Andrea Monacchi
Matr.-Nr.: 1261444

Assisting Energy Management in Smart Buildings and Microgrids

Dissertation

zur Erlangung des akademischen Grades
Doktor der technischen Wissenschaften

Alpen-Adria-Universität Klagenfurt
Fakultät für Technische Wissenschaften

Begutachter:

Univ.–Prof. Dr. -techn. Wilfried Elmenreich
Institut für Vernetzte und Eingebettete Systeme
Alpen-Adria-Universität Klagenfurt

Klagenfurt, im Juni 2016

Declaration of honor

I hereby declare in lieu of an oath that

- the submitted academic paper is entirely my own work and that no auxiliary materials have been used other than those indicated,
- I have fully disclosed all assistance received from third parties during the process of writing the paper, including any significant advice from supervisors,
- any contents taken from the works of third parties or my own works that have been included either literally or in spirit have been appropriately marked and the respective source of the information has been clearly identified with precise bibliographical references (e.g. in footnotes),
- to date, I have not submitted this paper to an examining authority either in Austria or abroad and that
- the digital version of the paper submitted for the purpose of plagiarism assessment is fully consistent with the printed version.

I am aware that a declaration contrary to the facts will have legal consequences.

Signature	Date/Place

Abstract

The increasing exploitation of renewable energy sources for power generation introduces a significant instability into the power grid, which has to be addressed with appropriate management strategies. Energy storage is a costly and inefficient solution. Demand-side control mechanisms can help mitigating the unbalance between available supply and demand. This includes both direct and indirect control, depending on the degree of controllability of demand-side loads. In the latter, congestion on the shared resource is managed using a price signal, exchanged throughout the power grid and reflecting the resource availability. This requires the timely exchange of information between energy consumers and producers, namely power and phase measurements to be used for the resource pricing. Furthermore, more fine-grained usage data is progressively becoming available to utilities thanks to the deployment of smart meters. Such an information is also relevant to facility managers and users, to become aware of the energy footprint of daily activities and seek a more efficient usage process.

This thesis deals with different applications of high-resolution power usage data for energy management in smart microgrids. To this end, the first stage included a measurement campaign in selected households in Italy and Austria. The resulting dataset, named GREEND, contains more than 1 year consumption data at 1 Hz. GREEND was released to the research community for open use, as well as used throughout the thesis.

We elaborate on the design of a data infrastructure capable of collecting data from heterogeneous data sources in highly dynamic environments. Specifically, architectural requirements are identified to achieve interoperability at the level of electrical devices as well as exchanged data. The proposed solution offers a single interface to query for status changes, which eases the application development process. In addition, we propose an ontology modeling both static and dynamic information of household appliances. This allows for the full integration of smart and non-smart devices, whose behavior can be tracked and recorded in a sort of datasheet to be exchanged across the network.

The availability of energy usage data allows for the provisioning of value-added services to both end-users and utilities. To this end, we investigate on the possibility of an interactive system to timely inform users on their energy usage, in order to promote an efficient use of local resources. In particular, advices are returned to consumers based on their usage behavior and building occupance. Using the GREEND, this solution alone was quantified as potentially yielding up to 34% of savings.

However, the effectiveness of demand response programs is greatly affected by the possibility to automate specific devices. Towards this vision, we introduced the HEMS market simulator, which allows for training appliance controllers. Because of the strictly competitive setting, pure market mechanisms do not offer a complete solution for automatic load management. Accordingly, competition is limited to a specific trading day, and has the potential effect of yielding service interruption. To solve this issue, we propose a microgrid power broker that acts as a retailer of available supply. The broker seeks profit by forecasting the price of different power provisioning durations.

The three different approaches are independent and give an individual contribution to the research community. The results provide the basis for future research in the field of energy management systems for microgrids and smart buildings.

Zusammenfassung

Die wachsende Nutzung von erneuerbaren Energien stellt das Stromnetz vor neue Herausforderungen die geeignete Maßnahmen für die stabile Energieversorgung im Stromnetz erfordern. Der massive Einsatz von Speichermöglichkeiten ist hierbei eine vergleichsweise teure und ineffiziente Lösung. Mechanismen zur Nachfragesteuerung können das Ungleichgewicht zwischen Energieangebot und -nachfrage verbessern. Dies inkludiert sowohl direkte als auch indirekte Regelmechanismen, die abhängig von der Steuerbarkeit der Lasten eingesetzt werden können. In letzterem Fall werden Netzengpässe und Überschüsse durch einen veränderlichen Strompreis behandelt. Das erfordert den rechtzeitigen Informationsaustausch zwischen Anbietern und Verbrauchern, konkret Daten zu Leistungsflüssen und Phasenlage als Grundlage zur Preisbildung. Dazu stehen den Energieversorgern durch den Einsatz von intelligenten Stromzählern zunehmend genauere Strommessdaten zur Verfügung. Solche Daten sind auch für das Gebäudemanagement und für Endkunden von Interesse, um den täglichen Energieverbrauch zu erfassen und auf dieser Grundlage Effizienzverbesserungen zu planen.

Diese Arbeit befasst sich mit verschiedene Anwendungen von hochauflösenden Stromverbrauchsdaten für das Energiemanagement in Smart Microgrids. Zu diesem Zweck wurde in der ersten Phase eine Messkampagne in ausgewählte Haushalten in Italien und Österreich durchgeführt. Der daraus entstehende Datensatz, kurz GREEND, enthält mehr als 1 Jahr an detaillierten Verbrauchsdaten die im Sekundenabstand gemessen wurden. Der GREEND-Datensatz wurde veröffentlicht und für die Forschungsgemeinschaft freigegeben. Er wurde ebenfalls durchgehend in dieser Dissertation verwendet.

Diese Arbeit stellt außerdem eine Dateninfrastruktur vor, welche das Zusammenführen von Daten aus unterschiedlichen Quellen in dynamischen Umgebungen erlaubt. Im Besonderen werden Architekturanforderungen identifiziert welche die Interoperabilität auf Geräte- und Datenebene ermöglicht. Die vorgestellte Lösung bietet eine einheitliche Schnittstelle um Datenänderungen zu verfolgen, was die Umsetzung in konkrete Anwendungen erleichtert. Zusätzlich wird ein Ontologiemodell vorgestellt welche statische und dynamische Information zu Haushaltsgeräten modellieren kann. Dies ermöglicht die volle Integration von intelligenten und konventionellen Geräten, so dass das Verhalten eines Geräts allgemein in einer Art elektronisches Datenblatt über das Netz kommuniziert werden kann.

Die Verfügbarkeit von Energieverbrauchsdaten ermöglicht das Bereitstellen von Mehrwertdiensten für Endverbraucher und Versorger. Dazu wird die Möglichkeit eines interaktiven Systems untersucht, welches Verbraucher über ihren Energie-

verbrauch zeitnah informieren kann um eine effiziente Nutzung von vorhandenen Ressourcen zu fördern. Konkret werden hier basierend auf dem Verbrauchsverhalten und der Wohnungsbelegung Vorschläge an den Kunden zurückgegeben. Unter Verwendung von GREEND konnte für diesen Ansatz ein Einsparungspotential von bis zu 34% ermittelt werden.

Die Effektivität von Nachfragesteuerungsprogrammen hängt jedoch stark von automatisiert steuerbaren Geräten ab. Zur Umsetzung dieser Vision stellt diese Arbeit im letzten Teil eine HEMS-Marktsimulation vor, welche das automatische Trainieren von Gerätesteuerungen ermöglicht.

Da dieser Ansatz auf einem kompetitiven Ansatz beruht liefert er nur bedingt kooperative Lösung zur Gerätesteuerung. Dementsprechend erfolgt die Optimierung nur jeweils für einen bestimmten Zeitraum was potentiell zu Betriebsunterbrechungen führen kann. Um dieses Problem zu lösen, wird ein Microgrid Power Broker vorgeschlagen, welcher als Vertriebsstelle für die verfügbare Energie agiert. Der Broker versucht dabei den Preis für eine gegebene Leistung über verschiedene Betriebszeiten vorherzusagen und dementsprechend zu handeln.

Die drei vorgestellten Ansätze sind unabhängig voneinander einsetzbar und liefern verschiedene Beiträge zum Stand der Forschung. Die Ergebnisse bieten eine Basis für zukünftige Forschung im Bereich Energiemanagement und in der intelligenten Gebäudetechnik.

Abstract

Lo sfruttamento crescente di risorse energetiche a carattere rinnovabile introduce una considerevole instabilità nella rete elettrica, la quale deve essere risolta mediante appropriate strategie di gestione. L'accumulo energetico offre una soluzione ancora inefficiente e costosa. I meccanismi di controllo lato consumatore (demand-side) possono mitigare il divario tra domanda ed offerta disponibili. Si distingue in particolare tra controllo diretto ed indiretto, in base alla possibilità di controllare direttamente i carichi dell'utenza. In quest'ultimo l'accesso alla risorsa è gestito tramite un segnale indicante il prezzo energetico, condiviso nella rete in modo da riflettere la disponibilità energetica. L'efficace calcolo del prezzo energetico richiede quindi un continuo scambio di informazioni tra consumatori e produttori, in termini di misure di potenza e fase. Inoltre l'installazione di contatori digitali rende sempre più disponibili dati ad alta frequenza. Tale informazione è utile a utenti e responsabili di strutture per divenire consapevoli del costo di attività quotidiane, e tentare quindi un utilizzo più efficiente.

Questa tesi concerne diverse applicazioni di dati di potenza ad alta risoluzione per la gestione energetica di smart microgrids. A tal fine la prima parte include una campagna di misura in abitazioni selezionate in Italia ed Austria. Il dataset risultante, chiamato GREEND, contiene più di 1 anno di dati di consumo ad 1 Hz. Il GREEND è stato rilasciato pubblicamente alla comunità di ricerca ed utilizzato all'interno di questa tesi.

Al fine di poter raccogliere dati da fonti eterogenee (diversi standard dati e di comunicazione) è stata progettata una infrastruttura dati. In particolare, requisiti architetturali sono stati identificati per poter raggiungere interoperabilità a livello di dispositivi elettrici e dati scambiati. La soluzione proposta offre una interfaccia singola per poter interrogare cambiamenti di stato. Ciò semplifica il processo di sviluppo di applicazioni. In aggiunta proponiamo una ontologia modellante informazioni statiche e dinamiche degli elettrodomestici. Ciò permette l'integrazione completa di elettrodomestici smart e comuni, il cui comportamento può essere monitorato e memorizzato in un profilo condivisibile nella rete.

La disponibilità di dati energetici permette la fornitura di servizi a valore aggiunto a utenti finali e fornitori di servizio. A tal fine, investighiamo la possibilità di sistemi interattivi di informare prontamente gli utenti sul loro utilizzo energetico, in modo da promuovere un utilizzo efficiente delle risorse locali. In particolare l'occupazione dell'edificio e il comportamento energetico vengono considerati per formulare e mostrare suggerimenti su come migliorare le performance. Usando il GREEND questa soluzione da sola è stata quantificata portare fino al 34% di risparmio.

L'efficacia dei programmi di demand response è influenzata dalla possibilità di automatizzare specifici dispositivi. Verso questa visione abbiamo introdotto il simulatore di mercato HEMS, che permette la progettazione di controllori per elettrodomestici. A causa dell'ambiente strettamente competitivo, i meccanismi di mercato non offrono una soluzione completa alla gestione automatica dei carichi. La competizione è infatti limitata ad una specifica giornata di negoziazione, il che può potenzialmente causare interruzioni di servizio. Per risolvere tale problema proponiamo un broker che funge da rivenditore della potenza disponibile. Il broker ricerca profitto stimando il prezzo per forniture di servizio aventi diversa durata.

I tre diversi approcci forniscono un contributo individuale alla comunità di ricerca. I risultati forniscono la base per ricerca futura nel campo dei sistemi di gestione energetica per edifici intelligenti e microgrids.

Acknowledgement

This work would never have been possible without the funding provided by the MONERGY project. The Alpen-Adria-Universität Klagenfurt and the state of Austria provided an exceptional environment to pursue research. I am indebted to my supervisor Prof. Wilfried Elmenreich for having selected me for such a project, as well as the Lakeside Labs research cluster, in the person of Ms. Claudia Prüggler. I was lucky to find a project matching my interest for smart homes, and I was even luckier to be given the freedom to coltivate my interests and apply the methods I found most powerful.

Living in Austria for 3 years was a great experience, also because of the different culture, food and landscape. At NES I have met wonderful colleagues with whom discussion in front of a beer was more than welcome. Pasquale Grippa, Wasif Masood and Nikolaj Marchenko had a special seat in this trip. My colleagues from the Smart Grid group also had a special role in this experience, most of all Dominik Egarter, Manfred Rabl-Pöchacker, Istvan Fehervari and Sergii Zhevzhyk. I would like to also thank Mrs Kornelia Lienbacher, who was always there to assist when a problem was popping up. Special thank also for Jacopo Gaggioli, a guy coming from my hometown Gubbio, who I have accidentally met on a Saturday at Neuer Platz, in Klagenfurt. Over time, it turned to a nice friendship.

A special thank goes to my parents and my sister Elisa, who convinced me not to give up in many occasions. Living away from the family is never too nice. It is only thanks to their presence and support that I could pursue this PhD work. Special thank also goes to my grandmas Gina and Virginia, as well as my aunts Silvana and Nadia, and their families. Finally I would like to thank the people at the public library of Gubbio, in which the last part of this dissertation was developed. They have set up a relaxing environment with a breathtaking view of the old town.

Last but not least, a special thank goes to those people that spent time to provide feedback: Prof. Gerhard Leitner and Prof. Sebastian Lenhoff.

List of Acronyms

3LN Three-Layered Network

6LoWPAN IPv6 over Low power Wireless Personal Area Networks

ADC Analog-to-Digital Converter

ANN Artificial Neural Network

API Application Program Interface

BLH Battery-based Load Hiding

CAR Carinthia

COAP Constrained Application Protocol

CSV Comma Separated Value

DER Distributed Energy Resources

DPWS Device Profile for Web Services

EMS Energy Management System

EXI Efficient XML Interchange

FHMM Factorial Hidden Markov Model

FMN Fully-Meshed Network

FSM Finite-State Machine

FVG Friuli-Venezia Giulia

GUI Graphical User Interface

HAVi Home Audio/Video interoperability

HEMS Home Energy Management System

HMM Hidden Markov Model

HSMM Hidden semi-Markov Model

ILM Intrusive Load Monitoring

JSON Javascript Simple Object Notation

LLH Load-based Load Hiding

MTBF Mean Time Between Failures

MTTR Mean Time To Recover

NAT Network Address Translator

NILM Non-Intrusive Load Monitoring

OSGi Open Services Gateway initiative

OWL Web Ontology Language

PLC Powerline Communication

PR-OWL Probabilistic OWL

RMI Java Remote Method Invocation

RNN Recurrent Neural Network

RPC Remote Procedure Call

RDF Resource Description Framework

REST Representational State Transfer

SDK Software Development Kit

SFTP Secure File Transfer Protocol

SOA Service-Oriented Architecture

SNR Signal-to-Noise Ratio

SP Smart Plug

SPARQL SPARQL Protocol and RDF Query Language

SPIN SPARQL Inferencing Notation

SPI Serial Peripheral Interface

UPnP Universal Plug and Play

Contents

List of Figures

List of Tables

CHAPTER 1

Introduction

"Simplicity is the ultimate sophistication" – Leonardo Da Vinci

"Μέγα βιβλίον, μέγα κακόν - Mega biblìon, mega kakòn" – Kallímachos

1.1 Motivation

Over the last decades, the growing concern towards global warming has raised massive investments on sustainability research. Fossil fuels still account for the highest share of utilized source, with consequences on pollution and climate change. According to the world wide fund for nature (WWF) fossil sources still account worldwide for more than 81% of sources[1]. This is similarly indicated by the international energy agency (IEA) in Fig. 1.1, which shows the proportion of energy supply in millions of tons of oil equivalent (Mtoe) for the period 1971 to 2013. Coal is the oldest and most polluting of those, being heavily responsible for higher emmisions of greenhouse gases and other detrimental substances (e.g., fine dusts, mercury, cadmium) than for instance natural gas. For instance, emissions of CO_2 from coal power plants are 30% higher than oil-powered ones, and 70% higher than those powered by natural gas. The Legambiente remarked in its dossier[2] that in Italy the 12 Coal-based power plants producing only 13% of supply are responsible for 30% of CO_2 of all Italian thermo-electric power plants. This translates into 36 Mt (millions of tons) of CO_2, out of the total 122 Mt. In Italy, installed coal-powered production amounts to 121 762 MW, against a demand peak of 59 126 MW. In any perspective, coal is far from being a sustainable solution, given also that Italy imports practically its entire demand of coal. Moreover, the market is often being compromised by statal incentives towards fossil sources. While nations have recognized the

[1] http://stopcarbone.wwf.it

[2] http://www.legambiente.it/sites/default/files/docs/dossier_carboneritornoalpassato.pdf

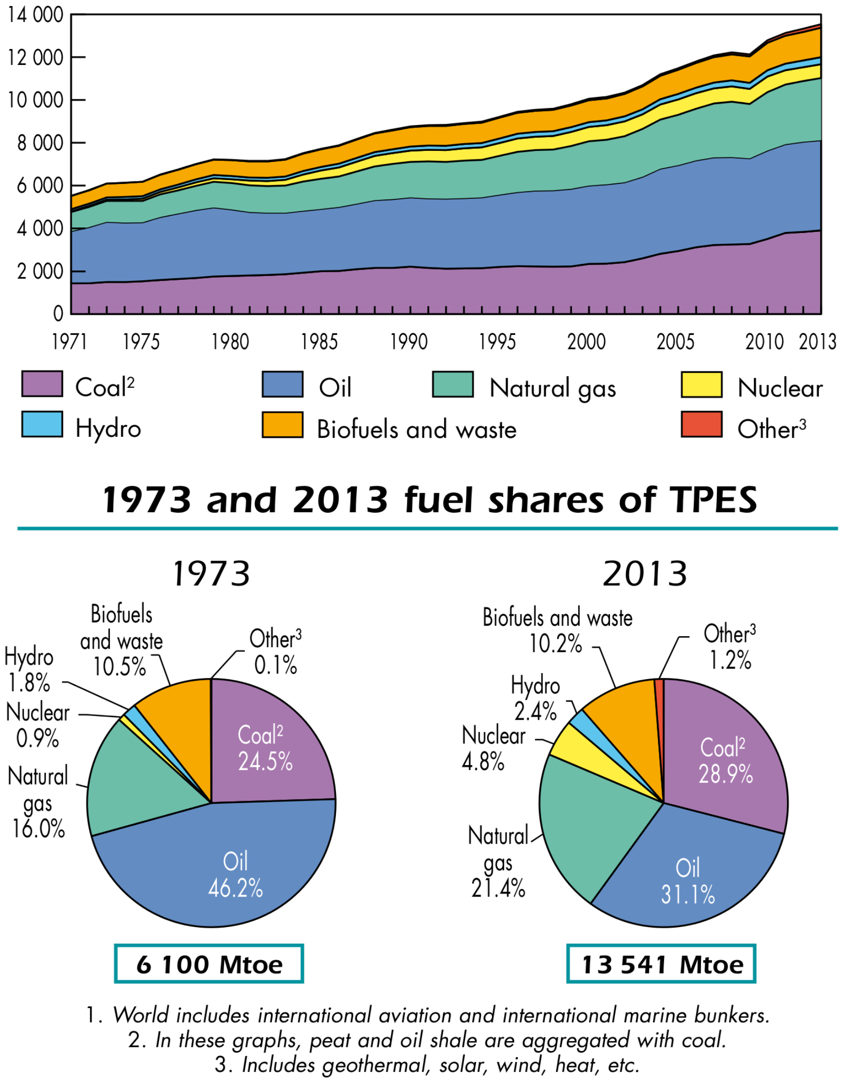

Figure 1.1: Worldwide total primary energy supply in Mtoe [OEC15]

problem and set a limit to their emissions, such as in the European 2020 20% reduction plan [Eur10] and the Kyoto protocol [Uni88], technology is still far from providing a sustainable solution to energy production, distribution and usage. The availability of fossil sources is destined to shrink and a diversification of sources is thus essential for the future. Fig. 1.2 shows the proportion of different energy sources for electricity production in the period 1971 to 2013. Energy production and its related effects are constantly rising, especially in developing countries, seen the increasing standards of living. For a community, access to electricity is the propeller of economic growth and quality of life. In its world energy outlook [IEA15], the IEA reports 1.2 billion people without access to electricity, mostly in rural areas of sub-Saharan Africa (635 millions)

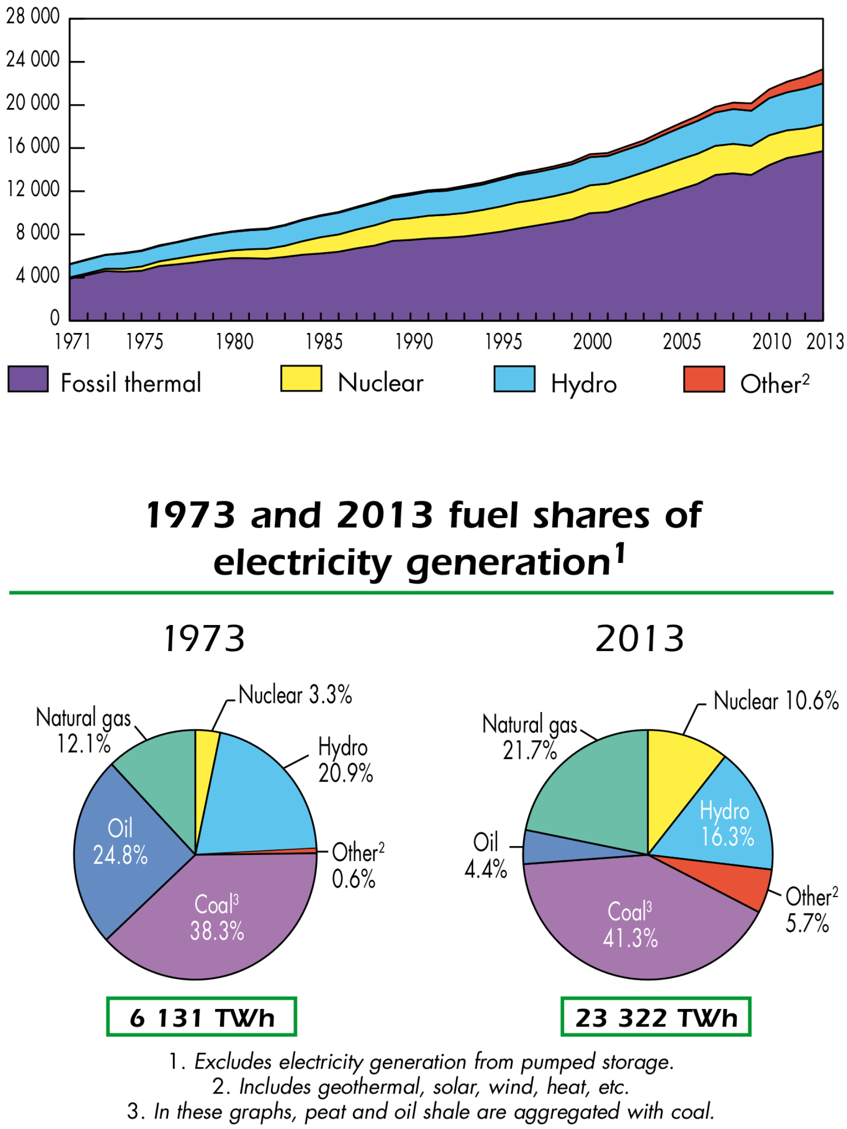

Figure 1.2: Electricity generation by fuel in TWh [OEC15]

and developing Asia (526 millions). These communities will soon claim their right for a decent life standard. Given the high installation and mantainance costs, the old energy supply model based on big power plants and a ubiquitous power distribution network provides an infeasible solution to such a demand. Power electronics will play a crucial role towards cleaner energy generation and more efficient storage and utilization, with an estimated saving potential of 20% [Bos10]. On the other hand, information and communication technologies have a core role in the transition towards a smarter power grid. The availability of an information channel connecting utilities, distribution network providers and customers can accordingly provide a coordination means for a greater efficiency. This is especially important with the increased complexity introduced by the use of renewable energy sources. Being dependent on weather conditions, generation becomes highly stochastic and yields unbalance between supply and demand.

Microgrids represent a bottom-up approach in coping with such a complexity, by dividing the grid in sub-systems that are easier to manage. This creates networks of small generators and consumers, such as households, industry sites and villages. Accordingly, energy can be generated locally (e.g., using a photovoltaic system) and its usage can be reduced or shifted in order to optimize use of local resources [Pal11]. Microgrids might operate in complete autonomy in what is called island mode, which also allows for better detecting and isolating system faults. As such, microgrids provide a cost-effective solution to the problem of electrification in developing countries.

However, the transition to a more sustainable energy system takes place at various levels, including the consumer side. Demand response allows for sharing with end users the uncertainty resulting from the employment of renewable sources. Accordingly, a price signal is shared to reflect the balance between available supply and current demand. This is made possible by the increased temporal measurement resolution provided by modern digital meters. On one hand this information can be used to improve billing and better act in the wholesale energy market. On the other hand the information can be used by users to improve energy usage. Accordingly, demand response exploits the flexibility offered by certain devices to react to grid instability. A typical example is the one of dwellings, which account for a significant share of the overall energy usage, being quantified as the 23% by [Mcm02] in 2010. Such a consumption is expected to grow further, given the increase of 10.8% registered for the EU25 in the period 1999 to 2004 [Ber07]. In particular, the study shown in [Car13], shows the main contributors to residential consumption in the US. Using national average penetration rates, 12 appliance types are shown responsible for 80% consumption. In particular, white goods such as the fridge, the dishwasher and the washing machine demand a higher amount of energy than brown goods (e.g., TV). A possibility is to better inform users of the footprint of their daily life activities, by embedding measurement units in electrical devices. Raised awareness can lead to more informed decisions and increased efficiency. A relevant aspect is thus the investigation of feedback means to aid decision making. However, demand response via feedback mechanisms relies on the timely reaction of human decision makers. More automatic systems are necessary to assist the scheduling of electrical loads based on the current and forecasted availability of energy. The solution should be decentralised in the sense that a mediator should not be necessary for the coordination of the whole grid. This is demanded by another requirement, that of scalability, as the system should be working at both household and microgrid scale, to the whole power grid.

Multi agent systems offer a straight solution to those requirements of decentralization and automatism [Woo09]. The main concern is thus on the selection of appropriate coordination mechanisms, as well as on techniques for modeling users' preferences and power generation. This thesis addresses the problem of microgrid energy management by undertaking the design and use of energy management systems. The main objective is to provide means for an increased efficiency and reliability in presence of renewable energy sources. As we will see, this poses several challenges that need further research.

1.2 Electrical energy management

An Energy Management System (EMS) is a system of computing components that can be employed for optimizing energy resources in building environments. Typically, an EMS should be able to collect consumption information of devices, as well as monitoring local production from renewable energy sources such as photovoltaics (Fig. 1.3). Typical building blocks of energy management systems are:

- *Smart meters* offer a means to collect high resolution energy data. Such a higher resolution opens to dynamic prices, which can better reflect the available energy to keep it balanced. This information can be used by both facility managers and consumers, who can can get a better understanding of usage across activities. Utilities can use demand data to improve billing, as well as to monitor the grid and better plan future investements.

- *Smart appliances* embed a computing unit and a network interface to interact with users and other appliances. Smart appliances are aware of consumed power based on local measurement units or built-in profiles [Elm12]. To interoperate with other devices in the network, smart devices need to provide a machine-readable description of their features and properties. Applications can thus be realized controlling distributed digital sensors and actuators, which can dynamically join and leave the network.

- *Legacy electrical devices* Energy management systems need to consider the presence of non-smart devices. A possible solution is to connect sensing units to each loads to track their consumption. So-called smart outlets and smart plugs form a network of distributed sensing nodes, which normally provide also the possibility to remotely switch them (on/off). Since current market solutions do not support identification of connected loads, any processing of consumption data has to be done at application level.

- *A Gateway* is used in automation systems to supervise the whole system and bridge the local network to the wide-area network. As such, the gateway connects the private network and the internet, therefore playing a crucial role in ensuring security and privacy. Moreover, the gateway is also the point where interconnection and interoperation between heterogeneous technologies takes place. Sub-networks using specific technologies, such as automation fieldbuses and Zigbee networks, can be managed from the gateway in order to provide a uniform interface to access resources.

- *Interfaces* such as mobile terminals or specific displays are commonly employed to control the whole system.

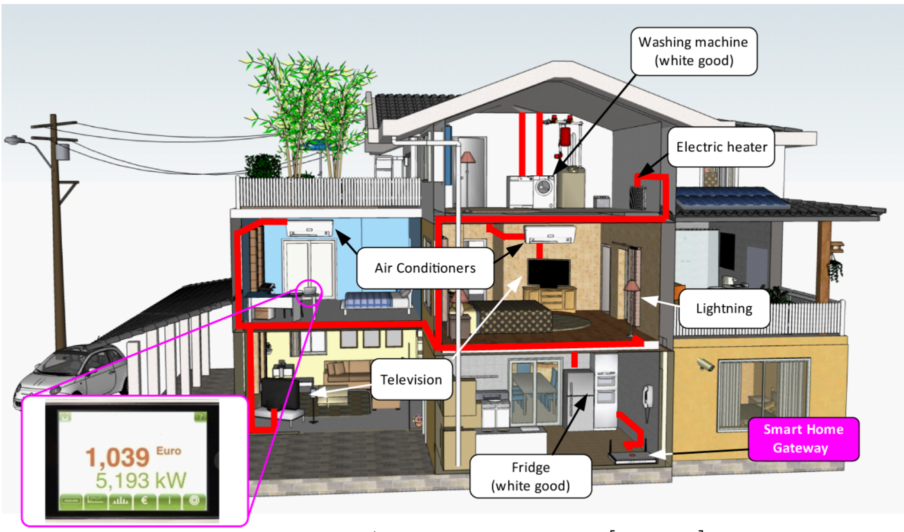

Figure 1.3: A smart building [Lak13]

1.3 Research questions

In the context of smart microgrids and buildings, this thesis undertakes three different paths, in particular:

1. **How can device and data interoperability be achieved towards the vision of a microgrid energy market?**
   This includes the identification of technologies to annotate device capabilities and exchanged data, as well as approaches for the integratation of those devices detected through load disaggregation.

2. **How can energy awareness be effectively boosted towards a more conscious use of energy resources?**
   This includes the identification of possible conservation strategies to be applied in the experiment regions. An estimation or real measurements of yielded savings should be provided to support their effectiveness.

3. **How can the design of appliance controllers be automated to optimize the use of local energy resources?**
   This includes the design of appliance controllers and the selection of appropriate market mechanisms.

1.4 Outline

The thesis started with **Chapter 1** which introduced the topic of energy management in buildings and microgrids, as well as the challenges undertaken.
Chapter 2 describes the background by overviewing existing technologies and previous work presented in the direction of this dissertation. This includes platforms for the collection and processing of energy data, as well as approaches towards the achievement of architectural interoperability. At a wider scale, interoperability can enable multiple self-* properties, according to the vision of autonomic computing. The availability of energy data opens to several other applications, such as detection and profiling of electrical loads. While this information is useful to all stakeholders: utilities, distribution network providers and consumers, this raises privacy concerns, which demand suitable solutions. On the other hand, this information should be exploitable locally to build models of demand and supply, which autonomous agents can use to minimize human intervention. We thus conclude the chapter with an overview of existing market mechanisms which can be employed for power trading.
Chapter 3 reports in detail of the measurement campaign carried out in Austria and Italy, which eventually yielded the GREEND dataset.
Chapter 4 documents our work towards architectural interoperability for building energy management systems. We introduce a multi-layer architecture, using both networked devices and the use of load disaggregation to collect energy information. Full integration of electrical devices is achieved by annotating information according to a shared semantic model. This allows for the application of logic queries on annotated data. **Chapter 5** analyzes the collected energy data to identify points for intervention. It introduces the Mjölnir web-based dashboard, which is a framework to process aggregated and disaggregated (i.e., appliance-level) consumption data. An advisor widget is designed to provide tailored feedback to energy consumers. The saving potential of the widget are finally estimated to up to 34%. Acceptance of the widget is finally evaluated in

a user test.
Chapter 6 documents our work towards automatic load management. We design a tool for the automatic design of controllers for energy prosumers. We show that because of the selected coordination mechanism, the approach might still lead to suboptimal solutions. A power broker is then introduced for pricing different power provisioning agreements. Agreements act as multiple provisioning services, constrained to different levels of quality of service.
Finally, **Chapter 7** summarizes the contribution of this work and lists aspects deserving further investigation for the future.

CHAPTER 2 Background and related work

"Verba volant, scripta manent"

– Caius Titus

The Smart Grid is a cyber-physical system whose smartness relies on the possibility to efficiently connect consumers and producers through an information channel. As such, its success heavily depends on the possibility to collect measurement data from distributed heterogeneous sources and to timely react to status changes. Demand side management promises a more efficient power grid through a better coordination of electrical loads. Accordingly, [Pal11] distinguishes in:

- energy efficiency and conservation, in which users are provided with an energy feedback to increase their awareness [Mon13c];
- time-of-use tariffs, in which the retail energy price is predefined in hourly or half-hourly intervals to meet expected peak periods;
- demand response, which includes direct control of certain customer processes by utilities, as well as voluntary response to emergency signals and price changes (i.e., indirect control). The latter can be implemented using more dynamic pricing schemes, such as critical peak pricing and real-time pricing schemes, which better reflect the wholesale market prices into the retail prices.

In this chapter, we provide a comprehensive literature work on energy monitoring systems (Sect. 2.1), namely looking at available datasets (Sect. 2.2), interoperability issues (Sect. 2.3), data management and analysis approaches (Sect. 2.4), as well as techniques to raise users' awareness (Sect. 2.5) and automate energy management (Sect. 2.6).

2.1 Metering systems for energy management

Building energy management is achieved through the collection of energy consumption and production data. A digital meter is an electronic device recording energy usage at regular intervals, to be processed for improving decision making of both humans and autonomous controllers. Before going to further detail, it is important to firstly recall the physical quantities a meter deals with: i) the voltage espressed in Volts, ii) the current in Amperes (A) and iii) the phase shift between them ϕ_{ui}. To handle those measures, the voltage can be lowered with a voltage divider and simply fed into an Analog-to-Digital Converter (ADC). The current (i.e., quantity of charge per second) can be measured using a Hall-effect sensor or a current transformer, which base their functioning on the magnetic field generated on the conducting wire. The phase shift $\phi_{vi} = \phi_v - \phi_i$ is the time shift between the measured voltage and current and can be estimated through numerical methods. From these parameters we can distinguish in three different quantities: the active power in Watts (W), the reactive power expressed in VAR and the apparent power in Volt-Amperes (VA) (see Fig. 2.1):

$$P = V_{RMS} \cdot I_{RMS} \cdot cos(\phi_{vi}) \quad \text{(active power)}$$

$$Q = V_{RMS} \cdot I_{RMS} \cdot sin(\phi_{vi}) \quad \text{(reactive power)}$$

$$S = V_{RMS} \cdot I_{RMS} \quad \text{(apparent power)}$$

The root mean square (RMS) of a time-varying function, such as voltage and

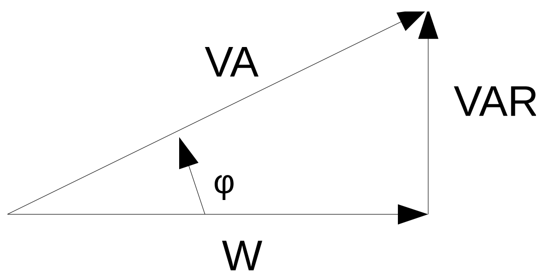

Figure 2.1: Relationship between power quantities

current, can be computed by dividing the peak value by the the crest factor. The crest factor is a signal-specific property. For instance, for a sinusoidal signal it is $\sqrt{2}$. Data collection can be carried out directly accessing digital meters or employing networks of distributed measurement units. In this dissertation we focus on the second approach, for which various solutions are available in the market (See Table 2.1). The main aspect to consider is the kind of features measurable and the highest sampling frequency achievable. According to the Nyquist-Shannon sampling theorem, in order to reconstruct a sampled signal and avoid aliasing, the sampling frequency should be not less than twice the highest frequency (i.e., $f_s > 2f_{max}$). Another aspect to remark from Table 2.1 is the availability of a Software Development Kit (SDK) or a documented Application Program Interface (API), as most of commercial solutions provide closed-source software applications to manage the system. Most of those employ a 2.4 GHz Zigbee connection, while only one, to the best of our knowledge, uses Powerline Communication (PLC) technology. In [D'A14] available solutions are discussed and network constraints are derived for a building energy management system.

Table 2.1: Commercially available measurement systems [D'A14]

Manufacturer	**Comm.tech.**	**Sampling freq.**	**SDK/API**	**SW**	**Price/SP**
Pikkerton	Zigbee 2.4 GHz	some secs	yes	no	100 €
Plugwise	Zigbee 2.4 GHz	1 h	no	yes	30 €
4-noks	Zigbee 2.4 GHz	-	yes	no	40 €
ThinkEco Modlet	Zigbee 2.4 GHz	-	yes	yes	100 €
FlexGrid	Zigbee 2.4 GHz	-	no	yes	85 €
CurrentCost	Wireless Proprietary 433 MHz	1 m	no	yes	18 €
SLSEnergy	Powerline Proprietary 115-132 kHz	30 s	no	yes	38 €

2.2 Energy-usage datasets

To evaluate solutions working for real scenarios, research on energy and sustainability relies on publically available datasets. Table 2.2 classifies energy-usage datasets using as classification attributes the sampling frequency and the characteristics of the signal being measured, such as active power (P), reactive power (Q), apparent power (S), energy (E), frequency (f), phase angle (Φ), voltage (V) and current (I). Collected data reflects the differences required on target applications. As visible, certain datasets (e.g., REDD [Kol11] and BLUED [And12]) only monitor a small number of households at a high sampling frequency. In load disaggregation, more representative features can be captured when sampling at a higher frequency [Arm13]. Another aspect is the location. Many datasets were collected in the USA, where voltage is $120V$ rather the 230V commonly used in Europe. The setting in which collection takes place greatly affects behavior and should therefore be as much representative as possible. The appliance tracks collected in TraceBase [Rei12] and ACS-F1 [Gis13] do not provide any relationship to the consumption scenario, which makes any behavioral analysis impossible. Besides, location determines weather and climate. Another aspect is the type and number of monitored devices and households. Statistical analyses are only possible with higher number of households, such as in HES and OCTES. Moreover, seasonal behaviors can only be captured by long term campaigns. Also, certain datasets monitor households over shifted time windows, which makes a comparison impossible.

Table 2.2: Existing datasets for energy consumption in households

Dataset	Location	Duration	#Houses	#Sensors (per house)	Features	Resolution
ACS-F1 [Gis13]	Switzerland	1 hour session (2 sessions)	N/A	100 devices in total (10 types)	I, V, Q, f, Φ	10 secs
AMPds [Mak13]	Greater Vancouver	1 year	1	19	I, V, pf, F, P, Q, S	1 min
BLUED [And12]	Pittsburg, PA	8 days	1	Aggregated	I, V, switch events	12 kHz
GREEND	Austria, Italy	1 year	8	9	P	1 Hz
HES	UK	1 month (255 houses) - 1 year (26 houses)	251	13-51	P	2 min
iAWE [Bat13]	India	73 days	1	33 sensors (10 appliance level)	V, I, f, P, S, E, Φ	1 Hz
IHEPCDS[1]	France	4 years	1	3 circuits	I, V, P, Q	1 min
OCTES[2]	Finland, Iceland, Scotland	4-13 months	33	Aggregated	P, Energy price	7 secs
REDD [Kol11]	Boston, MA	3 - 19 days	6	9-24	Aggregate: V, P; Sub-metered: P	15 kHz (aggr.), 3 sec (sub)
Sample dataset[3]	Austin, TX	7 days	10	12	S	1 min
Smart* [Bar12]	Western Massachussets	3 months	1 Sub-metered +2 (Aggregated + Sub-metered)	25 circuits, 29 appliance monitors	P, S (circuits), P (sub-metered)	1 Hz
Tracebase [Rei12]	Germany	N/A	15	158 devices in total (43 types)	P	1-10 sec
UK-DALE [Kel14a]	UK	499 days	4	5 (house 3) - 53 (house 1)	Aggregated P, Sub P, switch-status	16 kHz (aggr.), 6 sec (sub.)

[1] http://tinyurl.com/IHEPCDS
[2] http://octes.oamk.fi/final/
[3] http://www.pecanstreet.org/projects/consortium/

2.3 Device and data interoperability

The interoperability problem involves the whole smart grid system, which is lacking widely accepted standards [Gun11]. An effective coordination of energy consumers through demand response demands coping with the heterogeneity of such networked components. In particular, the system must deal with distributed resources, built by different manufacturers using different technologies. For instance, a building energy management system can include smart appliances [Elm12] using different communication protocols and data formats, as well as devices with no connectivity at all.

2.3.1 Distributed computing paradigms

In distributed computing communication between distributed components can be commonly distinguished in Remote Procedure Call (RPC) and Representational State Transfer (REST). Remote procedure call enables processes to call procedures located on different address spaces (e.g., remote servers). This is possible thanks to a *stub* component that converts and serializes requests (marshalling) in order to hide network mechanisms and structure to application developers. SOAP-based webservices follow this approach. In RPC services the sender perform a request along with parameters specified as payload. The server receives the request and performs the procedure call using the given parameters. This means that the service client sends part of its state to the server. The Java Remote Method Invocation (RMI) implements this approach for object-oriented environments. On the contrary, REST services are collections of representations of resources, which are addressable using unique identifiers (i.e., URI). A client can request resources to a server using a standardized interface (e.g., HTTP).

2.3.2 Device modeling and discovery

In service-oriented architectures (SOA), component-level interoperability and configurability can be achieved as a composition of loosely-coupled computing components. Specifically, this takes place by separating the component interface from implemented functionalities. Interoperability is firstly achieved at the message level by relying on standardized data formats such as XML and JSON. Binary variants have also been proposed to further reduce the overhead introduced for parsing those formats on memory-constrained embedded devices. For instance, the Efficient XML Interchange (EXI)[1] was shown providing the

[1]http://www.w3.org/TR/exi/

highest compression and efficiency to process XML messages [Sak09]. For JSON, its counterpart is the binary JSON format[2].

This allows smart devices to provide a machine-readable description of their features, that can be made available to the other devices across the system. Services are described in terms of their I/O interface: i) possible operations, ii) constraints on data given and iii) communication protocol. As indicated, this allows for the encapsulation of the service complexity and simpler management mechanisms. Applications can be built as composition of those distributed services, using i) service orchestration or ii) service choreography.

An orchestrator controls the dependencies between each specific service by means of a process workflow. The Business Process Execution Language for Web Services (BPEL4WS) is a standard for XML/SOAP services. On the contrary, in a choreography the global workflow is split down to local rules to be implemented directly by the individual services. While coordination can be directly hard-coded in the service definition, this approach does not scale to larger networks. The Web Services Choreography Interface (WSCIspec) and the Web Services Choreography Description Language (WS-CDL) are standardised choreography definition languages that can be used to automatically generate dependency rules given the overall business process.

