheid als ich.	☐	☐	☐	☐	☐
Ich bin bei technologischen Neuentwicklungen in meinem Interessengebiet immer auf dem Laufenden.	☐	☐	☐	☐	☐
Es macht mir Spaß, neue Technologien auszuprobieren.	☐	☐	☐	☐	☐

C8. Wie oft haben Sie die folgenden Aktionen IM LETZTEN JAHR durchgeführt?
Bitte kreuzen Sie an.

Im letzten Jahr habe ich…	nie				sehr oft
…überlegt, wie man Dinge wiederverwenden kann	☐	☐	☐	☐	☐
…Zeitungen recycled	☐	☐	☐	☐	☐
…Dosen und Flaschen recycled	☐	☐	☐	☐	☐
…Freunde und Familie ermutigt, zu recyceln	☐	☐	☐	☐	☐
…Produkte mit recycelbaren Verpackungen gekauft	☐	☐	☐	☐	☐
…fremden Müll aufgehoben	☐	☐	☐	☐	☐
…Essensabfälle kompostiert	☐	☐	☐	☐	☐
…Benzin gespart, indem ich gelaufen oder mit dem Fahrrad gefahren bin	☐	☐	☐	☐	☐

7 Bibliography

[1] W. Roediger, "Personalized Energy Measurements," Bachelor Thesis, Faculty for Computer Science (TU Munich), Institute for Pervasive Computing (ETH Zurich), TU Munich, Munich, Germany, 2009.

[2] X. Jiang, S. Dawson-Haggerty, J. Taneja, P. Dutta, and D. Culler, "Creating Greener Homes with Ip-Based Wireless Ac Energy Monitors," in Proceedings of the 6th ACM Conference on Embedded Network Sensor Systems, Raleigh, NC, USA, 2008, pp. 355-356.

[3] A. Helfenstein, "Smart Energy Detection - an Algorithm for Device Recognition Based on the Electrical Load Curve," Master Thesis, Institute for Pervasive Computing (Distributed Systems Group), ETH Zurich, Zurich, Switzerland, 2010.

[4] U.S. Energy Information Administration (EIA). (2011). Residential Energy Consumption Survey (Recs) 2009. Available: http://www.eia.gov/consumption/residential/reports/electronics.cfm

[5] International Energy Agency (IEA). (2008). Worldwide Trends in Energy Use and Efficiency - Key Insights from Iea Indicator Analysis. Available: http://www.iea.org/publications/free_new_Desc.asp?PUBS_ID=2026

[6] J. Froehlich, T. Dillahunt, P. Klasnja, J. Mankoff, S. Consolvo, B. Harrison, and J. A. Landay, "Ubigreen: Investigating a Mobile Tool for Tracking and Supporting Green Transportation Habits," in Proceedings of the 27th International Conference on Human Factors in Computing Systems, Boston, MA, USA, 2009, pp. 1043-1052.

[7] C. A. Björkskog, G. Jacucci, L. Gamberini, T. Nieminen, T. Mikkola, C. Torstensson, and M. Bertoncini, "Energylife: Pervasive Energy Awareness for Households," in Adjunct Proceedings of the 12th ACM International Conference on Ubiquitous Computing, Copenhagen, Denmark, 2010, pp. 361-362.

[8] International Energy Agency (IEA). (2010). Key World Energy Statistics 2010. Available: http://www.iea.org/textbase/nppdf/free/2010/key_stats_2010.pdf

[9] International Energy Agency (IEA). (2010). World Energy Outlook 2010. Available: http://www.iea.org/w/bookshop/add.aspx?id=422

[10] U.S. Environmental Protection Agency, 2011 U.S. Greenhouse Gas Inventory Report - Inventory of U.S. Greenhouse Gas Emissions and Sinks: 1990 - 2009, 2011.

[11] U.S. Energy Information Administration (EIA). (2010). Annual Energy Review. Available: http://www.eia.gov/totalenergy/data/annual/

[12] European Environment Agency (EEA). (2009). Available: http://www.eea.europa.eu/data-and-maps/figures/final-electricity-consumption-bysector-eu-27-1

[13] International Energy Agency (IEA). (2009). Gadgets and Gigawatts: Policies for Energy Efficient Electronics. Available: http://www.iea.org/w/bookshop/add.aspx?id=361

[14] Office of Energy Canada, Energy Efficiency Trends in Canada, 1990 to 2008, 2010.

[15] S. Firth, K. Lomas, A. Wright, and R. Wall, "Identifying Trends in the Use of Domestic Appliances from Household Electricity Consumption Measurements," Energy and Buildings, vol. 40, pp. 926-936, 2008.

[16] Energy Policy Branch and Energy Forecasting Division, Canada’s Energy Outlook, 1996-2020: Natural Resources Canada, 1997.

[17] A. Kirchner, A. Kemmler, P. Hofer, M. Keller, M. Jakob, and G. Catenazzi, Analysis of Energy Consumption in Switzerland 2000–2009. Bern, Switzerland: Swiss Federal Office of Energy, 2010.

[18] S. Z. Attari, M. L. DeKay, C. I. Davidson, and W. B. de Bruin, "Public Perceptions of Energy Consumption and Savings," Proceedings of the National Academy of Sciences, vol. 107, pp. 16054-16059, 2010.

[19] J. Pierce, D. J. Schiano, and E. Paulos, "Home, Habits, and Energy: Examining Domestic Interactions and Energy Consumption," in Proceedings of the 28th International Conference on Human Factors in Computing Systems Atlanta, GA, USA, 2010, pp. 1985-1994.

[20] F. Mattern, "The Vision and Technical Foundations of Ubiquitous Computing," Upgrade European Online Magazine, vol. 2, pp. 2-6, 2001.

[21] M. Weiser, "The Computer for the 21st Century," Scientific American, vol. 265, pp. 94-104, 1991.

[22] F. Mattern and C. Floerkemeier, "From the Internet of Computers to the Internet of Things," in From Active Data Management to Event-Based Systems and More. vol. 6462, K. Sachs, I. Petrov, and P. Guerrero, Eds. Springer Berlin / Heidelberg, 2010, pp. 242-259.

[23] E. Fleisch, "What Is the Internet of Things?," When Things Add Value. Auto-ID Labs White Paper WP-BIZAPP-053, Auto-ID Lab St. Gallen, Switzerland, 2010.

[24] R. Adelmann, M. Langheinrich, and C. Floerkemeier, "A Toolkit for Bar-Code-Recognition and-Resolving on Camera Phones - Jump Starting the Internet of Things," in Workshop on Mobile and Embedded Interactive Systems (MEIS'06) at Informatik 2006. GI Lecture Notes in Informatics Series (LNI), Dresden, Germany, 2006.

[25] A. Ilic, J. Ng, P. Bowman, and T. Staake, "The Value of RFID for Rti Management," Electronic Markets, vol. 19, pp. 125-135, 2009.

