Energy Reports 6 (2020) 94–106

Contents lists available at ScienceDirect

Energy Reports

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

Research paper

Understanding the practical issues of deploying energy monitoring and eco-feedback technology in the wild: Lesson learned from three long-term deployments

Lucas Pereira *, Nuno Nunes

ITI, LARSyS, Técnico Lisboa, 1049-001 Lisbon, Portugal

ARTICLE INFO

Article history:
Received 26 May 2019
Received in revised form 23 September 2019
Accepted 14 November 2019
Available online xxxx

Keywords:
Non-intrusive energy monitoring
Eco-feedback
Real-world deployments
Technological challenges
Social challenges
Financial challenges

ABSTRACT

This paper reports on the different engineering, social and financial challenges behind the building and deploying electric energy monitoring and eco-feedback technology in real-world scenarios, which despite being relevant to the research community are seldom reported in the literature. The objectives of this paper are two-fold: First, discuss the technical and social constraints of real-world deployments. This includes, for example, hardware and software requirements, and issues related to security and intrusiveness of the monitoring solutions. Second, identify and understand the costs associated with developing and deploying such systems. These include hardware costs and consumed energy. To this end, we rely on over five years of experience developing and improving a non-intrusive energy monitoring research platform to enable the deployment of long and short-term studies of eco-feedback technology. During this time, two versions of that platform were deployed in 50 homes for periods that lasted between 6 and 18 consecutive months. By iteratively developing and deploying our sensing and eco-feedback infrastructures, we managed to build upon previous findings and lessons learned to understand how to create, deploy, and maintain such systems. Concurrently, we gained insights regarding what are some of the most relevant costs associated with running such experiments.

1. Introduction

The global demand for electric energy has been experiencing a steady increase since 1990, emerging as the second most used end-form of energy with a 17.7% share, only behind oil with 40.8% (International Energy Agency, 2014). One of the factors leading to the growth in electricity consumption in the last years is the notion of well-being based on personal ownership and mass consumption. As more people in developing countries have access to higher levels of comfort it is expected (US Energy Information Administration, 2016) that the world demand for electricity will continue to increase in the next decades.

Nevertheless, the improvements in the quality of life enabled by electricity do not come without environmental costs. In fact, evidence shows that the carbon dioxide emissions from fuel combustion used to generate electric energy, have been steadily increasing since 1990, and are set to grow 46% by 2040 (International Energy Agency, 2014). Hence, the importance of domestic electric energy in the global context of energy over-consumption as outlined in Pacala and Socolow (2004), where the authors mention that the residential sector holds the potential for achieving one of the seven wedges required to stabilize carbon dioxide emissions by 2054.

Many studies suggest that providing householders with real-time and historical information about their consumption can lead to potential savings between 5% and 10% (Parker et al., 2006; Fischer, 2008; Jain et al., 2012), especially in the cases where the feedback is enhanced with individual appliance consumption information (Carrie Armel et al., 2013). This is commonly known as eco-feedback technology and is defined as the technology that provides feedback on individual or group behaviors with the goal of reducing environmental impact (Froehlich et al., 2010).

The underlying assumption behind eco-feedback technology is that people will be able to change their actions and consequently reduce their consumption if they can understand which appliances are responsible for their overall energy consumption. This effect was especially noticed in Parker et al. (2006), where Parker and colleagues evaluated two low-cost monitoring systems and found that users quickly discovered that by merely examining the differences in the overall demand by turning appliances ON and OFF they could easily approximate the energy usage of each electric device.

Therefore, there has been a substantial effort to create monitoring solutions that can provide the consumption figures of individual appliances. Including, for instance, electrical sub-metering

* Corresponding author.
E-mail address: lucas.pereira@tecnico.ulisboa.pt (L. Pereira).

https://doi.org/10.1016/j.egyr.2019.11.025

(i.e., installing sensors in each device) and the development of non-intrusive load monitoring (NILM) techniques that can sense and disaggregate energy consumption from a single sensing location in the distribution grid (Hart, 1985).

Even though the reported results are mostly positive regarding improved awareness and achieving savings in energy consumption (Fischer, 2008), it has been reported that after an initial period of exposure to this technology, the tendency is towards a decrease in the attention given to the feedback leading to behavior relapse (Peschiera et al., 2010; Pereira et al., 2013). This effect is defined in the literature as *response-relapse* and it suggests that to accurately assess the effectiveness of eco-feedback as a tool for promoting sustained energy savings, future studies should be carried for more extended periods.

Furthermore, and despite the abundance of literature in energy eco-feedback (e.g., Darby, 2006; Ehrhardt-Martinez et al., 2010; Vine et al., 2013; Karlin et al., 2015) and NILM (e.g., Zeifman and Roth, 2011; Zoha et al., 2012; Esa et al., 2016; Nalmpantis and Vrakas, 2018), it was only recently that the community saw the first publications regarding the value proposition of NILM as a tool to reduce energy consumption (Kelly and Knottenbelt, 2016; Batra et al., 2016), or trying to educate the research community about the practical issues of deploying and maintaining such systems in real-world conditions (Mayhorn et al., 2016; Kosonen and Kim, 2016), which we believe are of crucial importance to the large-scale adoption of such technologies in years to come.