To enable resource sharing, service-oriented architectures provide the means for resource naming [Tan06]. [Dar09] distinguishes in three generations of naming systems, i) *name services* (e.g., DNS) which resolve names to entities and possibly basic attributes, ii) *directory services* (e.g., LDAP) which also provide more complex attribute-based queries so as to retrieve entities satisfying certain attributes, and iii) *service discovery systems*. The first two can work fine in stable environments, such as wired networks, as they require human configuration and management. Contrarily, service discovery can reduce human intervation by dynamically detecting nodes joining or leaving the network. To a certain extent, this provides self-configuration and self-healing properties [Kep03]. In particular, discovery can be implemented using broadcast and multicast (e.g., ARP, UPnP), using shared service registries (e.g., DNS, LDAP, UDDI) or exploiting logical overlays such as with distributed hash tables (e.g., Chord, Kademlia) [Dar09].

It is important to remark that varius communication standards can be used for building automation, opearing with powerline (e.g., X10[3], HomePlug[4] and

[2] http://bjson.org/
[3] http://www.x10.com
[4] http://www.homeplug.org

LonWorks[5]) or wireless (e.g., ZigBee[6], Z-Wave[7], Bluetooth[8]) communication. However, a clear transition to IP-based technologies is ongoing [Vil13]. For instance, IETF Zeroconf[9] and the service location protocol (SLP)[10] are service discovery mechanisms for IP networks. Component-oriented interoperability has been advocated in various middlewares such as Jini, Open Services Gateway initiative (OSGi)[11], Home Audio/Video interoperability (HAVi), Universal Plug and Play (UPnP)[12] and Device Profile for Web Services (DPWS)[13].

2.3.3 Data modeling

While device description and discovery can be achieved with the use of service-oriented architectures, an open issue to be discussed is the readability of exchanged data. To this end, [Hea11], [Pfi11] and [Mon13a] suggest to semantically annotate sensor data using principles from the semantic web and the linked data initiatives. The core concept is to describe data in terms of their relationships, by using the Resource Description Framework (RDF) data model. Thereby, the basic information unit is a ⟨*subject*, *predicate*, *object*⟩ triple. In this way, data can be related to concepts defined in shared vocabularies, so that all entities sharing the same definition can interpret the data the same way.

For instance a washing machine could be described in Turtle[14] as:

```
@prefix ns: <http://myrepository.com/houses/12345/> .
@prefix en: <http://example.com/ontologies/appliances.owl#>
    .

ns:washing-machine en:model          "XYZ123456" ;
                   en:manufacturer   "Bob Inc." ;
                   en:type           en:washing-machine ;
                   en:consumption    "800" .
```

In order to achieve interoperability at this abstraction level, a possibility is to store the description on a server running on the networked device. In this way, when accessing the appliance at a specific URI (e.g., “http://myrepository.com/houses/12345/washing-machine”) a human-readable

[5] http://www.echelon.com
[6] http://www.zigbee.org
[7] http://www.z-wave.com
[8] http://www.bluetooth.com
[9] http://www.zeroconf.org
[10] https://www.ietf.org/rfc/rfc2608.txt
[11] https://www.osgi.org
[12] http://www.upnp.org
[13] http://docs.oasis-open.org/ws-dd/ns/dpws/2009/01
[14] Terse RDF Triple Language, a RDF serialization format. `http://www.w3.org/TeamSubmission/turtle/`

or a machine-readable representation can be returned. This mechanism relies on headers[15] to retrieve different representations of the same resource. To reduce latencies due to the resolution of URIs and the retrieval of descriptions, another possibility is to collect offline all descriptions in a centralised storage, such as a triple store [Mon13a]. Various technologies have been proposed towards this vision. The IPv6 over Low power Wireless Personal Area Networks (6LoWPAN)[16] aims at bringing IPv6 connectivity to resource-constrained embedded devices. Similarly, the Constrained Application Protocol (COAP) offers a lightweight version to HTTP, by employing UDP along with the CoRE link format (RFC6690)[17] to describe resources by their links. In detail, a node can send a POST for a link to its provided resources to the index "/.well-known/core" mantained on the selected directory node. This way, any client accessing the index can discover resources in the subnetwork.

The definition of shared vocabularies is thus central to the interpretation of exchanged data. Ontologies offer a formal conceptualization of a specific domain in axiomatic terms. Moreover, they can be integrated and extended to support larger domains. For instance, an energy management system for microgrids and buildings might require the definition of:

- *Building information* includes modeling of the dwelling, such as building geometry and insulation information [Kof13a].
- *Building automation* and device description includes service orientation [Sta12, Pre04] and building automation [Bon08].
- *User information and preferences* includes living processes in terms of appliance operation and system settings (e.g., for thermal comfort) [Kof13b].
- *Energy management* includes modeling of energy generation, management (cf. [Rei11]) and optimization through rule-based reasoning [Tom10].
- *Weather and climate modeling* includes modeling of both climate and weather conditions, as well as weather forecasts [Kof12].
- *Measurement units and sensors* include sensor modeling [Com12], physical phenomena [Ras04], measurement units, as well as geographic information.

The main benefit from using such knowledge engineering techniques is the possibility to built an abstraction over data collected from diverse sources. In addition, specific tools for retrieval of such data became available. The SPARQL

[15]http://www.w3.org/TR/cooluris/
[16]http://tools.ietf.org/wg/6lowpan/
[17]http://tools.ietf.org/html/rfc6690

Protocol and RDF Query Language (SPARQL) is a query language that can handle data modeled using the RDF framework. In detail, SPARQL allows for the addition, deletion, and modification of data triples [DuC11]. SPARQL can also be used to perform complex event processing by defining rules and constraints over data using the SPARQL Inferencing Notation (SPIN). While SPARQL was designed for static networks, such as the web, data collected in real environments tend to be highly dynamic and demand for different query languages able to tackle such a volatily. Various alternatives have been proposed: C-SPARQL [Bar10], SPARQLstream [Cal10], EP-SPARQL [Ani11], and CQELS [LP11].

2.4 Meter data management

The higher measurement resolution made available by the progressive rollout of smart meters allows for data analysis. On one hand this can benefit both utilities and distribution network providers towards a more efficient network. On the other hand customers can receive value-added services on top of their power provisioning plan.

2.4.1 Load disaggregation and detection

Monitoring and integration of energy usage data can take place by either measuring individual devices with a network of distributed monitoring units or detecting operating devices from aggregated measurements. Consequently, we distinguish Intrusive Load Monitoring (ILM) and Non-Intrusive Load Monitoring (NILM).

ILM requires monitoring each device with a sensing unit, such as a smart outlet or plug, with higher setup and mantainance costs. Device identification can either be solved by humans by indicating the type of each connected device, as well as inferred from collected data. Identifying loads requires modeling the operation dynamics of electrical devices, which can be either statically provided or extracted online from measurements [Don13]. For instance, [Par12] uses a general appliance model of refrigerators.

Non-intrusive load monitoring is a single-meter approach, as it aims at identifying running loads from the overall power consumption data. The approach, firstly introduced by Hart in [Har92], recognizes operating devices according to specific patterns on their power profile. Load disaggregation can be seen as an optimization problem where given the total power consumption and a database of known power profiles a composition is found to best best approximate the measured overall power [Lia10, Ega13, Suz08]. Disaggregation algorithms can be distinguished in supervised and unsupervised approaches, depending on the necessity to use labeled data to train the classifier. In particular, supervised techniques include Bayesian approaches [Zei12], artificial neural networks [Sri06] and support vector machines [Lin10]. The main disadvantage is the need of labeled data during a training phase. This implies greater effort and development costs. More recent unsupervised load disaggregation algorithms, such as those based on k-means clustering [Gon11], Factorial Hidden Markov Model (FHMM), and its variants [Kol12, Zai10, Kim11, Zoh13], promise to overcome this issue. The effectiveness of existing NILM algorithms strongly dependends on the employed sampling frequency. Accordingly, higher sampling frequency can provide more representative characteristics, which results in a more accurate appliance classification. [Arm13] suggest that approximately 10 different appliances can be detected with a sampling frequency of seconds. In particular,[Zei12] indicates the following requirements:

- power measurements at $1Hz$;
- a minimum acceptable accuracy between 80 and 90%;
- no training should be necessary;
- real-time appliance detection;
- possibility to detect between 20 to 30 appliances;
- should handle various appliance types, such as binary and multi-state devices, continuous appliances, as well as permanently operating ones [Zei11].

The design of a load disaggregation component challenges the research community with multiple questions, among which:

1. the dependence on the Signal-to-Noise Ratio (SNR) of power measurements;
2. the sensitivity with respect to the sampling frequency;
3. the sensitivity with respect to the given appliance models, which includes: similarity between different device types, as well as differences between differently-manufactured ones;
4. the necessity of an appliance dataset for modeling purposes, and the sensitivity of the resulting models to the completeness of the dataset;

5. the sensitivity to different usage modalities for the same devices, such as differences on time and duration of the operation;

6. the computational complexity and hardware costs entailed by the selected algorithm.

A complexity measure for load disaggregation is further discussed in [Ega15b] and [Poc15].

2.4.2 Usage profiling

The availability of fine-grained energy usage data allows for mining usage profiles, which can be used for instance to improve control strategies and feedback.

Lately a few works were proposed to infer occupancy from energy consumption data, as in [Che13, Kle13]. The main concept is to monitor consumption of user-driven devices to determine presence. In [Che13] threshold values are computed during inactivity periods, i.e. when only baseline loads such as fridges are operating. Most commonly, the thresholds are computed during the night when people are sleeping. This has the disadvantage of assuming inactivity as unoccupation and does not consider that smart appliances can be automatically scheduled to run over the night. In [Mon14a] we applied the algorithm on two months of power measurements (February and March 2014) to build weekdays and weekend occupancy models. We observed a good agreement between the estimated probability of occupancy and the habits of residents during the day. However, as expected the occupancy model shows inactivity during midnight and 6.30 AM.

Appliance usage mining concerns the extraction of models describing the use of electrical devices from event logs. This has been achieved through various techniques, such as association rule mining [Kan12], artificial neural networks (ANNs) [Ayd02], episode-generating Hidden Markov models (EGH) [Tru13] and Bayesian networks [Agu11]. In [Mon14a] we showed that the starting probability for a coffee machine can be easily modeled from a log of starting events (See Fig. 2.3). A good agreement with the residents activities was noticed, especially concerning the wake-up patterns which appear delayed in the weekend days.

2.4.3 Privacy aspects

The use of cyber-physical systems raises privacy concerns due to their effects on the physical world. Specifically, entities operating on user's behalf might be manipulated to operate differently.

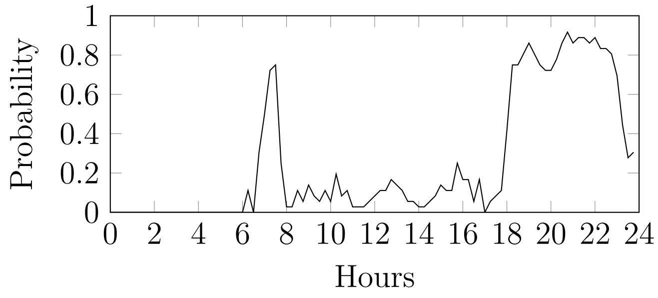

(a) Weekday occupancy probability

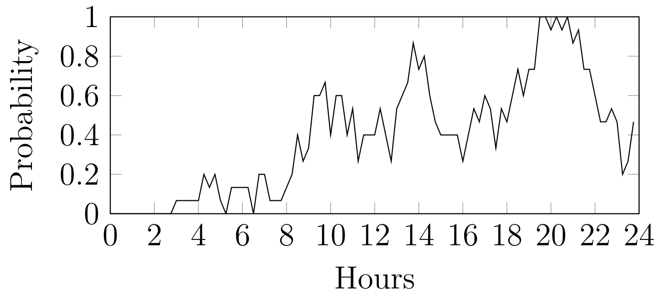

(b) Weekend occupancy probability

Figure 2.2: Occupancy model example [Mon14a]

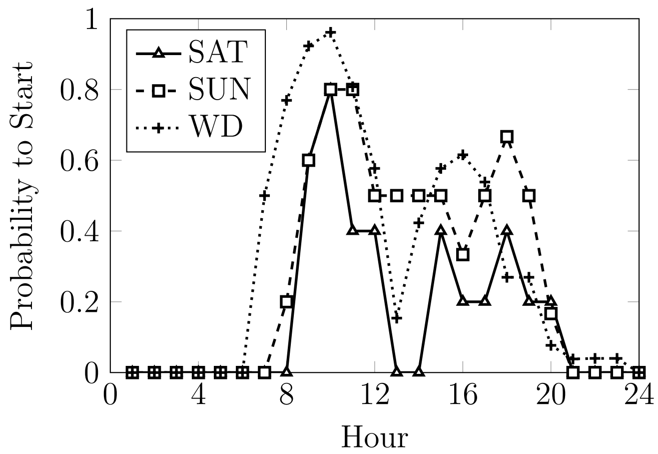

Figure 2.3: Usage forecasting for a coffee machine [Mon14a]

Also, handling energy information might result in theft or alteration. Accordingly, the availability of energy usage data indicate time and modality of use of devices, which can allow for the extraction of usage patterns [Ngu13, Lis10]. This has effects on multiple stakeholders such as utilities, marketing agencies, press and even criminals [Sko12]. For instance, [Gre12] showed the possibility of detecting the TV show being watched based on power consumption. Load hiding techniques emerged as a solution to obfuscate the overall power demand in order to ensure privacy. The concept is to operate controllable batteries or energy-intensive loads (e.g., water boilers) to reshape the measurable overall power demand of the building. Therefore, we distinguish between Battery-based Load Hiding (BLH) and Load-based Load Hiding (LLH) [Ega14].

2.5 Energy awareness

Energy awareness denotes the ability to understand the monetary, social and ecological impact behind the operation of electrical devices. While energy invoices return consumption information, they are normally sent out with significant delay from the actual energy use. This way, received consumption information is a too coarse-grained and late feedback to have any effect on the consumer's decision making. Prepaid billing is a possible way to increase the feedback resolution, and it was shown leading to average savings of 11% in UK, regardless of disconnections from the grid [Ozo13]. Greater efficiency can generally be achieved by i) replacing devices with more efficient ones, ii) improving the building efficiency, and iii) optimizing energy usage. To spot energy hogs, a possibility is to conduct an audit of energy usage, by means of surveys and interviews with facility managers or end users. However, smart meter data is going to provide higher resolution data, therefore offering the means for large scale automated energy audits. For instance, [Bec13] showed traces from more than 3000 households being used to extract specific customers' properties.

In demand response programmes, dynamic pricing schemes are used to incentivize load operation when the demand is lower. The effectiveness of this mechanism depends greatly on the possibility to timely inform users on their energy usage, in order to promote a better exploitation of local resources. Ambient interfaces, such as the power-aware cord [Gus05], are often employed as an unobtrusive feedback means. Another aspect is the selection of tariff plans that best suit with the usage behavior, as in the AgentSwitch [Ram13]. An evaluation on 10 users carried out for 3 months showed the system being effective in finding cheaper tariffs for most users [Fis13].

[Dar06] classifies feedback in two categories:

- *indirect*, when it provides consumption information after it occurred;
- *direct*, when the feedback concerns the amount of energy in use

Darby shows also that real-time consumption information can effectively raise user awareness and lead to up to 15% usage reduction. Indirect information is instead necessary to enable learning mechanisms towards a long-term change. Similarly, [Bon12] identify *antecedent* and *consequent* strategies. Antecedent strategies aim at preventing certain behaviors, for instance using goal-setting and advices, while consequent strategies concern direct and indirect feedback, which also includes monetary and social rewarding. However, studies have also shown that in spite of awareness, the effectiveness of these systems in making people responsible depends on their sensitivity and motivation [Str11]. The analysis in [EM10] analyzes 36 studies carried out between 1995 and 2010 to show that consumption information at device level can lead to the highest energy savings (see Fig. 2.4). Most effective feedback mechanisms were estimated

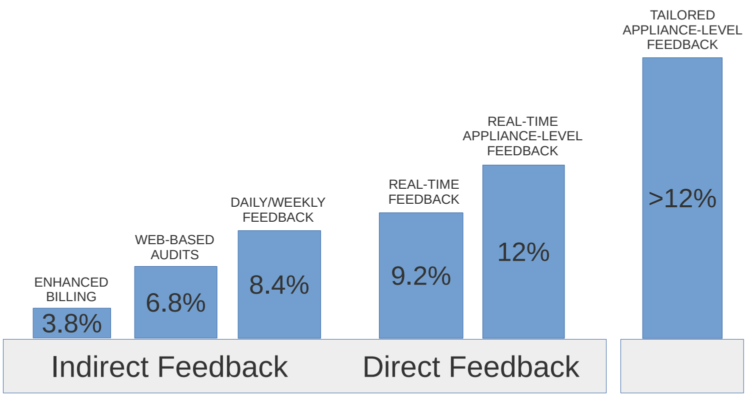

Figure 2.4: Effectiveness of feedback [EM10, CA13]

leading to around 20% savings [CA13], by exploiting a user model for offering tailored services, such as advices.

2.6 Autonomous systems

The smart grid is a system of systems, dealing with decentralised resources situated in highly dynamic environments. This comes with increased complexity in management and mantainance. A solution is to devolve a certain level of autonomy in order to reduce human intervention and minimize costs and risks, as well as increase security and availability. Autonomy is in fact present in nature. Typical examples are certain chemical processes, animal swarms and neural networks. In so called self-organizing systems, an overall order or coordination emerges from the local interaction of smaller components, without any external control [Ash47]. Being fully decentralized, these systems benefit from multiple properties, for example robustness to single-point failures, recoverability to damages and perturbations, scalability and adaptability.

2.6.1 Organic and autonomic computing

Biologically-inspired approaches, such as IBM's autonomic computing [Kep03] and organic computing [MS04], intend to tackle such a complexity by imitating natural processes of adaptation to endogenous and exogenous change. For instance, autonomic computing derives its principle from the autonomous nervous system, which uses multi-level feedback loops to react to stimuli while seeking higher-level objectives [Bru09]. Awareness of the monitored environment is achieved by periodically collecting sensor data. Self-awareness is achieved to adapt the controller's behavior according to the availability of internal resources and information.

A self-managing system exhibits one or more of so called self-* properties [Dar09]:

- *Self-configuration:* the system is able to adjust its internal parameters and behavior to adapt to dynamic environment changes. This property allows for continuous adaptation to unpredictable conditions, which is the case in stochastic environments. System administrators can thus specify high level policies or objectives without having to worry about specific actions undertaken by the system.
- *Self-optimization:* the system is able to optimize its operation in order to meet predefined goals while adapting to changes in available resources. The system can thus strive for the best result given its available resources and goals.
- *Self-healing:* the system is able to detect failures in its components and repair them in order to maximize its availability.
- *Self-protection:* the system is able to identify and prevent malicious accesses in order to ensure integrity, privacy and security.

This is a typical ability of living organisms, as they are able to adapt to different habitats by regulating the activity of internal organs, in order to maintain certain vital parameters and ensure the organism's survival. Beside these four core properties also called self-chop properties, researchers proposed further characteristics that autonomous systems should endow (Table 2.3). In autonomic computing, a controller is implemented as a loop of four tasks: monitoring, analyzing, planning and action execution, which result in the thus called MAPE-k architecture (Fig. 2.5). A shared knowledge base is used to both infer context information from heterogeneous sources and store system policies, which guide the planning process.

Table 2.3: Self-* properties [Pos09]

Self-* property	**Description**
Self-management	System manages itself without external intervention (Self-*)
Self-description	System can be understood by humans and other systems without further explanation
Self-configuration	System components automatically adapt to achieve high-level goals
Self-interest	System aims at pursuing its own goals
Self-monitoring	The system can retrieve information of its components and global status
Self-awareness	System knows its internal components and resources
Self-organisation	The system is formed via the the decentralised assembly of self-contained components and its structure is driven by certain models
Self-creation	System's members participate at the creation of the system according to social models
Self-optimisation	System monitors itself to improve its performance or efficiency
Self-protection	System detects and defend itself against malicious attacks
Self-healing	System detects and repairs hardware and software failures
Self-learning	System uses learning techniques to improve its performance
Self-regulation	System mantains a certain parameter
Self-evolution	System shows an emergent behaviour that arises from local interactions between its components

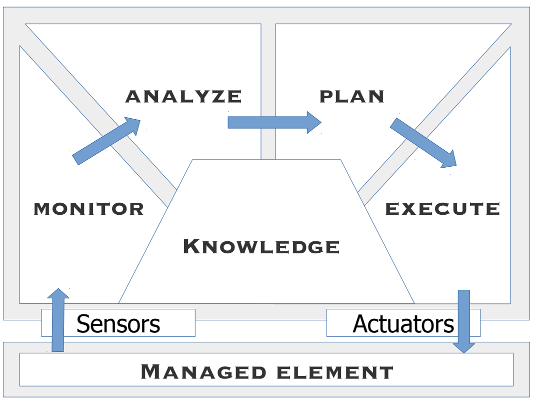

Figure 2.5: The MAPE-k architecture [Kep03]

2.6.2 Designing self-organizing systems

In spite of the advantages provided by self-organizing systems, addressing the design through a classic top-down approach is often not possible, because global properties emerge from bottom-up processes. A way to overcome this is to observe natural self-organizing systems to acquire insights that can be transferred to the technical domain [Flo08]. Specific examples are the ant colony optimization for solving the traveling-salesman problem [Dor96], the stigmergic pheromone laying behavior of ants being used for replica distribution [Sob13], as well as the firefly clock synchronization algorithm [Buc88].

However, imitating natural processes can not always provide solutions to technical problems. As also identified by [Feh13], a more general approach is to follow the process that led to the emergence of a certain pattern in natural systems: evolution. Evolutionary algorithms work on a pool of candidates, each described in terms of a genotype encoding a specific behavior [Eib03]. Each candidate of the population is a different solution to the targeted optimization problem. A fitness measure is used to assess the quality of candidates and rank them. The evolution is actuated by selecting best candidates for further generations, by applying elitarism, mutation and recombination operators. The main advantage is their universality, as they can be applied to a very wide set of problems and candidate representations. No problem-specific knowledge is necessary for the candidates, except for the fitness function evaluating the task outcome. Entire fields of research, such as evolutionary robotics [Nol00], base their functioning on artificial evolution. As indicated by [Tri11], the four necessary elements for the evolvability of robot controllers are: the ecology (i.e., the environment model), the sensory-motor system, the genotype-to-phenotype mapping, and the fitness function.

2.6.3 Artificial neural network controllers

An Artificial Neural Network (ANN) is a network consisting in nodes, called neurons, connected by weighted links called synapses. Synapses are used to propagate signals. In particular, we commonly distinguish in an input and an output layer of neurons, respectively receiving and providing signals from the environment. Internal neurons who are not directly connected to the environment constitute the hidden layers.

In its simplest model proposed by [McC43], the output of a neuron y_i is a sum function Φ of all incoming signals x_j, each weighted according to the connection strength w_{ij}. This translates into $y_i = \Phi(\sum_j^N w_{ij}x_j)$. An activation function θ_i is also included in the neuron function. Typical activation functions are:

$$\Phi(x) = kx \qquad \text{(linear function)}$$

$$\Phi(x) = \begin{cases} 1, & \text{if } x > \theta \\ 0, & \text{otherwise} \end{cases} \qquad \text{(step function)}$$

$$\Phi(x) = \frac{1}{1 + e^{-kx}} \qquad \text{(logistic function)}$$

$$\Phi(x) = tanh(x) = \frac{2}{1 + e^{-2x}} \qquad \text{(TanH function)}$$

$$\Phi(x) = max(0, x) \qquad \text{(rectifier function)}$$

$$\Phi(x) = \frac{e^{x_i}}{\sum_j e^{x_j}} \qquad \text{(softmax function)}$$

indicating the linear function, the step function, the logistic or sigmoid function, the hyperbolic tangent, the ramp function and the softmax. The linear function does not change the neuron output and is usually used for regression problems. On the contrary, the softmax is commonly employed in the output layer for classification problems. The softmax performs a sort of normalization over the output layer, consequently providing the probability associated to each class. The sigmoid and the hyperbolic tangent are a very common choice for feedforward neural networks, with the latter outputting values in $[-1, 1]$. However, latest research in deep architectures (i.e., with more than 3 hidden layers) tend to employ the rectified linear unit (ReLU) as the activation function of choice for hidden layers. The improved training results can be partially due to the linear and non-saturating nature of the function [Glo11]. Depending on the direction of propagation of signals, we can distinguish in feedforward and recurrent architectures. The simpler feedforward networks only allow for the learning of input-output reactive functions. On the contrary, recurrent architectures allow the existence of loops (i.e., connections from later layers to previous ones or from the same layer), thus creating bidirectional propagation of information. Having a much more complex temporal dynamics, the learning of those networks is also more complex [Pas13]. However, a Recurrent Neural Network (RNN) with sigmoidal activation function was shown in [Sie91] being Turing complete. In particular, RNNs can retain an internal state or context, which makes them especially effective in processing sequential input [Lip15].

Learning appropriate input-output mappings takes place by adjusting the weights on the synapses. This can be automated by continuously updating the weights while a sample set of patterns is presented to the network, such as in the backpropagation algorithm.

2.6.4 Neuroevolution

Evolutionary algorithms are often used in evolutionary robotics to learn artificial-neural networks. The synaptic weights are directly encoded as the genotype for each candidate and progressively adjusted according to the application of mutation and recombination operators. The main advantage of evolutionary algorithms (EA) over other supervised learning methodologies, such as back-propagation, is the possibility to train a network in absence of a training set. Accordingly, EAs work on a pool of candidates, each described in terms of a genotype encoding a specific behavior [Eib03]. The population is initially filled with random candidates. Each candidate represents a different solution to an optimization problem. By running each candidate in the population, it is possible to compute a fitness representing its quality for solving the problem. Elite candidates are those having the highest fitness and can be kept for further generations. To ensure diversity, the next generation also has space for some randomly selected candidates. To possibly find new and better solutions, mutation is applied to the selected candidates. This takes place by randomly modifying the agent's genotype. In addition, recombination is used as a reproduction mechanism which combines pairs of candidate solutions into a set of offspring candidates.

There are several reasons to prefer ANNs over other representations [Nol00]:

- they act as non-linear function approximators that can deal with both discrete and continuous signals to interact with the environment;
- they are robust to noise, as oscillations in the input signals do not drastically affect the network behavior;
- they offer a smooth search space, as gradual changes to their structure should correspond to gradual changes in behavior;
- evolutionary algorithms can be applied at different abstraction levels, such as the synaptic weights or the coordination of multiple networks
- they represent a metaphor for biological learning processes.

While this sounds nice in theory, designers will have to face the complexity of real scenarios, by copying with the following issues:

- **genotype-to-phenotype mapping**, for which changes to the controller representation does not imply a 1:1 change to the controller's behavior. In practice, this will render the fitness landscape less smooth.

- **selection of an appropriate fitness function**, which has a crucial influence on the final task. The selection is difficult because of the various constraints to consider in real world problems;
- **time and memory efficiency**. Being a population-based method that simulates repeated interaction with the environment, the learning process can be very computational intensive and last long time.
- **gap between simulation and real environment.** While the learned controllers might work well for the learning scenario, the end environment will certainly behave differently, because of its great stochasticity. Clearly, one can make the evolved controller less sensitive to certain features by varying them during the evolution. Typical approaches for producing more general solutions are incremental evolution and re-adaptation (see [Nol00]).
- **scarce understanding of internal working.** Being learned through a self-organizing process, the overall network behavior can not be inferred by analyzing the individual connections.

2.7 Automating energy management

The vision of a smart grid combining multiple microgrids raises questions related to the control of such networks. In particular, the availability of a dynamic price signal demands for agent-based solutions capable of negotiating resources to coordinate the networks.

An early work on scheduling household appliances using computational markets was presented by Ygge and Akkermans [Ygg96]. Towards this vision, [Pal11] surveys demand response solutions, while open challenges towards the employment of intelligent agents are discussed in [Ram12, Kok05, Lam10]. Energy trading has been implemented using various mechanisms [Saa12], such as cooperative games [Ala13], as well as based on cost-minimization and non-cooperative games [Kok05, Adi14, Cha14b, MR10] especially double-sided auctions [Vyt10, Ili12, Wan14]. As opposed to centralised scheduling based on optimization, market-based approaches can better deal with self-interested agents competing for scarce resources [Cle96]. A shared price balances the demand and supply in the system, therefore acting as congestion management mechanism that coordinates the individual agents. Agents can keep their preferences private and act solely based on their local view of the environment. Before going to further detail, it is important to distinguish in two different types of markets: i) a wholesale electricity market in which generators compete to supply their output to retailers and ii) a retail electricity market in which end-use customers can select their supplier from a pool of competing retailers. The "AMES Wholesale Power Market Test Bed" is a market simulator, in which energy markets are held to allocate hourly energy intervals [Li09]. As for retail markets, latest research has also targeted the consideration of appliance usage behavior and preferences. To plan unobtrusive control strategies, usage preferences are considered to keep the produced discomfort low. In particular, [Bap11] introduces a system to determine preferred time of use of appliances to minimize running costs and activity disruptions. Later works, such as iDR [Cha14a], DR-Sim [Wij13] and the HEMS simulator [Mon14c], analyze consumption behavior and preferences to reduce the discomfort produced by control strategies.

2.7.1 Auctions

Electronic markets provide a framework for the allocation of limited resources in communities of self-interested agents [Cle96]. In particular, pricing allows for adaptive control of distributed resources, leading to the emergence of global coordination of competitive entities. The result is optimal locally, i.e., according to the individual agent's utility, as well as global, i.e. in terms of social welfare. Market-based allocation takes place through auctions. Generally, auctions can be classified in [Par11]:

single-dimensional and multi-dimensional In a single-dimensional auction the offered price is the sole attribute used for ranking the bidders, whereas in multi-dimensional ones other quality measures are employed.

one-sided and two-sided One-sided auctions are characterised by only a type of bidders, be that buyers or sellers, and a central auctioneer determining the winner. Auctions with multiple sellers are called procurement auctions. In two-sided auctions, the auctioneer matches offers coming from multiple buyers and sellers.

open-cry and sealed bid In open-cry auctions all bidders are aware of every other bid, whereas in sealed-bid auctions only the auctioneer has global visibility on offers.

first price and kth price In first price auctions the winner pays the price of the winning bid, whereas in the kth price auction the bidder pays the price of the bid ranked kth. The kth pricing mechanism is used to implement incentive compatibility, such as in the Vickrey auction in which the winner pays the 2nd highest bid. The so called Vickrey-Clarke-Grove (VCG) mechanism bases its functioning on the absence of communication among agents to achieve preference elicitation. While the mechanism is sensitive to collusion between agents, it is this way possible to show a dominant strategy for agents to bid truthfully.

single-unit and multi-unit In a single-unit auction bids can be made for a single good or individual bundles of goods, whereas multi-unit auctions allows bidders to express their preference on multiple units of the same product.

single-item and multi-item (i.e., combinatorial) As opposed to single-item auctions, combinatorial auctions allow bidders for expressing preference over bundles of goods, as well as select goods of different type or quality. In particular, the preference expressed on bundles has a superadditive property, by which only when sold together multiple products can maximize user's preference. For instance, the agent's utility is maximized only when frequency in area A and B are sold together.

The most common types of auctions are the English and Dutch auctions, which are distinguished respectively for the presence of an ascending and a descending price. These are single-dimensional, one sided, open-cry first price auctions. The Vickrey auction is a second price sealed bid auction, which contrarily to the english and the dutch auction is incentive compatible. As anticipated, a dominant strategy is shown on bidding according to the agent's private preference. However, as anticipated, the auctioneer is shown extracting very low profit in presence of collusion or very low resource valuation among the bidders. Multi-unit variants are also available for the English and Dutch auctions. In that case, the bidders report the quantity to purchase at the current price and the auctioneer updates the price to match the demand and supply curves. For instance in a multi-unit english auction, the price is increased as long as demand exceeds available supply. Given that bidders are acting rationally they will not increase the amount for a price more than their utility, thus ultimately meeting the available supply. Buy-side and sell-side auctions match many sellers to a single buyer and vice versa many buyers to one seller. Two-sided auctions, also called double auctions, aim at matching multiple buyers and sellers in one place. We normally distinguish in discrete-time double auctions and continuous-time double auctions, respectively using a uniform and a discriminative pricing mechanism. In the former, also called periodic or clearing-house markets, bidders are entitled to make an offer throughout a time interval called the trading day. At the end of the day the auctioneer closes the market and collected demand and supply curves to compute a clearing price. The clearing price is the price at which available supply matches demand, i.e., there is no leftover supply or demand. In continuous-time double auctions, the auctioneer immediately matches offers as soon as compatible ones are received. Compatible offers are those in which the supply has higher quantity and lower price than the demanded one. Stock markets are examples of continuous-time double auctions in which blocks of shares are sold at certain prices, with buyers seeking to buy subset of those at a different price. In the context of energy management double-side auctions have been widely used. In the NOBEL project [Ili12], a discrete-time double auction is used for a district energy market, where traders allocate their energy at a 15-minute scale. Similarly, a day-ahead continuous-time double auction is used in [Vyt10] to balance supply and demand in a power network. A balancing mechanism is implemented by charging deviations from the forecasted demand and production amounts.

In general, multi-unit auctions have been less researched than single-unit ones, being harder to analyze. In fact, multi-unit auctions can be considered as a uniform-price auction, in which all bidders are charged the clearing price. However, the uniform-price auction is shown having a set of possible equilibria yielding inefficient outcome. The main reason for such an inefficiency is that uniform pricing incentivizes “demand reduction”, i.e., bidders will bid less than their utility value for a unit in order to lower the price payed for all units [Aus97]. In Vickrey’s multi-unit auction, truthful bidding is ensured by making the winner’s price dependent on the “opportunity cost” for the allocation, rather than the bid or the clearing price [Cra06]. In detail, the allocation of k objects to bidder i is charged $\sum_{i=0}^{k} k - u_i$, i.e., for the first unit with the price of the k^{th} highest rejected bid, for the second unit with the price of the $(k-1)^{st}$ highest rejected bid, and so forth. To achieve this, each bidder reports his valuation for all possible packages, so that items can be allocated efficiently (i.e., to maximize social welfare). Clearly, the enumeration alone has complexity $\mathcal{O}(2^n)$. In spite of its theoretical properties, this makes it often not applicable to practical contexts. Another kind of multi-unit auction is Ausubel’s [Aus97], which operates as a price-ascending auction in which each winning bidder is assigned the quantity demanded at the clearing price, although it is charged for the price at which he “clinched” his opponents. The auction is shown yielding the same outcome of a sealed-bid Vickrey auction. According to the clinching rule, for each price p being selected, the auctioneer determines whether for any bidder i the aggregation of demand of all other bidders for p is less than the available supply M. In that case, the difference is defined as “clinched” and is allocated to bidder i at price p. For instance [Aus97], three bidders initially bid for respectively 2, 1 and 1 units of the same resource. We assume 2 units of supply available. The price gets increased until price p, when the third bidder reduces its quantity from 1 to 0, thus dropping out the auction. Now, the opponents of the first bidder together demand only 1 unit. Consequently, such a demand for 1 unit is lower than the overall available supply of 2 units. In this case, the first bidder “clinches” 1 unit at the current price p. The auction continues as long as supply is still available. The mechanism is shown eliminating any incentive for demand reduction. Moreover, in presence of sealed-bids a dominant strategy is shown on truthful bidding [Aus97], thus it yields the same outcome of the Vickrey auction. The main advantage over Vickrey's is the possibility of agents in ascending-price auctions to retain their demand curve private, while acting truthfully. Also, the iterative ascending-price update provides price discovery, while still allowing for the retain of the individual bidder's utility private.

As previously discussed in the auction classification, there exists auctions for multiple objects. The simultaneous ascending auction is one of this kind [Cra06], in which bidders send sealed bids for each available item. Given the complementarity between items (i.e., $u(A+B) > u(A) + u(B)$), losing an item in a later auction would render the iniatial winning useless. Therefore, items are auctioned simultaneously in separated auctions rather than sequentially. The auction takes place in multiple rounds, at the end of which the sealed bids are taken to identify and expose the currently winning bidder. This allows the other bidder to adjust their future bid, by reducing amount for those products being evaluated more than their utility. A monotonicity activity rule makes sure that bidders do not decide to wait to observe others' behavior before bidding strategically. Accordingly, in order to be allowed for bidding each bidder has to indicate the demand for each item beforehand. Each bidder can only bid a lower or equal amount than initially stated.

As a matter of fact, the complementarity between items in multi-item auctions can be addressed only with combinatorial auctions. Combinatorial auctions allow bidders for expressing preference directly on combinations of goods. However, the main drawback is the increased complexity for bidding, as well as the NP-hardness in solving the winner determination problem. Moreover, combinatorial auctions suffer of the "freerider" problem. Assuming three agents, respectively interested in the item A, B and the bundle A and B. The third agent is in competition with the two others and loses whenever $u_1(A) > u_3(A+B)$ or $u_2(B) > u_3(A+B)$, with u_i evaluation for agent i. Thus, winning depends on the evaluation that either agent 1 or 2 attributes to A and B. Automatically, one of the two becomes a "free rider", as regardless of his evaluation he can be allocated the good. Various types of combinatorial auctions are available, such as proxy auctions, in which a central proxy agent is used to ultimately allocate items depending on their complementarities. We refer to [Cra06] for a thorough discussion.

2.7.2 Efficiency measures

Market mechanisms are commonly evaluated using the allocation efficiency, according to which efficient mechanisms are those that maximize social welfare, i.e., the sum of utilities delivered to traders in a certain allocation. Consequently, the efficiency is maximized when all possible profit is extracted from the traders, and can be computed as:

$$\epsilon = \frac{\Pi^a}{\Pi^e} = \frac{\Pi^a_b + \Pi^a_s}{\Pi^e_b + \Pi^e_s} \tag{2.1}$$

that is the ratio between the surplus of all traders Π^a and the maximum possible surplus Π^e that would be obtained in a centralized and optimum allocation. The profit Π_b for a buyer j is given by $\sum_{i\in\mathbb{N}}(\psi_{b_j} - p_{ij})q_{ij}$ with ψ_{b_j} sensitivity price and q_{ij} quantity bought from seller i at the unit price p_{ij}. The profit Π_s for a seller i is given by $\sum_{j\in\mathbb{N}}(p_{ij} - \psi_{s_i})q_{ij}$ with ψ_{s_i} reservation price and q_{ij} quantity sold to buyer j at the unit price p_{ij}. The actual overall profit Π^a is given by the sum of the actual profits of all buyers and sellers, computed as difference between the agent's private value and the actual unit price paid. The equilibrium profit Π^e is given by the sum of equilibrium profits of all buyers and sellers, computed as difference between the agent's private value and the market equilibrium price p_0. The equilibrium price is at the intersection of the supply and demand curve and can be computed as the price of an auction in which agents declare their private value for the good to trade [Phe08]. Market efficiency is related to Pareto optimality, as in any zero-sum game an agent can not get any better condition without worsening someone else [Par11].

In [Chu02], responsiveness of services is considered to provide a user-centric performance metric for job schedulers in clusters. Each request for service operation is associated to a utility capturing the value of the resource. The valuation of service provisioning considers its responsiveness, by decaying the delivered utility over time. For a buyer $b_j \in B$, $V_j(r) = \delta \cdot U_j(r)$, with δ expressing the discomfort received from the delayed allocation over the utility U_j to receive r amount of resources. The user-centric efficiency is computed as the overall value delivered to users at the allocation time, that is $\sum_{j=1}^{n} V_j$.

2.8 Summary

In this chapter, we provided a comprehensive literature work on energy monitoring systems (Sect. 2.1), specifically on existing energy datasets (Sect. 2.2), current interoperability issues (Sect. 2.3), data management and analysis approaches (Sect. 2.4), interactive systems and their effectiveness in fostering behavioral change (Sect. 2.5), as well as means to design autonomous systems (Sect. 2.6) and market mechanisms for the coordination of self-interested agents (Sect. 2.7).

CHAPTER 3 Collecting energy usage data

"In God we trust, all others bring data."

– W. Edwards Deming

This thesis relies on the GREEND dataset for most of its evaluations. This chapter describes the developed measurement platform and discusses the measurement campaign undertaken for its collection.

3.1 The measurement platform

The employed measurement platform consists of a combination of off-the-shelf solutions. The measurement is carried out through the Plugwise Basic[1] kit, consisting of a Zigbee network of 9 Smart Plugs (SPs), each collecting active power measurements from the connected load (see Fig. 3.1). Through a USB dongle, the network is managed by an ARM-based mini computer, namely a RaspberryPi[2] or a Beagle Bone[3] board. Outages are prevented on the system using the Anker Astro E5 15000mAh external battery[4].