[26] B. Nath, F. Reynolds, and R. Want, "RFID Technology and Applications," Pervasive Computing, IEEE, vol. 5, pp. 22-24, 2006.

[27] A. Ilic, T. Staake, and E. Fleisch, "Using Sensor Information to Reduce the Carbon Footprint of Perishable Goods," Pervasive Computing, IEEE, vol. 8, pp. 22-29, 2009.

[28] A. Dada, T. Staake, and E. Fleisch, "Reducing Environmental Impact in Procurement by Integrating Material Parameters in Information Systems: The Example of Apple Sourcing," in AMCIS 2010 Proceedings, Lima, Peru, 2010, p. 10.

[29] M. Lehtonen, T. Staake, and F. Michahelles, "From Identification to Authentication - a Review of RFID Product Authentication Techniques," Networked RFID Systems and Lightweight Cryptography, pp. 169-187, 2008.

[30] M. O. Lehtonen, F. Michahelles, and E. Fleisch, "Trust and Security in RFID-Based Product Authentication Systems," Systems Journal, IEEE, vol. 1, pp. 129-144, 2007.

[31] D. Guinard, V. Trifa, F. Mattern, and E. Wilde, "From the Internet of Things to the Web of Things: Resource-Oriented Architecture and Best Practices," in Architecting the Internet of Things, D. Uckelmann, M. Harrison, and F. Michahelles, Eds. Springer Berlin Heidelberg, 2011, pp. 97-129.

[32] S. Helal, W. Mann, H. El-Zabadani, J. King, Y. Kaddoura, and E. Jansen, "The Gator Tech Smart House: A Programmable Pervasive Space," Computer, vol. 38, pp. 50-60, 2005.

[33] J. P. Vasseur and A. Dunkels, Interconnecting Smart Objects with Ip: The Next Internet, 1 ed.: Morgan Kaufmann Publishers Inc., June 2010.

[34] M. Kovatsch, M. Weiss, and D. Guinard, "Embedding Internet Technology for Home Automation," in Proceedings of the 15th IEEE International Conference on Emerging Techonologies and Factory Automation (ETFA 2010), Bilbao, Spain, 2010, pp. 1-8.

[35] C. Floerkemeier and F. Mattern, "Smart Playing Cards: Enhancing the Gaming Experience with RFID," in Third International Workshop on Pervasive Gaming Applications - PerGames 2006 at PERVASIVE 2006, Dublin, Ireland, 2006, pp. 27-36.

[36] S. Hinske, M. Lampe, N. Yuill, S. Price, and M. Langheinrich, "Kingdom of the Knights: Evaluation of a Seamlessly Augmented Toy Environment for Playful Learning," in Proceedings of the 8th International Conference on Interaction Design and Children, Como, Italy, 2009, pp. 202-205.

[37] L.-d. Chen and R. Nath, "A Framework for Mobile Business Applications," International Journal of Mobile Communications, vol. 2, pp. 368-381, 2004.

[38] F. Mattern, T. Staake, and M. Weiss, "ICT for Green: How Computers Can Help Us to Conserve Energy," in Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking, Passau, Germany, 2010, pp. 1-10.

[39] A. Schmidt, A. Friday, H. W. Gellersen, and F. Mattern, "Ubiquitous Computing for Sustainable Energy (Ucse2010)," in Proceedings of the 12th ACM International Conference Adjunct Papers on Ubiquitous Computing, Copenhagen, Denmark, 2010, pp. 495-496.

[40] A. Krioukov, S. Dawson-Haggerty, L. Lee, O. Rehmane, and D. Culler, "A Living Laboratory Study in Personalized Automated Lighting Controls," in Buildsys 2011, Seattle, WA, USA, 2011.

[41] 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 (SenSys 2010), Zurich, Switzerland, 2010, pp. 211-224.

[42] W. Kleiminger, S. Santini, and M. Weiss, "Opportunistic Sensing for Smart Heating Control in Private Households," in Proceedings of the 2nd International Workshop on Networks of Cooperating Objects, Chicago, IL, USA, 2011.

[43] Y. Agarwal, B. Balaji, S. Dutta, R. K. Gupta, and T. Weng, "Duty-Cycling Buildings Aggressively: The Next Frontier in HVAC Control," in Proceedings of the 10th International Conference on Information Processing in Sensor Networks (IPSN), 2011, pp. 246-257.

[44] F. C. Schweppe, R. D. Tabors, J. L. Kirtley, H. R. Outhred, F. H. Pickel, and A. J. Cox, "Homeostatic Utility Control," IEEE Transactions on Power Apparatus and Systems, vol. PAS-99, pp. 1151-1163, 1980.

[45] J. W. Black and R. C. Larson, "Strategies to Overcome Network Congestion in Infrastructure Systems," Journal of Industrial and Systems Engineering, vol. 1, pp. 97-115, 2007.

[46] D. Livengood and R. Larson, "The Energy Box: Locally Automated Optimal Control of Residential Electricity Usage," Service Science, vol. 1, pp. 1-16, 2009.

[47] M. Ilic, J. W. Black, and J. L. Watz, "Potential Benefits of Implementing Load Control," in IEEE Power Engineering Society Winter Meeting, 2002, pp. 177-182.

[48] I. Lampropoulos, G. M. A. Vanalme, and W. L. Kling, "A Methodology for Modeling the Behavior of Electricity Prosumers within the Smart Grid," in IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT Europe), 2010, pp. 1-8.

[49] R. A. Verzijlbergh, Z. Lukszo, J. G. Slootweg, and M. D. Ilic, "The Impact of Controlled Electric Vehicle Charging on Residential Low Voltage Networks," in 2011 IEEE International Conference on Networking, Sensing and Control (ICNSC), 2011, pp. 14-19.

[50] M. D. Galus and G. Andersson, "Demand Management of Grid Connected Plug-in Hybrid Electric Vehicles (Phev)," in Proceedings of the IEEE Energy 2030 Conference, 2008, pp. 1-8.

[51] L. AlAbdulkarim and Z. Lukszo, "Impact of Privacy Concerns on Consumers' Acceptance of Smart Metering in the Netherlands," in 2011 IEEE International Conference on Networking, Sensing and Control (ICNSC), 2011, pp. 287-292.

[52] M. Lisovich and S. Wicker, "Privacy Concerns in Upcoming Residential and Commercial Demand-Response Systems," IEEE Proceedings on Power Systems, vol. 1, 2008.

[53] C. Fischer, "Feedback on Household Electricity Consumption: A Tool for Saving Energy?," Energy Efficiency, vol. 1, pp. 79-104, 2008.

[54] M. Chetty, D. Tran, and R. E. Grinter, "Getting to Green: Understanding Resource Consumption in the Home," in Proceedings of the 10th International Conference on Ubiquitous Computing, Seoul, Korea, 2008, pp. 242-251.