This paper discusses the practical issues of deploying NILM and eco-feedback systems in domestic environments. More concretely, we report on the many technical, social and financial challenges behind the building and deploying such technologies in real-world scenarios, which despite being relevant to the research community are seldom reported in the literature. To do this, we rely on more than 5 years of experience developing and improving a research platform that relies on non-intrusive energy meters to enable the quick deployment of eco-feedback technology studies, and at the same time serve as a research platform to evaluate NILM algorithms.

By iteratively developing and deploying our sensing and eco-feedback infrastructures we managed to build upon previous findings and lessons learned to gain a deeper understanding of how to create, deploy and maintain such systems. Concurrently, we gained valuable insights regarding what are some of the most relevant costs associated with running such experiments. When taken together, the different insights and lessons learned from the three deployments represent an advancement in the state of the art in the live deployment of non-intrusive energy monitoring and eco-feedback systems, and can certainly apply to researchers planning similar deployments.

The remaining of this paper is organized as follows: First, we provide an overview of the related literature. Second, we describe our research platform, and provide details of the three live deployments that were conducted. Third, we provide a comprehensive discussion of the many practical considerations of deploying and maintaining such systems. Finally, we conclude this paper highlighting the implications and limitations of this work, and briefly describing how some of the lessons learned were applied in a different research project.

2. Related literature

Literature on the practical issues of real-world deployments is not vast, with most of the published works referring to the early days of ubiquitous computing (Hansen et al., 2006; Konomi and Roussos, 2007; Huang et al., 2007), public displays (Ojala et al., 2012), and wireless sensor networks (Tateson et al., 2005; Langendoen et al., 2006) in the first decade of the 21st century. Overall, the issues reported in the different papers fall under four broad categories: (i) technical, (ii) infrastructural, (iii) organizational, and (iv) end-users.

Concerning the technical issues, one of the main challenges reported is the integration of the different hardware and software components, as well as countless communication issues (Tateson et al., 2005; Langendoen et al., 2006). Another challenge that is reported in the need for energy sources to power the devices. Such an issue is particularly relevant in the case of wireless sensor deployments in the wild, as reported in Langendoen et al. (2006). Ultimately, the failure to set up proper communication and energy sources results in tremendous data losses. For example, in Langendoen et al. (2006) communication issues lead to the quick dry-out of the sensor-node batteries leading to o total failure of the sensor network.

Regarding software, the literature reveals two main concerns. First, the possibility of debugging in case of software malfunction, and second, how easy it is to deploy updates remotely (Langendoen et al., 2006; Hansen et al., 2006). Likewise, in deployments that require integration with existing infrastructures or third-party systems, it is necessary to take into consideration the constraints of the environment where the technology will be deployed (e.g., the existence of firewalls) (Hansen et al., 2006; Huang et al., 2007). Another concern that appears reported in the literature are the potential safety issues in case of hardware malfunctioning, and if such issues would interfere with the systems already in place. Such concerns are particularly important when the deployments happen in critical scenarios like hospitals (Hansen et al., 2006).

With respect to the infrastructural challenges, it is important to keep space limitations in mind (Hansen et al., 2006). For example, in indoor deployments, it may be difficult to find enough space to install all the necessary hardware (e.g., sensing and feedback devices). On the other hand, for outdoor deployments, it is necessary to take into consideration the security of the environment (e.g., the risk of theft and vandalism) and extreme weather constraints that cannot be captured in laboratory studies. For example, in Tateson et al. (2005) a change in the weather prevented the research team from recovering the deployed sensors on the expected day, representing a delay of 10 days.

Another challenge that was reported in the literature is related to the organizational rules. For example, in Langendoen et al. (2006), the project timeline was restricted to the duration of the growing season of the crops being monitored, between mid-May and mid-August. In the real world, security and privacy are major issues. For example, in Hansen et al. (2006), the deployed system handled real patient data as well as the medical staff's location information. As such, the system had to be shielded by the existing firewall, and outside accesses could only be made using virtual private networks.

Finally, literature also reveals some challenges with regard to the end-users. First and foremost, its very likely that in the real-world the end-users have very different priorities and motivations, which are not possible to capture in laboratory settings (Huang et al., 2007; Ojala et al., 2012). Likewise, literature also reveals the challenges of preparing the staff to deal with the changes in their working routines. This was particularly evident in settings where end-user training needs to be performed during business as usual operation (Konomi and Roussos, 2007; Huang et al., 2007). Finally, there were some concerns regarding the loss of privacy, but surprisingly this was not among the main concerns of the end-users. In fact, different studies reveal that when there is a perceived benefit, people are willing to trade-off some of their privacy (Hansen et al., 2006; Konomi and Roussos, 2007).

3. Non-intrusive energy monitoring and eco-feedback deployments in the real-world

The research platform described here is part of a series of sustainability research projects, which involved a team of multidisciplinary researchers looking at using sensing, social networking, and context awareness to understand and motivate people to reduce their energy consumption in the residential sector.