Table 3.1: Specifications of selected boards

Model	Power	CPU	RAM	Connectivity	Price
Raspberry Pi B	700 mA (3.5W)	Broadcom BCM2835 @ 700MHz	512 MB	10/100 Ethernet	35$
BeagleBone Black	210–460 mA (2W)	AM335x 1GHz ARM Cortex-A8	512 MB	10/100 Ethernet	45$

[1] http://www.plugwise.com
[2] http://www.raspberrypi.org
[3] http://beagleboard.org/bone
[4] http://www.ianker.com/support-c1-g228.html

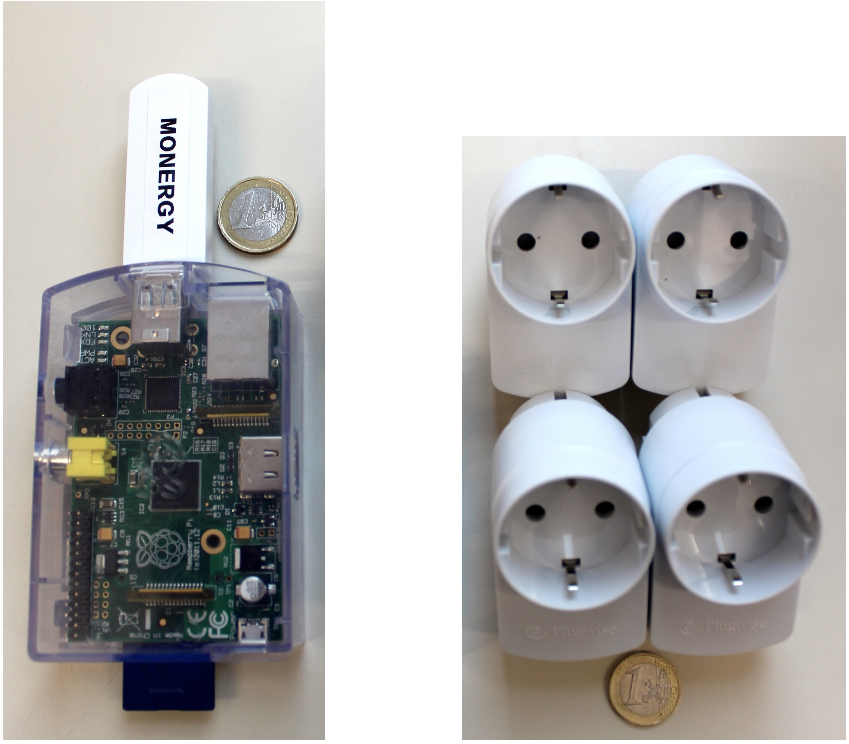

(a) A RaspberryPi (b) Plugwise smart plugs

Figure 3.1: The measurement platform

3.2 Data collection infrastructure

To guarantee a certain degree of reliability in our campaign we implemented fault recovery in our software daemon. The overall code base is freely available as a Sourceforge project[5]. The daemon consists of a collector and a manager script (See Fig. 3.2). The manager keeps the collector up to date, by periodically checking our servers for newer versions. The manager starts and monitors the collector. As soon as a newer version becomes available it replaces the old with the new one. All requests to our servers are authenticated through an API key, specifically assigned to each householder beforehand. This information is provided to the manager through a configuration file. The collector initializes the hardware measurement units, from which collection take place over epochs. In each epoch, power measurements are retried from each meter using the open source *python-plugwise*[6] library. In order to guarantee a uniform epoch duration, the collector detects and skips nodes which are temporarily not available. Failures can result from two different sources: i) incorrectly read values due

[5] http://sourceforge.net/projects/monergy

[6] https://bitbucket.org/hadara/python-plugwise/wiki/Home

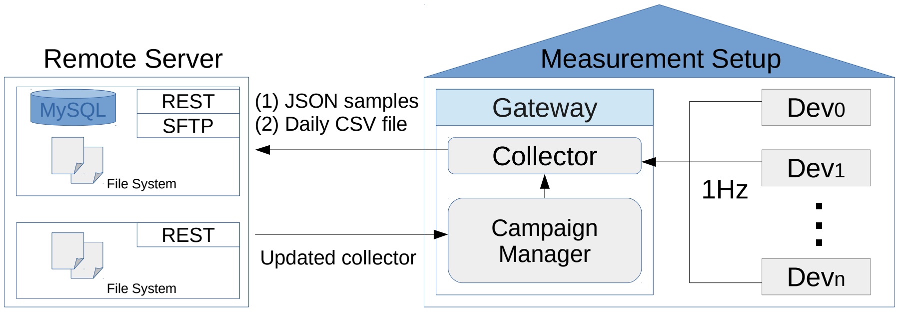

(a) Data management

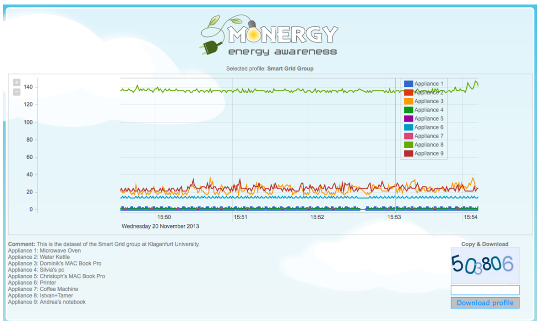

(b) The control interface

Figure 3.2: The measurement infrastructure

to communication failures or faulty nodes, as well as ii) timeouts due to node disconnections. For the former, retrial is performed only when remaining time is enough to guarantee the epoch duration. Unplugged nodes are handled using a blacklist, specifically to prevent their interrogation for a certain backoff time. While this allows for the prevention of future timeouts, it yields missing values for those nodes only temporarily unreachable. As for data storage, multiple possibilities exist. Collected measurements are collected in a window, whose size can be specified along with other settings in the configuration file. Upon completion of the window size, the data can be either sent as a Javascript Simple Object Notation (JSON) message sent to a REST interface. In this case, the data is remotely stored in a MySQL database, whose content can be visualized and further downloaded by the campaign managers (Fig. 3.2b). Another possibility is to save locally the samples as a daily Comma Separated Value (CSV) file. This solution was used for a building lacking internet connection. In addition, we provide the possibility to periodically upload collected CSV files using a Secure File Transfer Protocol (SFTP) connection to our servers. This is normally done at the end of the day, when a new CSV file is being created.

3.3 The GREEND dataset

GREEND was designed to overcome the limits of by existing datasets, as presented in Sect. 2.2. In particular, we based our requirements on those of load disaggregation [Zei12], as they are generally stricter than other kind of analysis. We collected active power measurements at 1 Hz, as according to [Arm13] this allows for the identification of 8 devices or more. Measurements were gathered from the energy hogs identified in Sect. 5.1.1 and [Mon13c]. We seeked diversity in both device types and users demography, and favoured a long campaign to allow observing seasonal patterns. The campaign lasted more than 1 year, with the first house being monitored in December 2013 and the last version being released in June 2015. In particular, this included:

- *House #0* a detached house with 2 floors in Spittal an der Drau (AT). The residents are a retired couple, spending most of time at home.
- *House #1* an apartment with 1 floor in Klagenfurt (AT). The residents are a young couple, spending most of daylight time at work during weekdays, mostly being at home in evenings and weekend.
- *House #2* a detached house with 2 floors in Spittal an der Drau (AT). The residents are a mature couple (1 housewife and 1 employed) and an employed adult son (28 years).
- *House #3* a detached house with 2 floors in Klagenfurt (AT). The residents are a mature couple (1 working part-time and 1 full time), living with two young kids.
- *House #4* an apartment with 2 floors in Udine (IT). The residents are a young couple, spending most of daylight time at work during weekdays, although being at home in evenings and weekend.
- *House #5* a detached house with 2 floors in Colloredo di Prato (IT). The residents are a mature couple (1 housewife and 1 employed) and an employed adult son (30 years).
- *House #6* a terraced house with 3 floors in Udine, (IT). The residents are a mature couple (1 working part-time and 1 full time), living with two young children.
- *House #7* a detached house with 2 floors in Basiliano (IT). The residents are a retired couple, spending most of time at home.

The device configurations for the selected households are shown in Table 3.2.

Table 3.2: Device configurations in the monitored households

House	Devices
0	Coffee machine, washing machine, radio, water kettle, fridge w/ freezer, dishwasher, kitchen lamp, TV, vacuum cleaner
1	Fridge, dishwasher, microwave, water kettle, washing machine, radio w/ amplifier, dryier, kitchenware (mixer and fruit juicer), bedside light
2	TV, networked-attached storage (NAS), washing machine, drier, dishwasher, notebook, kitchenware, coffee machine, bread machine
3	Entrance outlet, Dishwasher, water kettle, fridge w/o freezer, washing machine, hairdrier, computer, coffee machine, TV
4	Total outlets, total lights, kitchen TV, living room TV, fridge w/ freezer, electric oven, computer w/ scanner and printer, washing machine, hood
5	Plasma TV, lamp, toaster, stove, iron, computer w/ scanner and printer, LCD TV, washing machine, fridge w/ freezer
6	Total ground and first floor (including lights and outlets, with whitegoods, air conditioner and TV), total garden and shelter, total third floor.
7	TV w/ decoder, electric oven, dishwasher, hood, fridge w/ freezer, kitchen TV, ADSL modem, freezer, laptop w/ scanner and printer

3.4 Summary

In this chapter, we described the data collection infrastructure designed for a long-term measurement campaign we carried out in Italy and Austria. We addressed the reliability issues of the employed measurement platform using a best-effort collection strategy. The outcome of the campaign was the release of the GREEND dataset, which targets researchers in sustainability and is used in the following chapters of this dissertation.

CHAPTER 4

Achieving device and data interoperability

"If a picture is worth 1000 words, a prototype is worth 1000 meetings"

– D. Kelley

4.1 Architectural requirements

Figure 4.1 sketches an architectural model of an EMS that integrates a load disaggregation unit to detect legacy appliances. Specific driver components allows for the detection and management of sub-networks, thus acting as a proxy to integrate networked devices, smart devices and legacy devices. The model is implemented over the following 5 layers:

1. *Electric layer:* Electrical devices are connected to a common local power distribution network. This layer allows devices to deal with electrical power measurements. A classic meter works at this level.

2. *Network layer:* This layer provides network connectivity to embedded devices. A typical example is given by automation field-buses and wireless sensor networks, such as building automation systems and wireless smart outlets (e.g., ZigBee and WiFi). Management of the sub-network requires a specific driver to interface it to the Home Energy Management System (HEMS).

3. *Service layer:* In order to be automatically usable by other devices in the network, smart devices are required to provide a machine-readable description of their features and properties. The service layer includes the mechanisms by which devices can describe and advertise their features, so that functionalities can be discovered and exploited by other devices [Jam05].

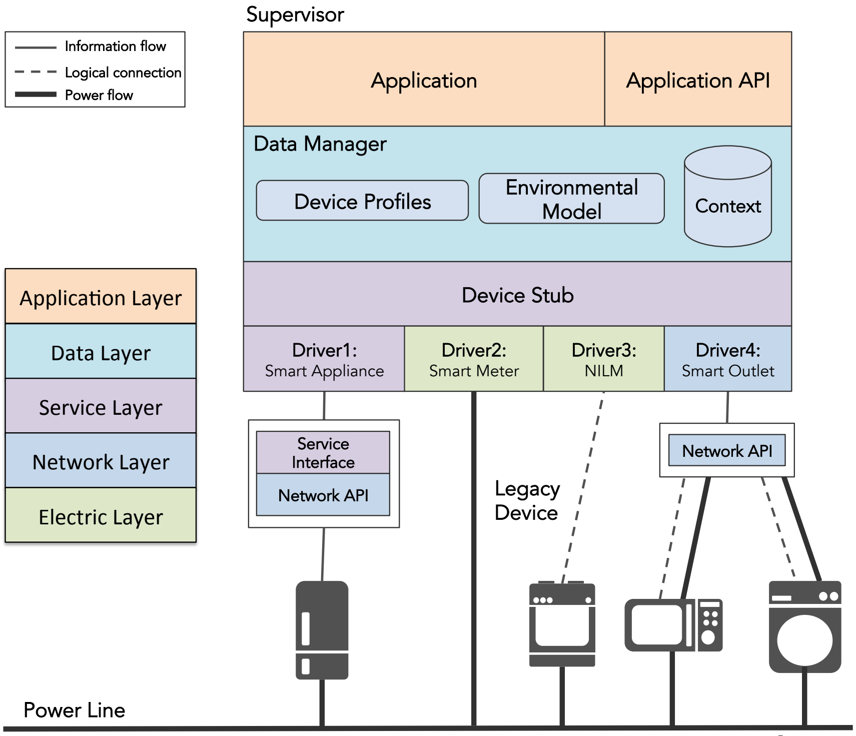

Figure 4.1: Architecture for microgrid energy management [Ega15a]

4. *Data layer:* The data layer provides an abstract representation of data and functionalities managed by the individual drivers, by providing a homogeneous interface to access this resource. This also includes the management of the device profiles, which are datasheets reporting static information of devices (e.g., sensor accuracy, type) and could be stored on the manufacturer's servers. In addition, a data model could be employed to perform a basic processing of raw data collected from the drivers, so as to produce a more abstract context representation, which can be stored in a knowledge base.

5. *Application layer:* User applications are run in the application layer. For instance a decision maker might rely on the context representation stored in the data layer to react to environment changes. A query engine provides an interface between the data and the application layer. On the other hand, the network API provides application-level interoperability to the architecture, thus representing a communication interface between applications running on different computing environments.

With exception of the electric layer, all layers could be implemented either locally to the building environment or on remote servers. For instance, a smart appliance is a device that embeds a computing unit and a network interface. In order to be integrated in a HEMS a smart appliance should implement the first three layers (electric, network and service), although the device could scale to the application layer in case data management and decision making at appliance level were necessary.

4.2 Integrating smart and legacy devices

A *device stub* can deal with physical devices, acting as a proxy for networked devices, which can be reached using different drivers. A driver interfaces the remote device using a specific communication technology, such as a ZigBee network of smart outlets. On the other hand, load disaggregation can be used to detect operating devices. The stub keeps a local representation of remote objects, to describe their characteristics and status. This information is made available to the data layer, so that applications can combine it to other data sources (e.g., environmental data) for providing specific services.

4.2.1 Detection of legacy devices

Modern EMS can exploit ILM or NILM to extract status information of electrical loads. For instance, Table 4.1 shows which information can be extracted using NILM and ILM, as compared to those made available by smart appliances. This gives the possibility to track device operation and build profiles that can

Table 4.1: Available appliance information for smart and legacy appliances [Ega15a]

Parameter	Smart	NILM	ILM
ID	✓	✓	✓
Type	✓	~	✓
Controllable	✓	✗	~
Current power	✓	~	✓
Energy per day	✓	~	✓
Appliance usage	✓	~	✓

be made available throughout the system. This information can be of great value when planning control strategies, as effective scheduling of smart devices should also consider the presence and behavior of non-schedulable ones. For instance, [Car13] indicates white goods as critical for the success of demand response programmes. Towards the same vision, the authors of [Ega15a] list important appliances that should be necessarily identified and integrated in a

Table 4.2: Relevant devices to be integrated in a HEMS [Ega15a]

Type	Control.	User-dr.	Tested in [Ega15a]
Fridge	✓	-	✓
Lighting	✓	✓	-
Dishwasher	✓	✓	✓
Oven	-	✓	-
Microwave	-	✓	-
Hob	-	✓	-
Washing mach.	✓	✓	✓
TV	-	✓	✓
Computer	-	✓	-
Water Kettle	-	✓	✓
Coffee Mach.	-	✓	✓
Vacuum Clea.	-	✓	✓

HEMS, as in Table 4.2. User-driven devices require the presence of users while operating. Examples are water kettles or vacuum cleaners. As such, they are not good candidates for automatic load control. Controllability of detected devices is also an important aspect. Contrarily to smart appliances, ILM and NILM do not provide means for controlling and thus scheduling loads. For instance, smart outlets can switch on/off loads, although load operation can't be correctly paused and continued afterwards.

4.2.2 Management and representation of appliance data

As previously presented in Sect. 2.3.3, ontologies provide shared vocabularies of a specific domain and allow for interoperability across networked entities. To achieve this, it is necessary to model:

- *measurement values* i.e., measurement data collected from the physical environment. In this case constraints arise related to the sampling frequency and the representation of gathered timeseries. Required memory is proportional to the sampling accuracy, therefore alternative representations such as edge or event-based should be considered.
- *appliance profiles* which include appliance properties. This information can be provided by manufacturers as datasheets describing the device operation modalities. Alternative approaches include having a certified trustworthy entity providing the profiles, or having the profiles filled by users. Datasheets can include static information such as as manufacturer,

type, energy rating and user controllability. Being the device operating in the physical environment, it offers a set physical services. Each service carries out a specific task, such as the washing cycle for a washing machine. For the service, a profile includes its power signature, the demanded energy and the current status (see Fig. 4.2). We distinguish in devices with permanent consumption, such as fire alarms, and devices operating over multiple states. A state is defined in terms of a peak active power (in Watts) and a tolerance to power variations, as well as a state duration and two discomfort factors. A delay sensitivity (in seconds) defines the responsiveness of the device, whereas the interruption sensitivity (in seconds) denotes the tolerance to interruptions of the state. For instance, the start of a coffee machine should not be postponed from its request because of an extremely low delay sensitivity, while its water heating state should not be interrupted because of a low interruption sensitivity. The status of a service describes the operation (i.e., on, off, or paused), as well as its progress in terms of start time and elapsed duration. A virtual service provides the possibility to control a physical service or retrieve information (e.g., temperature values) through a network interface. The virtual service is described as a machine-readable interface. In this way, the appliance can have multiple interfaces to describe the same service. This information can be automatically extracted and filled out by a load disaggregation unit, which allows applications for seamlessly access device information for both smart and legacy devices.

- *appliance identification models* i.e., algorithm-dependent model describing the electrical operation of devices through a set of observable states. While appliance profiles describe the operation for energy management purposes, identification models target load detection. For instance, this can be defined using state-based representations such as Finite-State Machine (FSM) and Hidden Markov Model (HMM), which model the appliance dynamics as a trajectory of state transitions over time. Specifically, each device observation is described by a set of features, such as current and power, as well as outgoing transitions to other observable states. Normally, state features are distinguished in steady state and transient features. Certain models also express the typical duration of a device observation (e.g., Hidden semi-Markov Model (HSMM)). Device dynamics are expressed by associating a transition probability to each edge connecting observations. Fig. 4.3 reports the ontology for electrical appliances, showing both the device profile and the load identification model.

- *appliance usage models.* As previously stated in Sect. 2.4.2, usage data can be used to extract models describing appliance usage. In order to

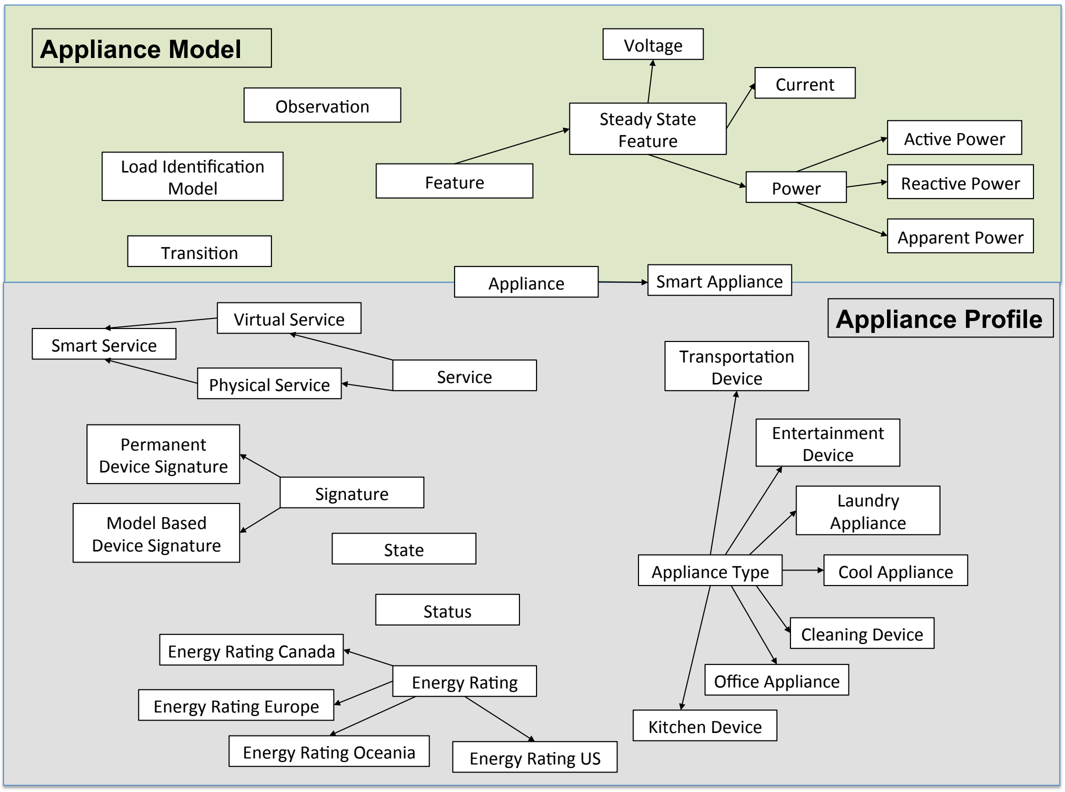

Figure 4.2: Taxonomy of appliance description and model [Ega15a]

allow applications for exploiting such models, common formats should be identified. For instance, the Probabilistic OWL (PR-OWL)[1] is an extension to the Web Ontology Language (OWL) that allows for modeling probabilistic semantic networks, i.e., bayesian networks.

To implement the appliance profile and the load identification model, we used the the open source tool Protégé[2] to build models in the OWL. The resulting ontology was released for open use on the MONERGY project webpage[3]. Fig. 4.4 shows an example profile for a water kettle. The device is user driven and has a physical service to heat water. The service demands 0.03 kWh and is currently in the OFF status. The service takes place over one state, requiring 1800 W with 5% tolerance being insensitive to interruption and start delay. Fig. 4.5 reports the load identification model for the water kettle. To identify the device, this model describes OFF and ON observations, using active power as a feature. As noticeable, device dynamics are captured using transition probabilities. Following this example, we further released on Github[4] a small

[1] http://www.pr-owl.org
[2] http://protege.stanford.edu
[3] http://www.monergy-project.eu/appliance-ontology/

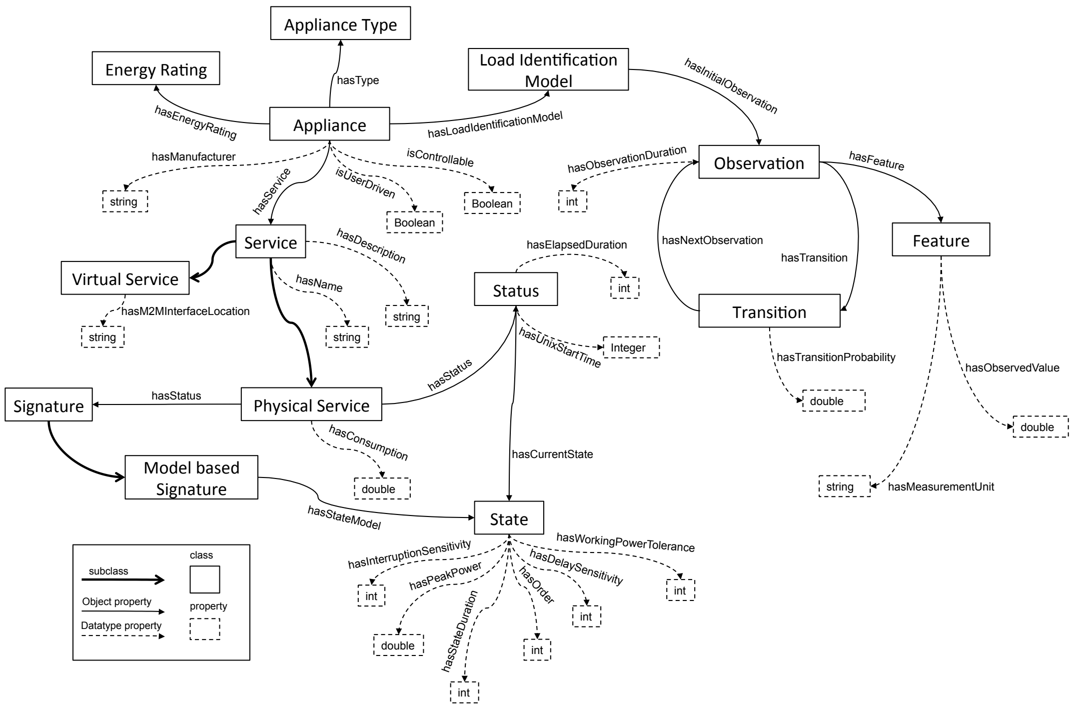

Figure 4.3: Ontology for appliance description [Ega15a]

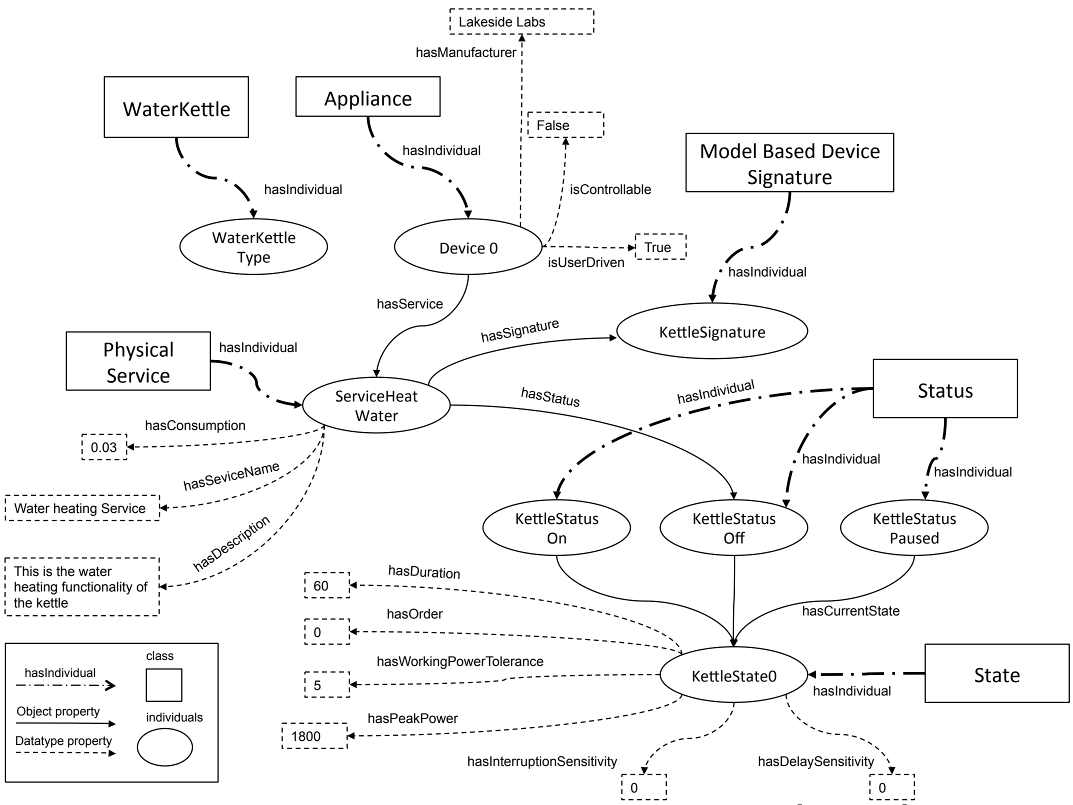

Figure 4.4: Device profile for the water kettle [Ega15a]

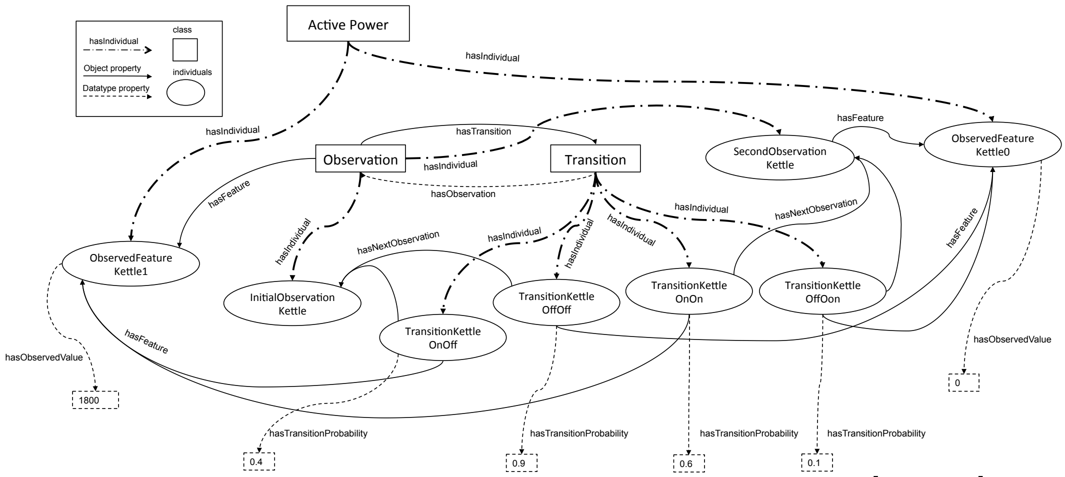

Figure 4.5: Load identification model for the water kettle [Ega15a]

library based on the Python RDFLib. The library allows for the management of the ontology and the knowledge base, as well as allows for the retrieval of information from the semantic network using a SPARQL query engine.

In the example, a gateway is started to manage the household namespace `http://www.monergy-project.eu/houses/1234/`. The gateway consists of a knowledge base (i.e., a graph) and a simulated load disaggregator, as well as a web interface exposing a SPARQL endpoint. The gateway role is to retrieve device status from both the network and the load disaggregation unit. For instance, a water kettle can be modeled as in the Listing 4.1.

Listing 4.1: Modeling a water kettle

```
# create an appliance in the given household
a = SmartAppliance()
a.set_appliance_attributes("Lakeside_Labs_GmbH", "WK300", APPS.WaterKettle, True)

# Add supported technologies for the device
a.add_M2M_technology("COAP")
a.add_M2M_technology("DPWS")

# define a service for the water kettle
kettleStateZero = State(order=0, peak_power=1800.0,
                        state_duration=60, power_tolerance=5,
                        delay_sensitivity=0, interruption_sensitivity=0)
kettleSignature = ModelBasedDeviceSignature()
kettleSignature.add_state(kettleStateZero)
kettleStatusOff = Status("Off", current_state=kettleStateZero)
waterHeatingService = PhysicalService(name="waterHeatingService",
            description='This_service_describes_the_operation_of_a_water_kettle',
            signature=kettleSignature, status=kettleStatusOff, consumption=0.03)
a.add_physical_service(waterHeatingService)

self.data_manager.add_appliance(a)
```

The *add_appliance* method converts all appliance attributes into RDF triples that are added to the main semantic network.

[4] https://github.com/pilillo/EMSDataManager

Listing 4.2: The semantic network in the n3 format

```
@prefix apps: <http://www.monergy-project.eu/ontologies/appliances.owl#> .
@prefix base: <http://www.monergy-project.eu/houses/12345/> .
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
@prefix xml: <http://www.w3.org/XML/1998/namespace> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .

<http://www.monergy-project.eu/houses/12345/appliances/4bdf9953a5880b0eb943360b8d90bd0b46330babcaad11a44f7b536a96110bc0#Appliance> a
      apps:Appliance ;
    apps:hasEnergyClass "UnknownRating" ;
    apps:hasManufacturer "Lakeside Labs GmbH" ;
    apps:hasManufacturerProductID "WK300" ;
    apps:hasService <http://www.monergy-project.eu/houses/12345/services/4
        bdf9953a5880b0eb943360b8d90bd0b46330babcaad11a44f7b536a96110bc0_waterHeatingService#PhysicalService> ;
    apps:hasType apps:WaterKettle ;
    apps:implementsM2MTechnology "COAP", "DPWS" ;
    apps:isControllable false ;
    apps:isUserDriven true .

<http://www.monergy-project.eu/houses/12345/appliances/811ea636d7ae213a3d5120e80070dbc9598d6673d9c838b5a3988983c7407c57#Appliance> a
      apps:Appliance ;
    apps:hasEnergyClass "UnknownRating" ;
    apps:hasManufacturer "Lakeside Labs GmbH" ;
    apps:hasManufacturerProductID "WK300" ;
    apps:hasService <http://www.monergy-project.eu/houses/12345/services/
                                                        811ea636d7ae213a3d5120e80070dbc9598d6673d9c838b5a3988983c7407c57_waterHeatingService#
                                                            PhysicalService> ;
    apps:hasType apps:WaterKettle ;
    apps:implementsM2MTechnology "COAP",
        "DPWS" ;
    apps:isControllable false ;
    apps:isUserDriven true .

<http://www.monergy-project.eu/houses/12345/services/
                            4bdf9953a5880b0eb943360b8d90bd0b46330babcaad11a44f7b536a96110bc0_waterHeatingService#PhysicalService> a
                                apps:PhysicalService ;
    apps:hasConsumption 3e-02 ;
    apps:hasDescription "This service describes the operation of a water kettle" ;
    apps:hasServiceName "waterHeatingService" ;
    apps:hasSignature <http://www.monergy-project.eu/houses/12345/signatures/
                                                                        4bdf9953a5880b0eb943360b8d90bd0b46330babcaad11a44f7b536a96110bc0#
                                                                            ModelBasedDeviceSignature> .

<http://www.monergy-project.eu/houses/12345/services/
                            811ea636d7ae213a3d5120e80070dbc9598d6673d9c838b5a3988983c7407c57_waterHeatingService#PhysicalService> a
                                apps:PhysicalService ;
    apps:hasConsumption 3e-02 ;
    apps:hasDescription "This service describes the operation of a water kettle" ;
    apps:hasServiceName "waterHeatingService" ;
    apps:hasSignature <http://www.monergy-project.eu/houses/12345/signatures/
                                                                        811ea636d7ae213a3d5120e80070dbc9598d6673d9c838b5a3988983c7407c57#
                                                                            ModelBasedDeviceSignature> .

<http://www.monergy-project.eu/houses/12345/signatures/4bdf9953a5880b0eb943360b8d90bd0b46330babcaad11a44f7b536a96110bc0#
```

```
      ModelBasedDeviceSignature>
      apps:hasStateModel <http://www.monergy-project.eu/houses/12345/states/4
           bdf9953a5880b0eb943360b8d90bd0b46330babcaad11a44f7b536a96110bc0_0#State> .

<http://www.monergy-project.eu/houses/12345/signatures/811ea636d7ae213a3d5120e80070dbc9598d6673d9c838b5a3988983c7407c57#
      ModelBasedDeviceSignature>
                apps:hasStateModel
                <http://www.monergy-project.eu/houses/12345/states/811ea636d7ae213a3d5120e80070dbc9598d6673d9c838b5a3988983c7407c57_0#
                     State> .

<http://www.monergy-project.eu/houses/12345/states/4bdf9953a5880b0eb943360b8d90bd0b46330babcaad11a44f7b536a96110bc0_0#State> a
      apps:State ;
      apps:hasDelaySensitivity 0 ;
      apps:hasInterruptionSensitivity 0 ;
      apps:hasOrder 0 ;
      apps:hasPeakPower 1.8e+03 ;
      apps:hasStateDuration 60 ;
      apps:hasWorkingPowerTolerance 5 .

<http://www.monergy-project.eu/houses/12345/states/811ea636d7ae213a3d5120e80070dbc9598d6673d9c838b5a3988983c7407c57_0#State> a
      apps:State ;
      apps:hasDelaySensitivity 0 ;
      apps:hasInterruptionSensitivity 0 ;
      apps:hasOrder 0 ;
      apps:hasPeakPower 1.8e+03 ;
      apps:hasStateDuration 60 ;
      apps:hasWorkingPowerTolerance 5 .
```

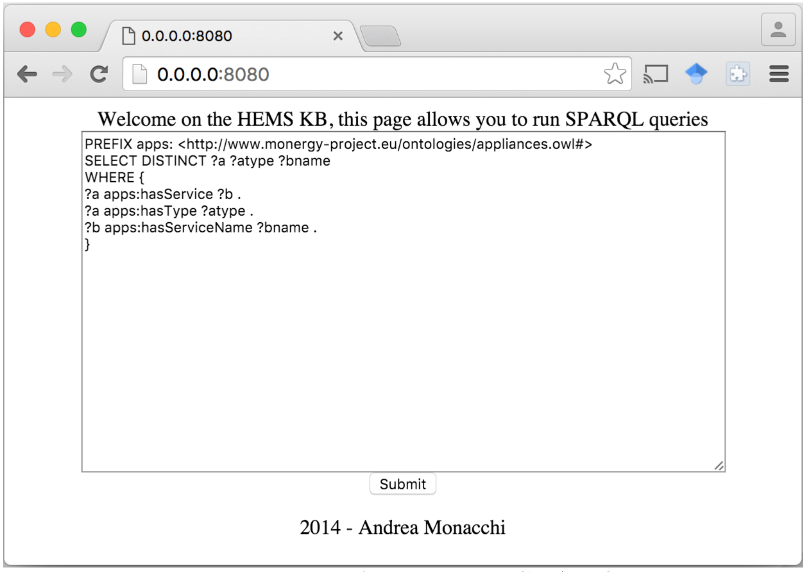

Figure 4.6: Web interface to run SPARQL queries

This makes the appliance already integrated in the system. Fig. 4.6 shows a web interface from which SPARQL queries can be run. The example query returns the URI of all appliances that provide a service, along with the device type (See Fig. 4.7).

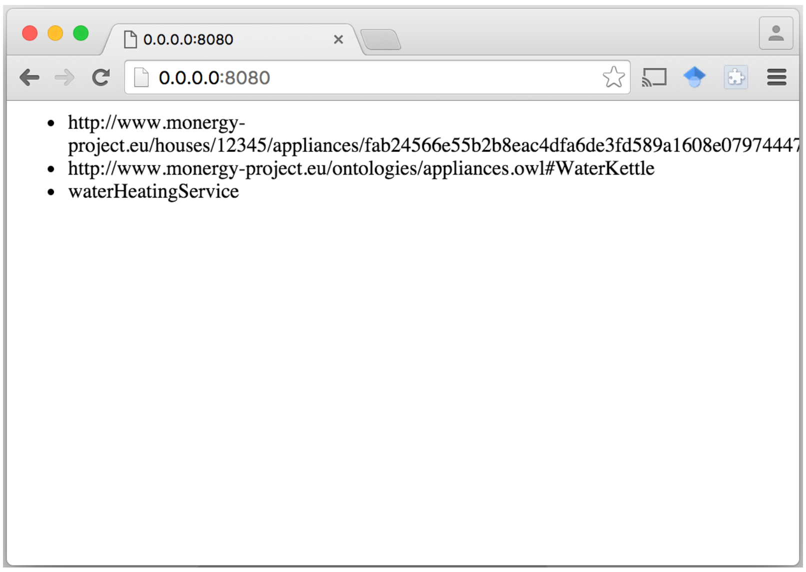

Figure 4.7: Web interface showing the result of the query

4.3 Summary

In this chapter we discussed interoperability issues in EMS. We introduced a multi-layer architecture for energy management systems and identified as a core challenge the integration of legacy and smart devices. To this end, we advocate an ontology representation of the domain which describes both device dynamics and models for load disaggregation. On one hand this allows for sharing appliance profiles with applications, which can this way consider legacy devices to better shedule resources. We provide as example the modeling of a water kettle and show how information can be managed and retrieved using existing RDF libraries. While the actual detection of devices goes out of the scope of this chapter, we expect the future integration with load disaggregation frameworks, as the nilm toolkit [Bat14]. The availability of appliance information can also potentially improve the load disaggregation process, by making models available at a larger scale. SPARQL can be used for performing both queries and updates on the semantic network. The integration with specific fieldbus standards for device control is yet to be addressed.

CHAPTER 5

Improving energy awareness

"The first step toward change is awareness. The second step is acceptance."

– Nathaniel Branden

In this chapter, we introduce and assess the effectiveness of feedback means in fostering energy conservation. We firstly describe the regions under experiment and identify potentially applicable solutions. The ultimate goal is to provide a working management system where feedback means can be freely implemented and evaluated.

5.1 Energy usage in Austria and Italy

Households account for a relevant portion of the overall energy consumption. The MONERGY project aimed at lowering such consumption by fostering efficiency through behavioral change. The project focused on the Italian region Friuli-Venezia Giulia (FVG) and the Austrian federal state Carinthia (CAR) and had as partners the University of Klagenfurt, the Lakeside Labs and the Italian company WiTikee. Therefore, the results hereby presented are the outcome of such collaboration and the ownership of individual results is explicitly stated.

5.1.1 Scenarios

The first stage included the analysis of consumption scenarios in the regions [Mon13c]. In particular, we carried out a web-survey on our project website. Citizens older than 18 were asked to provide characteristics of the building, as well as electrical devices and their usage. Advertisement was done via mailing lists (e.g., companies and universities).

The survey required an average of 15 minutes to be completed, and included 43 questions grouped in 5 categories:

- household information;
- use of electrical devices;
- sensitivity towards energy consumption and renewable energy generation;
- sensitivity and expectations towards technology;
- demographic information.

A total of 340 full responses out of 397 participants were collected, with a completion rate of 85.64%. This included respectively 186 responses from Carinthia (i.e., 96 female and 90 male) and 139 from FVG (i.e., 63 female and 76 male). The study showed a greater use of electrical devices for cooking and heating purposes in Carinthia. For FVG lower electricity costs can be accounted due to a more developed gas distribution network. Installed renewable energy generation is limited, with photovoltaic having the highest diffusion (7.91% in FVG and 2.69% in Carinthia). Differences on electrical devices are mostly due to climate differences. Residents in Friuli tend to use air conditioners (45.19% compared to the 2.16% of Austrian respondents) and are billed according to time-of-use tariff plans, which is possible due to the already available digital meters. As a consequence, householders in FVG declared to already exploit more favouravle pricing conditions when operating their washing machine (62.59%), lights (24.46%), iron (22.3%), electric oven (21.58%), dryer (10.79%), conditioner (10.07%), and dishwasher (9.35%). For Carinthia the only available countermeasure to increase efficiency is device replacement, as done by 67.20% respondents in the last 4 years in Carinthia. Nevertheless, householders expressed their willingness to exploit time-of-use pricing schemes to operate their washing machine (48%), electrical boiler (23%) and dryer (20%).

An estimation of energy usage in residential settings followed in [Kha14], along with an assessment of residents' attitude towards demand response and energy management systems. Similarly, the identified consumption scenarios, i.e., the involved electrical devices and building characteristics, were used to sketch requirements for the communication infrastructure [D'A14].

5.2 Increasing the feedback resolution

As shown in Sect. 5.1.1, in Carinthia the lack of high resolution consumption data makes hard to further analyze data at the meter level. However, as shown in

Sect. 2.5 the most effective conservation strategies are feedback approaches that provide device-level information. For such reason, we introduced in [Mon13b] the concept of pay-as-you-go electrical devices. This combines prepaid billing and device-level information. The main concept is the possibility to associate a credit to each electrical device. In this way, by having the credit decreased and shown for any usage of the appliance the user can receive a device-level feedback. To this end, we used a off-the-shelf monitoring solution to collect power consumption at device level and we implemented a mobile Android application[1]. A central gateway based on a Raspberry Pi gathers the samples from the networked measurement units, detects starting events and enhances such information with further contextual information. Ultimately, usage events of type $\langle start, duration, consumption, energyprice \rangle$ are collected by a remote Google AppEngine webservice[2] for each usage. Intuitively, differences on the energy price provides a non-technical measure of the costs for operating the device. This provides broken-down understanding of expenses, which might eventually trigger the replacement of the devices responsible for higher consumption and expense. The mobile application offers a quick gateway to such information, and displays event notifications (See Fig. 5.1). In particular, the Google Cloud

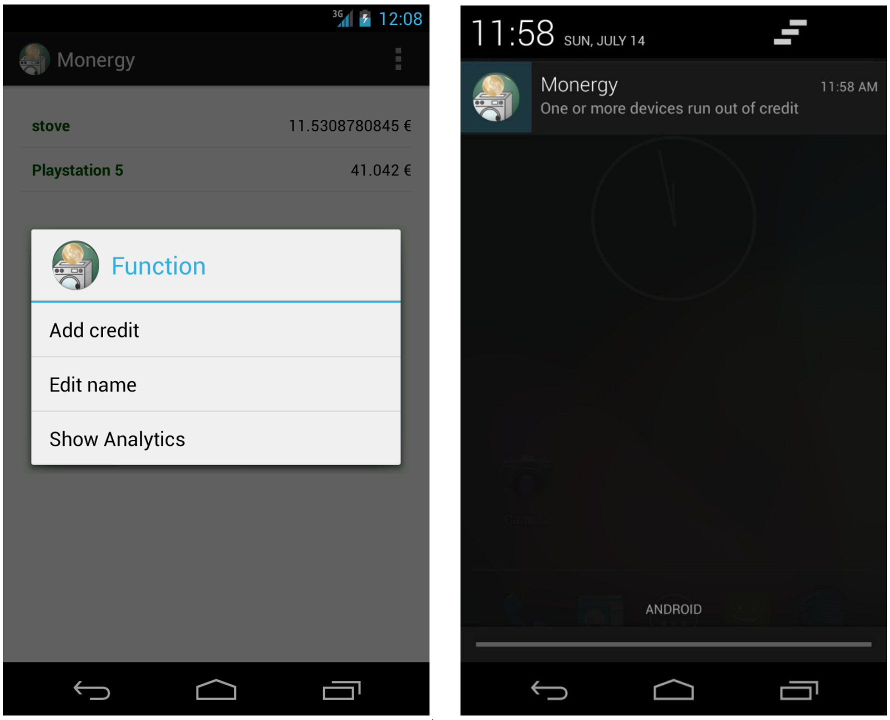

Figure 5.1: The Android application

messaging infrastructure was used to notify all users' terminals whenever a

[1]https://play.google.com/store/apps/details?id=at.aau.monergyapp
[2]https://appengine.google.com/

device runs out of credit (See Fig.5.2). The approach offers a reliable eventing

Figure 5.2: Event notification

mechanism while ensuring low battery consumption and network usage.