[55] J. Froehlich, L. Findlater, and J. Landay, "The Design of Eco-Feedback Technology," in Proceedings of the 28th International Conference on Human Factors in Computing Systems, Atlanta, Georgia, USA, 2010, pp. 1999-2008.

[56] W. Abrahamse, L. Steg, C. Vlek, and T. Rothengatter, "A Review of Intervention Studies Aimed at Household Energy Conservation," Journal of Environmental Psychology, vol. 25, pp. 273-291, 2005.

[57] D. Parker, D. Hoak, and J. Cummings, "Pilot Evaluation of Energy Savings from Residential Energy Demand Feedback Devices," FSEC, Rpt: FSEC-CR-1742-08, 2008.

[58] J. Thogersen and F. Ölander, "Spillover of Environment-Friendly Consumer Behaviour," Journal of Environmental Psychology, vol. 23, pp. 225-236, 2003.

[59] H. A. He, S. Greenberg, and E. M. Huang, "One Size Does Not Fit All: Applying the Transtheoretical Model to Energy Feedback Technology Design," in Proceedings of the 28th International Conference on Human Factors in Computing Systems, Atlanta, Georgia, USA, 2010, pp. 927-936.

[60] R. T. Fielding and R. N. Taylor, "Principled Design of the Modern Web Architecture," ACM Trans. Internet Technol., vol. 2, pp. 115-150, 2002.

[61] M. Weiss and W. Kleiminger, "Smart Residential Energy Systems – How Pervasive Computing Can Be Used to Conserve Energy," in Smart Energy Strategies Conference 2011, Zurich, Switzerland, 2011.

[62] M. Weiss, F. Mattern, T. Graml, T. Staake, and E. Fleisch, "Handy Feedback: Connecting Smart Meters with Mobile Phones," in Proceedings of the 8th International Conference on Mobile and Ubiquitous Multimedia (MUM09), Cambridge, United Kingdom, 2009, pp. 1-4.

[63] M. Weiss, T. Staake, F. Mattern, and E. Fleisch, "PowerPedia: Changing Energy Usage with the Help of a Community-Based Smartphone Application," Personal and Ubiquitous Computing, Springer London, pp. 1-10, 2011.

[64] D. Abdurachmanov, "Embedded Gateway: Implementation of eMeter Restful Gateway for Gumstix Overo Air/Fire Module," Lab Report, Distributed Systems Group, Institute for Pervasive Computing, ETH Zurich, Zurich 2011.

[65] S. C. Hayes and J. D. Cone, "Reducing Residential Electrical Energy Use: Payments, Information, and Feedback," Journal of Applied Behavior Analysis, vol. 10:3, pp. 425-435, 1977.

[66] R. Kohlenberg, T. Phillips, and W. Proctor, "A Behavioral Analysis of Peaking in Residential Electrical-Energy Consumers," Journal of Applied Behavior Analysis, vol. 9, pp. 13-18, 1976.

[67] L. McClelland and S. W. Cook, "Energy Conservation Effects of Continuous in-Home Feedback in All-Electric Homes," Journal of Environmental Systems, vol. 9, pp. 169-173, 1979.

[68] R. J. Sexton, N. B. Johnson, and A. Konakayama, "Consumer Response to Continuous-Display Electricity-Use Monitors in a Time-of-Use Pricing Experiment," Journal of Consumer Research, pp. 55-62, 1987.

[69] S. Darby, "The Effectiveness of Feedback on Energy Consumption," A Review for DEFRA of the Literature on Metering, Billing and Direct Displays, 2006.

[70] A. Faruqui, S. Sergici, and A. Sharif, "The Impact of Informational Feedback on Energy Consumption: A Survey of the Experimental Evidence," Energy, vol. 35, pp. 1598-1608, 2010.

[71] C. Seligman and J. M. Darley, "Feedback as a Means of Decreasing Residential Energy Consumption," Journal of Applied Psychology, vol. 62, pp. 363-368, 1976.

[72] K. Ehrhardt-Martinez, K. A. Donnelly, and S. Laitner, "Advanced Metering Initiatives and Residential Feedback Programs: A Meta-Review for Household Electricity-Saving Opportunities," American Council for an Energy-Efficient Economy 2010.

[73] I. Ayres, S. Raseman, and A. Shih, "Evidence from Two Large Field Experiments That Peer Comparison Feedback Can Reduce Residential Energy Usage," in 5th Annual Conference on Empirical Legal Studies Paper, 2009.

[74] J. Schleich, M. Klobasa, M. Brunner, S. Gölz, and K. Götz, "Smart Metering in Germany and Austria: Results of Providing Feedback Information in a Field Trial," Working Papers, Fraunhofer Institute for Systems and Innovation Research (ISI), 2011.

[75] E. S. Geller, R. A. Winett, P. B. Everett, and R. C. Winkler, Preserving the Environment: New Strategies for Behavior Change: Pergamon Press Elmsford, NY, 1982.

[76] J. Froehlich, K. Everitt, J. Fogarty, S. Patel, and J. Landay, "Sensing Opportunities for Personalized Feedback Technology to Reduce Consumption," in In the CHI 2009 Workshop on Defining the Role of HCI in the Challenge of Sustainability, Boston, MA, USA, 2009.

[77] A. S. Kulkarni, K. C. Welch, and C. K. Harnett, "A Review of Electricity Monitoring and Feedback Systems," in Proceedings of 2011 IEEE Southeastcon, 2011, pp. 321-326.

[78] D. Petersen, J. Steele, and J. Wilkerson, "Wattbot: A Residential Electricity Monitoring and Feedback System," in Proceedings of the 27th International Conference Extended Abstracts on Human Factors in Computing Systems, Boston, MA, USA, 2009, pp. 2847-2852.

[79] T.-J. Yun, "Investigating the Impact of a Minimalist in-Home Energy Consumption Display," in Proceedings of the 27th International Conference Extended Abstracts on Human Factors in Computing Systems, Boston, MA, USA, 2009, pp. 4417-4422.

[80] J. E. Petersen, V. Shunturov, K. Janda, G. Platt, and K. Weinberger, "Dormitory Residents Reduce Electricity Consumption When Exposed to Real-Time Visual Feedback and Incentives," International Journal of Sustainability in Higher Education, vol. 8, pp. 16-33, 2007.

[81] L. van Rensburg, "Energy Demand Research Project: Review of Progress for Period 04/2008 - 08/2008," TR Ofgem 2009.

[82] H. S. Matthews, L. Soibelman, M. Berges, and E. Goldman, "Automatically Disaggregating the Total Electrical Load in Residential Buildings: A Profile of the Required Solution," in Intelligent Computing in Engineering, Plymouth, UK, 2008, pp. 381-389.

[83] M. Baranski and J. Voss, "Detecting Patterns of Appliances from Total Load Data Using a Dynamic Programming Approach," in Fourth IEEE International Conference on Data Mining (ICDM'04), Brighton, United Kingdom 2004, pp. 327-330.