Under the umbrella of these projects, we developed a hardware and software platform to simplify and reduce the costs of deploying and maintaining energy monitoring and eco-feedback solutions in the real-world during long periods. The failure to find a viable commercial solution led to the development of the two custom end-to-end unobtrusive energy monitoring and eco-feedback platforms, which are described in this section.

During these projects, the two platforms were deployed in three long-term energy monitoring and eco-feedback deployments that lasted between six and 18 consecutive months. The three deployments involved a total of 50 different households that had the system installed and running continuously in their houses to acquire energy consumption and user-interaction data. During that period, the system was continuously monitored and perfected. Several eco-feedback studies, including qualitative interviews and surveys, were conducted to understand how people react and adopt energy monitoring and eco-feedback technologies (Viana, 2011; Barreto, 2014; Quintal, 2016).

Here, we focus on the practical issues of building, deploying, and maintaining such systems for long periods, which, as we mentioned previously, are very seldom reported in the literature.

Next, the two energy monitoring and eco-feedback research platforms, and the three deployments are briefly described. Additional information such as pictures of the deployment process and translations of the User Interfaces is provided as supplementary material.

3.1. Deployments timeline

The two research platforms were successfully deployed in three long-term energy monitoring and eco-feedback deployments that took place in Funchal, the capital city of Madeira Island in Portugal. Altogether, the three deployments involved 50 different households.

In the first deployment, the sensors and the eco-feedback device had to be installed directly in the main fuse box of the houses. To this end, and since this involved specialized work, certified electricians from the local electricity provider did all the installations.

Altogether, considering each day with at least one installation, we needed 16 days to complete the installation of the 23 energy monitors that comprised the first deployment. This deployment lasted for 658 consecutive days between the end of July 2010 when the first device was installed, and mid-May 2012 when the last one was removed.

Regarding the second and third deployments, since the monitoring platform had to be installed in the building main electric panel the sample was recruited from several apartment buildings in Funchal. Furthermore, due to the considerable complexity of the building's main electrical panels, all the installations were also performed by qualified electricians from the local electricity provider.

The second deployment started in the beginning of August 2012 and lasted until the end of May 2013 in a total of 298 consecutive days. As for the third deployment, it happened between the beginning of August 2013 to the end of April 2014, when the platform was removed from the building. Overall, this deployment lasted 268 consecutive days.

3.2. Single-house energy monitoring and eco-feedback platform

The requirements for a single-house energy monitoring and eco-feedback system with energy disaggregation capabilities can easily reach hundreds of Euros (each element costs between 50 to 100 Euros without the cost of integration) (Pereira et al., 2012).

As such, after several attempts with custom hardware, it was decided to use a netbook since it provided all the required elements in a compact package that would cost between 200 and 300 Euros. The sound-card served as the data acquisition module (two channels, one for current and another for voltage) using the built-in Analog to Digital Converter (ADC). The mini display and the speakers provided the feedback, while the Wi-Fi and Ethernet cards enabled communication over the Internet. Lastly, the built-in camera and microphone would act as low-cost sensors for human activity sensing.

In this version of the platform (used in the first deployment) the energy sensors were installed in the main power feed of the houses, hence covering the entire consumption from a single monitoring point. This platform is referred to as *single-house* in the sense that every house needs to have its own energy monitor.

3.2.1. Data acquisition and load monitoring

The current waveforms were sensed using standard non-invasive split-core (clamp-on) AC sensors. The voltage was measured with a custom-made voltage transformer that steps down the 230 V input voltage to 0.5 V such that it could be correctly sampled by the netbook sound-card. The two sensors were connected to the sound-card using a 3.5 mm TRS splitter.

The netbook and the sensors were installed in the main power feed, thus covering the entire house consumption and eliminating the need for additional sensing locations. The digitized waveforms were processed and transformed into power metrics representative of the energy consumption (e.g., apparent, real, and reactive power). The power metrics were calculated at a rate of 50 samples per second (i.e., the mains frequency in Portugal) and subsequently used for power event detection. All the power measurements were stored in a local database (aggregated at one sample per minute) along with the detected power events for feedback and future data analysis purposes.

It is important to remark that this solution is only viable in single-phase installations, where only two sensors are required to measure the load demand, i.e., one sensor for current and another for voltage.

3.2.2. Energy eco-feedback

The energy eco-feedback was provided on-site using the built-in display of the netbook through different custom-made applications that provided historical and real-time information on energy consumption and power events.

The first interface consisted mostly of traditional column charts to display the consumption information. The system shows a column chart with the total energy consumption over the current day and that of all the previous days. It is also possible to compare the electricity usage of the current week with the week earlier based on a daily average. Fig. 1 (left) provides an example of the daily consumption in a column chart, where each column represents the different hours of the day.