To assess the early acceptance of the notification mechanism we run a user testing session. Accordingly, 7 users were asked to interact with the mobile application. Specifically, we asked them to operate a water kettle and think aloud while interacting with the application. The objective was to observe their perception of the smart notification mechanism, namely concerning their understanding of costs for operating electrical devices and the intrusiveness of the interface. In particular, the first operation of the water kettle would only decrease the credit while the second one would reset it and consequently trigger the notification mechanism.

All users declared to have noticed the difference in credit terms for operating the water kettle. Nevertheless, the main negative comment was the necessity of refreshing the device list to have credit changes visible. The reason is that usage events are solely notified to the remote server, and are thus not propagated to the mobile terminals. As for the second operation, the credit gets decreased to a value below zero and triggers the notification as in Fig. 5.2. Opening the application causes the automatic refresh and a red-colored credit. Users declared the second process as more intuitive than the former.

In conclusion, the session allowed us for the identification of interface flaws, mainly related to the use of a RESTful backend. A more sophysticated transport technology providing full-duplex connectivity, i.e. websockets, should be employed in future versions. In particular, the web application messaging protocol (WAMP) [3] provides a publish-subscribe abstraction over websockets. WAMP creates a shared event bus where messages are delivered in a soft real-time fashion to coordinate loosely coupled distributed resources.

5.2.1 Room for intervention

To further analyze energy usage in the regions, the following step was to carry out a measurement campaign in actual households. As previously shown in Sect. 3.3, the main outcome was the GREEND dataset.

In [Lak15], we discuss the energy consumption in the monitored sites. In all sites, the fridge is the most consuming device as it accounts for 40% to 47% of monitored consumption. A significant fraction is also accounted to the dryer,

[3]https://tools.ietf.org/html/draft-oberstet-hybi-tavendo-wamp

the dishwasher and the washing machine. Lighting has also a considerable share, especially in site S1 where multiple incandescent lightbulbs are present. The bedside lamp alone accounts for 2% of monitored consumption. In site S2, we remark the high share given by the plasma TV and the stand-by consumption of consumer electronics devices (i.e., uninterruptible power supply, network attached storage, game console, personal computers). Similarly, in site S3 the desktop computer accounts for about 22% of monitored consumption, which translates into more than 12 kWh every month. The Italian deployments present a similar situation, with the fridge being responsible for the largest share, between 24% and 46% of the total monitored consumption. Televisions have a considerable impact, with respectively 20%, 25% and 39% of consumption in S4, S5 and S7. In S4 and S5 the consumption of televisions is higher than the washing machine, which accounts for only 5% and 10% of the total monitored in September. As shown in [Lak15], the sites under experiment employ a tariff plan divided in 2 time slots. Italian users were shown being aware of the incentive, as monitored devices are operated mostly during the cheaper T2. In particular, residents of S5 better exploit T2 than others, given the larger consumption spread between T1 and T2, expecially to operate the washing machine and the iron. However, it is important to remark that not all devices can be postponed. Also, the incentive of the Italian tariff plan can only yield limited savings. For instance, in S4 operating the washing machine in T2 only would save about 0.12 €.

Together with our colleagues from WiTikee, we identified in [Lak15, Mon15] multiple approaches to improve efficiency:

1. **lighting** promoting replacement of incandescent bulbs with energy saving ones;

2. **device diagnostics** promoting replacement of old appliances with more energy efficient ones, expecially regarding white goods but also involving consumer electronics (e.g., LCD/LED TV in place of a plasma TV);

3. **shedding of standby losses** promoting shed of consumer-electronic devices when people are not likely to be at home;

4. **device shifting** promoting postponement of energy demanding devices to off-peak periods, so as to operate them in cheaper time periods.

To estimate device diagnostics we considered two old fridges installed in site S5 (See Fig. 5.3). Given the coverage problems in the site, this required specific measurements for over a week. The measurements revealed an energy consumption of about 47.7 Wh and 28.6 Wh, for a total amount of 56 kWh per month and 668 kWh per year. This can be reduced to below 258 kWh per year by replacing the two freezers with new appliances having a A+++ energy class[4]. The resulting energy saving would be of 34 kWh per month that corresponds to the 11% of the total energy consumption of site S5.

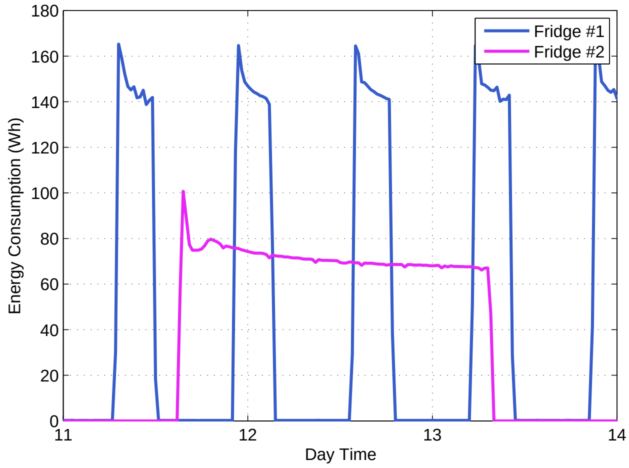

Figure 5.3: Power consumption of an old fridge and an old freezer in site S5 for 3 hours [Mon15]

The standby is commonly present on consumer electronic devices (e.g., players, televisions) to ensure a prompt reaction to users while operating in an idle state. In site S7 the analysis showed the television and the decoder being always in such a status. Fig. 5.4 shows the measured consumption for multiple days. Power consumption is approximately 6.57 W, which results in 57.57 kWh per year, namely the 1.4% of the total of site S7 (4099 kWh). If we observe that several devices can be in standby mode, this proportion can increase significantly. For instance, 10 devices with stand-by mode would already mean that 14% of total consumption is wasted this way. One device of this kind is the ADSL modem, which is being used for a few hours a day but is often left on all the time. Typical consumption of an ADSL modem with WiFi and Ethernet is of about 30 W[5], which leads to 263 kWh per year. For instance, a solution might be to shed this load for 3 hours a day and during the entire weekend. In this case, consumption would be of 98 kWh per year, i.e. about 37% less energy consumption than when was left on all time. Clearly, this strategy is effective when a model of device usage and occupancy are available.

Device shifting and curtailment are classic strategies of demand response. It consists in postponing or reducing the time for which particularly inefficient or energy-demanding devices are used. As example we consider the consumption

[4]The calculation assumes that the energy efficiency index (EEI) is equal to 22, the volume of the freezer is equal to 302 liters, and the appliance category is the 7th.

[5]goo.gl/IxWTiO

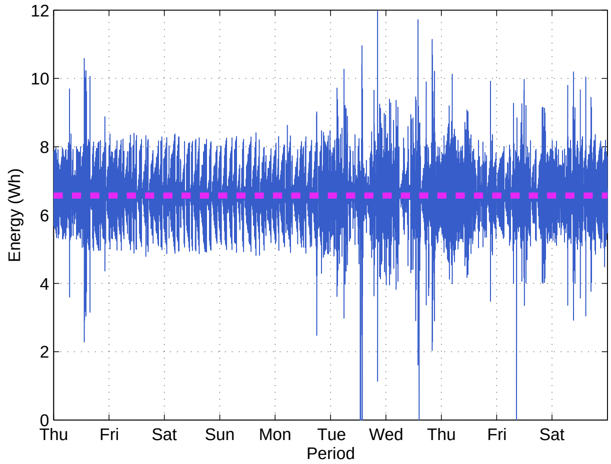

Figure 5.4: Measured and mean energy absorbed by the TV + decoder of site S7 [Mon15]

of the plasma TV (42") and the LCD TV (37") of site S4 and S5 respectively, for one day (Fig. 5.5). The plasma consumes significantly more energy than the

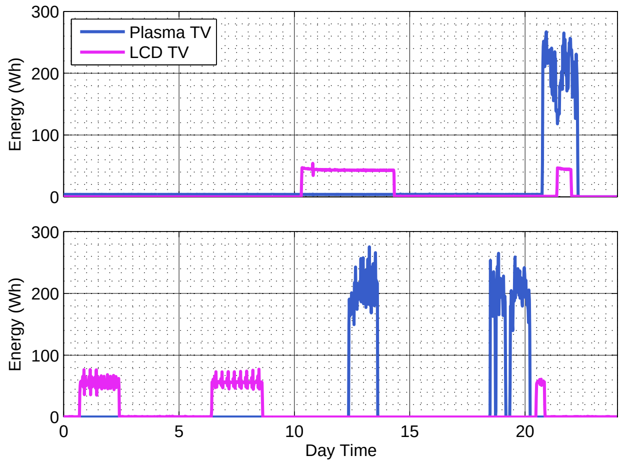

Figure 5.5: Energy consumption comparison of the plasma and the LCD TVs of site S4 (on top) and site S5 (on bottom) [Mon15]

LCD TV, with an estimated hourly consumption of respectively 200 Wh/hour and 80 Wh/hour. Beside mere replacement, a possibility to reduce consumption is to swap the two devices. In detail, the LCD can be used in those rooms whose occupancy is higher during the day (e.g., kitchen), while the plasma would be used in rooms occupied in off-peak periods (e.g., bedroom). Savings can be calculated by considering the time of use of those devices, which for site S4 is respectively 421 and 148 hours of operation for the plasma and the LCD TV. In site S5, 771 and 404 hours were estimated respectively for the plasma and

LCD TV. Swapping the devices in terms of location results in 34% and 23% lower energy consumption for site S4 and S5, respectively.

5.3 Mjölnir: a web-based energy dashboard

Following the potential room for intervention identified in Sect. 5.2.1, we developed a web-based energy management system, capable of analyzing energy consumption and production data resulting from both circuit-level and device-level measurements. The framework, named Mjölnir, is implemented in PHP 5, and its default working DBMS is MySQL, although it can be easily extended to connect to others, such as MongoDB. The project was initiated by myself and the student Manuel Herold, who took care of the front-end side.

5.3.1 Interface

The interface is implemented in CSS 3 and Javascript. In particular, we use Twitter's Bootstrap[6] to display seamlessly the dashboard on both mobile terminals and computers. The dashboard is organized in pages and cells. A public page is provided to place information visible to other peers, while private page report the results of data analysis. In particular, the cells can contains widgets, each performing a different type of analysis on data. This provides both interface modularity and allows users for adapting the system to display things in the way that is most meaningful to them, such as by placing things next to each other and only restricting the field of analysis.

Currently available widgets are:

timeserie showing collected circuit-level power measurements in comparison to the same hour interval recorded over the previous days;

production and consumption report showing daily energy information over the last month;

calendar view comparing daily energy use to show anomalous usage patterns;

cost report showing daily aggregated energy cost over the last month;

room-based cost report showing energy consumption and cost per individual rooms;

[6] http://getbootstrap.com

production and consumption gauges showing energy use for the current day;

energy estimation showing an estimation of energy production and consumption for the current day, as based on the previous days;

device itemization showing the consumption and cost per device, over the current day, week and year;

timeline showing energy usage events over a timeline, each described by their device, consumption and cost;

tariff switch showing the cost and use of devices over the available energy tariffs, in order to foster use in off-peak periods;

energy advisor returning tips to increase efficiency depending on usage behavior;

appliance usage showing the usage probability of user-driven devices, computed as frequency counting over monitored days;

occupancy model showing the building occupancy probability based on all extracted appliance usage models;

5.3.2 Modeling energy price and building information

To facilitate description of activities and processes, the system models buildings, rooms and individual devices. In particular, this is used to restrict the applicability of widgets to specific data sources. The description exploits the large vocabulary introduced in [Kel14b], specifically by indicating the device type (e.g., fridge), mobility and room, curtailability, autonomy (i.e., user control) and stand-by mode. In addition, we associate a credit to each monitored device, which is decreased upon device usage as in Sect. 5.2 and [Mon13b]. All cost analyses rely on an energy price model, which can be expressed as energy tariffs (e.g., priced time intervals).

5.3.3 Appliance usage and occupancy modeling

To tailor control strategies and feedback to end users, it is necessary to build models of appliance usage and building occupancy. In Mjölnir, usage of devices is modeled for weekdays, Saturdays and Sundays using a frequentist approach. Modeling is on a hourly base although this can be easily reduced to 15 minute intervals. In particular, based on received usage events we store the number of operations for each interval. For the interval, the usage probability is then computed as $P(e) \approx \frac{n_e}{n_t}$, that is the ratio between the number of usages occurred over the number of recorded events (i.e., the number of days the system is running). Accordingly, $P(e) = \lim_{n_t \to \infty} \frac{n_e}{n_t}$.

Similarly, we distinguish in occupancy in weekdays, saturdays and sundays. Occupancy probability is then calculated by selecting for each interval the usage probability of the device with the highest usage. While this allows for selecting the source of information of which we have highest amount of data, this has the clear drawback of not considering the underlying interdependence between devices.

5.3.4 Providing tailored efficiency advice

An advisor widget displays a list of advices to foster behavioral change, namely by displaying the efficiency policies identified in Sect. 5.2.1. In a first stage candidate advices are formulated and are ranked as based on previous user's acceptance or rejection (See Fig. 5.6). The following list reports the pseudocode

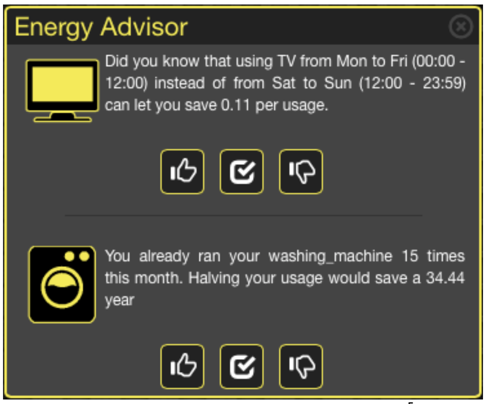

Figure 5.6: The advisor widget [Mon15]

for generating candidate advices, as resulting from Sect. 5.2.1:

- **Device diagnostics** advices replacement of appliances and it is thus useful to improve non-user-driven devices (e.g., fridge)
  1. Select non-user-driven devices
  2. Compute average consumption for each device type for all users[7]
  3. Retrieve devices whose average consumption is higher than the one for the device type of a certain threshold τ_1 (e.g., 30%) and suggest replacement

[7]Can be done periodically and cached in a separate location

- **Device shifting**
  1. Select user-driven devices
  2. Rank devices by their average consumption (according to consumption events)
  3. Rank tariffs by cost in order to select the best and worst tariffs available
  4. Suggest to use the device in the cheapest tariff and report the potential savings computed as $s = (l \cdot t) - (l \cdot c)$, respectively with l average consumption for the device, c and t cheapest and most espensive energy tariffs.
- **Shedding of standby losses** suggests to switch-off devices in standby mode (such as displays, decoders, DVD players, battery chargers without load, air-conditioning systems) in periods of not use (e.g., night). The advice can be returned to all devices with a standby mode, based on the building occupancy model. However, higher effectiveness can be achieved by also exploiting available device usage models.
- **Device curtailment and moderate usage**
  1. Select user-driven devices;
  2. Rank devices by their positive deviation from the average number of usage for the device type and cost;
  3. Suggest to reduce the amount of times the device is being used and compute the yearly savings by multiplying the running cost spent for the current month;

After their formulation, the most effective advices are selected to limit the information displayed to the user. In particular, we indicate with conversion of an advice into a behavior when the user explicitly accepts the recommendation. A conversion causes the disabilitation of the advice, in order to minimize user's discomfort. A feedback to an advice can be formalized through the tuple: $(user, advice\_type, device\_type, action, time)$. In this way, it is possible to omit advices which were previously converted into a behavior (i.e., goal) or involving device types and advice types with low acceptance (i.e., negative feedback). This is based on the results presented in [Bee13], which showed that displaying multiple times the same recommendation does not improve the conversion rate unless the user has a big opinion drift.

Consequently, the following 3-item Likert scale was used: "Ok thanks", "I'm already doing it", "No thanks" (See Fig. 5.6). Each advice can accordingly be formalized through the tuple ($user, advice\_type, device\_type, enabled, score$). A feedback of kind "I'm already doing it" directly causes the deactivation of the advice. A usefulness score is then computed for active advices using the votes resulting from "Ok thanks" and "No thanks". Such value is used to rank the advices, while randomness is used to order advices with same usefulness value. Positive feedback reinforces the advice by increasing its score, whereas negative feedback can result from a reluctance in operating the device or a mistrust in the specific advice type. Upon clicking on "No thanks", the user is asked to select one of the two causes. Based on this information, we decrease the score of all advices of the same type, that is, they either involve the same advice type or device type.

5.4 Acceptance of the advisor widget

As it was previously discussed, the estimated policies confirm that savings can be detected throughout the automatic analysis of energy usage. However, a relevant aspect to be addressed is how such opportunities should be presented to users in a compelling way that can foster behavioral change. To validate the usability of the advisor widget we decided to carry out a validation test on actual users. In particular, we were interested in assessing the effectiveness of the widget in informing and persuading users, as well as their satisfaction towards the means.

The target audience included householders of any age capable of using basic functionalities of computers, invited without the promise of a credit for the participation. We ran a total of 7 participants, all between 25 and 60 years. All participants had normal or corrected-to-normal vision. In particular, 3 subjects wore glasses during the study. None of them reported eye disorders, such as color disfunctions. It is important to remark that there is no optimal number of participants for this kind of tests. Whilst Virzi suggested 5 users being enough to spot 80% of usability problems [Vir92], 15 users is normally being considered as the upper number for this purpose [Nie93].

All subjects were initially informed of the widget functioning prior to being positioned in front of a desktop computer displaying the interface. The subjects were asked to use the "think aloud" protocol while following specific instructions to interact with the interface [Jør89]. We used a synthetic setup with the following appliances: coffee machine, washing machine, dishwasher, playstation 4 and television. A satisfaction questionnaire was finally given to rate the attractiveness of the design. Specifically, we used a 5-point Likert scale, with "strongly agree" as a left anchor and "strongly disagree" as a right anchor.

5.4.1 Results

Fig. 5.7 shows the results of the questionnaire, with −2 associated to the negative anchor and +2 to the positive anchor. The entries correspond to the following

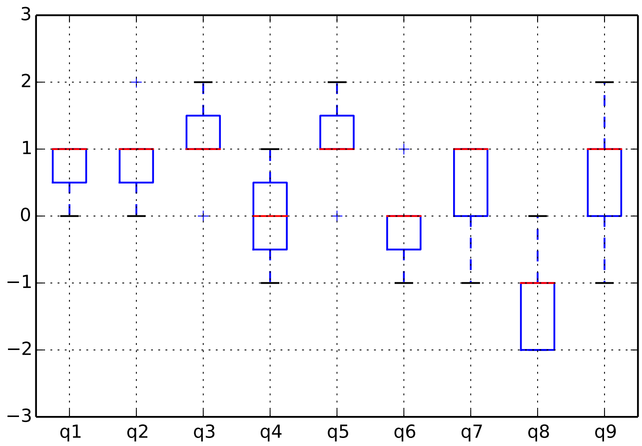

Figure 5.7: Results of the satisfaction questionnaire

questions: "it takes short time to learn the meaning of the buttons", "the position of the buttons is logical", "I understand what happens when I click the buttons", "the advices are unusual, inventive, original", "the advices are useful to improve energy efficiency", "The advices are doable", "I can learn something from the advices", "I would use this widget every day" and "I would use this widget again".

5.4.2 Discussion

All subjects immediately understood the functionality of the widget and could quickly determine which button to use depending on the meaningfulness of the displayed advice. However, it is important to remark that the advisor did exhibit a sort of cold start problem, as in some other types of recommender systems. Accordingly, given the initial absence of votes the advices are solely ranked on their estimated produced savings. However, we noticed that most users tried to get rid of all obvious advices by clicking on "I am already doing it" to remove them from the widget. The majority of the subjects commented the behavior as a curiosity to see what other advices they could learn from. Users commented the advisor as "useful" and in a couple of cases as "obvious" and declared not to be willing to use it every day. We remark that a mechanism is missing to keep users involved after the advices have been displayed and the initial learning phase has been overcome. A user commented the advice acting on stand-by consumption as "obvious" and "useless" and expressed the necessity for automatic means to shed such a consumption. The advice related device diagnostics is considered as the most useful, as it gives an estimation of possible savings that are not directly visible to householders. Another aspect that emerged from the guided interaction with the interface is also the lack of common sense in some advices. While advice applicability determines which devices a certain advice can cover, there is necessity for further mechanisms beyond the sole user's preference to select the candidates. An example is the advice "did you know that using the coffee machine from Sat to Sun ($00:00$ to $23:59$) instead of from Mon to Fri ($00:00$ to $23:59$) can let you save 0.03 € per usage?". While postponing certain devices to cheaper energy prices can yield savings, it should be possible to further diversify devices for their sensitivity to being shifted over time. In fact, rejection of the advice can prevent it from being applied to the same device type in future. However, the advice can be further improved by considering the actual usage model for covered devices.

5.5 Deployments

The 0.3 version of Mjölnir was deployed as a demonstrator at the premises of the Alpen-Adria-Universität Klagenfurt. In particular, the system consists of (See Fig. 5.8):

- a Carlo Gavazzi EM24 **aggregate power meter**, which can be interfaced using an industrial Modbus/RS485 bus.
- **an appliance-level monitoring system** based on the commercial Plugwise[8] kit.
- **a webserver running the dashboard**. The gateway is based on a Raspberry Pi and a Libelium RS485 hat[9], both cased in a DIN-Rail enclosure. The gateway runs a linux daemon[10] collecting and processing measurements from the meters, as well as listening for incoming device control events to be actuated on the network of smart outlets. In particular, real-time device monitoring and control is implemented using websockets. This allows a bidirectional communication between the dashboard and the server, as well as between the gateway and the server, even in presence of a Network Address Translator (NAT). In addition, the underlying blocking mechanism prevents continuous polling of devices for monitoring status changes.

[8] https://www.plugwise.com
[9] https://www.cooking-hacks.com/rs-485-modbus-shield-for-raspberry-pi
[10] http://sourceforge.net/projects/mjoelnir/files/Mjoelnir_Gateway.zip

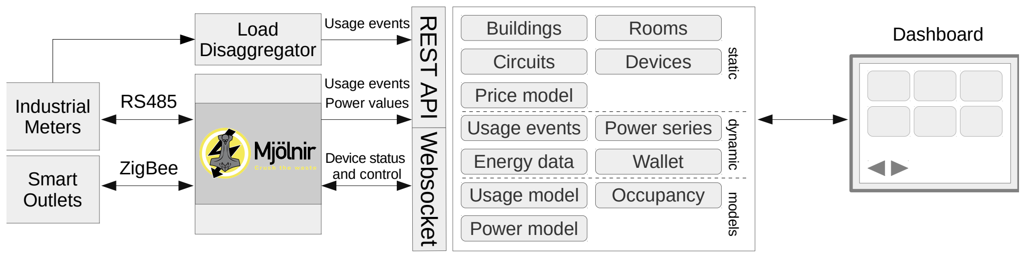

Figure 5.8: The deployed solution

The main purpose is the deployment of a demonstrator that could act as testbed for the assessment of new feedback mechanisms. As compared to residential settings, public spaces (e.g., universities) need to be addressed using different feedback means [Uta13, M1̈0]. Accordingly, in Table 5.1 we identified multiple roles and possible involvement strategies. Whilst the involvement of administrative personnel and employees is related to economical factors, the involvement of users has to be incentivized in multiple ways. In Mjölnir managers can display consumption at building or room (i.e., circuit level) and device level. In addition, the distinction between users and employees is done with public and private pages, where different kinds of widgets can be placed. The public page is displayed to users on public displays, while the private pages are meant to provide a full report to the managers.

Table 5.1: Use cases of energy management systems in public spaces

Role	**Involvement**	**Information**	**Action**
Owner (Administrator)	cost minimization	energy costs	personnel management
Energy Manager	salary	energy hogs	device replacement
Users (Students)	gamification and social competition	energy performance	informing manager and employee
Employees (Prof.)	performance evaluation (e.g., course feedback)	energy performance	device curtailment, selection of appropriate lecture classes and more efficient behavior

5.6 Summary

In this chapter we assessed consumption scenarios in Italy and Austria to identify potential room for intervention. One possibility to increase the feedback resolution in the regions is to implement prepaid billing at device level. As the most effective feedback is that tailored to the end-users' behavior we introduced an energy advisor. The advisor implements a list of common practices which were previously shown to have a potential of up to 34% of savings. The advisor can automatically process consumption data to return most meaningful policies for intervention. Accordingly, the information overload is kept low by considering the history of interaction with the users. The experiences led to the design and implementation of Mjölnir, a web-based dashboard where various feedback means can be displayed and assessed.

CHAPTER 6 Automating energy management

"The first rule of any technology used in a business is that automation applied to an efficient operation will magnify the efficiency. The second is that automation applied to an inefficient operation will magnify the inefficiency."

– Bill Gates

Chapter 5 assessed the room for improving awareness in domestic environments and proposed feedback means for the purpose. However, this might not provide a complete solution to the stabilization of the power grid, as it expects customers to actively react to changes on energy prices. Self-managing systems are therefore necessary to assist the scheduling of electrical loads.

This chapter deals with the design of controllers for small energy prosumers, which can communicate through an energy price and aim at minimizing operational costs. We identify devices that can be automatized and investigate possible controller representations. Towards this vision, we introduce a simulation tool to learn appliance controllers. While we initially apply market mechanisms used in the wholesale markets, we highlight the necessity of trading power in small power grids.

6.1 Microgrid modeling

A microgrid is a small power system built from the aggregation of local energy sources and small loads, and it is able to operate as an independent power island if necessary [Col09]. We therefore distinguish in (Fig. 6.1):

- a **smart meter**, which is a truth-telling agent metering the energy exchanged through the main power grid and the local grid. Its task is to expose an energy tariff (*get*) constrained to a power availability function, as well as a feed-in tariff (*fit*) and a power capability function.

- a set of **smart loads** bidding to allocate power on behalf of the residents. As discussed previously in Sect. 4.2.2, each appliance is a collection of services, each described by an *operation model* and a *usage model*. The operation model describes the coordinated operation of the system components in terms of a state sequence, in which a state σ_i is defined as a peak power level $P_i \in \mathbb{N}^+$ and a duration $d_i \in \mathbb{N}^+$ in seconds. Each state is associated to a *device start delay sensitivity* χ^b modeling the responsiveness in seconds, and a *state start delay sensitivity* χ^s modeling the sensitivity to a delayed start for intermediate states, as well as an *interruption sensitivity* χ^i defining the severity under which the state operation can be interrupted. Consequently, we can distinguish in: i) a *device begin discomfort* δ_b proportional to the overwaited time between the first offer and the beginning of the operation, ii) a *state begin discomfort* δ_s proportional to the overwaited time between the ending of a state and the beginning of the next and iii) an *interruption discomfort* δ_i having severe influence on the overall device operation. Based on their start delay sensitivity, appliances are classified in flexible and inflexible. A *usage model* defines the probability of desiring to operate an appliance at a specific time instant. In its simplest setting, we have a willingness $\omega^* \in [0,1]$ associated to a decay $\lambda \in \mathbb{R}$, which updates the probability based on concluded device operations. Given that $P_{use} = 1 - (P_{hold})^N$ is the probability for an OFF-ON event within a time interval of length N, $P_{use} = 1-(1-\omega^*)^N$ and consequently $(1-\omega^*)^N = 1-P_{use}$ and $\omega^* = 1 - \sqrt[N]{1-P_{use}}$. These values can be extracted from a consumption dataset using appropriate tools such as [Mon14b] (Sect. 6.2.3). A *price sensitivity* function ψ specifies the highest unit price users are willing to pay to operate the load at a given time.

- a set of **local generators** and **energy storage elements** to accumulate excessive energy. A *reservation price* function is used to model production costs, which depend on both technological and storage costs.

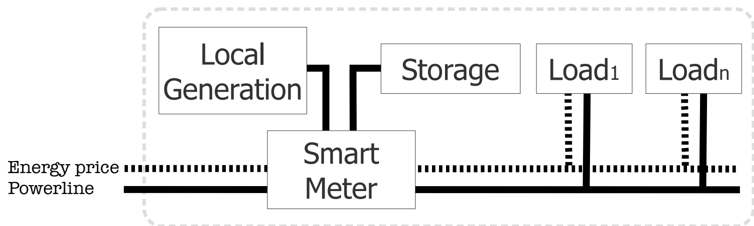

Figure 6.1: A smart microgrid

6.2 The HEMS simulator

To ease the design of smart controllers for energy prosumers we implemented a simulation tool, which is part of the FREVO evolutionary computing framework[1]. The called "home energy market simulator" (HEMS) tool targets both energy efficiency in terms of management policies, as well as comfort of inhabitants. The FREVO can handle both feed-forward and fully-connected networks, where each neuron is connected to every other and itself via several connections. Each connection is multiplied to a weight, and each neuron associated to a bias. Evolutionary algorithms are employed to optimize the weights (see Sect. 2.6.4).

6.2.1 Modeling a scenario

A scenario description is a JSON (Javascript simple object notation) dictionary specifying: i) a weather model, ii) grid connections, iii) local generators and iv) electrical loads (see [Mon14d] for a complete documentation). Models of weather, price and power series can be specified as a list of time intervals or as an external timeserie. Similarly, production of local generators can be computed using the weather model according to the techinques presented in [P1̈4], as well as directly read from external timeseries. For the available photovoltaic generators, weather models indicate the amount of sunlight intensity. This can be computed using internal sun models, depending on location and positioning of the generator (i.e., actual exposure), or can be directly collected from a weather station.

6.2.2 Simulation interface

A graphical user interface is provided in order to display the simulation status and to plot charts of the selected measures (see Fig. 6.2). In particular, the top left panel shows the logical topology of the power provisioning. Generators are represented as rectangles and loads as ellipses, which are marked green while inactive and red during operation. Edges indicate the direction of the power provision, and are annotated with the cost and amount transferred. The balance of each agent is reported on the top right panel. To characterize the ongoing simulation, multiple tabs at the bottom are used to display charts. The JFreeChart[2] library is used for the purpose. Each tab can be exported to an external file, as a picture, comma-separated value (CSV) file, as well as LaTeX TikZ. The *price* tab displays the average grid energy price and the local energy

[1] http://frevo.sourceforge.net
[2] http://www.jfree.org/jfreechart/

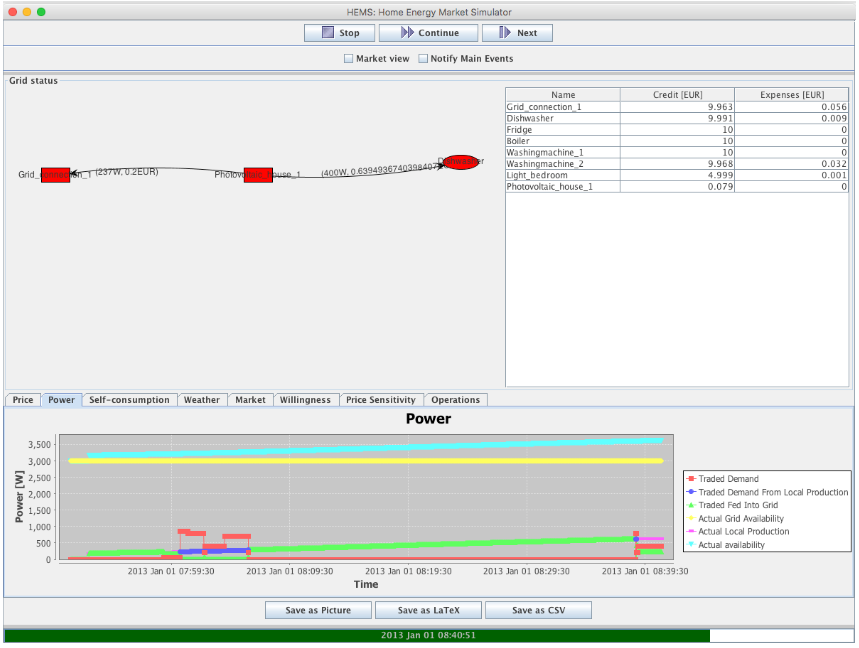

(a) The running simulation

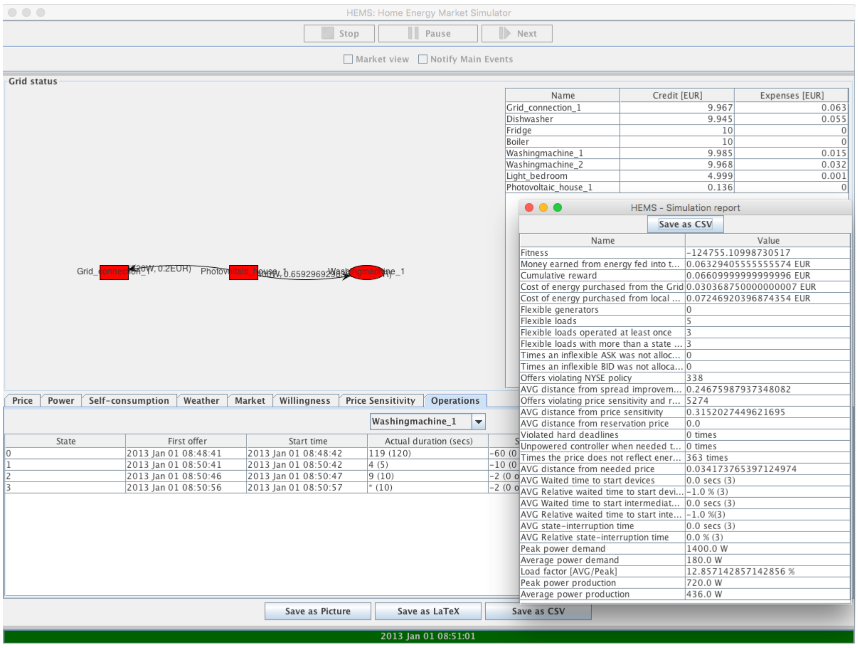

(b) The simulation report

Figure 6.2: Graphical user interface of the HEMS simulator

price, with the latter being computed as the average energy price for all running transactions. This equals the grid tariff when no energy is traded and the feed-in tariff when the whole local production is sold back to the grid. This shows the minimum price a device has to pay to start operating. The *power* tab displays aggregate power produced and demanded locally, while the *self-consumption* tab shows the exploitation of local generation. Weather conditions such as the sun factor are displayed in the *weather* tab. The *market* tab displays the market efficiency of the currently allocated devices. Since energy bought from the grid is charged under grid tariffs, the efficiency in presence of such transactions is 1 because $\Pi^a = \Pi^e$. Consequently, the efficiency measure only considers transactions between local generators and loads. The *willingness* tab shows the trading willingness of each agent, i.e., the tendency to buy or sell energy over time. We also show the price sensitivity and the reservation price of all agents in the *price* tab. The *operations* tab lists for each device all ran and running states. This serves to further analyze delays and state interruptions. Another view is available to analyze step-wise evolution of the market for the ongoing trading day (Fig. 6.3). The view shows the orderbook, including all ASK and

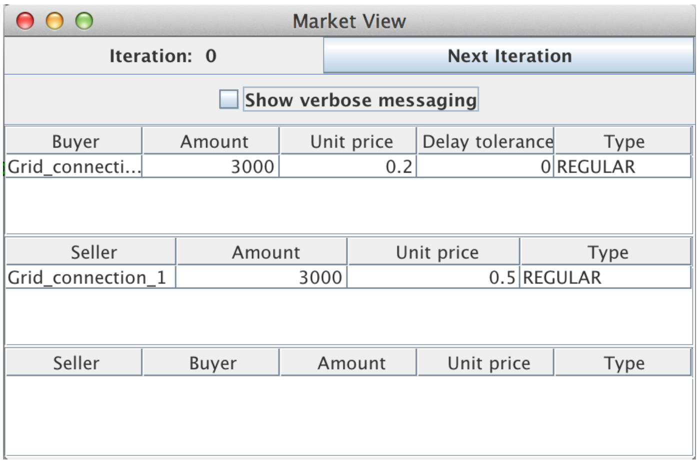

Figure 6.3: The market view

BID offers being received and matched over time. At the end of the simulation, a report prompts the overall result, consisting in a list of performance measures and the achieved fitness (Fig. 6.2).

6.2.3 The appliance-usage model manager

The appliance usage model manager (UMMA) is a tool for modeling device demand. Usage models can be extracted from energy datasets (see Fig. 6.4b) and used in simulations within the HEMS (see Fig. 6.4c). For the latter we use a TCP socket. This encapsulates the model management process into the manager, with the HEMS requesting the probability for a certain time interval via the socket. In particular, event datasets should be written in the CSV format as in Table 6.1. To ease the task of event extraction from power readings, a script is

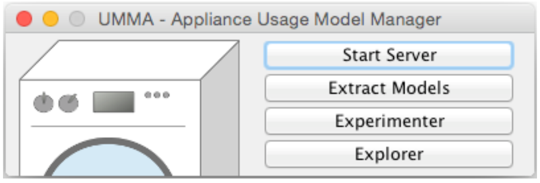

(a) The main window

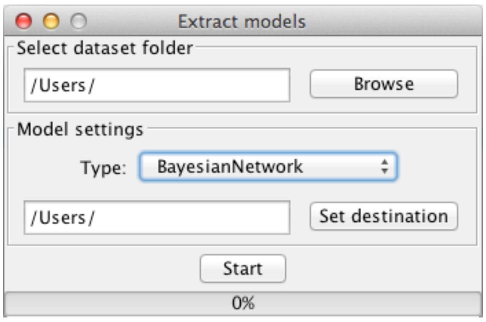

(b) The model extraction interface

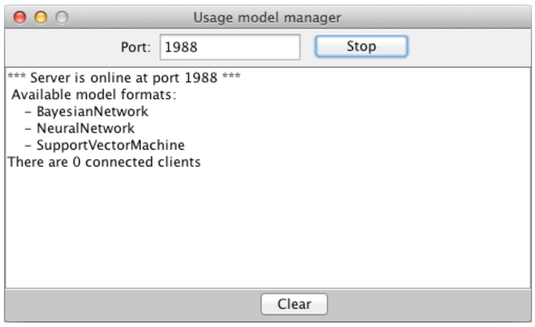

(c) The usage model server

Figure 6.4: The GUI

available on the UMMA project page [3]. Once parsed, values are normalized

Table 6.1: Event data format

Month	Hour	Weekday	ConcludedOperations	Starting
...	...	...	...	...
March	H7	Weekday	0	Start
March	H7	Weekday	1	Hold
March	H7	Weekday	2	Hold
March	H7	Weekday	3	Hold
March	H7	Weekday	4	Hold
...	...	...	...	...

(i.e., to a dataset in the range $[0, 1]$) using $x' = \frac{x - min(x)}{max(x) - min(x)}$. This prevents features with greater numeric ranges to dominate others. Interoperability for different technologies and formats is achieved at interface level. In particular, each model is provided with the current time and the number of operations already performed in the current hourly interval (see Fig. 6.5). This is important for those user-driven devices, such as coffee machines, which are generally used multiple times in the same interval. The tool supports the following technologies:

- Bayesian networks, for which we use the Netica Java library[4]. In particular, we model the network as in Fig. 6.6. Depending on the availability

[3] https://sourceforge.net/projects/umma/files
[4] http://www.norsys.com/netica-j/

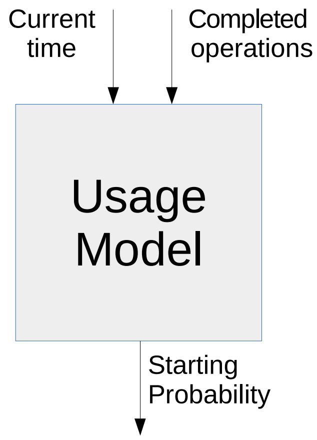

Figure 6.5: The device model interface

of longer-term measurements, the user can decide whether to include seasonal information (i.e., the month node). Learning of the network takes place using the Expectation Maximization (EM) algorithm, whereas the estimation of usage probability is done by means of the Junction tree algorithm [Kol09].

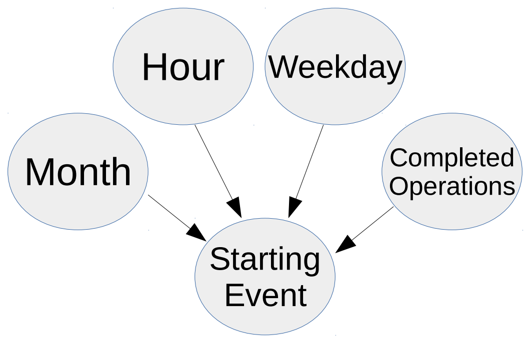

Figure 6.6: The Bayesian network

- artificial neural networks, for which we use the Neuroph Java library[5]. The ANN is a multilayer perceptron learned with dynamic backpropagation[6].
- support-vector machines, for which we use the libSVM java library[7] [Cha11].

It's important to remark that default settings were defined based on the GREEND dataset. Therefore use in different application scenarios will re-

[5] http://neuroph.sourceforge.net

[6] `http://neuroph.sourceforge.net/javadoc/org/neuroph/nnet/learning/DynamicBackPropagation.html`

[7] `http://www.csie.ntu.edu.tw/~cjlin/libsvm/`

quire further tuning of the models. To this end, the experimenter provides a benchmarking tool were the models can be assessed. The main experiment available is the k-fold cross validator, which extracts k randomised folds from the event dataset. The first k-1 folds are used for training and 1 fold for testing. The accuracy, mean-squared error, as well as the root-mean-squared error per each fold and on average are returned. In particular, the accuracy is defined as:

$$Acc = \frac{TP + TN}{TP + TN + FP + FN} \tag{6.1}$$

with TP and TN respectively as the number of true and negative cases correctly classified. Accordingly, FP and FN are respectively the number of true and negative cases incorrectly classified.

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - y)^2} \tag{6.2}$$

For an example application of the tool we refer to [Lak15], where we analyze

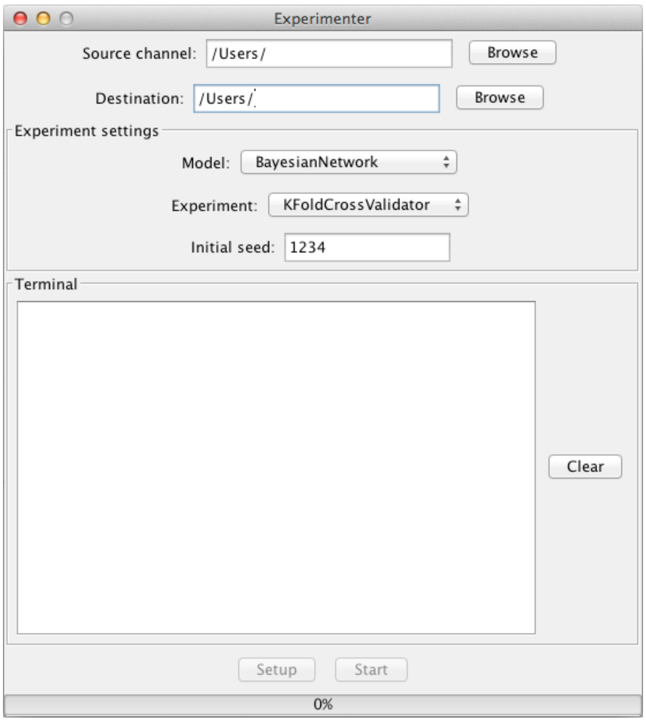

Figure 6.7: The usage model experimenter

appliance usage in the GREEND dataset.