[84] M. Baranski and J. Voss, "Genetic Algorithm for Pattern Detection in Nialm Systems," in 2004 IEEE International Conference on Systems, Man and Cybernetics, 2004, pp. 3462-3468.

[85] L. Farinaccio and R. Zmeureanu, "Using a Pattern Recognition Approach to Disaggregate the Total Electricity Consumption in a House into the Major End-Uses," Energy and Buildings, vol. 30, pp. 245-259, 1999.

[86] A. Prudenzi, "A Neuron Nets Based Procedure for Identifying Domestic Appliances Pattern-of-Use from Energy Recordings at Meter Panel," in Power Engineering Society Winter Meeting, 2002. IEEE, 2002, pp. 941-946.

[87] H. Lam, G. Fung, and W. Lee, "A Novel Method to Construct Taxonomy Electrical Appliances Based on Load Signatures," IEEE Transactions on Consumer Electronics, vol. 53, pp. 653-660, 2007.

[88] S. N. Patel, T. Robertson, J. A. Kientz, M. S. Reynolds, and G. D. Abowd, "At the Flick of a Switch: Detecting and Classifying Unique Electrical Events on the Residential Power Line," in Proceedings of the 9th International Conference on Ubiquitous Computing, Innsbruck, Austria, 2007, pp. 271-288.

[89] S. Gupta, M. S. Reynolds, and S. N. Patel, "Electrisense: Single-Point Sensing Using Emi for Electrical Event Detection and Classification in the Home," in Proceedings of the 12th ACM International Conference on Ubiquitous Computing, Copenhagen, Denmark, 2010, pp. 139-148.

[90] M. Zeifman and K. Roth, "Nonintrusive Appliance Load Monitoring: Review and Outlook," IEEE Transactions on Consumer Electronics, vol. 57, pp. 76-84, 2011.

[91] T. Ueno, R. Inada, O. Saeki, and K. Tsuji, "Effectiveness of an Energy-Consumption Information System for Residential Buildings," Applied Energy, vol. 83, pp. 868-883, 2006.

[92] J. Lifton, M. Feldmeier, Y. Ono, C. Lewis, and J. A. Paradiso, "A Platform for Ubiquitous Sensor Deployment in Occupational and Domestic Environments," in 6th International Symposium on Information Processing in Sensor Networks (IPSN 2007), Cambridge, USA, 2007, pp. 119-127.

[93] J. Paradiso, "Some Novel Applications for Wireless Inertial Sensors," in NSTI Nanotech 2006, 2006, pp. 431-434.

[94] X. Jiang, S. Dawson-Haggerty, P. Dutta, and D. Culler, "Design and Implementation of a High-Fidelity Ac Metering Network," in Proceedings of the 2009 International Conference on Information Processing in Sensor Networks, 2009, pp. 253-264.

[95] P. Dutta, J. Taneja, J. Jeong, X. Jiang, and D. Culler, "A Building Block Approach to Sensornet Systems," in Proceedings of the 6th ACM Conference on Embedded Network Sensor Systems, Raleigh, NC, USA, 2008, pp. 267-280.

[96] T. Schmid, D. Culler, and P. Dutta, "Meter Any Wire, Anywhere by Virtualizing the Voltage Channel," in Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings, Zurich, Switzerland, 2010, pp. 25-30.

[97] J. Lifton, M. Feldmeier, Y. Ono, C. Lewis, and J. A. Paradiso, "A Platform for Ubiquitous Sensor Deployment in Occupational and Domestic Environments," in Proceedings of the 6th International Conference on Information Processing in Sensor Networks, Cambridge, MA, USA, 2007, pp. 119-127.

[98] M. Weiss and D. Guinard, "Increasing Energy Awareness through Web-Enabled Power Outlets," in Proceedings of the 9th International Conference on Mobile and Ubiquitous Multimedia (MUM10), Limassol, Cyprus, 2010, pp. 1-10.

[99] M. Jahn, M. Jentsch, C. R. Prause, F. Pramudianto, A. Al-Akkad, and R. Reiners, "The Energy Aware Smart Home," in 5th International Conference on Future Information Technology (FutureTech), 2010, pp. 1-8.

[100] A. Reinhardt, D. Burkhardt, P. S. Mogre, M. Zaheer, and R. Steinmetz, "Smartmeter. Kom: A Low-Cost Wireless Sensor for Distributed Power Metering," in 6th IEEE International Workshop on Practical Issues in Building Sensor Network Applications, Bonn, Germany, 2011.

[101] Y. Kim, T. Schmid, Z. M. Charbiwala, and M. B. Srivastava, "Viridiscope: Design and Implementation of a Fine Grained Power Monitoring System for Homes," in Proceedings of the 11th International Conference on Ubiquitous Computing, Orlando, Florida, USA, 2009, pp. 245-254.

[102] V. Sundramoorthy, G. Cooper, N. Linge, and L. Qi, "Domesticating Energy-Monitoring Systems: Challenges and Design Concerns," Pervasive Computing, IEEE, vol. 10, pp. 20-27, 2011.

[103] A. Spagnolli, N. Corradi, L. Gamberini, E. Hoggan, G. Jacucci, C. Katzeff, L. Broms, and L. Jonsson, "Eco-Feedback on the Go: Motivating Energy Awareness," Computer, vol. 44, pp. 38-45, 2011.

[104] J. Lei, L. Jiaming, L. Suhuai, J. Jin, and S. West, "Literature Review of Power Disaggregation," in Proceedings of 2011 International Conference on Modelling, Identification and Control (ICMIC), 2011, pp. 38-42.

[105] G. Fitzpatrick. (2009) Technology-Enabled Feedback on Domestic Energy Consumption: Articulating a Set of Design Concerns. pp. 37-44. Available: http://doi.ieeecomputersociety.org/10.1109/MPRV.2009.17

[106] G. Wood and M. Newborough, "Energy-Use Information Transfer for Intelligent Homes: Enabling Energy Conservation with Central and Local Displays," Energy and Buildings, vol. 39, pp. 495-503, 2007.

[107] C. J. H. Midden and B. S. M. Ritsema, "The Meaning of Normative Processes for Energy Conservation," Journal of Economic Psychology, vol. 4, pp. 37-55, 1983.

[108] G. Jacucci, A. Spagnolli, L. Gamberini, A. Chalambalakis, C. Bjoerksog, M. Bertoncini, C. Torstensson, and P. Monti, "Designing Effective Feedback of Electricity Consumption for Mobile User Interfaces," Psychology Journal, vol. 7, pp. 265-289, 2009.

[109] Y. Riche, J. Dodge, and R. A. Metoyer, "Studying Always-on Electricity Feedback in the Home," in Proceedings of the 28th International Conference on Human Factors in Computing Systems, Atlanta, Georgia, USA, 2010, pp. 1995-1998.