The second version was designed based on feedback received from the deployment of the first version. This interface used a gauge analogy to display consumption information to the user. The interface displayed information for the hourly, weekly, monthly, and yearly consumption organized in a tabbed menu. The consumption levels were mapped using a color scale going from green to dark red. A cursor hover on the gauge would trigger the display of information regarding CO_2 emissions and

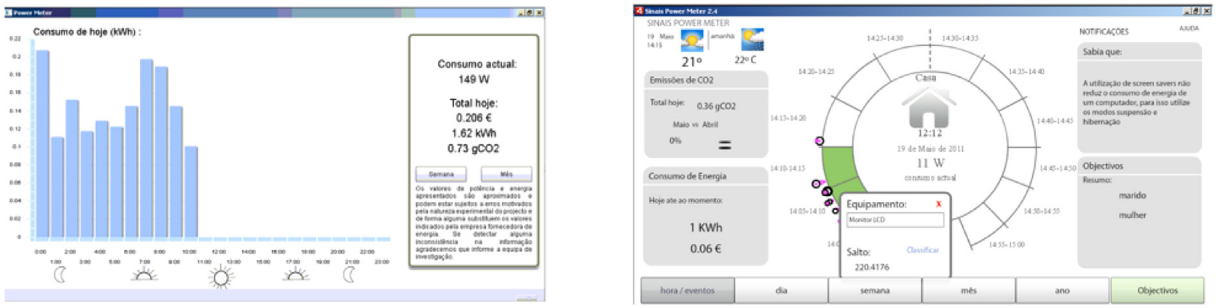

Fig. 1. Eco-feedback interfaces used in deployment one: version 1 (left), version 2 (right).

the cost associated with that time slot. Fig. 1 (right) provides a screenshot of the hourly consumption screen, where the dots represent power events (i.e., the instants when appliances change their operation status).

For additional information about the eco-feedback user interfaces used in the first deployment please refer to the following publications (Nunes et al., 2011; Pereira et al., 2012, 2013).

3.2.3. Installation and data integration

When installing the first version of the research platform, the current sensor had to be placed in the main fuse box. The voltage transformer and netbook had to be connected to a power source.

Regarding the data integration, each energy meter stored all the data locally using an SQLite[1] database. All the databases were synchronized using a Dropbox[2] folder linked to an account shared by all the houses in the deployment. The individual databases were later integrated into a single data warehouse using the SQL Server Integration Services[3] running in a machine also linked to the same Dropbox folder. Finally, Teamviwer[4] was installed in every computer to enable remote maintenance and updates.

Fig. 2 provides a general overview of the single-house platform installation and data integration process.

3.3. Multi-house energy monitoring and eco-feedback platform

The hardware and software platform evolved according to the limitations found on the first deployment. For example, in the initial setup, the sensing and eco-feedback were installed at the main breaker box, which raised some issues of limited accessibility for some household members (especially children). Because of these, the original monitoring platform has undergone some updates. The most significant one involved the replacement of the netbook sound-card with a more capable Data Acquisition System (DAQ).

In this version of the platform, the energy monitors were installed in the main lobby of the apartment buildings along with the electric utility meters. This solution enabled the monitoring of multiple homes from a single sensing location, thus reaching a new level of unobtrusiveness and security since no hardware had to be installed inside the participant houses. Furthermore, using a multi-channel DAQ enable the possibility of deploying the platform in two and three-phase electric systems. This platform is referred to as *multi-house* since multiple houses are monitored from a single energy monitor.

3.3.1. Data acquisition and load monitoring

The current and voltage signals for all the monitored houses are acquired from the building main electric panel and processed by a single computer using a dedicated DAQ board . The computer, referred to as the Energy Monitoring Base Station (EMBS), is also responsible for storing and providing remote access to consumption data. To this end, all the data is stored in a single MySQL[5] database and a layer of REST[6] web-services was implemented to enable easy access to the data.

Regarding the data acquisition hardware, the LabJack U6 DAQ[7] was used. This device can scan 14 analog input signals with a bit resolution up to 16 bit and a maximum sampling rate of 50 kHz (to be shared among all the active input channels). The LabJack DAQ connects to the EMBS via USB 2.0.

3.3.2. Energy eco-feedback

The multi-house energy monitoring and eco-feedback platform enabled householders to access the eco-feedback in different places of the house, or even outside the household premises provided there was an Internet connection available. As such, to take advantage of this feature, the eco-feedback was provided using custom-made mobile applications running on 7″ Android tablets (see Fig. 3 - left).

The eco-feedback application used in the second deployment involved two main modes of operation. If not used for two minutes the app would revert to the *Energy Awareness* mode, showing the consumption mapped as a digital illustration of the local endemic forest (see Fig. 3 — left). Once the user interacted with the tablet, by pressing the back softkey, the system went to the *Detailed Consumption* mode, showing daily, weekly and monthly information about the home energy consumption (see Fig. 3 — right).

As for the eco-feedback system used in the third deployment, the *Energy Awareness* mode was replaced with information about the electric energy generation in Madeira island. The developed application was composed of a set of tabs presenting the electric generation information, and summaries of the consumption on a daily, weekly, and monthly basis.

The energy generation view was the default mode of the app, and the system reverted to this visualization when no interaction happened during a pre-defined period. The electric energy generation was represented using a cumulative area chart of all the sources of energy used during the day, their quotas relative to each other. A forecast of the sources that would be available for the rest of the day was also available (Fig. 4 — left). The summary view (Fig. 4 — right) shown two charts representing the consumption of the current day, week, and month, as well as a comparison between homologous periods.