6.3 Traders for energy markets

In this section, we design neural controllers for smart prosumers using an evolutionary learning process. Part of this work was presented in [Mon14c]. For the purpose, we use a market-based approach typical of wholesale energy markets. Therein the allocation is decided for hourly intervals on a day-ahead basis, due to the physical limits for the actuation of generators. However, given the more stringent requirements of microgrids, our solution provides an infrastructure for power trading at a second resolution. This is important to avoid underuse of resources, as appliance states can demand provisioning over seconds rather than hours.

6.3.1 Controllers for smart prosumers

Our approach for the design of smart controllers follows the methodology of [Feh10], who applies evolutionary methods to train ANN controllers. In particular, a controller is designed to trade energy for both electrical loads and local generators (see Fig. 6.8). We use a fully-meshed ANN controller

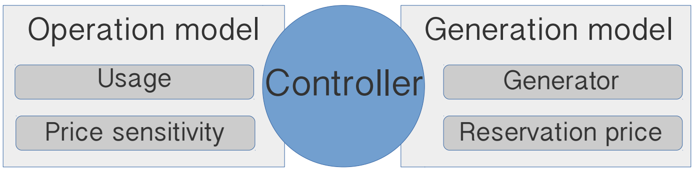

Figure 6.8: The agent structure

representation, which is already availeble in FREVO. In particular, a trading tendency $\tau \in [-1.0, 1.0]$ (with -1 to sell, 0 to skip the trade and $+1$ to buy) is computed to respectively reflect the availability and necessity of power to be traded. The trading tendency is computed using the leftover of local available power, which is always used to firstly satisfy local demand. The tendency is the most important input for prosumers, such as batteries, where it can assume continuous values to reflect the amount of charge. It is also important to remark that the starting willingness of a load is initially determined by its usage model (i.e., $\omega = \omega^*$). However, after making a first BID offer, the agent engages for the allocation (i.e., $\omega = 1.0$), which ensures rationality and commitment for pursuing the operation of the device.

An action is found using the controller each time the agent wants to buy or sell energy, according to the following structure:

- seller's inputs, which include the reservation price (ψ_{s_i}/p_{max}), the unit price of the outstanding ASK (a_{min}/p_{max}), the position in the ASK orderbook (with 1.0 denoting first and 0.0 the last), the percentage of already matched ASK offer ($P_i^{reserved}/P_i^{demanded}$);

- context information, including the time the decision is being taken, i.e., the hour (midnight is 0.0, 11 pm is 1.0), month (January is 0.0, December is 1.0) and weekday (sunday is 0.0, weekdays are 0.5, saturday is 1.0);
- trading tendency, which includes the offer importance (1.0 for inflexible and 0.0 for flexible) and trading tendency τ;
- buyer's inputs, which model the delayed start tolerance left (χ_l^b/χ^b , with χ_l^b initially equal to χ^b and progressively decreased), the price sensitivity (ψ_{b_j}/p_{max}), the unit price of the outstanding bid (b_{max}/p_{max}), the position in the BID orderbook, and the percentage of already matched BID offer.

All inputs are provided as relative values. Similarly, given that the controller outputs a real value between 0 and 1, we scale it to $[-p_{max}, +p_{max}]$ using $p = 2 \cdot p_{max} \cdot p_{output} - p_{max}$. A market threshold p^{th} is then used to decide whether to formulate a BID ($p > p^{th}$), ASK ($p < p^{th}$) or an opt out otherwise. The primary goal of the controller is to minimize costs, while selecting a price rationally reflecting the availability of local production and the willingness to start or complete an ongoing state. Since FREVO is using an absolute ranking-based selection, there was no need to normalize fitness to positive values or to squeeze fitness values into a given number range.

The formulated fitness function is:

$$F = R + (\delta_g \cdot I_{grid}) - C, \tag{6.3}$$

The reward R is the sum of the utility delivered to users upon completion of device operation, which is the price sensitivity multiplied to the duration and power of each state described in the device profile. The income I_{grid} from the energy fed into the grid is also considered. All costs are then subtracted (Eq. 6.4) after being weighted through various penalties δ.

$$\begin{aligned} C =& \delta_g \cdot C_{grid} + \delta_b \frac{1}{B_f^o} \textstyle\sum_{b_j \in B_f^o} d_{b_j} \\ & \delta_s \frac{1}{B_f^s} \textstyle\sum_{b_j \in B_f^s} d_{s_j} + \delta_i \frac{1}{B_f^o} \textstyle\sum_{b_j \in B_f^o} d_{c_j} + \\ & \delta_i (\sum_{j=1}^{B} v_{i_j} + \sum_{i=1}^{S} v_{i_i}) + \delta_m (\sum_{j=1}^{B} v_{m_j} + \sum_{i=1}^{S} v_{m_i}) + \\ & \delta_l (\sum_{j=1}^{B} v_{p_j} + \sum_{i=1}^{S} v_{p_i}) + \delta_n (\sum_{j=1}^{B} v_{n_j} + \sum_{i=1}^{S} v_{n_i}). \end{aligned} \tag{6.4}$$

The costs include: i) the energy purchased from the main energy grid C_{grid}, ii) the discomfort resulting from user interaction (i.e., average delayed device start d_b, average delayed state start d_s, and average interruption time within states d_c), iii) the cost v_m of violating the NYSE market policy and iv) the cost v_p for trading irrationally (i.e., with losses). Given that devices have have different settings, we normalize each average to the delay tolerance of each device. User discomfort results only from flexible loads B_f, which were operated B_f^o and have more than a state B_f^s. For inflexible services we consider all time-instants v_i in which offers were not allocated, which is then penalized through δ_i.

6.3.2 Results

To assess the designed controller we employed a uniform-price double auction (UCDA). Transactions between local agents are priced under a k-pricing scheme with $k = 0.5$ (i.e., equally distributing the profit between buyer and seller), whereas transactions involving the main power grid are charged under the given tariffs (i.e., fit and get).

The size of the allocation interval depends strictly on the size of present electrical loads. On one hand, we desire operating without service interruptions, as they affect both user comfort and the correct device operation. A 15-minute allocation interval as in [Vyt10] would reserve the resource for a longer period. We chose a smaller allocation interval to prevent loads from monopolizing the resource. In particular, we chose an allocation interval of 1 second, with a 1-second long trading day taking place over multiple duration-less iterations. The selected resolution guarantees a minimal delay between trading and allocation time, and acts as an interrupt mechanism to change scheduling plans to timely react to environment changes.

To assess the designed controller we evolved a fully-meshed artificial neural network over 200 generations with the simulation properties shown in Table 6.2 and 6.3. To favor competition, the appliances were set with same price sensitivity. The simulation scenarios are reported entirily in Appendix B. The first scenario consists of a sole generation unit, with no present demand. As visible in Fig. 6.9a, the controller can properly learn to sell its local power in less than 20 generations. We complicate further the scenario by adding a pool of loads, namely: a dishwasher, a fridge, a water boiler, two washing machines and a bedside light. Fig. 6.9b shows the fitness landscape, which although less steep shows convergence to the solution in about 40 generations. The performance of the learned controller are shown in Fig. 6.10. The controllers are able to strategically bid as based on their trading tendency and price models. As before, the generators are able to sell all of their produced energy. The operated loads learned to run properly, i.e., without service interruptions and excessive delay

Table 6.2: Evolutionary algorithm properties

Property	**Value**
Random Seed	12345
Generations	200
interXover frequency	10
Mutation probability	100%
Mutation severity	30%
Percentage of elite candidates	15%
Percentage of mutated candidates	40%
Percentage of randomly selected candidates	10%
Percentage of randomly generated candidates	5%
Percentage of generated through recombination	30%
Population size	60
ANN Type	Fully-meshed net
Activation function	SIGMOID
Hidden nodes	2
Iterations	2
Mutation rate	20%
Random bias	No
Variable mutation rate	No

before starting the states. So far the learned controllers properly handled their resources within the market. However, the complexity of the allocation task is directly proportional to the amount of congestion, i.e., the number of competing appliances. Moreover, local generation was managed by a single entity. In the third scenario, the local generation unit is split into two smaller generators, independently seeking profit maximization. On one hand this complicates its coordination, and makes the allocation of multiple packages of power to the same buyer more difficult. After evolving the controllers for 200 generations (see Fig. 6.9c), the generators can still properly sell their power back to the grid. However, with exception of the fridge (200 W) the appliances are not able to properly purchase power and prevent service interruption (see Fig. 6.11) In particular, the dishwasher and the washing machines are visibly being shifted over time (see Fig. 6.12). The selected settings produce suboptimal results in scenarios with multiple small generators, which would require further coordination to fully supply big loads. This is due to the independence of consecutive trading days and the pure competitive setting created by the market mechanism. The problem gets more accentuated in presence of congestion. In the worst-case scenario loads with very similar preferences and high budget engage in price wars which might ultimately yield service interruptions. The controllers

Table 6.3: Properties of the market simulation

Property	Value
Auction type	UCDA
Allocation size	1 sec
Auction iterations	4
Simulation beginning	2013 01 01, 07:51
Simulation duration	3600 secs
Market price threshold	0.02 €
Market limit price	1.0 €
reward completed operation	1000
reward device start	100
reward state start	100
penalty device start delay	100
penalty grid use	5
penalty inflexible offers	100000
penalty state interruption	1000
penalty state start delay	1000
penalty unnecessary trading	10000
penalty unpowered controller	10000
penalty unstarted device	1000
penalty violation market rules	10
penalty violating price sensitivity	10

might be further improved by providing information of the trading tendency of other loads, and better reflecting the "opportunity cost" for operating a load into the fitness function. However, the issue is a peculiarity of the selected market mechanism, which is sensitive to the statically-sized allocation interval. Sizing the allocation interval presents the problem of internal and external fragmentation of dynamic memory allocation [Sil08]. Similarly, the division of a shared resource in fixed partitions limits the size of the process and thus the degree of multiprogramming. A larger allocation interval can be used to suit any existing load state, although this might often lead to internal fragmentation, i.e., resource underuse. Contrarily, while smaller allocation intervals can allow for timely reacting to system changes, this might cause service interruption. In this section we used 1-second-based allocation intervals. This way the market reallocates power at each trading cycle, which can result in service interruptions with competing demand. To solve this problem, we investigate in the next section on the possibility to learn a power broker able to sell dynamically-sized provisioning durations.

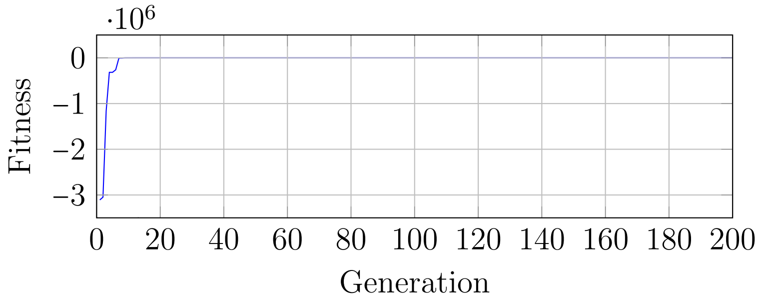

(a) Fitness development for the first scenario

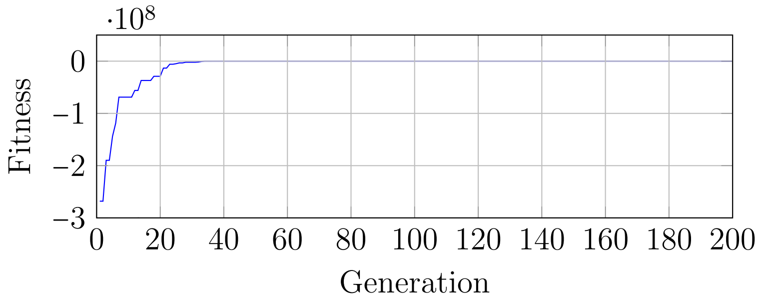

(b) Fitness development for the second scenario

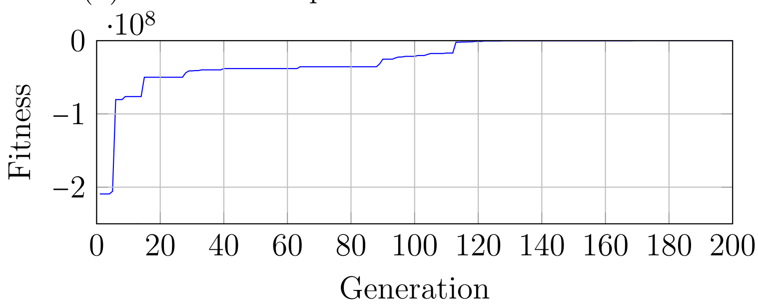

(c) Fitness development for the third scenario

Figure 6.9: Fitness development for the different scenarios

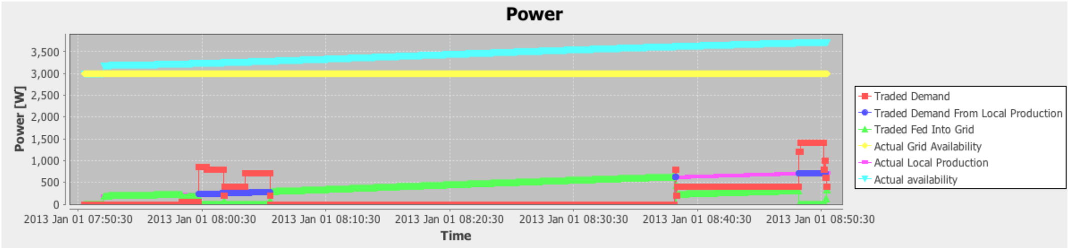

(a) The "Power" tab

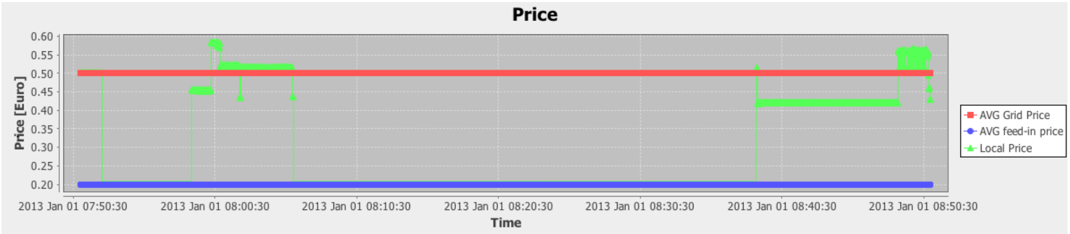

(b) Market price

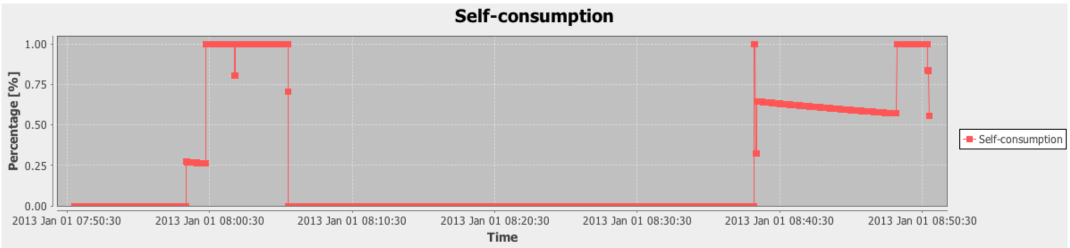

(c) Use of local production

Washingmachine_2

State	First offer	Start time	Actual duration (secs)	Start delay (secs)	State interruptions [secs]	Reward [EUR]
0	2013 Jan 01 08:00:14	2013 Jan 01 08:00:15	119 (120)	-60 (0 of 60)	[]	0.013361111111111088
1	2013 Jan 01 08:02:14	2013 Jan 01 08:02:15	4 (5)	-10 (0 of 10)	[]	1.6666666666666667E-4
2	2013 Jan 01 08:02:19	2013 Jan 01 08:02:20	9 (10)	-2 (0 of 2)	[]	5.555555555555556E-4
3	2013 Jan 01 08:02:29	2013 Jan 01 08:02:30	9 (10)	-2 (0 of 2)	[]	5.555555555555556E-4
4	2013 Jan 01 08:02:39	2013 Jan 01 08:02:40	9 (10)	-2 (0 of 2)	[]	5.555555555555556E-4
5	2013 Jan 01 08:02:49	2013 Jan 01 08:02:50	9 (10)	-2 (0 of 2)	[]	5.555555555555556E-4
6	2013 Jan 01 08:02:59	2013 Jan 01 08:03:00	9 (10)	-2 (0 of 2)	[]	5.555555555555556E-4
7	2013 Jan 01 08:03:09	2013 Jan 01 08:03:10	9 (10)	-2 (0 of 2)	[]	5.555555555555556E-4
8	2013 Jan 01 08:03:19	2013 Jan 01 08:03:20	9 (10)	-2 (0 of 2)	[]	5.555555555555556E-4
9	2013 Jan 01 08:03:29	2013 Jan 01 08:03:30	9 (10)	-2 (0 of 2)	[]	5.555555555555556E-4
10	2013 Jan 01 08:03:39	2013 Jan 01 08:03:40	9 (10)	-2 (0 of 2)	[]	5.555555555555556E-4
11	2013 Jan 01 08:03:49	2013 Jan 01 08:03:50	9 (10)	-2 (0 of 2)	[]	5.972222222222222E-4
12	2013 Jan 01 08:03:59	2013 Jan 01 08:04:00	119 (120)	-10 (0 of 10)	[]	0.011597222222222241
13	2013 Jan 01 08:05:59	2013 Jan 01 08:06:00	1 (2)	-10 (0 of 10)	[]	2.777777777777778E-5

(d) Performance for a washing machine controller

Light_bedroom

State	First offer	Start time	Actual duration (secs)	Start delay (secs)	State interruptions [secs]	Reward [EUR]
0	2013 Jan 01 07:58:52	2013 Jan 01 07:58:53	119 (120)	0 (0 of 0)	[]	0.0019999999999999944

(e) Performance for a light controller

Figure 6.10: Performance for a 1 hour simulation

Fridge

State	First offer	Start time	Actual duration (secs)	Start delay (secs)	State interruptions [secs]	Reward [EUR]
0	2013 Jan 01 07:58:05	2013 Jan 01 07:58:06	4 (5)	-60 (0 of 60)	[]	1.388888888888889E-4

(a) Fridge operation

Dishwasher

State	First offer	Start time	Actual duration (secs)	Start delay (secs)	State interruptions [secs]	Reward [EUR]
0	2013 Jan 01 07:59:15	2013 Jan 01 08:26:29	14 (5)	1573 (1633 of 60)	[2 1 4 3]	5.833333333333334E-4
1	2013 Jan 01 08:26:43	2013 Jan 01 08:26:44	4 (5)	-10 (0 of 10)	[]	1.666666666666667E-4
2	2013 Jan 01 08:26:48	2013 Jan 01 08:26:50	924 (600)	-1 (1 of 2)	[1 1 1 1 1 1 1 1 1 ...	0.03336111111111064
3	2013 Jan 01 08:42:14	2013 Jan 01 08:42:17	304 (120)	-8 (2 of 10)	[2 1 2 1 3 2 3 2 2 ...	0.009944444444444447
4	2013 Jan 01 08:47:21	2013 Jan 01 08:47:22	4 (5)	-10 (0 of 10)	[]	1.1111111111111112E-4

(b) Dishwasher operation

Washingmachine_1

State	First offer	Start time	Actual duration (secs)	Start delay (secs)	State interruptions [secs]	Reward [EUR]
0	2013 Jan 01 08:07:10	2013 Jan 01 08:26:30	708 (120)	1099 (1159 of 60)	[3 2 1 8 5 4 1 8 4 ...	0.013361111111111088
1	2013 Jan 01 08:38:18	2013 Jan 01 08:38:19	6 (5)	-10 (0 of 10)	[1 1]	1.666666666666667E-4
2	2013 Jan 01 08:38:25	2013 Jan 01 08:38:26	11 (10)	-2 (0 of 2)	[1 1]	5.555555555555556E-4
3	2013 Jan 01 08:38:37	2013 Jan 01 08:38:39	15 (10)	-1 (1 of 2)	[1 1 1 1 1 1]	5.555555555555556E-4
4	2013 Jan 01 08:38:54	2013 Jan 01 08:38:55	14 (10)	-2 (0 of 2)	[1 1 1 1 1]	5.555555555555556E-4
5	2013 Jan 01 08:39:09	2013 Jan 01 08:39:11	15 (10)	-1 (1 of 2)	[1 1 1 1 1 1]	5.555555555555556E-4
6	2013 Jan 01 08:39:26	2013 Jan 01 08:39:27	12 (10)	-2 (0 of 2)	[1 1 1]	5.555555555555556E-4
7	2013 Jan 01 08:39:39	2013 Jan 01 08:39:40	10 (10)	-2 (0 of 2)	[1]	5.555555555555556E-4
8	2013 Jan 01 08:39:50	2013 Jan 01 08:39:52	13 (10)	-1 (1 of 2)	[1 1 1 1]	5.555555555555556E-4
9	2013 Jan 01 08:40:05	2013 Jan 01 08:40:07	14 (10)	-1 (1 of 2)	[1 1 1 1 1]	5.555555555555556E-4
10	2013 Jan 01 08:40:21	2013 Jan 01 08:40:23	15 (10)	-1 (1 of 2)	[1 1 1 1 1 1]	5.555555555555556E-4
11	2013 Jan 01 08:40:38	2013 Jan 01 08:40:40	18 (10)	-1 (1 of 2)	[1 1 1 1 1 1 1 1 1]	5.972222222222222E-4
12	2013 Jan 01 08:40:58	2013 Jan 01 08:41:01	363 (120)	-8 (2 of 10)	[4 3 3 3 3 3 3 3 3 ...	0.011597222222222241
13	2013 Jan 01 08:47:04	2013 Jan 01 08:47:05	1 (2)	-10 (0 of 10)	[]	2.777777777777778E-5

(c) Washing machine #1 operation

Washingmachine_2

State	First offer	Start time	Actual duration (secs)	Start delay (secs)	State interruptions [secs]	Reward [EUR]
0	2013 Jan 01 08:02:46	2013 Jan 01 08:26:28	683 (120)	1361 (1421 of 60)	[4 2 3 8 6 5 4 4 4 ...	0.013361111111111088
1	2013 Jan 01 08:37:51	2013 Jan 01 08:37:52	4 (5)	-10 (0 of 10)	[]	1.666666666666667E-4
2	2013 Jan 01 08:37:56	2013 Jan 01 08:37:58	10 (10)	-1 (1 of 2)	[1]	5.555555555555556E-4
3	2013 Jan 01 08:38:08	2013 Jan 01 08:38:10	11 (10)	-1 (1 of 2)	[2]	5.555555555555556E-4
4	2013 Jan 01 08:38:21	2013 Jan 01 08:38:22	15 (10)	-2 (0 of 2)	[1 1 1 1 1 1]	5.555555555555556E-4
5	2013 Jan 01 08:38:37	2013 Jan 01 08:38:38	15 (10)	-2 (0 of 2)	[1 1 1 1 1 1]	5.555555555555556E-4
6	2013 Jan 01 08:38:53	2013 Jan 01 08:38:54	14 (10)	-2 (0 of 2)	[1 1 1 1 1]	5.555555555555556E-4
7	2013 Jan 01 08:39:08	2013 Jan 01 08:39:09	13 (10)	-2 (0 of 2)	[1 1 1 1]	5.555555555555556E-4
8	2013 Jan 01 08:39:22	2013 Jan 01 08:39:24	15 (10)	-1 (1 of 2)	[1 1 1 1 1 1]	5.555555555555556E-4
9	2013 Jan 01 08:39:39	2013 Jan 01 08:39:41	15 (10)	-1 (1 of 2)	[1 1 1 1 1 1]	5.555555555555556E-4
10	2013 Jan 01 08:39:56	2013 Jan 01 08:39:58	13 (10)	-1 (1 of 2)	[1 1 1 1]	5.555555555555556E-4
11	2013 Jan 01 08:40:11	2013 Jan 01 08:40:12	13 (10)	-2 (0 of 2)	[1 1 1 1]	5.972222222222222E-4
12	2013 Jan 01 08:40:25	2013 Jan 01 08:40:27	362 (120)	-9 (1 of 10)	[1 1 1 1 1 1 1 1 1 ...	0.011597222222222241
13	2013 Jan 01 08:46:29	2013 Jan 01 08:46:30	1 (2)	-10 (0 of 10)	[]	2.777777777777778E-5

(d) Washing machine #2 operation

Figure 6.11: Controller performance in the third scenario

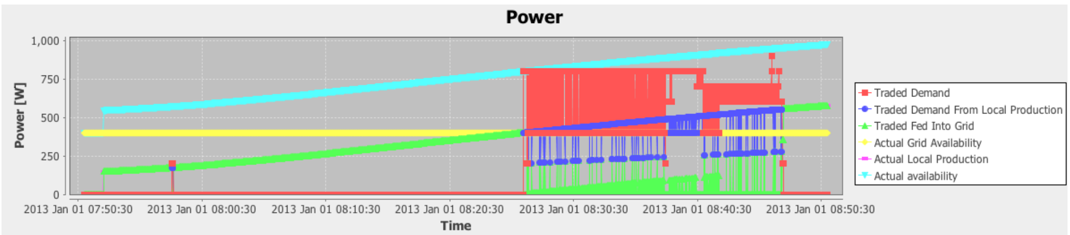

Figure 6.12: Service interruptions in the third scenario

6.4 Forward contracting in Microgrids

6.4.1 Problem statement

Energy brokerage is the problem of formulating a price describing the cost associated to the provisioning of currently available power.

In this part, we consider the power grid and the local generators as truth-telling agents, whose reservation price is given by the grid-energy tariff and the feed-in tariff, hereby represented as *get* and *fit*. For the broker, the cost to supply the local grid at a generic time t is thus given by the power drawn from the local generator and the one requested from the energy grid, as in Eq. 6.5.

$$p(P_s) = \begin{cases} fit & \text{if } P_s \leq P_{re} \\ \frac{P_{re} \cdot fit + (P_s - P_{re}) \cdot get}{P_s} & \text{if } P_s > P_{re} \end{cases} \tag{6.5}$$

where $P_s = P_{re} + P_{grid}$. Fig. 6.13 shows a scenario with 3000 W provided by

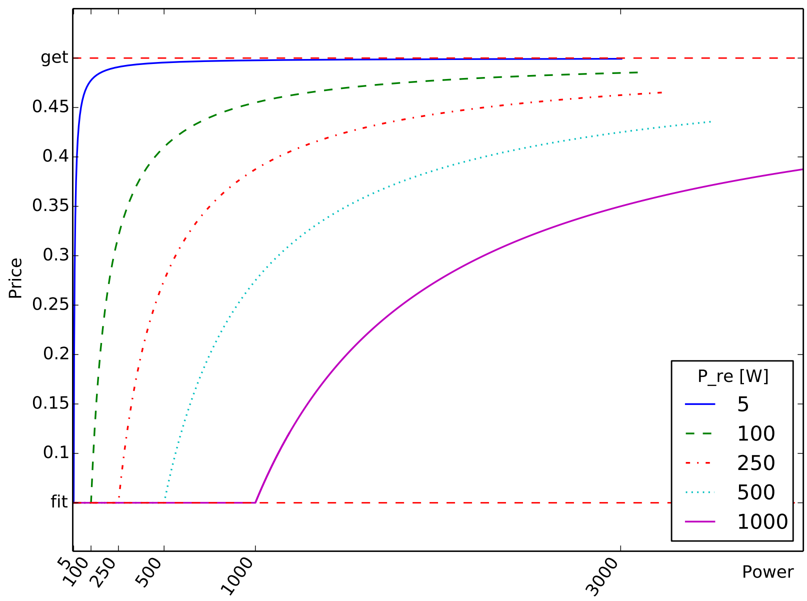

Figure 6.13: Cost function for different amounts of locally generated power P_{re}

the power grid, under a 0.05 €/kWh feed-in tariff and a 0.5 €/kWh grid energy tariff. The price is computed for different levels of P_{re}. As visible, higher amounts of P_{re} lowers the portion of P_{grid} being used, with a consequent lower price to supply the local grid. The broker is also required to minimize service interruptions by providing multiple provisioning durations or service-level agreements (SLA). Trading different durations as different products can better reflect demand differences for the service agreements. For instance, with a majority of loads with long states this would imply higher costs to purchase long-term service agreements, which would favour the purchase of short-term contracts and the mitigation of fragmentation. The example shows three loads

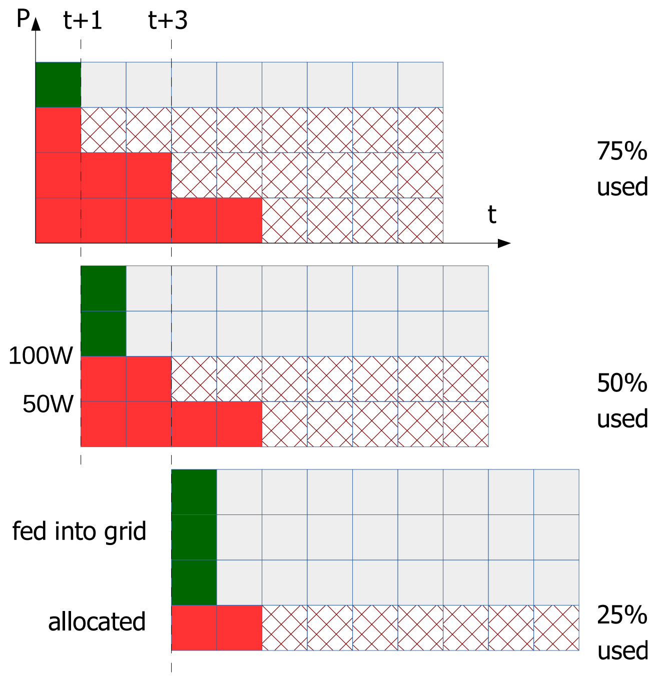

Figure 6.14: Providing multiple provisioning contracts

of same type, and consequently same demand (i.e., 50 W), competing for the allocation of power at time t. At time $t+1$ the running loads cause a 150 W demand, while the remaining power is fed into the grid. As time passes and no new allocations are matched, the service-level agreements are shifted to the left, thus resulting in the situation showed in the second and third example. The broker's objective is to maximize the profit, indicated as difference between its income and costs (i.e., to buy energy from the grid and the local generator). This includes a profit Π_{uGrid} resulting from power sold throughout the microgrid, as well as Π_{feedin} resulting from power injected back to the main power grid. Similarly, we distinguish in a procurement cost C_{supply} and a compensation cost $C_{reimbursement}$. Agreements that can not be satisfied due to insufficient supply are reimbursed. Specifically, the broker refunds involved loads with the supply cost for the remaining portion of the SLA. The overall broker profit Π is thus given as:

$$\Pi = (\Pi_{uGrid} + \Pi_{feedin}) - (C_{supply} + C_{reimbursement}) \tag{6.6}$$

6.4.2 Performance measures

For a performance comparison of different brokerage schemes we have identified the following metrics:

Peak-to-average ratio (PAR) is the ratio between the peak power and the average over the considered time window. The PAR indicates the proportion of power peaks over the overall demand and directly affects the loss of load probability. It is desirable to keep this value as low as possible, as higher PAR values denote a lower system reliability and consequently inefficiency [Liu14].

Service availability defined as the proportion of time in which a system is in a working condition. For clarity we distinguish in two more measures: the Mean Time Between Failures (MTBF) which models the average uptime between consecutive failures, and the Mean Time To Recover (MTTR) describing the average downtime due to a service recovery. The MTBF is directly related to the failure rate (i.e, the frequency of service interruption) as $\lambda = \frac{1}{MTBF}$. The availability is computed as $A = \frac{MTBF}{MTBF+MTTR}$, while the unavailability is $U = 1 - A = \frac{MTTR}{MTBF+MTTR}$.

System reactivity which describes the degree of responsiveness. In the context of energy management this can be defined for a load as the probability of having enough power to operate. We can distinguish in: i) CBP number of times the load could make an offer to get enough power and ii) $CNBP$ number of times the load could not make an offer to get enough power. These values can be collected for each load once for each trading day. Consequently, we can compute $R = \frac{CBP}{CBP+CNBP}$. Clearly, the presence of longer service agreements increases the proportion of $CNBP$ with respect to CBP, thus lowering R.

Economic profit (Π) The profit in economical terms is computed as difference between retail revenues and production costs. This is directly proportional to market power, the ability to raise the price of a good or service over its marginal cost. Market power is high in monopolies and oligopolies, and absent in perfectly competitive markets. Because the broker operates in a pure mopoly, by pricing available power to steer the system, the economic profit is a good quality measure of its performance.

6.4.3 Rule-based brokers

Given that costs are fixed by the *get* and *fit* plans, the broker seeks profit maximization by reflecting future supply costs and expected demand into the retail price. Accordingly, it is possible to distinguish in two different strategies: i) a *pessimistic* broker that charges SLAs proportionally to their duration, and ii) an *optimistic* broker that keeps the same price for all SLA durations.

The pessimistic broker matches mostly single-unit SLAs (i.e., the shortest available), which presents the problem of market competition and service interruption previously encountered. The optimistic broker attributes same uncertainty to all SLA durations. This can result in economic losses for the broker. Moreover, it favours the sale of long-terms SLAs, which reduce market competition and consequently the reactivity of the allocation mechanism.

The selected scenario is a small Austrian household with a pool of appliances and a photovoltaic power generator of 3.3 kWp [Mon13c]. The production depends on two different weather models (see Fig. 6.15): i) a clear-sky sunlight intensity computed solely according to the sun position [PÏ4], and ii) a 15-minute-resolution illuminance timeserie collected from a weather station at the University of Klagenfurt[8], Austria. Due to the lack of digital meters, the current

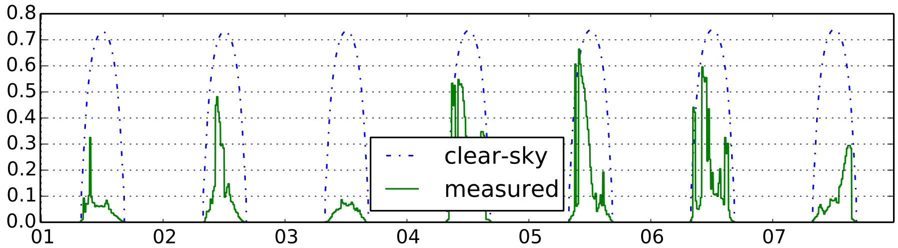

(a) January 2015

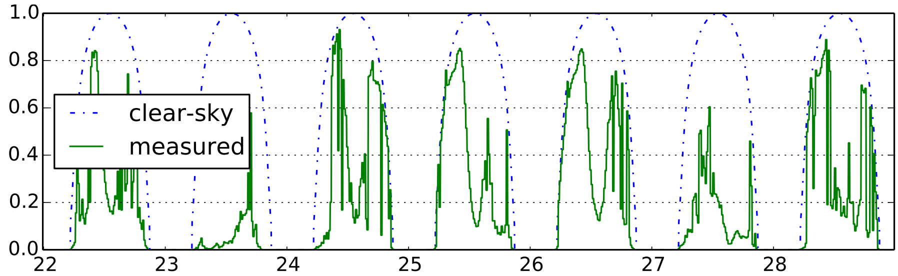

(b) June 2015

Figure 6.15: Sunlight intensity from the employed models

Austrian energy system does not implement time-based tariffs. Consequently, we assumed a pricing scheme similar to the Italian one. Accordingly, the energy exchanged with the main power grid *get* was set to 0.29 €/kWh in the interval 6 a.m. to 9 p.m and 0.15 €/kWh otherwise. The feed-in tariff fit is 0.04 €/kWh from 6 a.m. to 9 p.m. and 0.02 €/kWh otherwise. Consumption data of building 2 in the first week of 2015 (i.e., Jan 1st to 7th) is taken from

[8] http://wetter-cms.aau.at/info.php

[Mon14a] to model the usage behavior of a retired couple with an adult son. The following devices were selected: television, washing machine, dishwasher, tumble dryer and coffee machine. A low-pass filter was applied to the measurements and edge-detection techniques were employed to identify device starting events. Fig. 6.16 shows events being detected for January 1st, 2015. As visible in Fig. 6.16c the following edges were detected: for device #0 at $01:23:53$, for #7 at $07:39:18$, for #7 at $08:16:39$, for #0 at $11:59:09$, for #0 at $13:46:16$, for #0 at $16:07:55$ and for #4 at $16:25:11$. Finally, load operation was described as a sequence of states (Table 6.4) accompanied by an external usage timeserie. For simplicity, a fridge was modeled as an an additional periodic device with $\omega^* = 0.8$ and $\lambda = 0.5$. All models were implemented in the previously

Table 6.4: Simulation scenario

	Operation model $(P[kW], d[sec])$
TV	$(0.18, 3600)$
Dishwasher	$(2.1, 300), (0.1, 120), (0.3, 60), (0.1, 120), (2.1, 300)$
Dryer	$(2.5, 120)^{10}$
W.machine	$(2.1, 120), (0.3, 300), (0.2, 120), (0.6, 300), (0.2, 60)$
Fridge	$(0.2, 30), (0.16, 600)$
Coffee m.	$(2, 60)$

presented HEMS framework. The market-based allocation was disabled by setting the price sensitivity of all loads to 0.9 €/kWh. Selected provisioning durations are: a unitary agreement occupying 1 time instant, and respectively 10, 30, 60, 120, 600 and 1800 seconds. In the experiment the amount of P_{grid} available to the broker is varied: i) 0 kW, ii) 1.5 kW, iii) 3 kW in 6 a.m. to 6 p.m. and 1 kW otherwise, iv) 3 kW and v) 6 kW.

As visibile in Fig. 6.17, lowering P_{grid} causes the postponement of loads to off-peak periods. While this reduces the peak power demand, this also decreases the average power demand, thus leading to the PAR reported. The postponement of devices is not directly reflected on A, as this captures only the performance of operating loads. For instance, the fridge and the entertainment system are the only operating loads in the first scenario. Hence, lowering P_{grid} has the effect of preventing the operation of certain high-power demanding loads, which results in a lower profit Π_{uGrid}. Moreover, it is remarkable that availability and reliability are two opposite objectives to be optimized. Specifically, with a low P_{grid} the sale of long-term provisioning agreements results in resource monopolization (i.e., $R \simeq 0$). The absence of a connection to the main power grid, as in the first scenario, makes the system more sensitive to variations of P_{re}. The broker reimburses involved loads with the supply cost of the remaining portion of the SLA. The very low values are due to the pricing of the SLA,

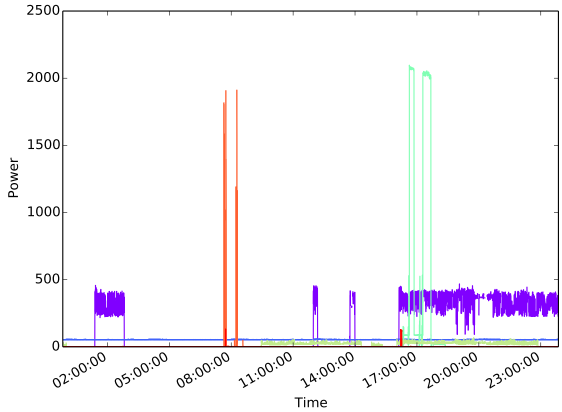

(a) Data as provided by the GREEND

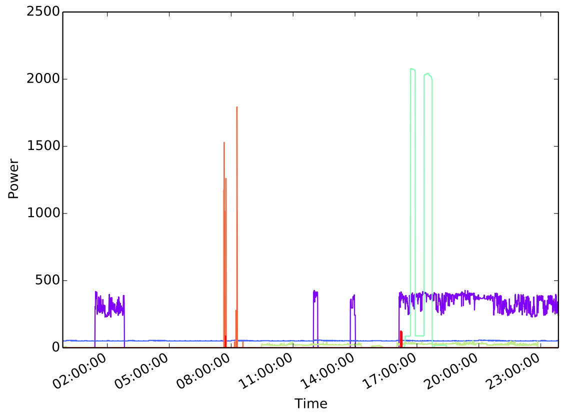

(b) Data after the preprocessing stage

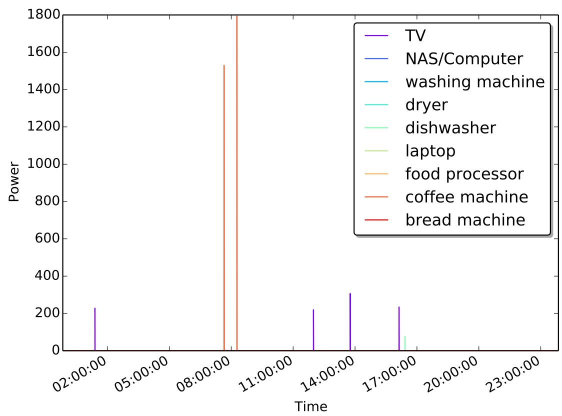

(c) Detected events

Figure 6.16: Usage behavior of January 1st 2015 for the selected household

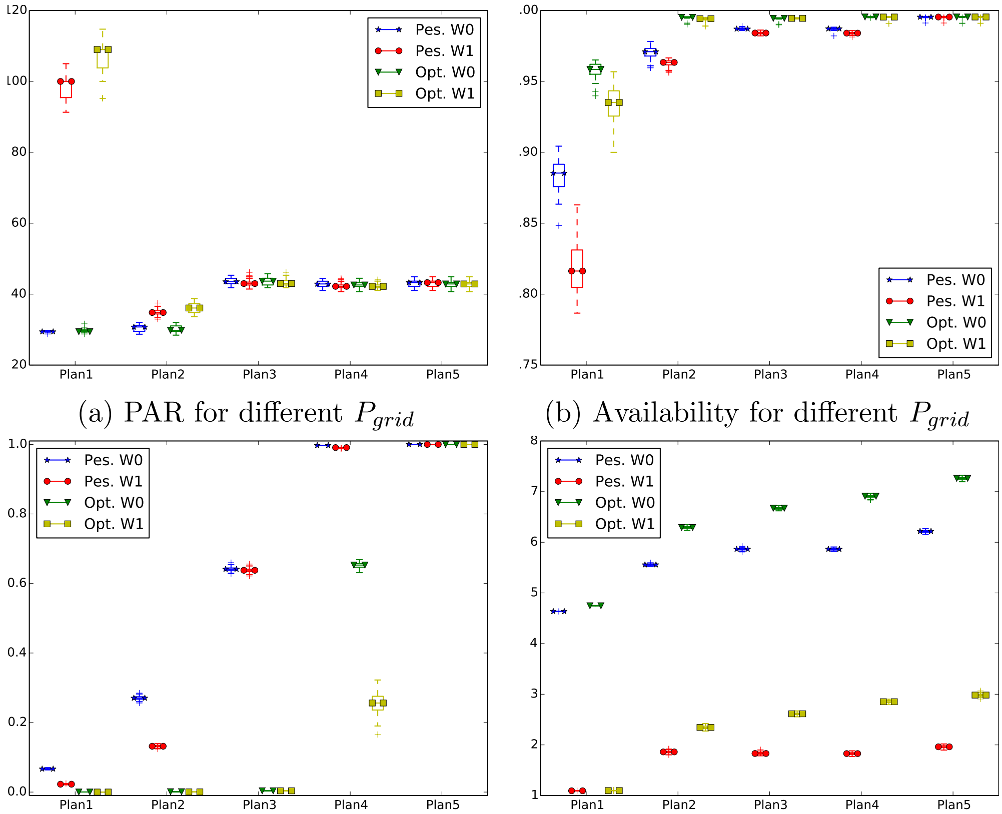

(a) PAR for different P_{grid} (b) Availability for different P_{grid}

(c) System reactivity for different P_{grid} (d) Broker's profit for different P_{grid}

Figure 6.17: Evaluation metrics for 100 evaluations

which in absence of P_{grid} is charged under the fit tariff (see Eq. 6.5). As visible from the overall profit and income values, this issue gets even more accentuated with more realistic weather models (i.e., $W1$). The weather stochasticity is reflected on the produced power, which causes higher reimbursement and supply costs to fulfill the SLAs. This demands approaches able to dynamically tune the amount of sellable provisioning agreements.