[110] W. F. Van Raaij and T. M. M. Verhallen, "A Behavioral Model of Residential Energy Use," Journal of Economic Psychology, vol. 3, pp. 39-63, 1983.

[111] R. B. Ammons, "Effects of Knowledge of Performance: A Survey and Tentative Theoretical Formulation," Journal of General Psychology, 1956.

[112] M. L. Dennis, E. J. Soderstrom, W. S. Koncinski, and B. Cavanaugh, "Effective Dissemination of Energy-Related Information: Applying Social Psychology and Evaluation Research," American Psychologist, vol. 45, p. 1109, 1990.

[113] L. A. Furrey, S. Nadel, and J. A. S. Laitner, "Laying the Foundation for Implementing a Federal Energy Efficiency Resource Standard," ACEEE, Washington D.C. Research Report E091, 2009.

[114] T. Luckenbach, P. Gober, S. Arbanowski, A. Kotsopoulos, and K. Kim, "Tinyrest-a Protocol for Integrating Sensor Networks into the Internet," in Workshop on Real-World Wireless Sensor Network, 2005.

[115] S. Dawson-Haggerty, X. Jiang, G. Tolle, J. Ortiz, and D. Culler, "Smap: A Simple Measurement and Actuation Profile for Physical Information," in Proceedings of SenSys 2010, Zurich, Switzerland, 2010.

[116] D. Guinard, M. Mueller, and J. Pasquier-Rocha, "Giving RFID a Rest: Building a Web-Enabled Epcis," in Internet of Things 2010, Tokyo, Japan, 2010, pp. 1-8.

[117] D. Yazar and A. Dunkels, "Efficient Application Integration in Ip-Based Sensor Networks," in Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings, Berkeley, California, 2009, pp. 43-48.

[118] D. Guinard, M. Weiss, and V. Trifa, "Are You Energy-Efficient? Sense It on the Web," in Adjunct Proceedings Pervasive09, Nara, Japan, 2009.

[119] J. Froehlich, E. Larson, S. Gupta, G. Cohn, M. S. Reynolds, and S. N. Patel, "Disaggregated End-Use Energy Sensing for the Smart Grid," IEEE Pervasive Computing, pp. 28-39, 2010.

[120] M. Weiss, C. M. Loock, T. Staake, F. Mattern, and E. Fleisch, "Evaluating Mobile Phones as Energy Consumption Feedback Devices," in 7th International ICST Conference on Mobile and Ubiquitous Systems (Mobiquitous 2010), Sydney, NSW, Australia, 2010.

[121] M. Weiss, A. Merkle, T. Staake, and E. Fleisch, "Towards a PowerPedia - a Collaborative Energy Encyclopedia," in UCSE2010 Workshop at Ubicomp2010 (1st Workshop on Ubiquitous Computing for Sustainable Energy), Kopenhagen, Denmark, 2010.

[122] O. Baecker, T. Ippisch, F. Michahelles, S. Roth, and E. Fleisch, "Mobile Claims Assistance," in Proceedings of the 8th International Conference on Mobile and Ubiquitous Multimedia, Cambridge, United Kingdom, 2009, pp. 1-9.

[123] R. Adelmann, "Mobile Phone Based Interaction with Everyday Products - on the Go," in Proceedings of the 2007 International Conference on Next Generation Mobile Applications, Services and Technologies (NGMAST '07), 2007, pp. 63-69.

[124] J. Landgren and U. Nulden, "A Study of Emergency Response Work: Patterns of Mobile Phone Interaction," in Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, San Jose, California, USA, 2007, pp. 1323-1332.

[125] A. Dada, F. G. Von Reischach, and T. Staake, "Displaying Dynamic Carbon Footprints of Products on Mobile Phones," in Adjunct Proceedings of the 6th International Conference on Pervasive Computing, Sydney, Australia, 2008.

[126] L. Bartram, J. Rodgers, and R. Woodbury, "Smart Homes or Smart Occupants? Supporting Aware Living in the Home." vol. 6947, P. Campos, N. Graham, J. Jorge, N. Nunes, P. Palanque, and M. Winckler, Eds. Springer Berlin / Heidelberg, 2011, pp. 52-64.

[127] T. G. Holmes, "Eco-Visualization: Combining Art and Technology to Reduce Energy Consumption," in Proceedings of the 6th ACM SIGCHI Conference on Creativity and Cognition, Washington, DC, USA, 2007, pp. 153-162.

[128] J. Froehlich, "Promoting Energy Efficient Behaviors in the Home through Feedback: The Role of Human-Computer Interaction," University of Washington, Seatle, USA , Technical Note #09-02-01, 2009.

[129] K. Sami, "Consumer Preferences for Feedback on Household Electricity Consumption," Energy and Buildings, vol. 43, pp. 458-467, 2011.

[130] S. L. Smith and J. N. Mosier, Guidelines for Designing User Interface Software: Userlab, 1986.

[131] J. Kjeldskov and C. Graham, "A Review of Mobile Hci Research Methods Human-Computer Interaction with Mobile Devices and Services," Lecture Notes in Computer Science, vol. 2795, pp. 317-335, 2003.

[132] F. von Reischach, F. Michahelles, D. Guinard, R. Adelmann, E. Fleisch, and A. Schmidt, "An Evaluation of Product Identification Techniques for Mobile Phones Human-Computer Interaction," in International Conference on Human-Computer Interaction (INTERACT 2009), Uppsala, Sweden, 2009, pp. 804-816.

[133] F. D. Davis, "Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology," MIS Quarterly, pp. 319-340, 1989.

[134] I. P. Levin, S. L. Schneider, and G. J. Gaeth, "All Frames Are Not Created Equal: A Typology and Critical Analysis of Framing Effects," Organizational Behavior and Human Decision Processes, vol. 76, pp. 149-188, 1998.

[135] C. Snyder, Paper Prototyping: The Fast and Easy Way to Design and Refine User Interfaces: Morgan Kaufmann Pub, 2003.

[136] R. Sefelin, M. Tscheligi, and V. Giller, "Paper Prototyping - What Is It Good For?: A Comparison of Paper- and Computer-Based Low-Fidelity Prototyping," in CHI03 Extended Abstracts, Ft. Lauderdale, Florida, USA, 2003, pp. 778-779.

[137] D. G. Saari, "The Optimal Ranking Method Is the Borda Count," Discussion Papers, 1985.

[138] D. Black, The Theory of Committees and Elections. London: University Press Cambridge, 1958.

[139] A. M. Lund, "Measuring Usability with the Use Questionnaire," STC Usability and User Experience SIG Newsletter, vol. 8(2), 2001.

[140] A. Parasuraman, "Technology Readiness Index (TRI)," Journal of Service Research, vol. 2, p. 307, 2000.

[141] W. Kempton and L. Montgomery, "Folk Quantification of Energy," Energy, vol. 7, pp. 817-827, 1982.