1 SQLite, https://www.sqlite.org.
2 Dropbox, https://www.dropbox.com.
3 SQL-IS, https://www.microsoft.com/en-us/sql-server/.
4 Teamviewer, https://www.teamviewer.com/.
5 MySQL, https://www.mysql.com/.
6 REST, http://www.ics.uci.edu/~fielding/pubs/dissertation/top.htm.
7 LabJack U6, https://labjack.com/products/u6.

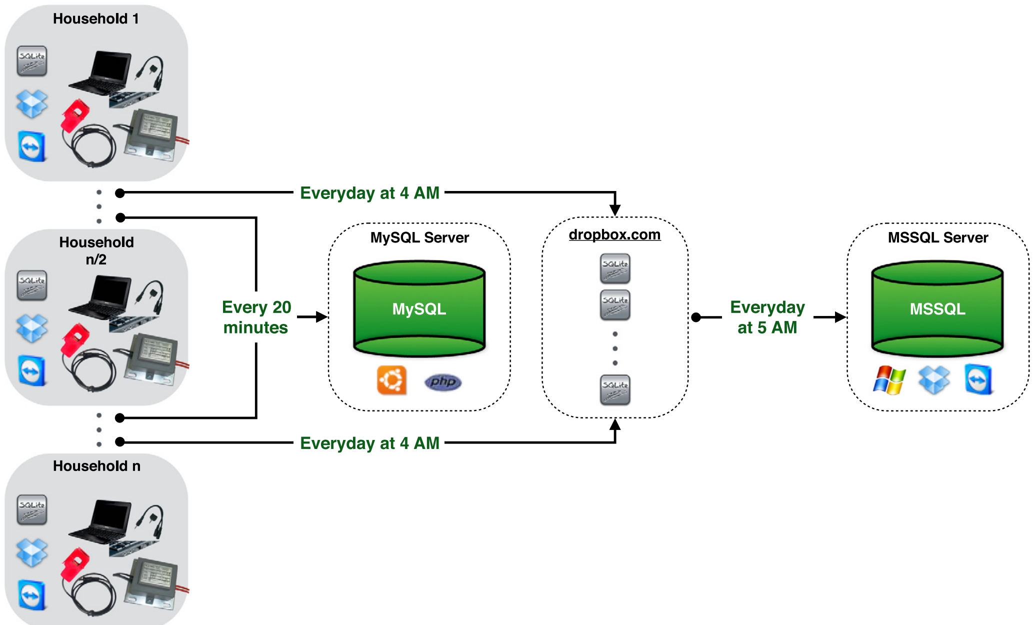

Fig. 2. General overview of the single-house version of the energy monitoring platform.

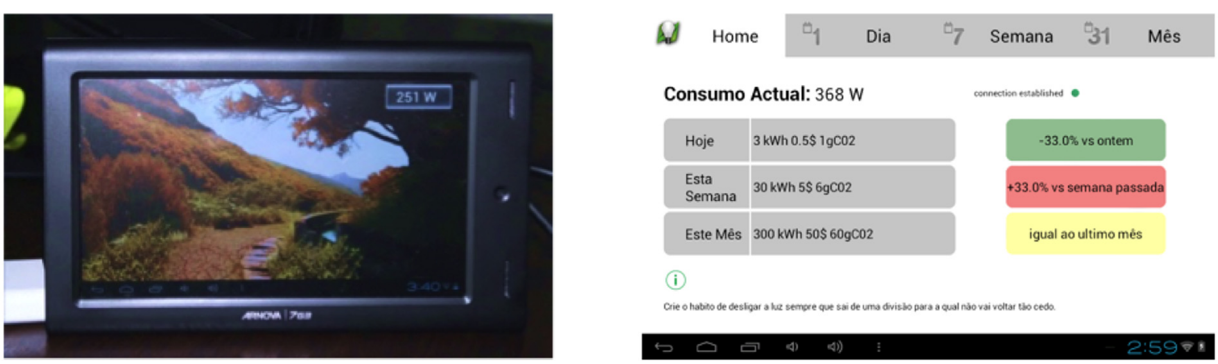

Fig. 3. Energy eco-feedback applications used in deployment two: energy awareness mode (left), detailed consumption mode (right).

Fig. 4. Energy eco-feedback applications used in deployment three: energy generation information (left), consumption summary (right).

3.3.3. Installation and data integration

The physical installation of the multi-house platform can be performed in different ways, depending on the number of houses to be monitored and the desired sampling rate. For example, a building with 11 apartments or less can be fully monitored using a single DAQ as depicted in Fig. 5. All the three voltage phases and the 11 current signals are sensed and digitized using the multi-port DAQ. The digitized waveforms are then fed to the energy monitoring software that computes the different power metrics by combining the corresponding voltage and current waveforms of each house.