6.4.4 Learning contract brokers

The broker's objective is the maximization of economic profit and the minimization of reimbursement costs. To this end, we employ an artificial neural network, trained using evolutionary algorithms according to the fitness function in Eq. 6.6. Given the predefined fit and get price models, the broker can seek profit maximization by modeling the expectation of future resource availability. We therefore further penalize the reimbursement cost $C_{reimbursement}$ by multiplying it to a reimbursement penalty δ_r. The broker's input layer includes P_{re}, fit, P_{grid} and get, whereas the output layer consists of a price

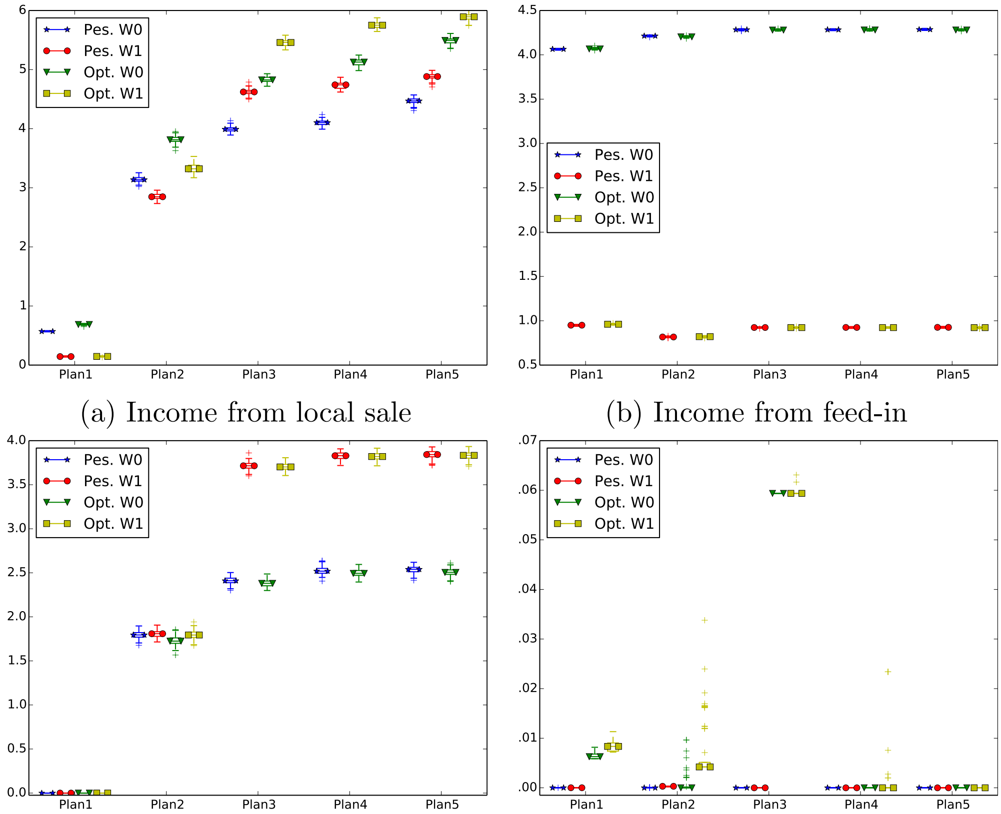

(a) Income from local sale (b) Income from feed-in

(c) Supply costs (d) Reimbursement costs

Figure 6.18: Profit components for 100 evaluations

for each provisioning duration: i.e., the unitary agreement and SLAs lasting respectively 10, 30, 60, 120, 600 and 1800 seconds. The first design choice concerns the selection of an input interface, which is related to weather conditions. In particular, for photovoltaics this is related to sunlight availability and air and cell temperature. A possibility to model seasonal patterns is thus to use: i) an input for the hour of the day and ii) one for the day of the year (Fig. 6.19a). This can be modeled using gaussian or sinusoidal functions, in order to reflect the higher availability of light in the central part of the day. Accordingly, for the ANN in Fig. 6.19a we used $\sin(\pi \cdot \frac{t}{t_{max}})$, with t used to indicate either the hour of the day or the day of the year. Another possibility is to directly use the sunlight intensity as an input to the broker (Fig. 6.19b). This can represent the expected light given the season and hour of the day, or the measured sunlight intensity (see Fig. 6.15). Therefore, the effect of different sensory input interfaces is an aspect that requires further evalutation. While artificial neural networks are universal function approximators, their ability to learn a function is in fact greatly affected by their topology. The number of neurons in the hidden layer affects the ability to generalize their

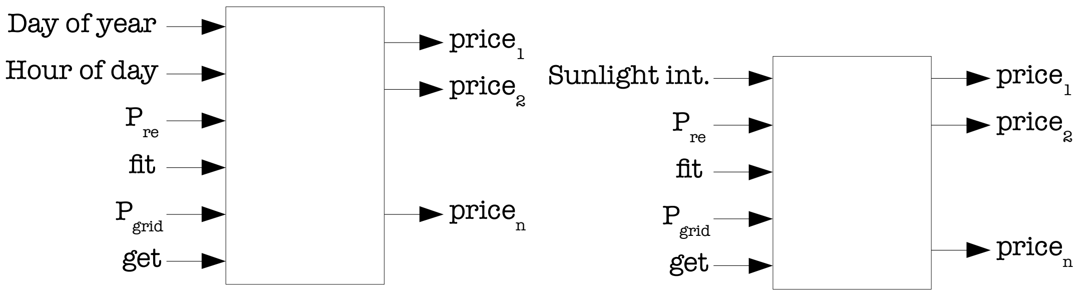

(a) Hour and day time inputs (b) Sunlight intensity input
Figure 6.19: Proposed ANN interfaces

experience, leading to overfitting when using fewer neurons and underfitting when using too many. The optimal number of hidden neurons depends on the complexity of the function to be approximated, and, therefore, indirectly on the number of input and output nodes [Zhe15]. There exist empirically derived rules-of-thumb for selecting the number of hidden neurons providing a range of possible configurations [Blu92, Swi96, Ber97, Bog97, Cau92]. In this work we selected the number of hidden neurons based on a experiments within the suggested range of 2,3, and 4 neurons, where 2 neurons showed best behavior. In addition, we applied two different representations: three-layered ANNs and fully connected ANNs. This allows for the assessment of both simple feedforward and more expressive recurrent structures. In particular, [Sie91] showed a RNN with sigmoidal activation function being Turing complete. In a fully-connected ANN each neuron is connected to itself and any other neuron, which allows for retaining a state or context, although this complicates their learning [Pas13]. Such an ability to model temporal dependencies makes RNNs especially effective for processing sequential input [Lip15]. In this study, we used both a linear and sigmoid activation function, respectively for the fully-meshed and the three-layered version.

Training the proposed brokers

For each of the scenarios used to assess the rule-based brokers, a neural network was trained using the NNGA evolutionary algorithm, with the parameters reported in Table 6.5. The settings include: ideal and real weather conditions, different season and different grid provisioning plan.

In particular, the networks were trained on a day-long simulation data (see Fig. 6.15) at 1 Hz resolution. We initially planned the training to 800 generations. However, since we noticed stabilization of the fitness already before 500 generations we shortened the simulation for time issues. The reimbursement cost $C_{reimbursement}$ is further penalized by a reimbursement penalty δ_r which we empirically set to 100000.

Table 6.5: Parameters of the evolutionary algorithm

Population size	50
Number of generations	500
Elite selection rate [%]	15
Mutation rate [%]	40
Crossover rate [%]	30
Random-creation rate [%]	5
Random-selection rate [%]	10

Results

In the first experiment, we assess the fitness landscape over different scenarios, namely: i) different grid energy provisioning plans and ii) different season and weather conditions. Fig. 6.20 shows the fitness landscape for the proposed model over the winter (Fig. 6.20a) and summer season (Fig. 6.20b). Differences in the availability of both renewable and grid-supplied energy determine different fitness for the broker. To compare the brokers against their rule-based counterpart, we used the selected performance measures. For each scenario, the best candidate network (i.e., in the last generation) is selected and used over respectively 1 day and 1 week time. Moreover, the brokers were placed in both ideal and real weather conditions (i.e., as in their simulation enviroment). Fig. 6.21 and 6.22 show the main performance metrics, from which we omit the legenda for a better clarity. For each grid plan, the points on the left and on the right to the label represent respectively results for the winter and the summer season. The symbols correspond to: i) ideal weather for 1 day (circle), ii) ideal weather for 1 week (star), iii) real weather for 1 day (plus) and iv) real weather for 1 week (triangle). Similarly to the rule-based brokers, higher amount of power resulting from the main grid or renewable sources allows for higher market volume and consequently profit. In Fig. 6.21a, the high PAR for the $Plan_0$ is due to the absence of P_{grid}, which makes not possible the operation of certain loads and results in $P_{max} = 2kW$ and $P_{avg} = 10W$. Similar differences are encountered for the summer weather as opposed to the winter weather, as well as the ideal weather with respect to the actual weather. A further remark is that the brokers correctly learned to minimize the reimbursement costs, as compared to their rule-based counterparts (see Fig. 6.24b).

To observe behavioral differences in the different scenarios we report in Table 6.6 the number of sold provisioning durations, respectively for 1 day and 7 days long simulation. The first part lists the results for the winter weather, whilst the second for the summer season. Given the aggressive trading attitude of the designed loads (i.e., independent from their price sensitivity), the brokers' pricing mechanism has more evident effects in the resource constrained settings, as in $Plan_0$ and $Plan_1$. Therein the effects of real weather conditions significantly affect the global power availability, which makes shorter SLAs favorite. Contrarily, in other provisioning plans even with very stochastic weather the power supplied by the main grid is normally enough to back the loads. Effects can therefore be observed on an economical basis, with the broker increasing the price for the resource proportionally to the grid energy tariff. Consequently, the pricing of SLAs depends strictly on the expectation of future demand. By setting the price sensitivity to 0.9 (€ / kWh) we simulated the worst possible congestion scenario, in which depending on the employed usage model all loads desire to operate regardless of the SLA pricing. In fact, users will assign different price sensitivity models to the loads, according to the delivered utility. This has the favourable effect of determining an ordering over the loads, and consequently more favourable conditions for the broker and the resulting SLA prices. The following simulation study employs the second proposed model, of which Fig. 6.25 reports the fitness landscape. In particular,

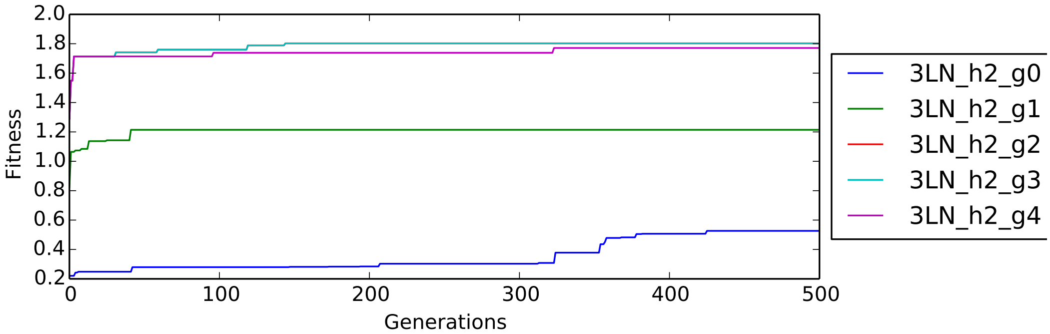

(a) Training on the winter day

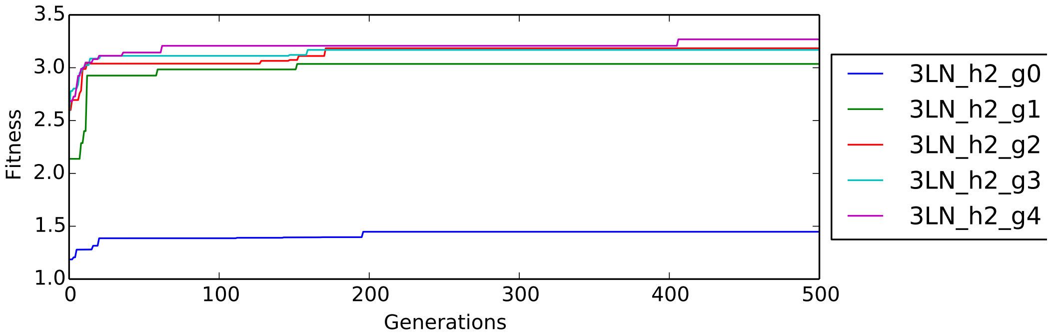

(b) Training on the summer day

Figure 6.20: Fitness landscape for the ANN_A 3LN

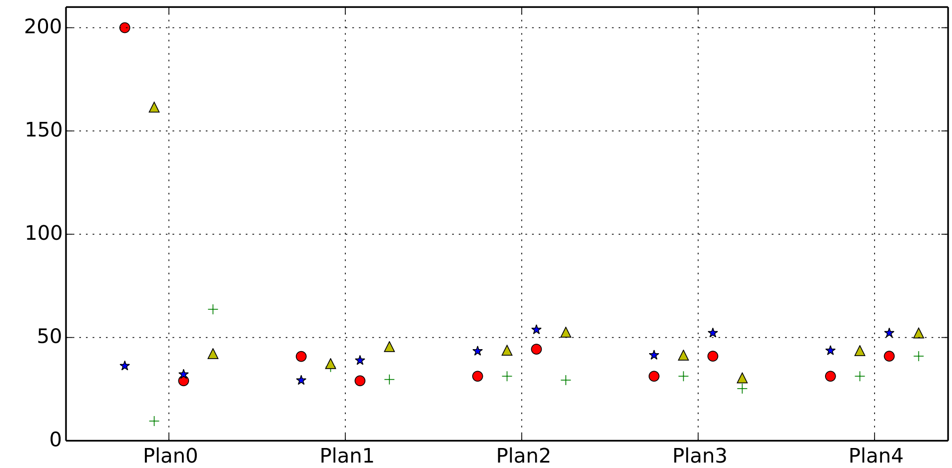

(a) PAR for different P_{grid} scenarios

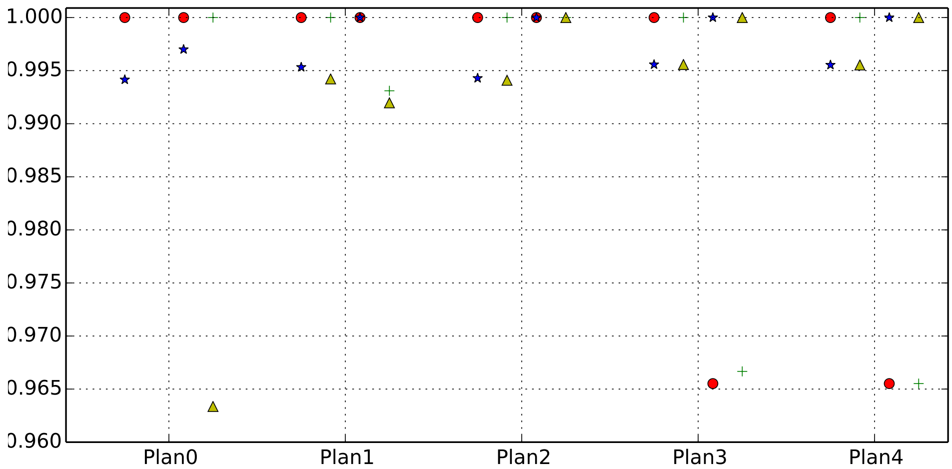

(b) Service availability for different P_{grid} scenarios

Figure 6.21: Evaluation metrics for the ANN_A 3LN

the sunlight intensity input provided to the ANN_B is the one measured in the area of Klagenfurt (Austria) (See Fig. 6.15).

Fig. 6.27 report the evaluation metrcis for the ANN_B, while Fig. 6.28 and 6.29 shows the economic profit and its components.

We now report the result of the trained brokers, using a fully-meshed neural network architecture. Fig.6.30 and 6.35 show the fitness for the fully-meshed version of the ANN_A and ANN_B. As visible, the Fully-Meshed Network (FMN) variants seem to lead to worse fitness values than their Three-Layered Network (3LN) counterparts.

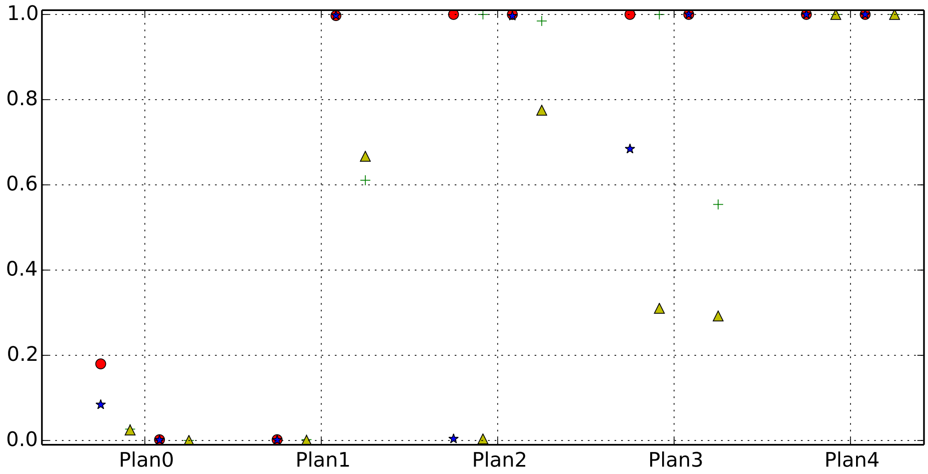

(a) System reactivity for different P_{grid} scenarios

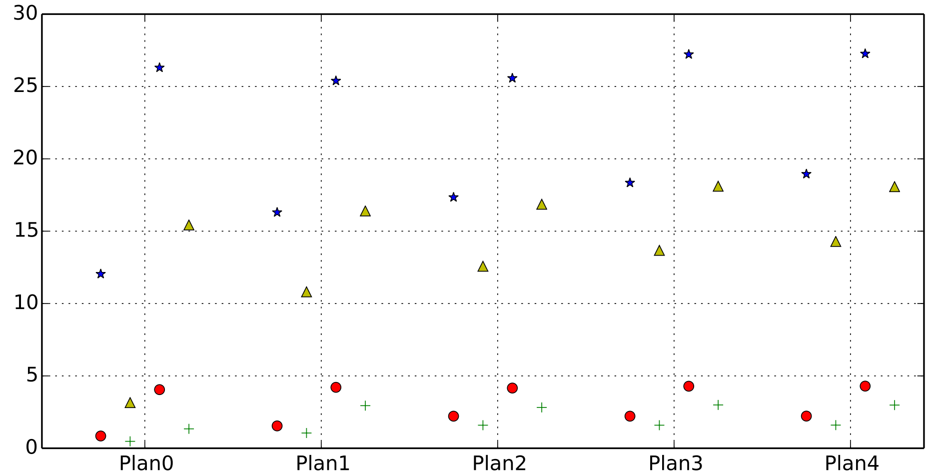

(b) Broker's profit for different P_{grid} scenarios

Figure 6.22: Evaluation metrics for the ANN_A 3LN

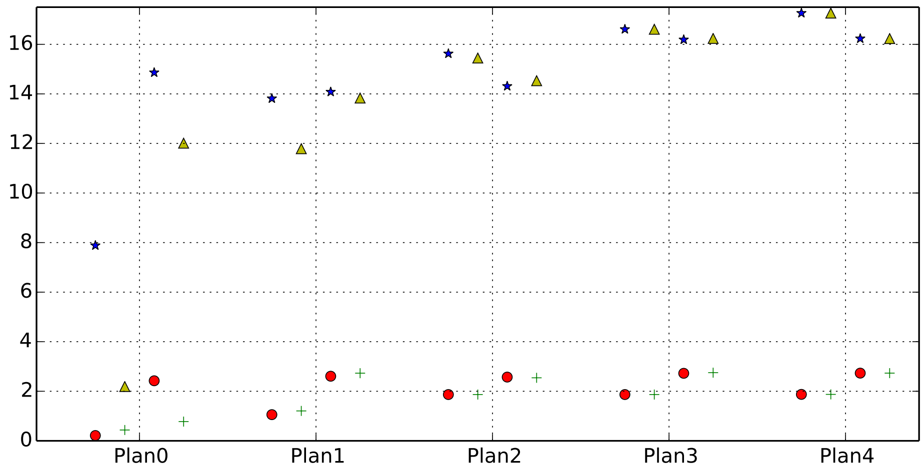

(a) Income from local sale for different P_{grid} scenarios

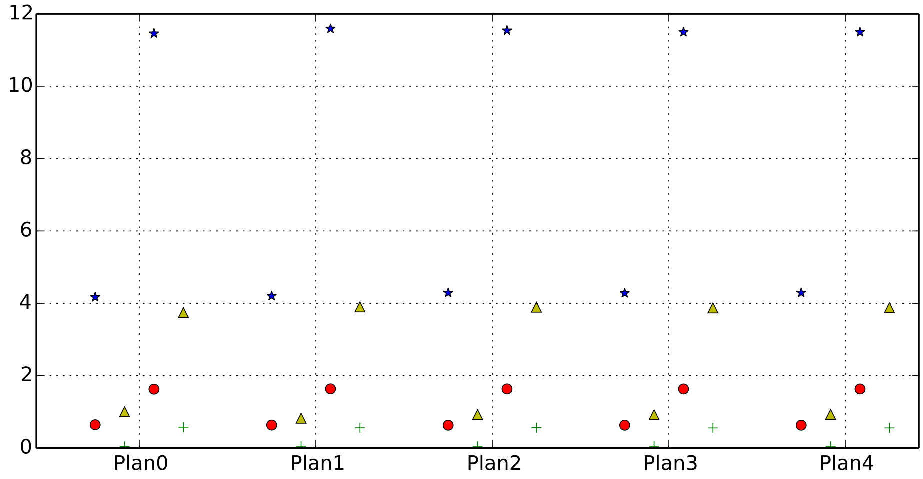

(b) Income from feed-in for different P_{grid} scenarios

Figure 6.23: Profit components for the ANN_A 3LN

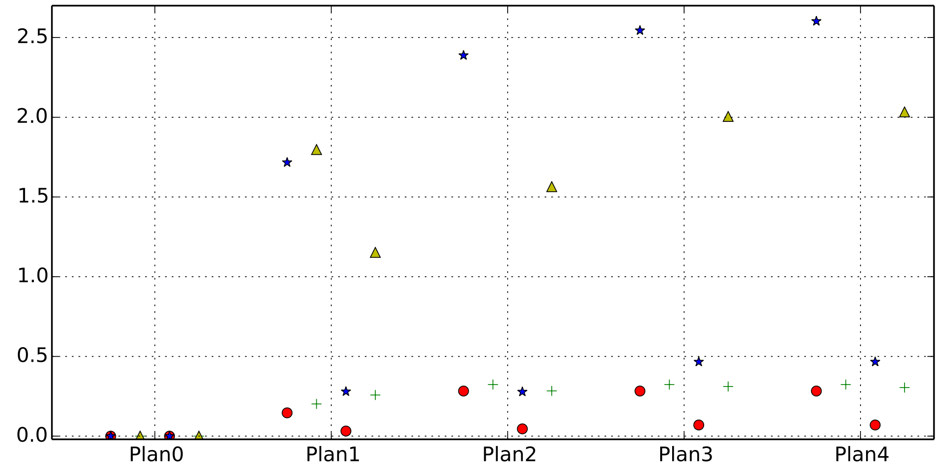

(a) Supply costs for different P_{grid} scenarios

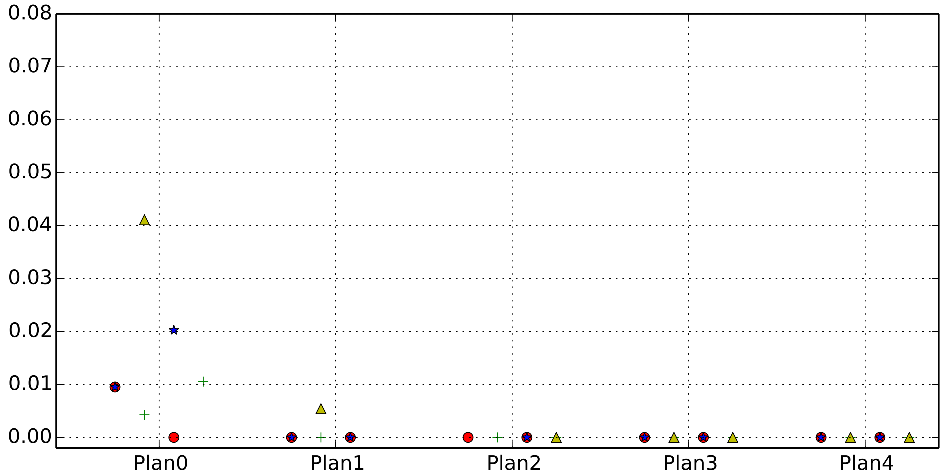

(b) Reimbursement costs for different P_{grid} scenarios

Figure 6.24: Profit components for the ANN_A 3LN

Table 6.6: Traded service-level agreements

Plan	Weather	SLA duration						
		0	10	30	60	120	600	1800
0	Ideal	0/0	0/0	8/51	1/24	0/98	8/89	0/3
	Real	0/0	0/1	5/41	0/3	0/5	11/55	2/5
1	Ideal	0/0	0/0	14/156	2/34	0/104	38/290	0/0
	Real	0/0	0/0	21/152	1/23	0/64	45/276	0/0
2	Ideal	0/0	0/0	21/144	3/43	2/106	47/280	0/0
	Real	0/0	0/0	21/137	3/43	2/106	47/273	0/0
3	Ideal	0/0	0/0	21/159	3/43	2/108	47/297	0/0
	Real	0/0	0/0	21/159	3/43	2/108	47/297	0/0
4	Ideal	0/0	0/0	21/157	3/43	2/108	47/295	0/0
	Real	0/0	0/0	21/157	3/43	2/108	47/295	0/0
0	Ideal	0/0	0/0	15/99	6/42	16/66	27/246	0/0
	Real	0/0	0/1	15/72	2/22	2/53	23/202	0/0
1	Ideal	0/0	0/0	22/158	6/42	16/66	28/184	0/0
	Real	0/0	0/1	21/145	6/36	18/70	27/170	0/0
2	Ideal	0/0	0/0	19/139	6/42	16/66	25/171	2/18
	Real	0/0	0/0	21/156	6/42	16/66	27/200	0/6
3	Ideal	0/0	0/0	20/150	6/42	16/66	38/296	0/0
	Real	0/0	0/0	21/152	6/42	16/66	39/298	0/0
4	Ideal	0/0	0/0	20/150	6/42	16/66	38/296	0/0
	Real	0/0	0/0	20/150	6/42	16/66	38/296	0/0