[142] A. Gustafsson, C. Katzeff, and M. Bang, "Evaluation of a Pervasive Game for Domestic Energy Engagement among Teenagers," Journal of Computer Entertainment, vol. 7, pp. 1-19, 2010.

[143] C. M. Loock, T. Graml, M. Baeriswyl, and T. Staake, "How to Motivate Energy Efficiency Online: ICT Can Help to Induce Behavior Change," in 20th International Conference on Management of Technology, Miami, FL, USA, 2011, p. 15.

[144] W. Abrahamse, L. Steg, C. Vlek, and T. Rothengatter, "The Effect of Tailored Information, Goal Setting, and Tailored Feedback on Household Energy Use, Energy-Related Behaviors, and Behavioral Antecedents," Journal of Environmental Psychology, vol. 27, pp. 265-276, 2007.

[145] G. Peschiera, J. E. Taylor, and J. A. Siegel, "Response‚Äìrelapse Patterns of Building Occupant Electricity Consumption Following Exposure to Personal, Contextualized and Occupant Peer Network Utilization Data," Energy and Buildings, vol. 42, pp. 1329-1336, 2010.

[146] T. Graml, C.-M. Loock, M. Baeriswyl, and T. Staake, "Improving Residential Energy Consumption at Large Using Persuasive Systems," in Proceedings of the 19th European Conference on Information Systems (ECIS), Helsinki, Finland, 2011, pp. 2267-2280.

[147] H. L. O'Brien and E. G. Toms, "What Is User Engagement? A Conceptual Framework for Defining User Engagement with Technology," Journal of the American Society for Information Science and Technology, vol. 59, pp. 938-955, 2008.

[148] H.-M. Chen, C.-W. Lin, S.-H. Hsieh, H.-F. Chao, C.-S. Chen, R.-S. Shiu, S.-R. Ye, and Y.-C. Deng, "Persuasive Feedback Model for Inducing Energy Conservation Behaviors of Building Users Based on Interaction with a Virtual Object," Energy and Buildings, vol. 45, pp. 106-115, 2012.

[149] A. Gustafsson, M. Bang, and M. Svahn, "Power Explorer: A Casual Game Style for Encouraging Long Term Behavior Change among Teenagers," in Proceedings of the International Conference on Advances in Computer Enterntainment Technology, Athens, Greece, 2009, pp. 182-189.

[150] A. Gustafsson, M. D. Johnson, and I. Roos, "The Effects of Customer Satisfaction, Relationship Commitment Dimensions, and Triggers on Customer Retention," Journal of Marketing, pp. 210-218, 2005.

[151] L. Gamberini, N. Corradi, L. Zamboni, M. Perotti, C. Cadenazzi, S. Mandressi, G. Jacucci, G. Tusa, A. Spagnolli, C. Bjorkskog, M. Salo, and P. Aman, "Saving Is Fun: Designing a Persuasive Game for Power Conservation," in Proceedings of the 8th International Conference on Advances in Computer Entertainment Technology, Lisbon, Portugal, 2011, pp. 1-7.

[152] M. Weiss, A. Helfenstein, F. Mattern, and T. Staake, "Leveraging Smart Meter Data to Recognize Smart Home Appliances," in IEEE International Conference on Pervasive Computing and Communications 2012 (PerCom 2012), Lugano, Switzerland, 2012.

[153] J. Liu, L. Nie, and Z. Xu, "The Input-Sensing Problem in Ternary Computing and Its Application in Household Energy-Saving," in 2011 IEEE/ACM International Conference on Green Computing and Communications (GreenCom), 2011, pp. 131-138.

[154] M. Parti and C. Parti, "The Total and Appliance-Specific Conditional Demand for Electricity in the Household Sector," The Bell Journal of Economics, pp. 309-321, 1980.

[155] D. W. Caves, J. A. Herriges, K. E. Train, and R. J. Windle, "A Bayesian Approach to Combining Conditional Demand and Engineering Models of Electricity Usage," The Review of Economics and Statistics, pp. 438-448, 1987.

[156] D. J. Aigner, C. Sorooshian, and P. Kerwin, "Conditional Demand Analysis for Estimating Residential End-Use Load Profiles," The Energy Journal, vol. 5, pp. 81-98, 1984.

[157] L. G. Swan and V. I. Ugursal, "Modeling of End-Use Energy Consumption in the Residential Sector: A Review of Modeling Techniques," Renewable and Sustainable Energy Reviews, vol. 13, pp. 1819-1835, 2009.

[158] R. Bartels and D. G. Fiebig, "Integrating Direct Metering and Conditional Demand Analysis for Estimating End-Use Loads," The Energy Journal, vol. 11, pp. 79-98, 1990.

[159] D. G. Fiebig, R. Bartels, and D. J. Aigner, "A Random Coefficient Approach to the Estimation of Residential End-Use Load Profiles," Journal of Econometrics, vol. 50, pp. 297-327, 1991.

[160] R. Bartels and D. G. Fiebig, "Residential End-Use Electricity Demand: Results from a Designed Experiment," The Energy Journal, vol. 21, pp. 51-82, 2000.

[161] D. L. Ryan and R. Liu, "Conditional Demand Analysis Revisited: Evaluating Residential End-Use Energy Consumption in Canada," Canadian Building Energy End-Use Data and Analysis Centre Research Report CBEEDAC, 2007.

[162] M. P. E. Lins, A. C. M. da Silva, and L. P. Rosa, "Regional Variations in Energy Consumption of Appliances: Conditional Demand Analysis Applied to Brazilian Households," Annals of Operations Research, vol. 117, pp. 235-246, 2002.

[163] D. Yi, D. Liang, L. Bin, R. Harley, and T. Habetler, "A Review of Identification and Monitoring Methods for Electric Loads in Commercial and Residential Buildings," in IEEE Energy Conversion Congress and Exposition (ECCE), 2010, pp. 4527-4533.

[164] G. W. Hart, "Nonintrusive Appliance Load Monitoring," Proceedings of the IEEE, vol. 80, pp. 1870-1891, 1992.

[165] L. Norford and N. Mabey, "Non-Intrusive Electric Load Monitoring in Commercial Buildings," in Eighth Symposium on Improving Building Systems in Hot and Humid Climates, Dallas, USA, 1992.

[166] S. R. Shaw, L. K. Norford, D. Luo, and S. B. Leeb, "Detection and Diagnosis of HVAC Faults Via Electrical Load Monitoring," HVAC&R Research, vol. 8, pp. 13-40, 2002.

[167] H. Pihala, "Non-Intrusive Appliance Load Monitoring System Based on a Modern kWh-Meter," Technical Research Centre of Finland, 1998.

[168] D. Luo, L. K. Norford, S. Leeb, and S. Shaw, "Monitoring HVAC Equipment Electrical Loads from a Centralized Location-Methods and Field Test Results," ASHRAE Transactions, vol. 108, pp. 841-857, 2002.