Regarding the data integration task, the MySQL database in each EMBS was available for remote access. This way, it was possible to remotely access and integrate the individual databases whenever new data was required. Moreover, to further optimize the interactions with the databases, we maintained summary tables on the database server (e.g., hourly, daily, weekly, and monthly energy consumption averages).

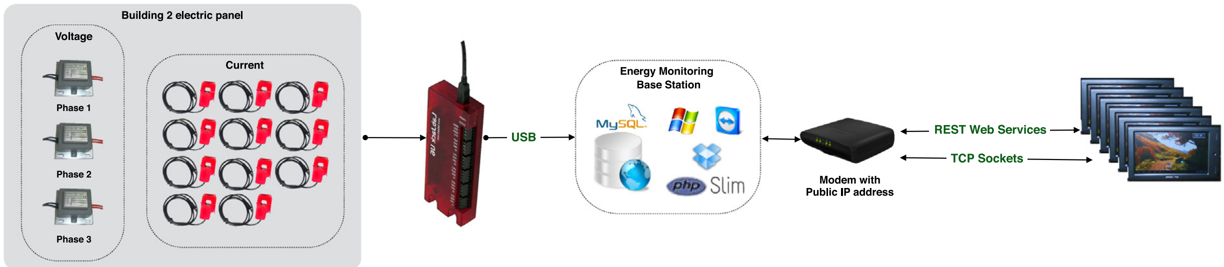

Fig. 5. Example of a possible configuration of the multi-house energy monitoring platform.

4. Practical deployment considerations

After describing the research platforms and the three live deployments, we now reflect on more than five years of experience in developing and deploying NILM systems. We include not only technical considerations but also the outcomes of several interviews that were conducted with the participants, to identify and clarify what we consider to be the practical issues behind deploying and maintaining NILM and eco-feedback systems in the wild.

First, we describe the different *technical challenges* faced by researchers, like physical installation constraints and data management issues. Secondly, we discuss the different *social challenges*, which involve for example maintaining a steady sample during the entire deployment. Finally, we discuss the *costs* associated with deploying energy monitoring systems for eco-feedback research purposes.

4.1. Technical considerations

Technological (or physical) issues refer to the different challenges that research teams are presented when developing and deploying this kind of systems. In our case, we have identified three main categories of technological issues, namely: (i) installation and maintenance; (ii) communication; and (iii) data management.

4.1.1. Installation and maintenance

Regarding the *installation* of the system, the main challenges are related to the location of the breaker boxes, particularly when deploying hardware inside the houses. For example, even though all the homes in the first deployment had the fuse box near the main door or in the kitchen, there were cases in which the fuse box was located inside a bedroom , thus making the whole process extremely intrusive. Likewise, it is also expected that in some older houses the breaker box will be in the basement or in the attic, which in any of the cases invalidates solutions with built-in eco-feedback.

Additionally, we should remark that most energy monitoring solutions require a constant power source. However, as we observed in our deployments, it was not very common to find power outlets near the breaker box. Consequently, when deploying single-house monitoring systems it is important to take into consideration that some extra work might be required to install all the necessary equipment.

With respect to the multi-house platform, the issues of accessing the breaker box were naturally avoided. However, installing the system in the building breaker box is by far a more challenging task and should always be conducted by experienced electricians. Nevertheless, the biggest challenge of deploying the second platform is also related to the actual physical location of the electrical panels and the circumstance that they are not prepared for the installation of this kind of systems. This fact was particularly evident in the third deployment where there was no space to store the necessary hardware.

Lastly, the *maintenance* of such long-term deployments was considerably challenging. In particular, the need to constantly monitor all the installations to ensure that everything was working smoothly. Moreover, it should be taken into consideration that constantly monitoring the status of the deployment will not necessarily mean that all the failures are detected and acted upon. As such, we argue in favor of following a *pro-active* maintenance strategy in which the monitoring solutions themselves are responsible for at least detecting and notifying the system administrators in case of failure.

4.1.2. Connectivity

In the wild research-platforms rely heavily on the availability of stable network connectivity, for several reasons including data transmission and system maintenance. This dependency was particularly evident in our second platform since everything was done remotely.

Yet, contrary to what one would expect, Internet connections and particularly Wi-FI are not widely available or easily accessible. This was the case in many of the homes monitored in the first deployment, which lead to the creation of a *had-hoc* Local Area Network to provide Internet to the participants. Likewise, Internet access was also a constraint in the second and third deployments, since we had to contract Internet connections from a local provider in order to connect our energy monitoring base stations to the Internet.

Consequently, when deploying this kind of systems, it is important to take into consideration that Internet connections may represent extra costs. This will become even more important if the deployments happen in remote places where the only available connections are mobile (e.g., 3G or 4G) since these are normally more expensive than traditional DSL or cable connections.

Furthermore, it is necessary to assume that the Internet will not be available all the time and therefore there is a need to account for long periods without an Internet connection. Such measures include, for example, storing the data locally when no connection is available and uploading to the server when an Internet collection is available without compromising the storage of real-time data.

4.1.3. Data management

One of the most important challenges in our deployments was to cope with the rate at which data was generated by the deployed energy monitors. Taking as an example the power readings that are stored at one sample per minute and considering the 50 households, there were 504 k records in the database after one week, and 2.160 M after 30 days.