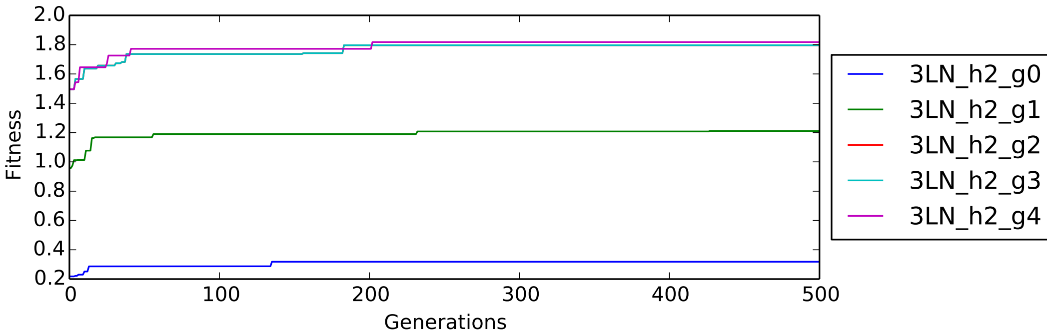

(a) Training on the winter day

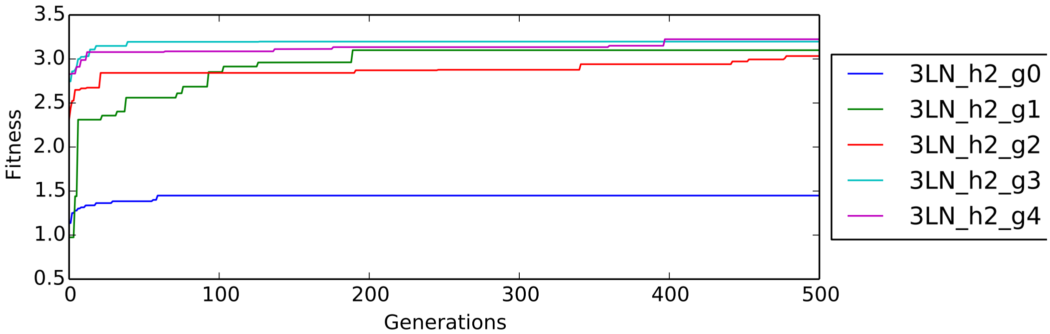

(b) Training on the summer day

Figure 6.25: Fitness landscape for the ANN_B 3LN

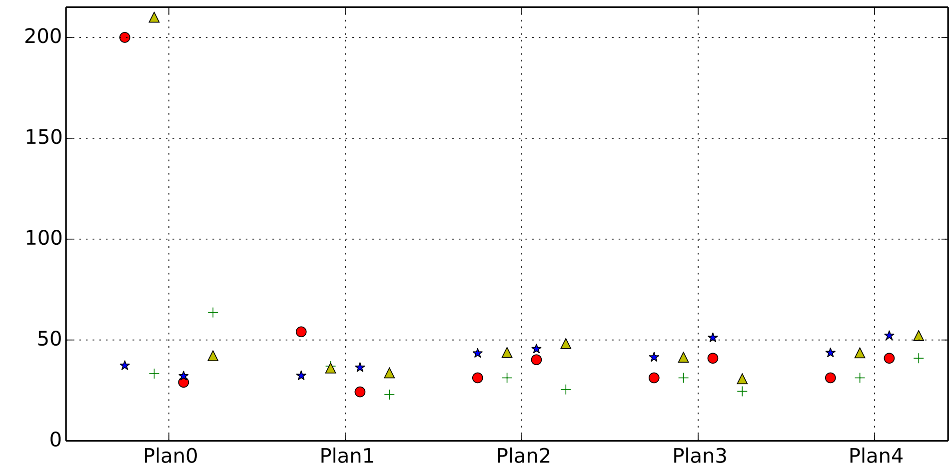

(a) PAR for different P_{grid} scenarios

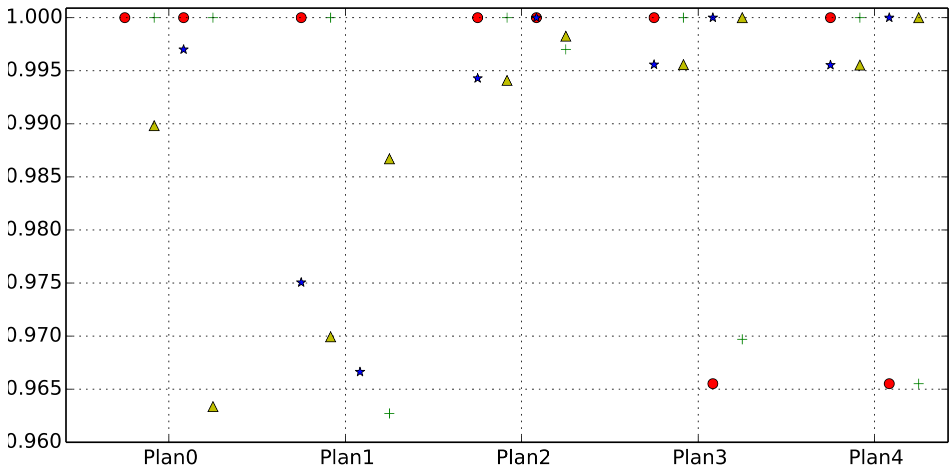

(b) Service availability for different P_{grid} scenarios

Figure 6.26: Evaluation metrics for the ANN_B 3LN

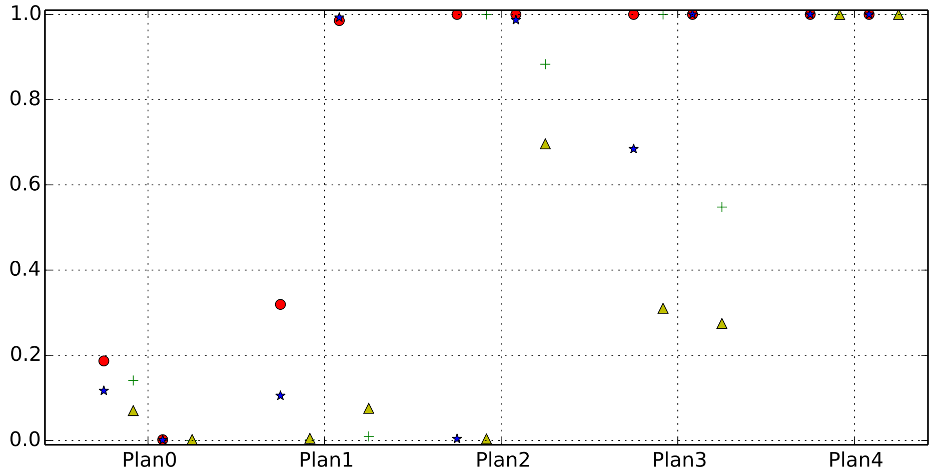

(a) System reactivity for different P_{grid} scenarios

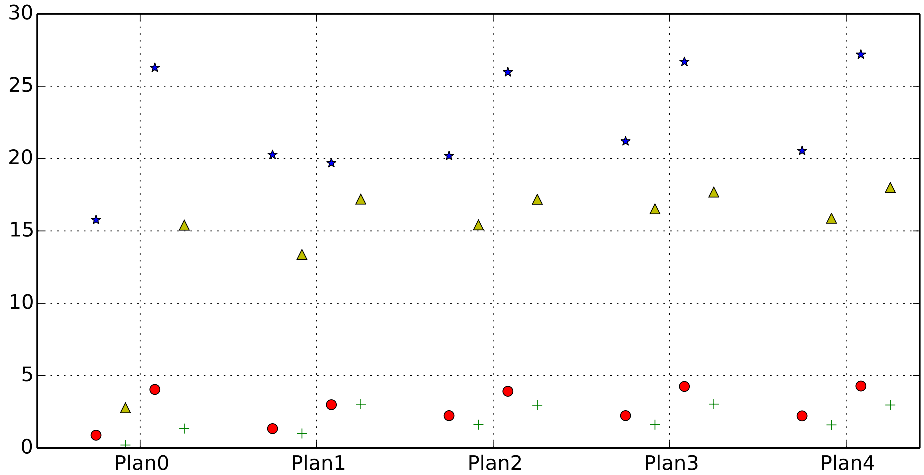

(b) Broker's profit for different P_{grid} scenarios

Figure 6.27: Evaluation metrics for the ANN_B 3LN

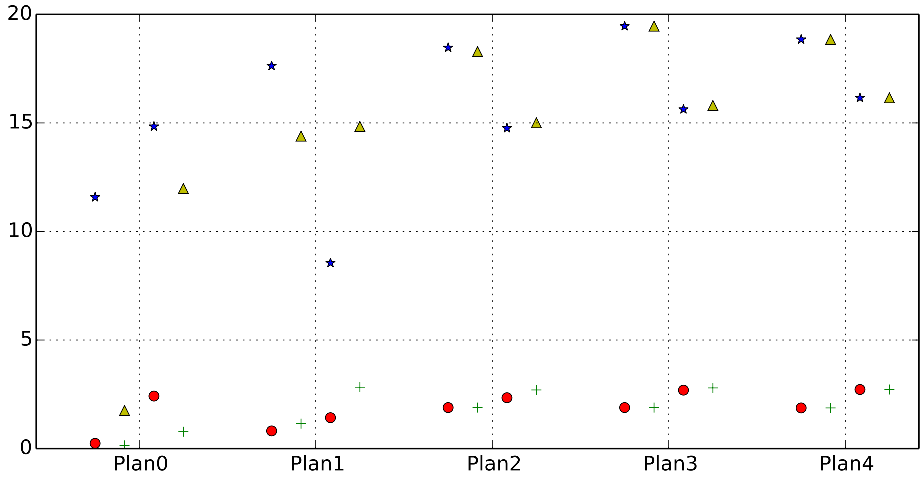

(a) Income from local sale for different P_{grid} scenarios

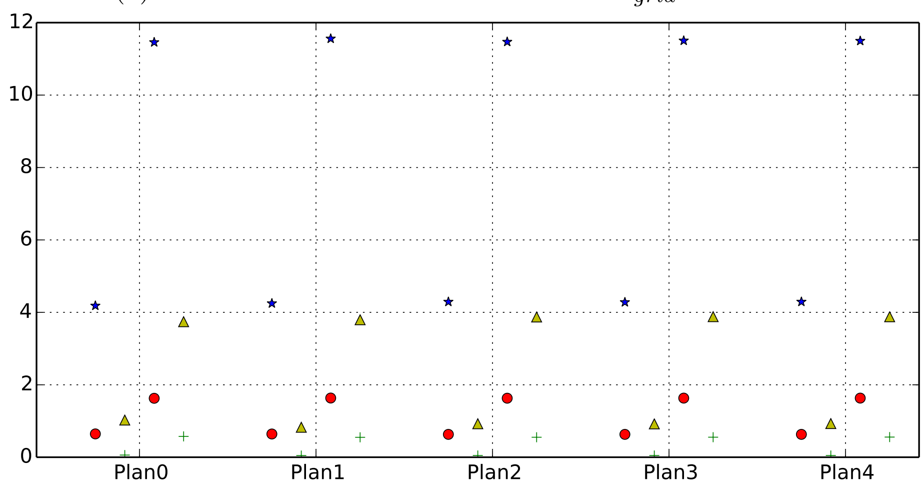

(b) Income from feed-in for different P_{grid} scenarios

Figure 6.28: Profit components the ANN_B 3LN

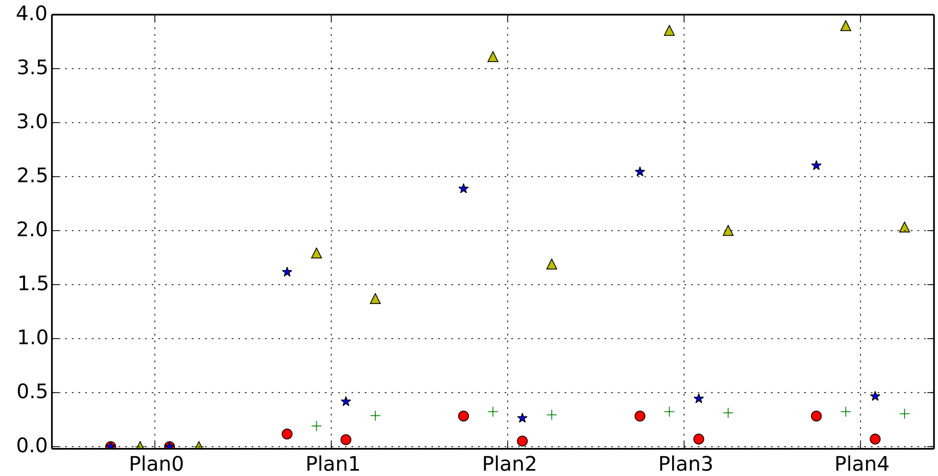

(a) Supply costs for different P_{grid} scenarios

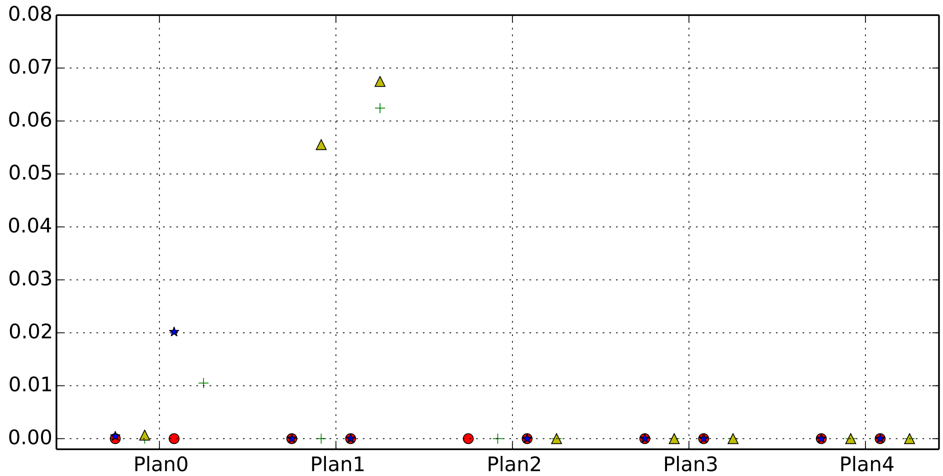

(b) Reimbursement costs for different P_{grid} scenarios

Figure 6.29: Profit components the ANN_B 3LN

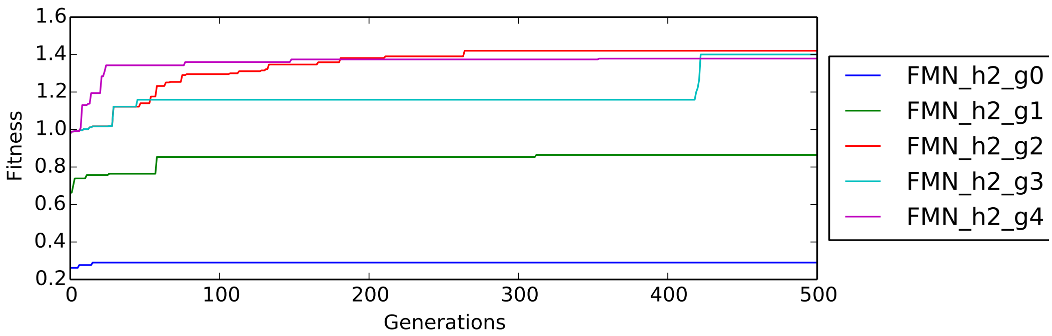

(a) Training on the winter day

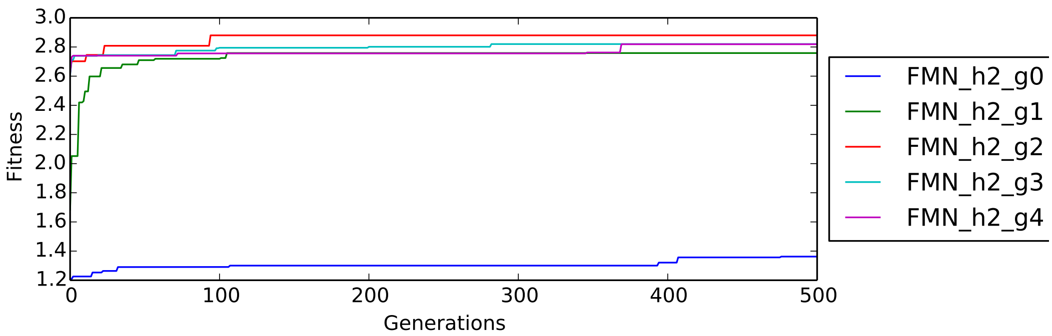

(b) Training on the summer day

Figure 6.30: Fitness landscape for the ANN_A FMN

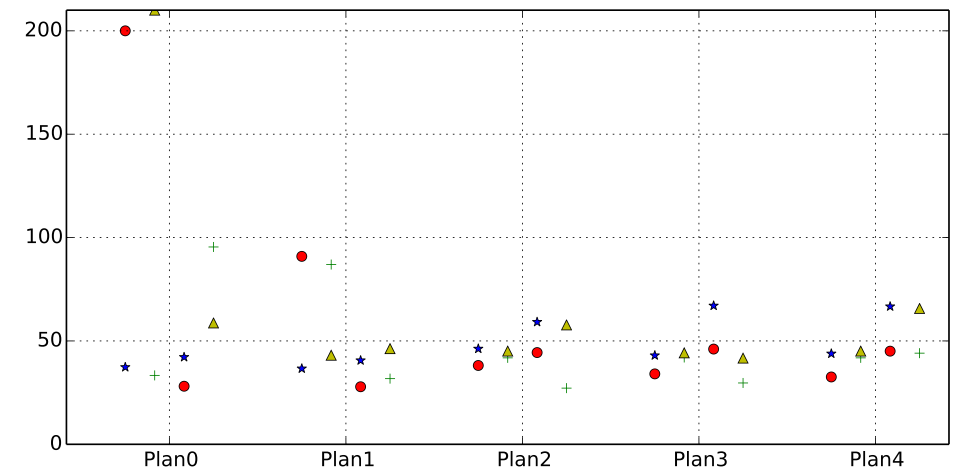

(a) PAR for different P_{grid} scenarios

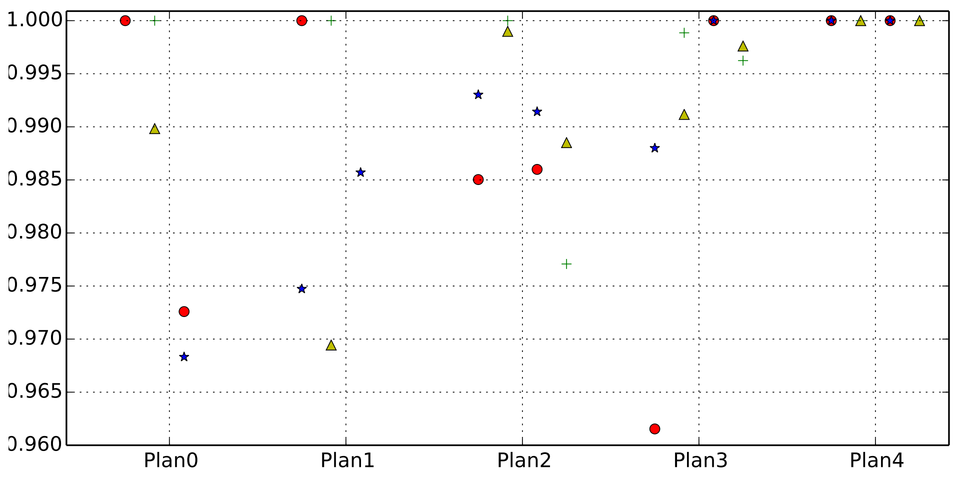

(b) Service availability for different P_{grid} scenarios

Figure 6.31: Evaluation metrics for the ANN_A FMN

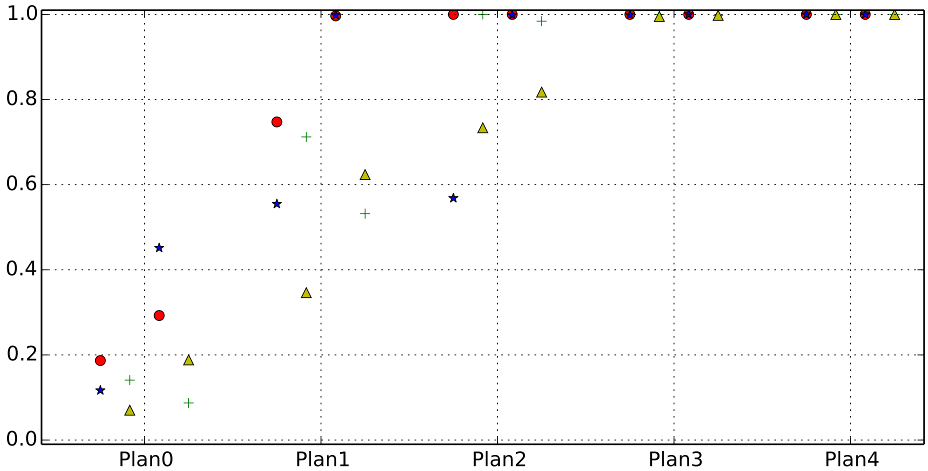

(a) System reactivity for different P_{grid} scenarios

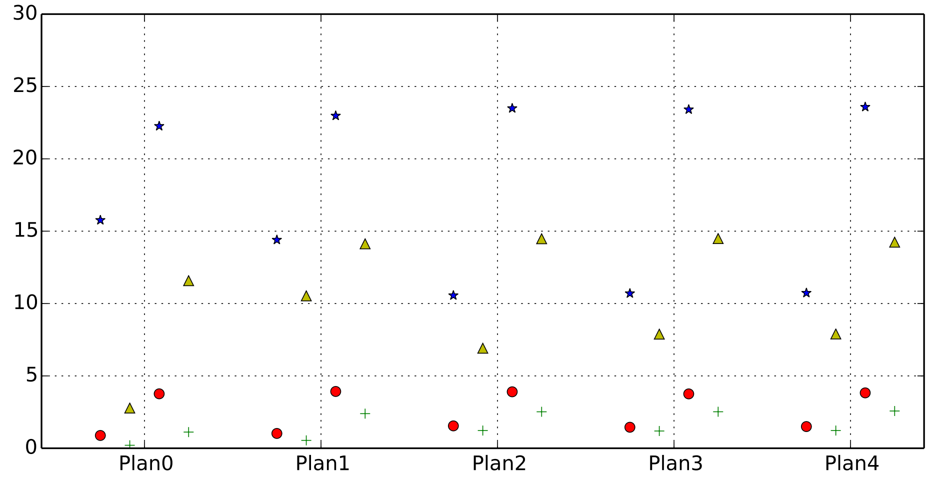

(b) Broker's profit for different P_{grid} scenarios

Figure 6.32: Evaluation metrics for the ANN_A FMN

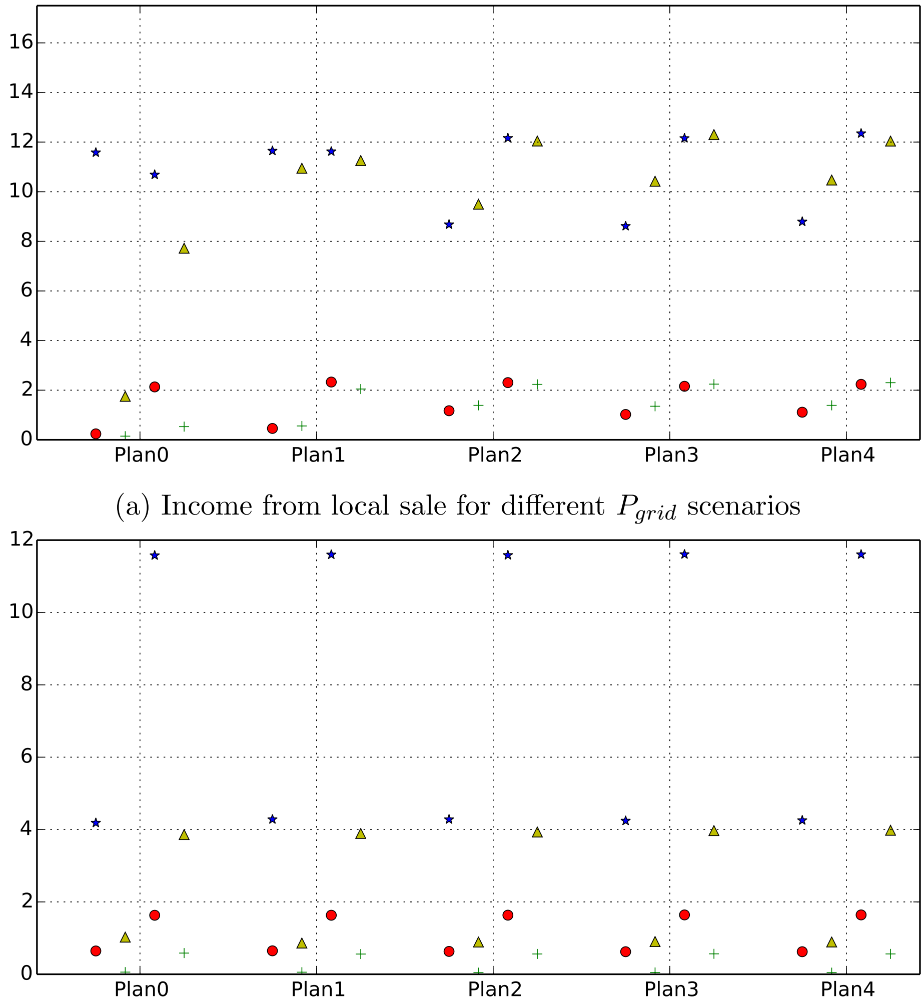

(a) Income from local sale for different P_{grid} scenarios

(b) Income from feed-in for different P_{grid} scenarios

Figure 6.33: Profit components for the ANN_A FMN

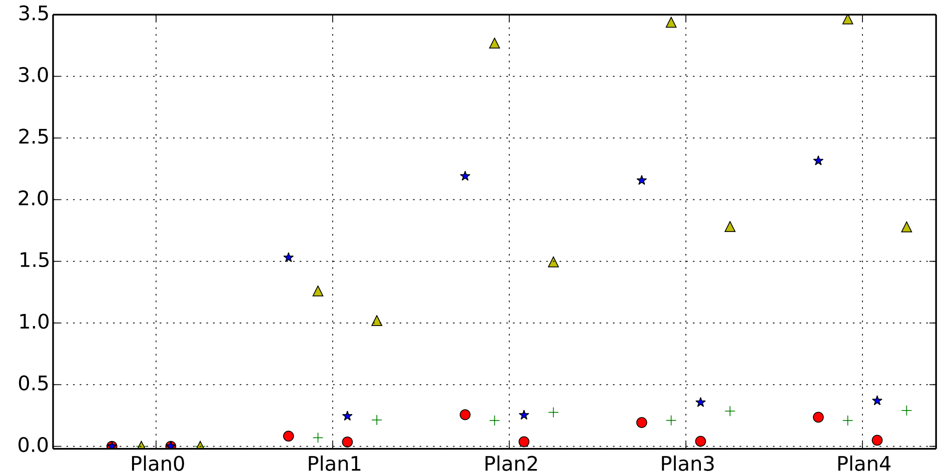

(a) Supply costs for different P_{grid} scenarios

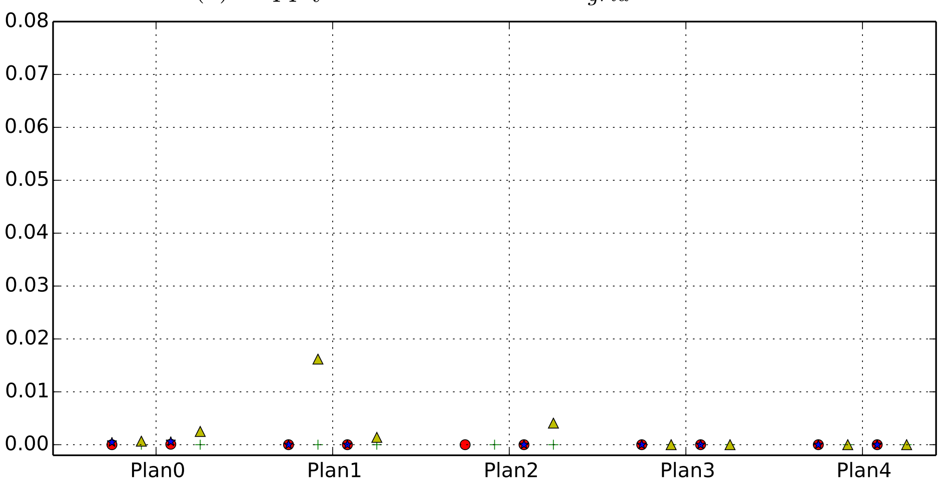

(b) Reimbursement costs for different P_{grid} scenarios

Figure 6.34: Profit components for the ANN_A FMN

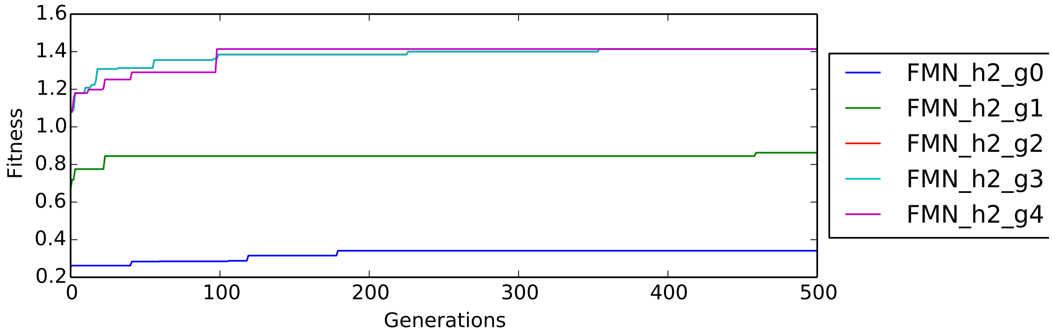

(a) Training on the winter day

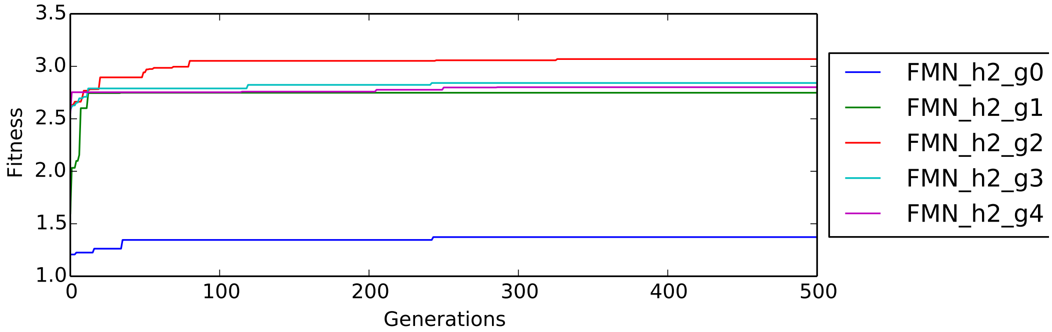

(b) Training on the summer day

Figure 6.35: Fitness landscape for the ANN_A FMN

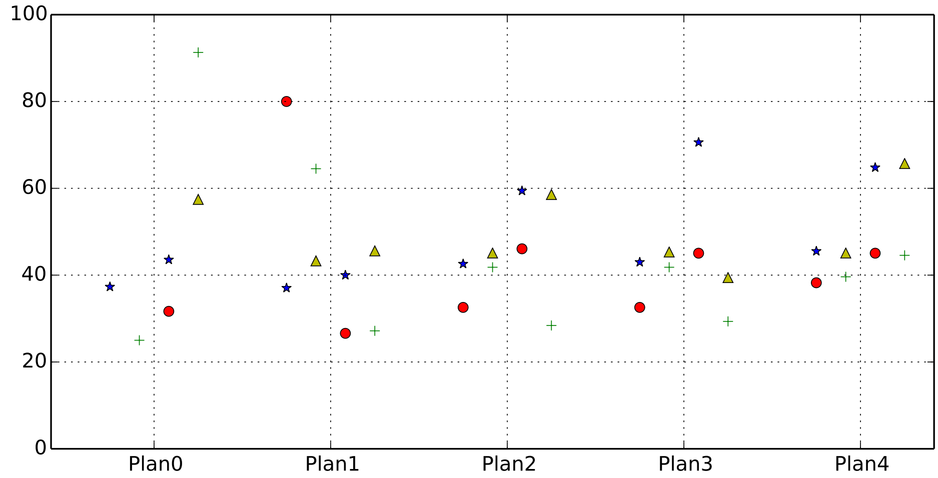

(a) PAR for different P_{grid} scenarios

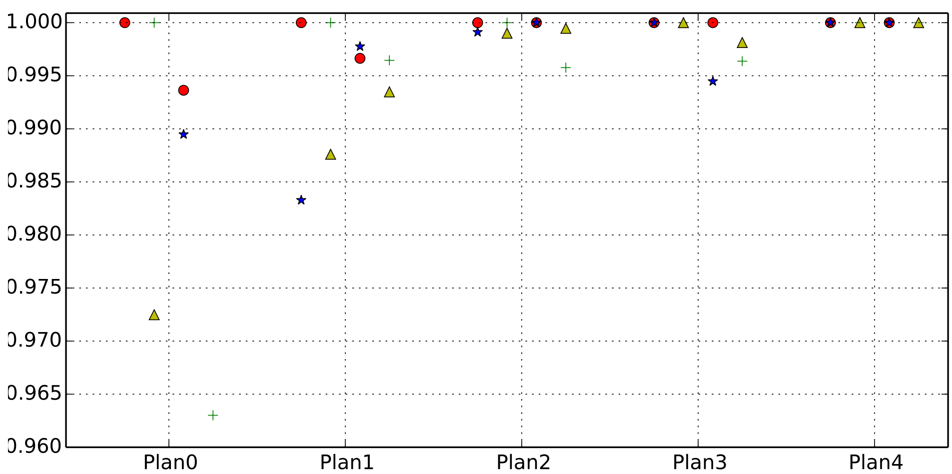

(b) Service availability for different P_{grid} scenarios

Figure 6.36: Evaluation metrics for the ANN_B FMN

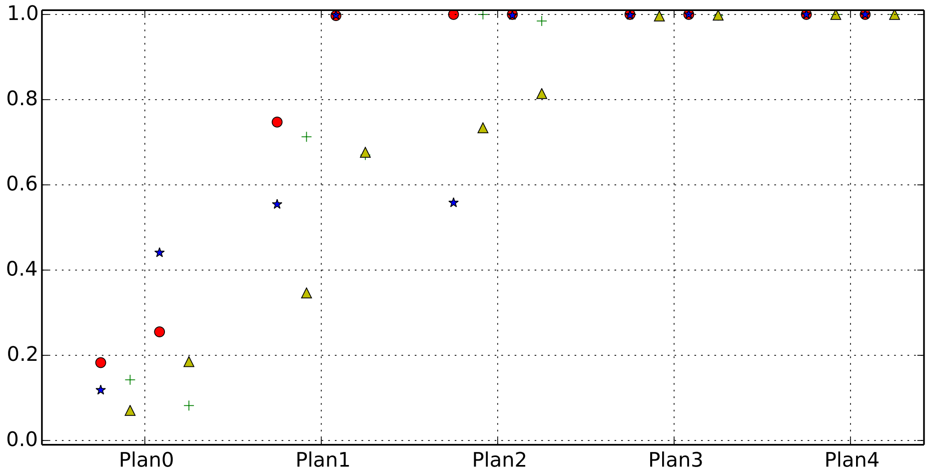

(a) System reactivity for different P_{grid} scenarios

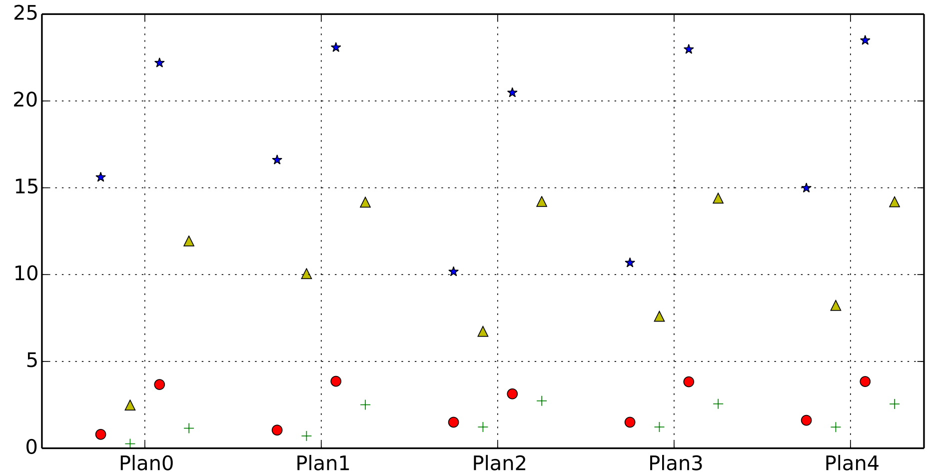

(b) Broker's profit for different P_{grid} scenarios

Figure 6.37: Evaluation metrics for the ANN_B FMN

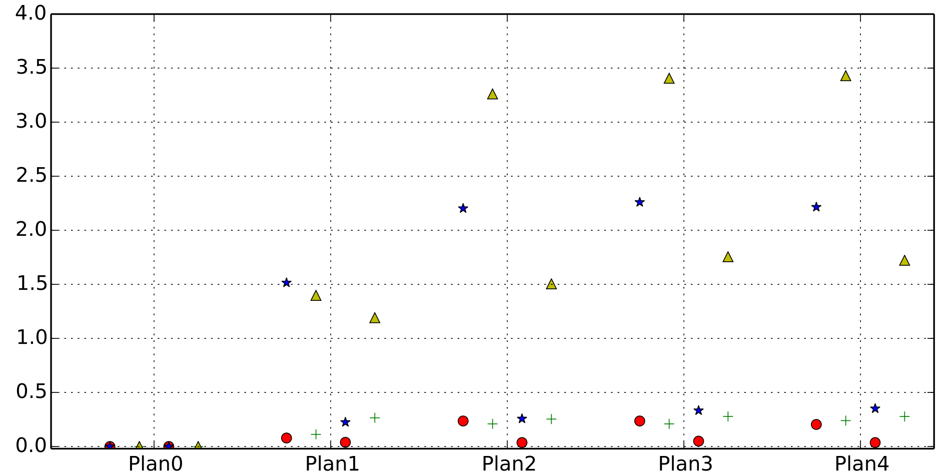

(a) Supply costs for different P_{grid} scenarios

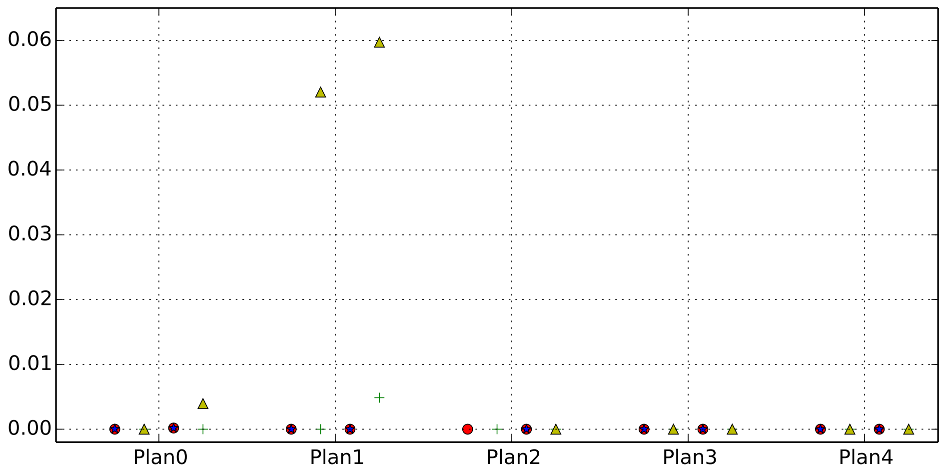

(b) Reimbursement costs for different P_{grid} scenarios

Figure 6.39: Profit components for the ANN_B FMN

Discussion

Results show the brokers effectively minimizing the reimbursement costs. By setting the price sensitivity to 0.9 (€ / kWh) we simulated the worst possible congestion scenario, in which depending on the employed usage model all loads desire to operate regardless of the SLA pricing.

Due to the designed trading attitude, loads seek operation regardless of their price sensitivity. The formulated prices have more evident effect in resource-constrained setups, such as $Plan_0$ and $Plan_1$. Accordingly, weather conditions determine the overall power availability, which makes shorter SLAs favourite. In presence of grid power connections, the stochastic nature of the weather is backed by the power provided by the main grid. This leads to higher supply costs, given the higher proportion of P_{grid} employed.

In actual scenarios, users would assign different price sensitivity to the loads, depending on the delivered utility. This has the favourable effect of determining an ordering over the loads, and consequently more favourable conditions for the broker and the resulting SLA prices.

6.5 Summary

This chapter dealt with the problem of automating energy management, by delegating software agents for the control of electrical loads. A smart microgrid is initially formalized in Sect. 6.1 and later implemented as a simulation framework (Sect. 6.2). The framework was used in Sect. 6.3 to show the possibility of learning controllers for smart prosumers, i.e., agents that embedd both a production and a generation module. While the tool can be used for real-world applications, a drawback of this approach is the size selected for the allocation interval. To solve this issue, we investigated in Sect. 6.4 on the possibility of learning a power broker, able to price power provisioning depending on the duration of service agreements.

CHAPTER 7

Conclusions

"Prediction is very difficult, especially about the future"

– Niels Bohr

In this chapter, we recap the provided contribution and discuss limitations that deserve further investigation in future. We further list published work related to the topic of this dissertation. We conclude the chapter with a list of open questions for future work.

7.1 Contributions

The main objective of the research presented in this dissertation is to address the problem of energy management in microgrids. In particular, we desire assisting users towards a more conscious and efficient use of local renewable sources, which being highly dependent on weather increase significantly the complexity for their management. To this end we have contributed with the following tools:

The GREEND dataset which we released to the public, contains more than 1 year power consumption data in selected household in Italy and Austria. The dataset was used in several intelligent energy applications such as load disaggregation and appliance usage modeling.

Ontology for integration of legacy and smart devices can be used to annotate information related to the functioning of electrical loads. The ontology allows for the seamless integration of smart and legacy devices, and opens to the use of other semantic web technologies, such as SPARQL for querying.

Mjölnir - the open energy advisor features state-of-the-art processing algorithms acting on production and consumption data, on both aggregated and disaggregated level. The tool can benefit both human-computer interaction scientists, as well as thinkerers and open-source enthusiasts.

The HEMS simulator framework to train energy prosumers, such as renewable energy generators, appliances, as well as batteries. The model proposed in this study is general and can be applied in any trading setup. The simulator can be used as a testbed for assessing different demand-response policies, as well as to automatically learn controllers that could be then run on physical devices.

The Smart-Microgrid broker serves as proof of concept towards the application of forward contracts for power trading in smart microgrids. We identified important metrics for the evaluation of power brokers and designed both basic rule-based and model-based ones. The model is general and can be further extended to other network architectures, e.g. deep networks and NEAT [Sta02]. The broker can be easily integrated in existing energy management tools, where it can assist users by displaying expected energy prices.

7.2 Limitations

Because of the selected techniques we also face multiple limitations:

- This dissertation focused on residential environments, by gaining insights from real households being monitored through surveys and a measurement campaign. Specifically, the GREEND offers insights into residential energy consumption. To tackle public and industrial settings different datasets are necessary.

- The proposed architecture allows for the full integration of smart and legacy devices, and ultimately for device and data interoperability. However, the released proof of concept serves only to demonstrate the possibility to describe and query device information. In particular, the selected RDFLib can rely on relational or file-based databases to store handled triples. Generally this means using a table with three columns (i.e., subject, predicate and object), which has clear scalability limitations for both storage and interrogation.

- The policies automatically generated by the advisor widget were estimated leading up to 34% of savings. This was calculated by adding up the individual contribution of each advice. This represents the optimal case in which users are exposed to all policies and decide to actuate them. In addition, the GREEND dataset provides a limited number of setups targeting the peculiarities of Austrian and Italian households. Thus, the reported findings do not have statistical significance and the actual effectiveness of the advisor is yet to be assessed through large-scale deployments.

- The designed prosumer controllers were trained in specific scenarios. Our study still lacks a comparison of possible controllers, as this would be dependent on the used dataset.
- The designed broker adapts to the given simulation environment, and thus, results are hardly generalizable. However, the main contribution is the used methodology, which can be generalized to other setups.

7.3 Related publications

During the course of this study various publications have been presented to the research community (see B.3). This dissertation includes material from the following papers:

- **Strategies for Domestic Energy Conservation in Carinthia and Friuli-Venezia Giulia** [Mon13c] Households account for a significant fraction of overall energy consumption. Energy usage can be reduced by improving the efficiency of devices and optimizing their use as well as by encouraging people to change their behaviour towards a more sustainable lifestyle. In this study, we investigate patterns of domestic energy use in Carinthia (Austria) and Friuli-Venezia Giulia (Italy). In particular, we report the results of an online survey about electrical devices and their use in households. We outline typical scenarios in the two regions and discuss possible strategies to reduce the consumption of energy in these regions.
- **GREEND: an energy consumption dataset of households in Italy and Austria** [Mon14a] Home energy management systems can be used to monitor and optimize consumption and local production from renewable energy. To assess solutions before their deployment, researchers and designers of those systems demand for energy consumption datasets. In this paper, we present the GREEND dataset, containing detailed power usage information obtained through a measurement campaign in households in Austria and Italy. We provide a description of consumption scenarios and discuss design choices for the sensing infrastructure. Finally, we benchmark the dataset with state-of-the-art techniques in load disaggregation, occupancy detection and appliance usage mining.

- **Integrating Households into the Smart Grid** [Mon13a] The success of the Smart Grid depends on its ability to collect data from heterogeneous sources such as smart meters and smart appliances, as well as the utilization of this information to forecast energy demand and to provide value-added services to users. In our analysis, we discuss requirements for collecting and integrating household data within smart grid applications. We put forward a potential system architecture and report state-of-the-art technologies that can be deployed towards this vision.

- **Integration of Legacy Appliances into Home Energy Management Systems** [Ega15a] The progressive installation of renewable energy sources requires the coordination of energy consuming devices. At consumer level, this coordination can be done by a home energy management system (HEMS). Interoperability issues need to be solved among smart appliances as well as between smart and non-smart, i.e., legacy devices. We expect current standardization efforts to soon provide technologies to design smart appliances in order to cope with the current interoperability issues. Nevertheless, common electrical devices affect energy consumption significantly and therefore deserve consideration within energy management applications. This paper discusses the integration of smart and legacy devices into a generic system architecture and, subsequently, elaborates the requirements and components which are necessary to realize such an architecture including an application of load detection for the identification of running loads and their integration into existing HEM systems. We assess the feasibility of such an approach with a case study based on a measurement campaign on real households. We show how the information of detected appliances can be extracted in order to create device profiles allowing for their integration and management within a HEMS.

- **An Open Solution to Provide Personalized Feedback for Building Energy Management** [Mon15] The integration of renewable energy sources increases the complexity in mantaining the power grid. In particular, the highly dynamic nature of generation and consumption demands for a better utilization of energy resources, which seen the cost of storage infrastructure, can only be achieved through demand-response. Accordingly, the availability of energy and potential overload situations can be reflected using a price signal. The effectiveness of this mechanism arises from the flexibility of device operation, which is nevertheless heavily reliant on the exchange of information between the grid and its consumers. In this paper, we investigate the capability of an interactive energy management system to timely inform users on energy usage, in order to promote an optimal use of local resources. In particular, we analyze data being collected in several households in Italy and Austria to gain insights into usage behavior and drive the design of more effective systems. The outcome is the formulation of energy efficiency policies for residential buildings, as well as the design of an energy management system, consisting of hardware measurement units and a management software. The Mjölnir framework, which we release for open use, provides a platform where various feedback concepts can be implemented and assessed. This includes widgets displaying disaggregated and aggregated consumption information, as well as daily production and tailored advices. The formulated policies were implemented as an advisor widget able to autonomously analyze usage and provide tailored energy feedback. The advisor is estimated leading to a potential of 34% of savings using measurement data from the GREEND dataset.

- **HEMS - A Home Energy Market Simulator** [Mon14c] Stability issues in the electric power grid originate from the rising of renewable energy generation and the increasing number of electric vehicles. The uncertainty and the distributed nature of generation and consumption demand for optimal allocation of energy resources, which, in the absence of sufficient control reserve for power generation, can be achieved using demand-response. A price signal can be exploited to reflect the availability of energy. In this paper, market-based energy allocation solutions for small energy grids are discussed and implemented in a simulator, which is released for open use. Artificial neural network controllers for energy prosumers can be designed to minimize individual and overall running costs. This enables a better use of local energy production from renewable sources, while considering residents' necessities to minimize discomfort.

- **Assisted energy management in Smart Microgrids** [Mon16] Demand response provides utilities with a mechanism to share with end users the stochasticity resulting from the use of renewable sources. Pricing is accordingly used to reflect energy availability, to allocate such a limited resource to those loads that value it most. However, the strictly competitive mechanism can result in service interruption in presence of competing demand. To solve this issue we investigate on the use of forward contracts, i.e., service-level agreements priced to reflect the expectation of future supply and demand curves. Given the limited resources of microgrids, service availability is an opposite objective to the one of system reactivity. We firstly design policy-based brokers and identify then a learning broker based on artificial neural networks. We show the latter being progressively minimizing the reimbursement costs and maximizing the overall profit.

7.4 Future work

- To achieve seamless integration of smart and legacy devices further work is necessary on the proposed architecture. This includes the integration of external load disaggregation libraries and frameworks, such as the NILM toolkit (NILMTK). Another relevant aspect to be considered is the scalability of the triple store used to handle the semantic network. A possibility is to rely on distributed database management systems, such as Apache Cassandra[1] and Hadoop HBase[2].

- In Mjölnir we foresee various developments. A better integration of available occupancy models and appliance usage models with external tools, such as the if-this-then-that (IFTTT[3]) service, would allow users for benefit of other smart-home applications. Moreover, a better integration with social networking sites would allow for new social features, such as comparisons with "friends" and follower-like lists, to be used for compelling goal-setting widgets. For instance, this might include competitions at microgrid or regional level, where performance of rooms and buildings would contribute to the neighborhood performance in a sort of crowd sensing tool (e.g., smartroadsense[4]). It is also important to remark that today's installations take place mostly in residential environments. Since the motivation of users in different settings is not clear, a deeper evaluation of potential of feedback systems in public buildings should thus be carried out in future. Specific widgets will thus be necessary, especially given that public displays provide different interaction modalities and resolution.

- To actually benefit of the possibility of automatically design controllers for energy prosumers, the HEMS Simulator should be further extended with the possibility of exporting learned controllers. An investigation of existing formats should be undertaken. Moreover, assessing the gap between simulated and physical setting is critical challenge for the effective operation.

- A comparison of possible structures for the broker should be carried out. This includes neuroevolution of augmenting topologies (i.e., NEAT [Sta02] and HyperNEAT [Sta09]) as well as the employment of deep architectures and different activation functions. In addition, further research is necessary towards the integration of the designed microgrid broker into energy management systems.

[1] http://cassandra.apache.org/
[2] http://hbase.apache.org/
[3] http://ifttt.com
[4] http://smartroadsense.it

Bibliography

[Adi14] C. Adika and L. Wang. Autonomous Appliance Scheduling for Household Energy Management. *IEEE Transactions on Smart Grid*, 5(2):673–682, March 2014.

[Agu11] P. Aguilera, A. Fernández, R. Fernández, R. Rumí, and A. Salmerón. Bayesian networks in environmental modelling. *Environmental Modelling & Software*, 26(12):1376 – 1388, 2011.

[Ala13] M. Alam, S. D. Ramchurn, and A. Rogers. Cooperative Energy Exchange for the Efficient Use of Energy and Resources in Remote Communities. In *Proc. of the 2013 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '13, pages 731–738, 2013.

[And12] K. Anderson, A. Ocneanu, D. Benitez, D. Carlson, A. Rowe, and M. Berges. BLUED: A Fully Labeled Public Dataset for Event-Based Non-Intrusive Load Monitoring Research. In *Proc. of the 2nd KDD Workshop on Data Mining Applications in Sustainability (SustKDD)*, 2012.

[Ani11] D. Anicic, P. Fodor, S. Rudolph, and N. Stojanovic. EP-SPARQL: a unified language for event processing and stream reasoning. In *Proceedings of the 20th international conference on World wide web*, pages 635–644. ACM, 2011.

[Arm13] K. C. Armel, A. Gupta, G. Shrimali, and A. Albert. Is disaggregation the holy grail of energy efficiency? The case of electricity. *Energy Policy*, 52(0):213 – 234, 2013.

[Ash47] W. R. Ashby. Principles of the Self-Organizing Dynamic System. *The Journal of General Psychology*, 37(2):125–128, 1947.

[Aus97] L. M. Ausubel. An efficient ascending-bid auction for multiple objects. *AMERICAN ECONOMIC REVIEW*, 94, 1997.

[Ayd02] M. Aydinalp, V. I. Ugursal, and A. S. Fung. Modeling of the appliance, lighting, and space-cooling energy consumptions in the residential sector using neural networks. *Applied Energy*, 71(2):87 – 110, 2002.

[Bap11] T. Bapat, N. Sengupta, S. K. Ghai, V. Arya, Y. B. Shrinivasan, and D. Seetharam. User-sensitive Scheduling of Home Appliances. In *Proceedings of the 2Nd ACM SIGCOMM Workshop on Green Networking*, GreenNets '11, pages 43–48, New York, NY, USA, 2011. ACM.

[Bar10] D. F. Barbieri, D. Braga, S. Ceri, and M. Grossniklaus. An execution environment for C-SPARQL queries. In *Proceedings of the 13th International Conference on Extending Database Technology*, pages 441–452. ACM, 2010.

[Bar12] S. Barker, A. Mishra, D. Irwin, E. Cecchet, P. Shenoy, and J. Albrecht. Smart*: An open data set and tools for enabling research in sustainable homes. *SustKDD, August*, 2012.

[Bat13] N. Batra, M. Gulati, A. Singh, and M. B. Srivastava. It's Different: Insights into Home Energy Consumption in India. In *Proc. of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, BuildSys'13, New York, NY, USA, 2013. ACM.

[Bat14] N. Batra, J. Kelly, O. Parson, H. Dutta, W. Knottenbelt, A. Rogers, A. Singh, and M. Srivastava. NILMTK: An Open Source Toolkit for Non-intrusive Load Monitoring. In *Proceedings of the 5th International Conference on Future Energy Systems*, e-Energy '14, pages 265–276. ACM, 2014.

[Bec13] C. Beckel, L. Sadamori, and S. Santini. Automatic Socio-Economic Classification of Households Using Electricity Consumption Data. In *Proceedings of the 4th International Conference on Future Energy Systems (ACM e-Energy '13). Berkeley, CA, USA*. ACM, May 2013.

[Bee13] J. Beel, S. Langer, M. Genzmehr, and A. Nürnberger. Persistence in Recommender Systems: Giving the Same Recommendations to the Same Users Multiple Times. In T. Aalberg, C. Papatheodorou, M. Dobreva, G. Tsakonas, and C. Farrugia, Editors, *Research and Advanced Technology for Digital Libraries*, volume 8092 of *Lecture Notes in Computer Science*, pages 386–390. Springer Berlin Heidelberg, 2013.

[Ber97] M. Berry and G. Linoff. *Data Mining Techniques: For Marketing, Sales, and Customer Support*. Database management / Wiley. Wiley, 1997.

[Ber07] P. Bertoldi and B. Atanasiu. Electricity Consumption and Efficiency Trends in the Enlarged European Union. Technical Report EUR 22753 EN, European Commission, Institute for Environment and Sustainability, 2007.

[Blu92] A. Blum. *Neural Networks in C++: An Object-Oriented Framework for Building Connectionist Systems*. Wiley, 1 edition, 5 1992.

[Bog97] Z. Boger and H. Guterman. Knowledge extraction from artificial neural network models. In *Systems, Man, and Cybernetics, 1997. Computational Cybernetics and Simulation., 1997 IEEE International Conference on*, volume 4, pages 3030–3035 vol.4, 1997.

[Bon08] D. Bonino and F. Corno. DogOnt - Ontology Modeling for Intelligent Domotic Environments. In *The Semantic Web - ISWC 2008*, volume 5318 of *Lecture Notes in Computer Science*, pages 790–803. Springer Berlin Heidelberg, 2008.

[Bon12] D. Bonino, F. Corno, and L. D. Russis. Home energy consumption feedback: A user survey. *Energy and Buildings*, 47(0):383 – 393, 2012.

[Bos10] B. Bose. Global Warming: Energy, Environmental Pollution, and the Impact of Power Electronics. *Industrial Electronics Magazine, IEEE*, 4(1):6–17, March 2010.

[Bru09] Y. Brun, G. Marzo Serugendo, C. Gacek, H. Giese, H. Kienle, M. Litoiu, H. Müller, M. Pezzè, and M. Shaw. Software Engineering for Self-Adaptive Systems. Chapter Engineering Self-Adaptive Systems through Feedback Loops, pages 48–70. Springer, 2009.

[Buc88] J. Buck. Synchronous Rhythmic Flashing of Fireflies. II. *The Quarterly Review of Biology*, 63(3):265–289, 1988.

[CA13] K. Carrie Armel, A. Gupta, G. Shrimali, and A. Albert. Is disaggregation the holy grail of energy efficiency? The case of electricity. *Energy Policy*, 52(C):213–234, 2013.

[Cal10] J.-P. Calbimonte, O. Corcho, and A. J. G. Gray. Enabling ontology-based access to streaming data sources. In *Proceedings of the 9th international semantic web conference on The semantic web - Volume Part I*, pages 96–111. Springer-Verlag, 2010.

[Car13] D. R. Carlson, H. S. Matthews, and M. Berges. One size does not fit all: Averaged data on household electricity is inadequate for residential energy policy and decisions. *Energy and Buildings*, 64(0):132 – 144, 2013.

[Cau92] M. Caudill and C. Butler. *Understanding Neural Networks; Computer Explorations.* MIT Press, Cambridge, MA, USA, 1992.

[Cha11] C.-C. Chang and C.-J. Lin. LIBSVM: A Library for Support Vector Machines. *ACM Trans. Intell. Syst. Technol.*, pages 1–27, April 2011.

[Cha14a] V. Chandan, T. Ganu, T. K. Wijaya, M. Minou, G. Stamoulis, G. Thanos, and D. P. Seetharam. iDR: Consumer and Grid Friendly Demand Response System. In *Proceedings of the 5th International Conference on Future Energy Systems*, e-Energy '14, pages 183–194, New York, NY, USA, 2014. ACM.

[Cha14b] P. Chavali, P. Yang, and A. Nehorai. A Distributed Algorithm of Appliance Scheduling for Home Energy Management System. *IEEE Transactions on Smart Grid*, 5(1):282–290, Jan 2014.

[Che13] D. Chen, S. Barker, A. Subbaswamy, D. Irwin, and P. Shenoy. Non-Intrusive Occupancy Monitoring Using Smart Meters. In *Proc. of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, BuildSys'13, 2013.

[Chu02] B. N. Chun and D. E. Culler. User-Centric Performance Analysis of Market-Based Cluster Batch Schedulers. In *Proc. of the 2Nd IEEE/ACM International Symposium on Cluster Computing and the Grid*, pages 30–39, 2002.