[169] S. Drenker and A. Kader, "Nonintrusive Monitoring of Electric Loads," Computer Applications in Power, IEEE, vol. 12, pp. 47-51, 1999.

[170] F. Sultanem, "Using Appliance Signatures for Monitoring Residential Loads at Meter Panel Level," IEEE Transactions on Power Delivery, vol. 6, pp. 1380-1385, 1991.

[171] L. K. Norford and S. B. Leeb, "Non-Intrusive Electrical Load Monitoring in Commercial Buildings Based on Steady-State and Transient Load-Detection Algorithms," Energy and Buildings, vol. 24, pp. 51-64, 1996.

[172] S. B. Leeb, S. R. Shaw, and J. L. Kirtley, Jr., "Transient Event Detection in Spectral Envelope Estimates for Nonintrusive Load Monitoring," IEEE Transactions on Power Delivery, vol. 10, pp. 1200-1210, 1995.

[173] C. Laughman, K. Lee, R. Cox, S. Shaw, S. Leeb, L. Norford, and P. Armstrong, "Power Signature Analysis," Power and Energy Magazine, IEEE, vol. 1, pp. 56-63, 2003.

[174] W. Wichakool, A. T. Avestruz, R. W. Cox, and S. B. Leeb, "Modeling and Estimating Current Harmonics of Variable Electronic Loads," IEEE Transactions on Power Electronics, vol. 24, pp. 2803-2811, 2009.

[175] K. D. Lee, S. B. Leeb, L. K. Norford, P. R. Armstrong, J. Holloway, and S. R. Shaw, "Estimation of Variable-Speed-Drive Power Consumption from Harmonic Content," IEEE Transactions on Energy Conversion, vol. 20, pp. 566-574, 2005.

[176] S. R. Shaw, S. B. Leeb, L. K. Norford, and R. W. Cox, "Nonintrusive Load Monitoring and Diagnostics in Power Systems," IEEE Transactions on Instrumentation and Measurement, vol. 57, pp. 1445-1454, 2008.

[177] A. Cole and A. Albicki, "Nonintrusive Identification of Electrical Loads in a Three-Phase Environment Based on Harmonic Content," in Proceedings of the 17th IEEE Instrumentation and Measurement Technology Conference (IMTC 2000), 2000, pp. 24-29.

[178] A. I. Cole and A. Albicki, "Algorithm for Nonintrusive Identification of Residential Appliances," in Proceedings of the 1998 IEEE International Symposium on Circuits and Systems (ISCAS98) 1998, pp. 338-341 vol.3.

[179] A. I. Cole and A. Albicki, "Data Extraction for Effective Non-Intrusive Identification of Residential Power Loads," in IEEE Conference Proceedings of the Instrumentation and Measurement Technology Conference (IMTC98), 1998, pp. 812-815.

[180] A. Marchiori, D. Hakkarinen, Q. Han, and L. Earle, "Circuit-Level Load Monitoring for Household Energy Management," IEEE Pervasive Computing, vol. 10, pp. 40-48, 2011.

[181] A. Marchiori and Q. Han, "Using Circuit-Level Power Measurements in Household Energy Management Systems," in Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings, Berkeley, California, 2009, pp. 7-12.

[182] J. Powers, B. Margossian, and B. Smith, "Using a Rule-Based Algorithm to Disaggregate End-Use Load Profiles from Premise-Level Data," Computer Applications in Power, IEEE, vol. 4, pp. 42-47, 1991.

[183] A. G. Ruzzelli, C. Nicolas, A. Schoofs, and G. M. P. O'Hare, "Real-Time Recognition and Profiling of Appliances through a Single Electricity Sensor," in Proceedings of the 7th Annual IEEE Communications Society Conference on Sensor Mesh and Ad Hoc Communications and Networks (SECON2010), 2010, pp. 1-9.

[184] M. L. Marceau and R. Zmeureanu, "Nonintrusive Load Disaggregation Computer Program to Estimate the Energy Consumption of Major End Uses in Residential Buildings," Energy Conversion and Management, vol. 41, pp. 1389-1403, 2000.

[185] H.-H. Chang, H.-T. Yang, and C.-L. Lin, "Load Identification in Neural Networks for a Non-Intrusive Monitoring of Industrial Electrical Loads," Computer Supported Cooperative Work in Design IV, vol. 5236, pp. 664-674, 2008.

[186] M. Akbar and D. Z. A. Khan, "Modified Nonintrusive Appliance Load Monitoring for Nonlinear Devices," in IEEE International Multitopic Conference (NMIC2007), 2007, pp. 1-5.

[187] W. L. Chan, A. T. P. So, and L. L. Lai, "Harmonics Load Signature Recognition by Wavelets Transforms," in Proceedings of the International Conference of Electric Utility Deregulation and Restructuring and Power Technologies, 2000, pp. 666-671.

[188] Norman, N. C. F. Tse, Z. Long, and L. L. Lai, "Wavelet-Based Algorithm for Nonstationary Power System Waveform Analysis," in International Conference on Wavelet Analysis and Pattern Recognition (ICWAPR08), 2008, pp. 729-735.

[189] W. Lee, G. Fung, H. Lam, F. Chan, and M. Lucente, "Exploration on Load Signatures," in International Conference on Electrical Engineering (ICEE), 2004.

[190] D. Srinivasan, W. S. Ng, and A. C. Liew, "Neural-Network-Based Signature Recognition for Harmonic Source Identification," IEEE Transactions on Power Delivery, vol. 21, pp. 398-405, 2006.

[191] L. Jian, S. Ng, G. Kendall, and J. Cheng, "Load Signature Study - Part I: Basic Concept, Structure, and Methodology," IEEE Transactions on Power Delivery, vol. 25, pp. 551-560, 2010.

[192] L. Jian, S. Ng, G. Kendall, and J. Cheng, "Load Signature Study Part II: Disaggregation Framework, Simulation and Applications," in Power and Energy Society General Meeting, 2010 IEEE, 2010, pp. 1-11.

[193] J. Z. Kolter, S. Batra, and A. Y. Ng, "Energy Disaggregation Via Discriminative Sparse Coding," in 24th Annual Conference on Neural Information Processing Systems, Vancouver, BC, Canada, 2010.

[194] J. Z. Kolter and M. J. Johnson, "Redd: A Public Data Set for Energy Disaggregation Research," in Workshop on Data Mining Applications in Sustainability (SustKDD 2011), San Diego, USA, 2011.

[195] O. Parson, S. Ghosh, M. Weal, and A. Rogers, "Using Hidden Markov Models for Iterative Non-Intrusive Appliance Monitoring," in Neural Information Processing Systems Workshop on Machine Learning for Sustainability, Sierra Nevada, Spain, 2011.