Consequently, making the right choice of database technology is a crucial step when deploying such systems. There are

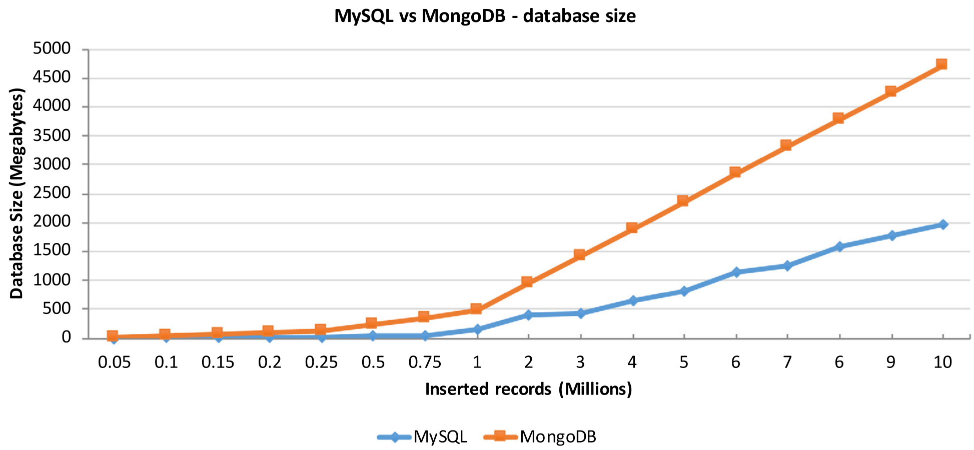

Fig. 6. MySQL vs. MongoDB: database physical size.

two aspects that we consider of great importance, namely the *query performance* and the *physical size* of the data. The former is expected to greatly affect the performance of any systems that rely on the stored data (e.g., eco-feedback applications), whereas the latter plays an important role regarding the selection of the hosting services.

Given the relevance of this issue in all the three deployments, we have decided to go beyond the theoretical guarantees of data storage technologies and performed a benchmark between SQL and NoSQL database management systems. Regarding the former, we selected MySQL, which is probably the most widely used database management system and normally the first option of most researchers, including us. As for the latter, we selected MongoDB[8] since it was one of the fastest growing NoSQL solutions at that time. This benchmark was performed using the SustData dataset that emerged from the three deployments (Pereira et al., 2014). For additional details about this benchmark, please refer to the following publications (Gonçalves, 2016; Pereira et al., 2018).

In one of the tests, we wanted to evaluate how much disk space would be necessary to store the same amount of data in both technologies. To this end, we performed the sequential insertion of 10 million records in MySQL and MongoDB and the results have shown that just after two million records (i.e., one month of power measurements for the 50 households) the size of the MongoDB database almost doubled the size of its MySQL counterpart (0.95 GB vs. 0.4 GB). Fig. 6 shows a graphical representation of the obtained results.

In another test, we wanted to measure how much time it would take to query the data in each engine. To this end, we selected the top five queries that were performed in our eco-feedback applications and executed each one of them using the same database sizes of the previous test. Fig. 7 shows a graphical representation of the average time it took to complete the five queries in each of the different database engines. As it can be observed, MongoDB clearly outperforms MySQL especially after there are half a million records in the database (7 days considering the same 50 households).

As we see from the results of the two tests, there is an obvious trade-off between the two database technologies. On the one hand, MySQL takes considerably less disk space, but just after 250 k records, the performance of the queries starts to degrade (4.71 s in MySQL against 1.78 s on MongoDB on average). On the contrary, MongoDB more than doubles the required disk space but manages to keep the query times in average 5 times faster than MySQL.

This said, at the end of the day, the most suitable database technology (or combination of technologies) is highly dependent on the type of application. For instance, in our case we are interested in providing the information to the user in the shortest period of time possible, thus the query time is much more relevant than the disk space.

4.2. Social considerations

By social issues, we refer to the different human–computer interaction related challenges that researchers face when conducting long-time research. More particularly, in our research we have identified two main categories of social issues, namely: (i) installing and maintaining a steady sample, and (ii) physical location and security of the deployed systems.

4.2.1. Installing and maintaining a steady sample

From our experience deploying these systems in real-world scenarios, we realize that one of the most interesting challenges was to deal with the very different agendas of everyone involved.

This fact became particularly evident during the installation phase of the first platform, which took 16 days to complete (and other 16 to remove) due to the difficulties in scheduling the visits to the houses. This was also evident when the different research teams had to interact with the householders to conduct the different eco-feedback research studies.

After exploring the evolution of the weekly user interactions with the eco-feedback (i.e., number of mouse clicks and screen touches) during the three deployments, we also found that with time participants tend to lose interest in the topic and stop using the devices.

The weekly interactions for each household during the deployment are shown in 8, where the week number represents the number of weeks since the system was installed in each household (i.e., week one means that the system was installed for an entire week, week two means that the system was installed for two weeks, and so on).