[Cle96] S. H. Clearwater, Editor. *Market-Based Control: A Paradigm for Distributed Resource Allocation.* World Scientific, River Edge, NH, 1996.

[Col09] C. Colson and M. Nehrir. A review of challenges to real-time power management of microgrids. In *Power Energy Society General Meeting, 2009. PES '09. IEEE*, pages 1–8, July 2009.

[Com12] M. Compton, P. Barnaghi, L. Bermudez, R. Garca-Castro, O. Corcho, S. Cox, J. Graybeal, M. Hauswirth, C. Henson, A. Herzog, V. Huang, K. Janowicz, W. D. Kelsey, D. L. Phuoc, L. Lefort, M. Leggieri, H. Neuhaus, A. Nikolov, K. Page, A. Passant, A. Sheth, and K. Taylor. The SSN Ontology of the W3C Semantic Sensor Network Incubator

Group. *Web Semantics: Science, Services and Agents on the World Wide Web*, 17(0), 2012.

[Cra06] P. Cramton, Y. Shoham, and R. Steinberg. *Combinatorial Auctions*. MIT Press, 2006.

[D'A14] S. D'Alessandro, A. M. Tonello, A. Monacchi, and W. Elmenreich. Home Energy Management Systems: Design Guidelines for the Communication Infrastructure. In *Proc. of the EEE International Energy Conference (ENERGYCON 2014)*, May 2014.

[Dar06] S. Darby. The effectiveness of feedback on energy consumption: a review for DEFRA of the literature on metering, billing and direct displays. Technical report, Environmental Change Institute, University of Oxford, 2006.

[Dar09] W. Dargie. *Context-Aware Computing and Self-Managing Systems*. Chapman & Hall/CRC, 1 edition, 2009.

[Don13] M. Dong, P. C. M. Meira, W. Xu, and C. Y. Chung. Non-Intrusive Signature Extraction for Major Residential Loads. *IEEE Transactions on Smart Grid*, PP(99):1 –10, 2013.

[Dor96] M. Dorigo, V. Maniezzo, and A. Colorni. Ant system: optimization by a colony of cooperating agents. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 26(1):29–41, Feb 1996.

[DuC11] B. DuCharme. *Learning SPARQL*. O'Reilly Media, Inc., 2011.

[Ega13] D. Egarter, A. Sobe, and W. Elmenreich. Evolving Non-Intrusive Load Monitoring. In *Proc. of 16th European Conference on the Applications of Evolutionary Computation*, Vienna, Austria, 2013. Springer.

[Ega14] D. Egarter, C. Prokop, and W. Elmenreich. Load hiding of household's power demand. In *IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pages 854–859. IEEE, 2014.

[Ega15a] D. Egarter, A. Monacchi, T. Khatib, and W. Elmenreich. Integration of legacy appliances into home energy management systems. *Journal of Ambient Intelligence and Humanized Computing*, pages 1–15, 2015.

[Ega15b] D. Egarter, M. Pöchacker, and W. Elmenreich. Complexity of Power Draws for Load Disaggregation. *CoRR*, abs/1501.02954, 2015.

[Eib03] A. E. Eiben and J. E. Smith. *Introduction to evolutionary computing*. Springer, 2003.

[Elm12] W. Elmenreich and D. Egarter. Design guidelines for smart appliances. In *Proceedings of the Tenth Workshop on Intelligent Solutions in Embedded Systems (WISES)*, pages 76 –82, 2012.

[EM10] K. Erhardt-Martinez, K. Donnelly, and J. A. S. Laitner. Advanced Metering Initiatives and Residential Feedback Programs: A Meta-Review for Household Electricity-Saving Opportunities. Technical report, American Council for an Energy-Efficient Economy, Washington, D.C., 2010.

[Eur10] European Commission. Europe 2020: Commission proposes new economic strategy, 2010.

[Feh10] I. Fehérvári and W. Elmenreich. Evolving neural network controllers for a team of self-organizing robots. *Journal of Robotics*, 2010.

[Feh13] I. Fehérvári. *On Evolving Self-organizing Technical Systems*. PhD Thesis, Alpen-Adria-Universität Klagenfurt, 2013.

[Fis13] J. E. Fischer, S. D. Ramchurn, M. A. Osborne, O. Parson, T. D. Huynh, M. Alam, N. Pantidi, S. Moran, K. Bachour, S. Reece, E. Costanza, T. Rodden, and N. R. Jennings. Recommending Energy Tariffs and Load Shifting Based on Smart Household Usage Profiling. In *International Conference on Intelligent User Interfaces*, pages 383–394, 2013.

[Flo08] D. Floreano and C. Mattiussi. *Bio-Inspired Artificial Intelligence: Theories, Methods, and Technologies*. The MIT Press, 2008.

[Gis13] C. Gisler, A. Ridi, D. Zujferey, O. Khaled, and J. Hennebert. Appliance consumption signature database and recognition test protocols. In *8th International Workshop on Systems, Signal Processing and their Applications (WoSSPA)*, May 2013.

[Glo11] X. Glorot, A. Bordes, and Y. Bengio. Deep Sparse Rectifier Neural Networks. In *Proc. of the Fourteenth International Conference on Artificial Intelligence and Statistics (AISTATS-11)*, volume 15, pages 315–323, 2011.

[Gon11] H. Goncalves, A. Ocneanu, and M. Berges. Unsupervised disaggregation of appliances using aggregated consumption data. In *Proc. of KDD Workshop on Data Mining Applications in Sustainability (SustKDD)*, 2011.

[Gre12] U. Greveler, B. Justus, and D. Loehr. Multimedia Content Identification Through Smart Meter Power Usage Profiles. In *Proceedings of the International Conference on Information and Knowledge Engineering IKE'12*, 2012.

[Gun11] V. Gungor, D. Sahin, T. Kocak, S. Ergut, C. Buccella, C. Cecati, and G. Hancke. Smart Grid Technologies: Communication Technologies and Standards. *IEEE Trans Ind. Informat.*, 7(4):529–539, Nov 2011.

[Gus05] A. Gustafsson and M. Gyllenswärd. The power-aware cord: energy awareness through ambient information display. In *CHI '05 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '05, pages 1423–1426, New York, NY, USA, 2005. ACM.

[Har92] G. Hart. Nonintrusive appliance load monitoring. *Proc. of the IEEE*, 80(12):1870–1891, 1992.

[Hea11] T. Heath and C. Bizer. *Linked Data: Evolving the Web into a Global Data Space*. Morgan & Claypool, 2011.

[IEA15] IEA. *World Energy Outlook 2015*. OECD Publishing, 2015.

[Ili12] D. Ilic, P. Da Silva, S. Karnouskos, and M. Griesemer. An energy market for trading electricity in smart grid neighbourhoods. In *Proc. of the 6th IEEE International Conference on Digital Ecosystems Technologies (DEST)*, pages 1–6, June 2012.

[Jam05] F. Jammes and H. Smit. Service-oriented paradigms in industrial automation. *IEEE Transactions on Industrial Informatics*, 1(1):62–70, Feb 2005.

[Jør89] A. H. Jørgensen. Using the Thinking-aloud Method in System Development. In *Proceedings of the Third International Conference on Human-computer Interaction on Designing and Using Human-computer Interfaces and Knowledge Based Systems (2Nd Ed.)*, pages 743–750, 1989.

[Kan12] K. J. Kang, B. Ka, and S. J. Kim. A service scenario generation scheme based on association rule mining for elderly surveillance system in a smart home environment. *Engineering Applications of Artificial Intelligence*, 25(7):1355 – 1364, 2012. Advanced issues in Artificial Intelligence and Pattern Recognition for Intelligent Surveillance System in Smart Home Environment.

[Kel14a] J. Kelly and W. Knottenbelt. "UK-DALE": A dataset recording UK Domestic Appliance-Level Electricity demand and whole-house demand. *ArXiv e-prints*, Apr. 2014.

[Kel14b] J. Kelly and W. Knottenbelt. Metadata for Energy Disaggregation. In *The 2nd IEEE International Workshop on Consumer Devices and Systems (CDS 2014)*, Västerås, Sweden, July 2014.

[Kep03] J. O. Kephart and D. M. Chess. The Vision of Autonomic Computing. *Computer*, 36:41–50, January 2003.

[Kha14] T. Khatib, A. Monacchi, W. Elmenreich, D. Egarter, S. D'Alessandro, and A. M. Tonello. European End-User's Level of Energy Consumption and Attitude Toward Smart Homes: A Case Study of Residential Sectors in Austria and Italy. *Energy Technology & Policy*, 1(1):97–105, 2014.

[Kim11] H. Kim, M. Marwah, M. F. Arlitt, G. Lyon, and J. Han. Unsupervised Disaggregation of Low Frequency Power Measurements. In *Proc. of the 11th SIAM International Conference on Data Mining*, 2011.

[Kle13] W. Kleiminger, C. Beckel, T. Staake, and S. Santini. Occupancy Detection from Electricity Consumption Data. In *Proc. of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, BuildSys'13, 2013.

[Kof12] M. J. Kofler, C. Reinisch, and W. Kastner. An Ontological Weather Representation for Improving Energy-Efficiency in Interconnected Smart Home Systems. In *Proceedings of Applied Simulation and Modelling/Artificial Intelligence and Soft Computing (ASC2012)*, Napoli, Italy, 2012.

[Kof13a] M. J. Kofler and W. Kastner. Towards an ontology representing building physics parameters for increased energy efficiency in smart home operation. In *Proceedings of the 2nd Central European Symposium on Building Physic (CESBP2013)*, Vienna, Austria, 2013.

[Kof13b] M. J. Kofler, F. I. Vázquez, and W. Kastner. An Ontology for Representation of User Habits and Building Context in Future Smart Homes. In *Proceedings of the 20th Workshop on Intelligent Computing in Engineering (EG-ICE2013)*, Vienna, Austria, 2013.

[Kok05] J. K. Kok, C. J. Warmer, and I. G. Kamphuis. PowerMatcher: Multiagent Control in the Electricity Infrastructure. In *Proc. of the Fourth*

International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS '05, pages 75–82, New York, NY, USA, 2005. ACM.

[Kol09] D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques - Adaptive Computation and Machine Learning*. The MIT Press, 2009.

[Kol11] J. Z. Kolter and M. J. Johnson. REDD: A Public Data Set for Energy Disaggregation Research. In *Proceddings of the SustKDD Workshop on Data Mining Applications in Sustainability*, 2011.

[Kol12] Z. Kolter and T. Jaakkola. Approximate Inference in Additive Factorial HMMs with Application to Energy Disaggregation. In *Proc. of the International Conference on Artifical Intelligence and Statistics*, 2012.

[Lak13] Lakeside Labs and WiTiKee. Monergy Technical Report: Definition of Requirements and System Architecture, September 2013.

[Lak15] Lakeside Labs and WiTiKee. Monergy Technical Report: Validation and Analysis of results, March 2015.

[Lam10] S. Lamparter, S. Becher, and J.-G. Fischer. An Agent-based Market Platform for Smart Grids. In *Proc. of the 9th International Conference on Autonomous Agents and Multiagent Systems: Industry Track*, AAMAS '10, pages 1689–1696, Richland, SC, 2010. International Foundation for Autonomous Agents and Multiagent Systems.

[Li09] H. Li and L. Tesfatsion. Development of open source software for power market research: The ames test bed. *Journal of Energy Markets*, pages 111–128, 2009.

[Lia10] J. Liang, S. Ng, G. Kendall, and J. Cheng. Load Signature Study, Part I: Basic Concept, Structure, and Methodology. *IEEE Trans. Power Del.*, 25(2):551–560, 2010.

[Lin10] G. Lin, S. Lee, J. Hsu, and W. Jih. Applying power meters for appliance recognition on the electric panel. In *Proc. of IEEE Conference on Industrial Electronics and Applications (ICIEA)*, 2010.

[Lip15] Z. C. Lipton. A Critical Review of Recurrent Neural Networks for Sequence Learning. *eprint arXiv:cs-lg/1506.00019*, 2015.

[Lis10] M. Lisovich, D. Mulligan, and S. Wicker. Inferring Personal Information from Demand-Response Systems. *IEEE Security Privacy*, 8(1):11–20, 2010.

[Liu14] Y. Liu, C. Yuen, S. Huang, N. Ul Hassan, X. Wang, and S. Xie. Peak-to-Average Ratio Constrained Demand-Side Management With Consumer's Preference in Residential Smart Grid. *IEEE Journal of Selected Topics in Signal Processing*, 8(6):1084–1097, Dec 2014.

[LP11] D. Le-Phuoc, M. Dao-Tran, J. X. Parreira, and M. Hauswirth. A native and adaptive approach for unified processing of linked streams and linked data. In *Proceedings of the 10th international conference on The semantic web - Volume Part I*, pages 370–388. Springer-Verlag, 2011.

[MÏ0] J. Müller, F. Alt, D. Michelis, and A. Schmidt. Requirements and Design Space for Interactive Public Displays. In *Proc. of the 18th ACM International Conference on Multimedia*, MM '10, pages 1285–1294, 2010.

[Mak13] S. Makonin, F. Popowich, L. Bartram, B. Gill, and I. V. Bajic. AMPds: A Public Dataset for Load Disaggregation and Eco-Feedback Research. In *Proceeding of the IEEE Conference on Electrical Power and Energy (EPEC)*, 2013.

[McC43] W. McCulloch and W. Pitts. A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5:115–133, 1943.

[Mcm02] A. H. Mcmakin, E. L. Malone, and R. E. Lundgren. Motivating Residents to Conserve Energy Without Financial Incentives. *Environment and Behavior*, 34(6):848–863, Nov. 2002.

[Mon13a] A. Monacchi, D. Egarter, and W. Elmenreich. Integrating households into the smart grid. In *Proc. of the IEEE Workshop on Modeling and Simulation of Cyber-Physical Energy Systems*, pages 48–53, Berkeley, California, 2013.

[Mon13b] A. Monacchi and W. Elmenreich. Poster Abstract: Insert-coin: Turning the Household into a Prepaid Billing System. In *Proc. of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, BuildSys'13, pages 26:1–26:2, New York, NY, USA, 2013. ACM.

[Mon13c] A. Monacchi, W. Elmenreich, S. D'Alessandro, and A. M. Tonello. Strategies for Domestic Energy Conservation in Carinthia and Friuli-Venezia Giulia. In *Proc. of the 39th Annual Conference of the IEEE Industrial Electronics Society*, Vienna, Austria, 2013.

[Mon14a] A. Monacchi, D. Egarter, W. Elmenreich, S. D'Alessandro, and A. M. Tonello. GREEND: an energy consumption dataset of households in Italy and Austria. In *Proc. of IEEE International Conference on Smart Grid Communications (SmartGridComm)*, Venice, Italy, Nov 2014.

[Mon14b] A. Monacchi and W. Elmenreich. Home Energy Market Simulator: The Appliance Usage Model Manager. Technical report, Alpen-Adria Universität Klagenfurt, November 2014.

[Mon14c] A. Monacchi, S. Zhevzhyk, and W. Elmenreich. HEMS: a home energy market simulator. *Computer Science - Research and Development*, pages 1–8, 2014.

[Mon14d] A. Monacchi, S. Zhevzhyk, and W. Elmenreich. Home Energy Market Simulator: The Scenario Definition Format. Technical report, Alpen-Adria Universität Klagenfurt, September 2014.

[Mon15] A. Monacchi, F. Versolatto, M. Herold, D. Egarter, A. M. Tonello, and W. Elmenreich. An Open Solution to Provide Personalized Feedback for Building Energy Management. *CoRR*, abs/1505.01311, June 2015.

[Mon16] A. Monacchi and W. Elmenreich. Assisted Energy Management in Smart Microgrids. *CoRR*, abs/1606.01949, June 2016.

[MR10] A.-H. Mohsenian-Rad, V. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia. Autonomous Demand-Side Management Based on Game-Theoretic Energy Consumption Scheduling for the Future Smart Grid. *Smart Grid, IEEE Transactions on*, 1(3):320–331, Dec 2010.

[MS04] C. Muller-Schloer. Organic computing - on the feasibility of controlled emergence. In *International Conference on Hardware/Software Codesign and System Synthesis*, pages 2–5, Sept 2004.

[Ngu13] T. A. Nguyen and M. Aiello. Energy intelligent buildings based on user activity: A survey. *Energy and Buildings*, 56(0):244 – 257, 2013.

[Nie93] J. Nielsen and T. K. Landauer. A Mathematical Model of the Finding of Usability Problems. In *Proc. of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems*, CHI '93, pages 206–213. ACM, 1993.

[Nol00] S. Nolfi and D. Floreano. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines.* MIT Press, Cambridge, MA, USA, 2000.

[OEC15] OECD. World Energy Statistics 2015. Technical report, IEA, 2015.

[Ozo13] M. Ozog. The Effect of Prepayment on Energy Use. Technical report, DEFG LLC, Washington DC, March 2013.

[P̈14] M. Pöchacker, T. Khatib, and W. Elmenreich. The microgrid simulation tool RAPSim: description and case study. In *Proc. of the 2014 IEEE Innovative Smart Grid Technologies (ISGT) Conference - Asia*, Kuala Lumpur, Malaysia, May 2014.

[Pal11] P. Palensky and D. Dietrich. Demand Side Management: Demand Response, Intelligent Energy Systems, and Smart Loads. *IEEE Transactions on Industrial Informatics*, pages 381–388, Aug 2011.

[Par11] S. Parsons, J. A. Rodriguez-Aguilar, and M. Klein. Auctions and Bidding: A Guide for Computer Scientists. *ACM Comput. Surv.*, 43(2):10:1–10:59, Feb. 2011.

[Par12] O. Parson, S. Ghosh, M. Weal, and A. Rogers. Non-intrusive load monitoring using prior models of general appliance types. In *Proc. of the Twenty-Sixth Conference on Artificial Intelligence (AAAI)*, 2012.

[Pas13] R. Pascanu, T. Mikolov, and Y. Bengio. On the difficulty of training recurrent neural networks. In *Proc. of the 30th International Conference on Machine Learning*, pages 1310–1318, 2013.

[Pfi11] D. Pfisterer, K. Romer, D. Bimschas, O. Kleine, R. Mietz, C. Truong, H. Hasemann, A. Kro andller, M. Pagel, M. Hauswirth, M. Karnstedt, M. Leggieri, A. Passant, and R. Richardson. SPITFIRE: toward a semantic web of things. *IEEE Communications Magazine*, pages 40 –48, 2011.

[Phe08] S. Phelps, K. Cai, P. McBurney, J. Niu, S. Parsons, and E. Sklar. Auctions, Evolution, and Multi-agent Learning. In *Adaptive Agents and Multi-Agent Systems III. Adaptation and Multi-Agent Learning*, volume 4865 of *Lecture Notes in Computer Science*, pages 188–210. Springer Berlin Heidelberg, 2008.

[Poc15] M. Pochacker, D. Egarter, and W. Elmenreich. Proficiency of Power Values for Load Disaggregation. *Instrumentation and Measurement, IEEE Transactions on*, PP(99):1–10, October 2015.

[Pos09] S. Poslad. *Ubiquitous Computing: Smart Devices, Environments and Interactions.* Wiley Publishing, 1st edition, 2009.

[Pre04] D. Preuveneers, J. V. D. Bergh, D. Wagelaar, A. Georges, P. Rigole, T. Clerckx, E. Berbers, K. Coninx, and K. D. Bosschere. Towards an extensible context ontology for ambient intelligence. In *Proceedings of the Second European Symposium on Ambient Intelligence*, 2004.

[Ram12] S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. R. Jennings. Putting the 'Smarts' into the Smart Grid: A Grand Challenge for Artificial Intelligence. *Commun. ACM*, 55(4):86–97, Apr. 2012.

[Ram13] S. Ramchurn, M. Osborne, O. Parson, T. Rahwan, S. Maleki, S. Reece, T. D. Huynh, M. Alam, J. Fischer, T. Rodden, L. Moreau, and S. Roberts. AgentSwitch: towards smart electricity tariff selection. In *12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, May 2013.

[Ras04] R. G. Raskin. Enabling semantic interoperability for earth science data. http://sweet.jpl.nasa.gov, 2004.

[Rei11] C. Reinisch, M. J. Kofler, F. Iglesias, and W. Kastner. ThinkHome Energy Efficiency in Future Smart Homes. *EURASIP Journal Embedded Syst.*, 2011:1:1–1:18, Jan. 2011.

[Rei12] A. Reinhardt, P. Baumann, D. Burgstahler, M. Hollick, H. Chonov, M. Werner, and R. Steinmetz. On the Accuracy of Appliance Identification Based on Distributed Load Metering Data. In *Proc. of the 2nd IFIP Conference on Sustainable Internet and ICT for Sustainability (SustainIT)*, 2012.

[Saa12] W. Saad, Z. Han, H. Poor, and T. Basar. Game-Theoretic Methods for the Smart Grid: An Overview of Microgrid Systems, Demand-Side Management, and Smart Grid Communications. *Signal Processing Magazine, IEEE*, 29(5):86–105, Sept 2012.

[Sak09] S. Sakr. XML compression techniques: A survey and comparison. *Journal of Computer and System Sciences*, 75(5):303–322, 2009.

[Sie91] H. T. Siegelmann and E. D. Sontag. Turing computability with neural nets. *Applied Mathematics Letters*, 4(6):77 – 80, 1991.

[Sil08] A. Silberschatz, P. B. Galvin, and G. Gagne. *Operating System Concepts*. Wiley Publishing, 8th edition, 2008.

[Sko12] F. Skopik. Security Is Not Enough! On Privacy Challenges in Smart Grids. *International Journal on Smart Grid and Clean Energy*, pages 7–14, 2012.

[Sob13] A. Sobe and W. Elmenreich. Replication and replacement in dynamic delivery networks. *Complex Adaptive Systems Modeling*, 1, 2013.

[Sri06] D. Srinivasan, W. S. Ng, and A. Liew. Neural-network-based signature recognition for harmonic source identification. *IEEE Trans. Power Del.*, 21(1):398–405, 2006.

[Sta02] K. O. Stanley and R. Miikkulainen. Evolving Neural Networks Through Augmenting Topologies. *Evol. Comput.*, 10(2):99–127, June 2002.

[Sta09] K. O. Stanley, D. B. D'Ambrosio, and J. Gauci. A Hypercube-based Encoding for Evolving Large-scale Neural Networks. *Artif. Life*, 15(2):185–212, April 2009.

[Sta12] T. G. Stavropoulos, D. Vrakas, D. Vlachava, and N. Bassiliades. BOnSAI: A Smart Building Ontology for Ambient Intelligence. In *Proceedings of the 2nd International Conference on Web Intelligence, Mining and Semantics*, 2012.

[Str11] Y. A. Strengers. Designing eco-feedback systems for everyday life. In *Proc. of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, pages 2135–2144, New York, NY, USA, 2011. ACM.

[Suz08] K. Suzuki, S. Inagaki, T. Suzuki, H. Nakamura, and K. Ito. Nonintrusive appliance load monitoring based on integer programming. In *Proc. of International conference on Instrumentation, Control, Information Technology and System Integration (SICE)*, 2008.

[Swi96] K. Swingler. *Applying Neural Networks: A Practical Guide*. Morgan Kaufmann, pap/dsk edition, 5 1996.

[Tan06] A. S. Tanenbaum and M. v. Steen. *Distributed Systems: Principles and Paradigms (2nd Edition)*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 2006.

[Tom10] S. Tomic, A. Fensel, and T. Pellegrini. SESAME Demonstrator: Ontologies, Services and Policies for Energy Efficiency. In *Proceedings of the 6th International Conference on Semantic Systems (I-SEMANTICS)*, 2010.

[Tri11] V. Trianni and S. Nolfi. Engineering the Evolution of Self-Organizing Behaviors in Swarm Robotics: A Case Study. *Artificial Life*, 17(3):183–202, July 2011.

[Tru13] N. C. Truong, L. Tran-Thanh, E. Costanza, and D. S. Ramchurn. Activity prediction for agent-based home energy management. In *Agent Technologies for Energy Systems (ATES 2013)*, May 2013.

[Uni88] United Nations. Kyoto protocol to the united nations framework convention on climate change, 1988.

[Uta13] H. Uta, C. Sheelagh, V. Nina, K. Kai, J. Giulio, and A. Moere. Guest Editors' introduction: interactive public displays. *IEEE Computer Graphics and Applications*, 33(2):25–27, 3 2013.

[Vil13] B. Villaverde, R. Alberola, A. Jara, S. Fedor, S. Das, and D. Pesch. Service Discovery Protocols for Constrained Machine-to-Machine Communications. *IEEE Communications Surveys Tutorials*, pages 1–20, 2013.

[Vir92] R. A. Virzi. Refining the Test Phase of Usability Evaluation: How Many Subjects is Enough? *Hum. Factors*, 34(4):457–468, Aug 1992.

[Vyt10] P. Vytelingum, S. D. Ramchurn, T. D. Voice, A. Rogers, and N. R. Jennings. Trading Agents for the Smart Electricity Grid. In *Proc. of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS '10)*, pages 897–904, 2010.

[Wan14] Y. Wang, W. Saad, Z. Han, H. Poor, and T. Basar. A Game-Theoretic Approach to Energy Trading in the Smart Grid. *IEEE Transactions on Smart Grid*, 5(3):1439–1450, May 2014.

[Wij13] T. Wijaya, D. Banerjee, T. Ganu, D. Chakraborty, S. Battacharya, T. Papaioannou, D. Seetharam, and K. Aberer. DRSim: A cyber physical simulator for Demand Response systems. In *Proc. of the IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pages 217–222, Oct 2013.

[Woo09] M. Wooldridge. *An Introduction to MultiAgent Systems.* Wiley Publishing, 2nd edition, 2009.

[Ygg96] F. Ygge and H. Akkermans. Power load management as a computational market. In *Second International Conference on Multiagent Systems (ICMAS)*, pages 393–400, Kyoto, 1996.

[Zai10] A. A. Zaidi, F. Kupzog, T. Zia, and P. Palensky. Load Recognition for Automated Demand Response in Microgrids. In *Proc. of the 36th IEEE Conference on Industrial Electronics IECON*, 2010.

[Zei11] M. Zeifman and K. Roth. Nonintrusive appliance load monitoring: Review and outlook. *IEEE Trans. Consum. Electron.*, 57(1):76 –84, february 2011.

[Zei12] M. Zeifman. Disaggregation of home energy display data using probabilistic approach. *IEEE Trans. Consum. Electron.*, 58(1):23–31, 2012.

[Zhe15] S. Zhevzhyk and W. Elmenreich. Comparison of Metaheuristic Algorithms for Evolving a Neural Controller for an Autonomous Robot. *Transactions on Machine Learning and Artificial Intelligence*, 2(6):62–76, 2015.

[Zoh13] A. Zoha, A. Gluhak, M. Nati, and M. Imran. Low-power appliance monitoring using Factorial Hidden Markov Models. In *Proc. of IEEE Conference on Intelligent Sensors, Sensor Networks and Information Processing*, 2013.

CHAPTER A Acceptance of the advisor widget

This section lists the questions used for the satisfaction questionnaire and the answers collected.

1. It takes short time to learn the meanings of the buttons
2. The position of the buttons is logical
3. I understand what happens when I click the buttons
4. The advices are unusual, inventive, original
5. The advices are useful to improve energy efficiency
6. The advices are doable
7. I can learn something from the advices
8. I would use this widget every day
9. I would use this widget again

Table A.1: Answers to the satisfaction questionnaire

	r1	r2	r3	r4	r5	r6	r7
q1	1	0	1	1	1	1	0
q2	1	0	2	1	1	0	1
q3	2	1	1	0	1	2	1
q4	0	1	0	0	-1	1	-1
q5	1	2	1	1	2	1	0
q6	0	1	0	-1	0	-1	0
q7	0	1	1	0	1	-1	1
q8	-1	-1	-2	-2	-1	-2	-2
q9	2	1	1	-1	0	1	0

CHAPTER B

Scenarios used for training appliance controllers

B.1 Scenario 1: learning to sell energy

In this scenario, a controller is learned for a photovoltaic generator having peak power of $4kW$ and located in the area of Klagenfurt (Austria) (List. B.1).

Listing B.1: First scenario


```
{"weather":{"type":"time-intervals",
            "intervals":[{"start":[12,21,0,0,0],"end":[3,21,0,0,0], "cloud-factor":0.6},
                         {"start":[3,21,0,0,0],"end":[6,21,0,0,0], "cloud-factor":0.4},
                         {"start":[6,21,0,0,0],"end":[9,21,0,0,0], "cloud-factor":0.2},
                         {"start":[9,21,0,0,0],"end":[12,21,0,0,0], "cloud-factor":0.9}]},
 "grid_connections":[{"name":"Grid_connection_1", "credit":10.0,
                      "price-model": {"tariff": {"type": "time-intervals",
                                                 "intervals" : [{"name":"day", "start":[6,0], "end":[20,0], "cost":0.50},
                                                                {"name":"night", "start":[20,0], "end":[6,0], "cost":0.30}]},
                                      "feed-in":{"type": "time-intervals",
                                                 "intervals" : [{"name":"day", "start":[6,0], "end":[20,0], "cost":0.20},
                                                                {"name":"night", "start":[20,0], "end":[6,0], "cost":0.05}]} },
                      "capability-model":{"power-availability": {"type": "time-intervals", "intervals" : [{"start":[0,0], "end":[23,59], "amount":600}] }
                           ,
                                          "power-capability": {"type" : "time-intervals", "intervals" : [{"start":[0,0], "end":[23,59], "amount":3000}] }
                                                }
                      }],
 "producers":[{"name":"Photovoltaic_house_1", "idle":0, "price-model":0.1,
               "type":"MODELED-PV", "payload":{"peakPower":4000.0, "efficiency":0.15, "latitude":46.6, "longitude":14.4, "height":0.446, "size":50.0}} ],
 "loads":[]
}
```

B.2 Scenario 2: learning a prosumer able to buy and sell energy

In this scenario, a pool of loads is added to assess the possibility to buy power depending on the provided synthetic usage models (List. B.2).

Listing B.2: Second scenario

```
{
 "weather":{"type":"time-intervals",
            "intervals": [{"start":[12,21,0,0,0],"end":[3,21,0,0,0], "cloud-factor":0.6},
                          {"start":[3,21,0,0,0],"end":[6,21,0,0,0], "cloud-factor":0.4},
                          {"start":[6,21,0,0,0],"end":[9,21,0,0,0], "cloud-factor":0.2},
                          {"start":[9,21,0,0,0],"end":[12,21,0,0,0], "cloud-factor":0.9}]},
 "grid_connections":[ {"name":"Grid_connection_1", "credit":10.0,
                       "price-model": {"tariff": {"type": "time-intervals",
                                                  "intervals" : [{"name":"day", "start":[6,0], "end":[20,0], "cost":0.50},
                                                                 {"name":"night", "start":[20,0], "end":[6,0], "cost":0.30}]},
                                       "feed-in":{"type": "time-intervals",
                                                  "intervals" : [{"name":"day", "start":[6,0], "end":[20,0], "cost":0.20},
                                                                 {"name":"night", "start":[20,0], "end":[6,0], "cost":0.05}]}},
                       "capability-model": {"power-availability": {"type": "time-intervals", "intervals" : [{"start":[0,0], "end":[23,59], "amount":3000}]},
                                            "power-capability": {"type" : "time-intervals", "intervals" : [{"start":[0,0], "end":[23,59], "amount":3000}]}}
                       }],
 "producers":[ {"name":"Photovoltaic_house_1", "idle":0, "price-model":0.15,
                "type":"MODELED-PV", "payload":{"peakPower":4000.0, "efficiency":0.15, "latitude":46.6, "longitude":14.4, "height":0.446, "size":200.0}}],
 "loads":[
    {"name":"Dishwasher", "credit":10.0, "idle":0, "offer_expiration":3600, "type":"RANDOM", "deferrable":true,
        "payload":{"power":[[800,5,60],[200,5,10],[400,600,2],[600,120,10],[200,5,10]],
                   "willingness":0.6, "willingness_decay":0.6, "sensitivity":0.25}},
    {"name":"Fridge", "credit":10.0, "idle":0, "offer_expiration":3600, "type":"RANDOM", "deferrable":true,
        "payload":{"power":[[200,5,60]],
                   "willingness":0.8, "willingness_decay":0.1, "sensitivity":0.6}},
    {"name":"Boiler", "credit":10.0, "idle":0, "offer_expiration":3600, "type":"RANDOM", "deferrable":true,
        "payload":{"power":[[800,5,60],[800,5,10],[800,5,10],[800,5,10]],
                   "willingness":0.5, "willingness_decay":0.5, "sensitivity":0.3}},
    {"name":"Washingmachine_1", "credit":10.0, "idle":0, "offer_expiration":3600, "type":"RANDOM", "deferrable":true,
        "payload":{"power":[[800,120,60],[200,5,10],[400,10,2],[400,10,2],[400,10,2],[400,10,2],[400,10,2],[400,10,2],
                            [400,10,2],[400,10,2],[400,10,2],[400,10,2],[700,120,10],[200,2,10]],
                   "willingness":0.6, "willingness_decay":0.5, "sensitivity":0.4}},
    {"name":"Washingmachine_2", "credit":10.0, "idle":0, "offer_expiration":3600, "type":"RANDOM", "deferrable":true,
        "payload":{"power":[[800,120,60],[200,5,10],[400,10,2],[400,10,2],[400,10,2],[400,10,2],[400,10,2],[400,10,2],
                            [400,10,2],[400,10,2],[400,10,2],[400,10,2],[700,120,10],[200,2,10]],
                   "willingness":0.6, "willingness_decay":0.5, "sensitivity":0.5}},
    {"name":"Light_bedroom", "credit":5.0, "idle":0, "offer_expiration":3600, "type":"RANDOM", "deferrable":false,
        "payload":{"power":[[60,120]],
                   "willingness":0.6, "willingness_decay":0.5}}
]}
```

B.3 Scenario 3: multiple demand and fragmented local production

In this scenario, local generation is split in two different photovoltaic plants, indipendently seeking profit maximization through their trade (List. B.3).

Listing B.3: Third scenario

```
{
 "weather":{"type":"time-intervals",
           "intervals": [{"start":[12,21,0,0,0],"end":[3,21,0,0,0], "cloud-factor":0.6},
                         {"start":[3,21,0,0,0],"end":[6,21,0,0,0], "cloud-factor":0.4},
                         {"start":[6,21,0,0,0],"end":[9,21,0,0,0], "cloud-factor":0.2},
                         {"start":[9,21,0,0,0],"end":[12,21,0,0,0], "cloud-factor":0.9}]},
 "grid_connections":[ {"name":"Grid_connection_1", "credit":10.0,
                      "price-model": {"tariff": {"type": "time-intervals",
                                                 "intervals": [{"name":"day", "start":[6,0], "end":[20,0], "cost":0.50},
                                                               {"name":"night", "start":[20,0], "end":[6,0], "cost":0.30}]},
                      "feed-in":{"type": "time-intervals",
                                 "intervals":[{"name":"day", "start":[6,0], "end":[20,0], "cost":0.20},
                                              {"name":"night", "start":[20,0], "end":[6,0], "cost":0.05}]}},
                      "capability-model": {"power-availability": {"type": "time-intervals", "intervals" : [{"start":[0,0], "end":[23,59], "amount":400}]
                           },
                                           "power-capability": {"type" : "time-intervals", "intervals" : [{"start":[0,0], "end":[23,59], "amount":3000}]
                                                }}
                      }],
 "producers":[
    {"name":"Photovoltaic_house_1", "idle":0, "price-model":0.1, "type":"MODELED-PV",
           "payload":{"peakPower":1600.0, "efficiency":0.15, "latitude":46.6, "longitude":14.4, "height":0.446, "size":50.0}},
    {"name":"Photovoltaic_house_2", "idle":0, "price-model":0.1, "type":"MODELED-PV",
           "payload":{"peakPower":1600.0, "efficiency":0.15, "latitude":46.6, "longitude":14.4, "height":0.446, "size":50.0}}],
 "loads":[
  {"name":"Dishwasher", "credit":10.0, "idle":0, "offer_expiration":3600, "type":"RANDOM", "deferrable":true,
           "payload":{"power":[[800,5,60],[200,5,10],[400,600,2],[600,120,10],[200,5,10]],
                      "willingness":0.6, "willingness_decay":0.6, "sensitivity":0.5}},
  {"name":"Fridge", "credit":10.0, "idle":0, "offer_expiration":3600, "type":"RANDOM", "deferrable":true,
          "payload":{"power":[[200,5,60]],
                     "willingness":0.8, "willingness_decay":0.1, "sensitivity":0.5}},
  {"name":"Boiler", "credit":10.0, "idle":0, "offer_expiration":3600, "type":"RANDOM", "deferrable":true,
          "payload":{"power":[[800,5,60],[800,5,10],[800,5,10],[800,5,10]],
                     "willingness":0.5, "willingness_decay":0.5, "sensitivity":0.3}},
  {"name":"Washingmachine_1", "credit":10.0, "idle":0, "offer_expiration":3600, "type":"RANDOM", "deferrable":true,
          "payload":{"power":[[800,120,60],[200,5,10],[400,10,2],[400,10,2],[400,10,2],[400,10,2],[400,10,2],
                              [400,10,2],[400,10,2],[400,10,2],[400,10,2],[400,10,2],[700,120,10],[200,2,10]],
                     "willingness":0.6, "willingness_decay":0.5, "sensitivity":0.5}},
  {"name":"Washingmachine_2", "credit":10.0, "idle":0, "offer_expiration":3600, "type":"RANDOM", "deferrable":true,
          "payload":{"power":[[800,120,60],[200,5,10],[400,10,2],[400,10,2],[400,10,2],[400,10,2],[400,10,2],
                              [400,10,2],[400,10,2],[400,10,2],[400,10,2],[400,10,2],[700,120,10],[200,2,10]],
                     "willingness":0.6, "willingness_decay":0.5, "sensitivity":0.5}},
  {"name":"Light_bedroom", "credit":5.0, "idle":0, "offer_expiration":3600, "type":"RANDOM", "deferrable":false,
          "payload":{"power":[[60,120]], "willingness":0.6, "willingness_decay":0.5}}
]}
```

List of Own Publications

During the course of this study several publications have been presented to the research community. In particular, we collaborated to the publication of 2 book chapters, 5 journal papers, 3 conference and 3 workshop papers.

Journal Publications

1. A. Monacchi, F. Versolatto, M. Herold, D. Egarter, W. Elmenreich and Andrea M. Tonello. An Open Solution to Provide Personalised Feedback for Building Energy Management. IOS Journal of Ambient Intelligence and Smart Environments. 2016 (to appear).

2. A. Monacchi, W. Elmenreich. Assisted Energy Management in Smart Microgrids. Springer Journal of Ambient Intelligence and Humanized Computing. 2016.

3. D. Egarter, A. Monacchi, T. Khatib and W. Elmenreich. Integration of Legacy Appliances into Home Energy Management Systems. *Journal of Ambient Intelligence and Humanised Computing*, 2015.

4. T. Khatib, A. Monacchi, W. Elmenreich, D. Egarter, S. D'Alessandro and A. M. Tonello. European end-user's level of energy consumption and current structural barriers for smart homes: A case study of residential sectors in Austria and Italy. *Energy Technology & Policy*, 2014.

5. A. Monacchi, S. Zhevzhyk and W. Elmenreich. HEMS - A Home Energy Market Simulator *Computer Science – Research and Development*, 2014.

Conference Publications

6. A. Monacchi, S. Zhevzhyk and W. Elmenreich. HEMS - A Home Energy Market Simulator. *Energieinformatik*, 2014.

7. A. Monacchi, D. Egarter, W. Elmenreich, S. D'Alessandro and A. M. Tonello. GREEND: an energy consumption dataset of households in Italy and Austria. *IEEE International Conference on Smart Grid Communication (SmartGridComm)*, 2014.

8. S. D'Alessandro, A. M. Tonello, A. Monacchi and W. Elmenreich. Home Energy Management Systems: Design Guidelines for the Communication Infrastructure. *IEEE International Energy Conference (ENERGYCON 2014)*, 2014.

9. A. Monacchi, W. Elmenreich, S. D'Alessandro and A. M. Tonello. Strategies for Domestic Energy Conservation in Carinthia and Friuli-Venezia Giulia. *39th Annual Conference of the IEEE Industrial Electronics Society (IECON)*, 2013.

Workshop Publications

10. A. Monacchi and W. Elmenreich. Poster Abstract: Insert coin: turning the household into a prepaid billing system. *5th ACM Workshop On Embedded Systems For Energy-Efficient Buildings (BuildSys 2013)* November 2013.

11. A. Monacchi, D. Egarter, and W. Elmenreich. Integrating Households into the Smart Grid. *Workshop on Modeling and Simulation of Cyber-Physical Energy Systems (at E-Energy conference)*, 2013.

12. A. Kercek, W. Elmenreich and A. Monacchi. Energieverbrauch in den Regionen Kärnten, Österreich und Friaul-Julisch-Venetien, Italien - Ein Vergleich. *13. Symposium Energieinnovation. Graz*, 2014.

Book Chapter

13. D. Egarter, A. Monacchi, M. Pöchacker, K. Schweiger and B. Steinwender: Smart Grid: Visionen & Herausforderungen, published in Energie - Interdisziplinäre Perspektiven auf eine knappe Ressource (Profil Verlag 2015; G. Getzinger, H. Gross)

14. W. Elmenreich, A. Monacchi. Configuration and Management of Networked Embedded Devices, published in Industrial Communication Technology Handbook, Second Edition (CRC Press 2014)

Andrea Monacchi

Personal	Citizenship:	Italian
	Date and place of birth:	Gubbio, September 13th, 1988

Industrial experience	**Big Data Engineer**, Data Reply GmbH, München, Germany	July 2016 - today
	IT Consultant, Power Reply GmbH, München, Germany	Feb. 2016 - June 2016
	Software Engineer, Townet Srl, Cagli, PU Italy	June - Sept. 2009
	Intern Technician, A.T.E. Elettronica Srl, Gubbio, PG Italy	July 2006
	Intern Web Developer, Lyn-x Snc, Gubbio, PG Italy	July - August 2004
Research experience	**Grant researcher**, Alpen-Adria-Universität Klagenfurt	Oct. 2015 - Jan. 2016
	Research assistant, Lakeside Labs GmbH, Klagenfurt, Austria	Oct. 2012 - Aug. 2015
	Student Researcher, CADIA lab, Reykjavík, Iceland	Aug. 2011 - June 2012
Education	**Ph.D. in Informationstechnik - Information Engineering** Institute of Networked Embedded Systems, Smart Grid group Alpen-Adria Universität Klagenfurt, Klagenfurt (Austria)	Oct. 2012 - today
	Meistaragráða í tölvunarfræði - M.Sc. in Computer Science Reykjavik University, Reykjavík (Iceland)	June 2012
	M.Sc. *cum Laude* in Computer Science University of Camerino, Camerino (Italy)	July 2012
	B.Sc. *cum Laude* in Applied Computer Science “Carlo Bo” University of Urbino, Urbino (Italy)	Oct. 2010
	Technical High School in Electronics and Telecommunications I.T.I. “M. L. Cassata”, Gubbio (Italy)	June 2007

Academic Honors	Italian Ministry of Education & Research	”Cooperlink” scholarship	2011
	Camerino University	Erasmus scholarship	2011
	Camerino University	Merit scholarship	2012
	Marche Region (Italy)	MS Thesis award	2014
	Alpen-Adria Universität Klagenfurt	Forschungsstipendium	2015