[196] H. Goncalves, A. Ocneanu, and M. Berges, "Unsupervised Disaggregation of Appliances Using Aggregated Consumption Data," in Workshop on Data Mining Applications in Stustainability (SustKDD 2011), San Diego, USA, 2011.

[197] H. Kim, M. Marwah, M. Arlitt, G. Lyon, and J. Han, "Unsupervised Disaggregation of Low Frequency Power Measurements," in SIAM Conference on Data Mining (SDM11), Mesa, USA, 2011.

[198] M. E. Berges, E. Goldman, H. S. Matthews, and L. Soibelman, "Enhancing Electricity Audits in Residential Buildings with Nonintrusive Load Monitoring," Journal of Industrial Ecology, vol. 14, pp. 844-858, 2010.

[199] M. Berges, E. Goldman, H. S. Matthews, and L. Soibelman, "Learning Systems for Electric Consumption of Buildings," in ASCE International Workshop on Computing in Civil Engineering, Austin, USA, 2009.

[200] M. Berges, E. Goldman, H. S. Matthews, L. Soibelman, and K. Anderson, "User-Centered Non-Intrusive Electricity Load Monitoring for Residential Buildings," Journal of Computing in Civil Engineering, vol. 1, 2011.

[201] G. Marubayashi, "Noise Measurements of the Residential Power Line," in International Symposium on Power Line Communications and its Applications (ISPLC), 1997, pp. 104-108.

[202] M. Ito, R. Uda, S. Ichimura, K. Tago, T. Hoshi, and Y. Matsushita, "A Method of Appliance Detection Based on Features of Power Waveform," in Proceedings of the International Symposium on Applications and the Internet, 2004, pp. 291-294.

[203] M. Baranski and J. Voss, "Nonintrusive Appliance Load Monitoring Based on an Optical Sensor," in IEEE Power Tech Conference, Bologna, Italy, 2003, pp. 8, Vol. 4.

[204] A. Schoofs, A. Guerrieri, D. T. Delaney, G. O'Hare, and A. G. Ruzzelli, "ANNOT: Automated Electricity Data Annotation Using Wireless Sensor Networks," in 7th Annual IEEE Communications Society Conference on Sensor Mesh and Ad Hoc Communications and Networks (SECON 2010), 2010, pp. 1-9.

[205] M. Berges, E. Goldman, H. S. Matthews, and L. Soibelman, "Training Load Monitoring Algorithms on Highly Sub-Metered Home Electricity Consumption Data," Tsinghua Science & Technology, vol. 13, pp. 406-411, 2008.

[206] M. Berges, L. Soibelman, and H. S. Matthews, "Building Commissioning as an Opportunity for Training Non-Intrusive Load Monitoring Algorithms," in 16th International Conference on Innovation in Architecture, Engineering and Construction (AEC), State College, USA, 2010.

[207] M. Roberts and H. Kuhns, "Towards Bridging the Gap between the Smart Grid and Smart Energy Consumption," in Proceedings of the 2010 ACEEE Summer Study on Energy Efficiency in Buildings, Pacific Grove, USA, 2010.

[208] W. R. Schucany, "Kernel Smoothers: An Overview of Curve Estimators for the First Graduate Course in Nonparametric Statistics," Statistical Science, pp. 663-675, 2004.

[209] H. C. McAlpine, B. Hicks, and S. Culley, "Improving Re-Use of Informal Information through the Creation of an Engineering Electronic Logbook (Eel): A Demonstrator," in ASME Conference Proceedings, 2009, pp. 843-850.

[210] J. Scott, A. J. B. Brush, J. Krumm, B. Meyers, M. Hazas, S. Hodges, and N. Villar, "Preheat: Controlling Home Heating Using Occupancy Prediction," in Proceedings of the 13th International Conference on Ubiquitous Computing, Beijing, China, 2011, pp. 281-290.

[211] W. Poortinga, L. Steg, C. Vlek, and G. Wiersma, "Household Preferences for Energy-Saving Measures: A Conjoint Analysis," Journal of Economic Psychology, vol. 24, pp. 49-64, 2003.

[212] F. Ölander and J. ThØgersen, "Understanding of Consumer Behaviour as a Prerequisite for Environmental Protection," Journal of Consumer Policy, vol. 18, pp. 345-385, 1995.

[213] Iowa Energy Center. (2010). Lower Energy Bills with a Set-Back Thermostat. Available: http://www.energy.iastate.edu/news/pr/energysavingideas/setbackt herm.htm

[214] L. M. Hilty, P. Arnfalk, L. Erdmann, J. Goodman, M. Lehmann, and P. A. Wager, "The Relevance of Information and Communication Technologies for Environmental Sustainability - a Prospective Simulation Study," Environmental Modelling & Software, vol. 21, pp. 1618-1629, 2006.

[215] A. Esser, A. Kamper, M. Franke, D. Möst, and O. Rentz, "Scheduling of Electrical Household Appliances with Price Signals," Operations Research Proceedings 2006, pp. 253-258, 2007.

[216] A. Faruqui and S. Sergici, "Household Response to Dynamic Pricing of Electricity: A Survey of 15 Experiments," Journal of Regulatory Economics, vol. 38, pp. 193-225, 2010.

[217] A. Moholkar, P. Klinkhachorn, and A. Feliachi, "Effects of Dynamic Pricing on Residential Electricity Bills," in Power Systems Conference and Exposition, 2004. IEEE PES, 2004, pp. 1030-1035.

Curriculum Vitae: Markus Weiss

Personal Data

Date of Birth	May 8, 1981
Birthplace	Würzburg, Germany
Citizenship	German

Education

03/2008 – 02/2012	ETH Zurich, Ph.D. Student at the Institute for Pervasive Computing
02/2011 – 11/2011	Massachusetts Institute of Technology (MIT), Visiting Scholar at the ESD/MIT Energy Initiative
12/2003 – 02/2008	Julius-Maximilians-University of Würzburg, M.Sc. (Diplom) in Computer Science
08/2005 – 06/2006	University of Texas at Austin, Visiting Graduate Student in Computer Science
10/2001 – 11/2003	Julius-Maximilians-University of Würzburg, B.Sc. Studies in Computer Science
09/1991 – 06/2000	Schönborngymnasium, High School (Gymnasium), Würzburg, Germany

Civil Service

09/2000 – 07/2001	Center for Handicapped, Würzburg, Germany

Experience

03/2008 – 05/2012	Research Assistant, Department for Computer Science, ETH Zurich, Zurich, Switzerland
03/2008 – 05/2012	Senior Researcher, Bits to Energy Lab (ETH Zurich/University of Sankt Gallen), Zurich, Switzerland
11/2006 – 10/2007	Master Thesis in Cooperation with DATEV eG, Nürnberg, Germany
03/2004 – 07/2005	Student Assistant, Department of Computer Science, Julius-Maximilians-University of Würzburg, Würzburg, Germany