As it can be easily observed, in the first three to four weeks of each deployment there was a considerable number of interactions with the feedback devices. However, as the number of weeks increase, the number of interactions drops considerably. Likewise, during the first deployment, it is also possible to observe a very significant increase in the number of interactions between weeks 24 and 27, which corresponded with the update to the eco-feedback user interfaces that mentioned in Section 3.2.

To further understand this effect, we also looked at the average of weekly interactions per household, which is depicted in Fig. 9.

[8] MongoDB, https://www.mongodb.com/.

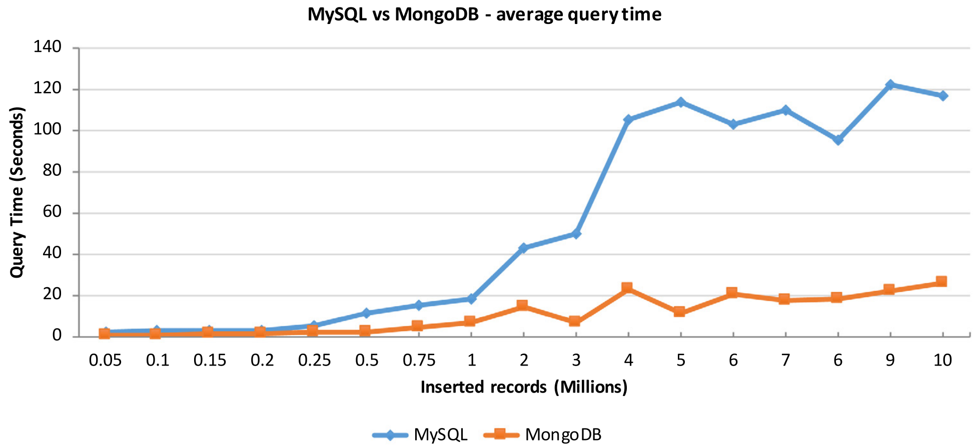

Fig. 7. MySQL vs. MongoDB: average query time.

As it can be seen, there are houses with very few interactions (e.g., house 11 with only one weekly interaction), and others that are way above average (e.g., houses 16, 18, and 19). After interviewing the participants, the research teams found that some of the participants did not have a particular interest in the research project, and only accepted to participate to be helpful. This was the case of house 11, that at some point ended up hiding the eco-feedback device behind a picture frame and house 4. On the other hand, in the families with a higher number of interactions, we found out that at least one of the members was highly motivated for this subject either for financial or environmental reasons.

While the reasons for this effect are out of the scope of this paper (interested readers should refer to Nunes et al., 2011; Pereira et al., 2013; Quintal et al., 2013b), this strongly suggests that field experiments are not just hard to implement from a technical stand-point, but also that any findings may be very difficult to validate and generalize, unless the sample size is large enough to account for caveats like the inability to conduct fully controlled trials (e.g., start and conclude an experiment with all participants at the same time), or the sample mortality (i.e., participants that opt to drop out, or that although remaining in the deployment do not relevant to the trials).

4.2.2. Physical location, security and intrusiveness

The fact that the first version of the monitoring platform had to be installed at the entrance of the monitored houses, presented some concerns. Firstly, the system was not easily accessible to all family members in particular children, as one of the mothers shared with us: *"She didn't reach it" (youngest daughter with 7 years old)*. Additionally, the location of the netbook near the main power feed made it harder for family members to interact with the eco-feedback as some users were afraid of either dropping it on the floor or damaging the equipment since they considered it to be very fragile (the computer was stuck to the wall with sticky Velcro) and they did not own the system (Pereira et al., 2013).

Likewise, some families also expressed concerns regarding the intrusiveness and safety of the system, even though it was properly and securely installed by qualified electricians. For instance, some families did not allow their kids to come nearby or interact with the devices, fearing the risk of electric shock.

Finally, with regards to the second platform, since all the measurements were taken from the main electrical panel of the building and the eco-feedback was provided using mobile applications, we did not observe any major concerns regarding the security and intrusiveness of the equipment.

4.3. Appropriation of the eco-feedback technology

From the extensive interviews, we have also learned that family members tend to have naturally defined roles. This included the role of checking and controlling the energy consumption.

Ultimately, this made other family members feel that they did not have to worry or use the eco-feedback systems, since someone (usually the husband or the person more comfortable around computers) was taking care of it, as shared by two spouses: *"There are certain things I leave for him to do and other things I take care of myself. I was curious to use it and I would use it but not as often as him"* and *"He would check more because he would be more curious (husband) and me I would let him give me the report of it. He would summarize the information"*.

4.3.1. Privacy and data protection concerns

Regarding privacy and data protection, our participants revealed little to no concerns, as long as their identities were not revealed. For example, when asked about possible additional features, most of the participants responded that they would like to know how their consumption compares with that of similar households, even if that meant having to share their data (Quintal et al., 2013a). Likewise, when inquired about the tracking of the mouse clicks and touches on the mobile applications, no concerns were raised whatsoever.

Finally, with regard to data ownership, most participants asked for a copy/summary of their data once the deployment was over, but did not raise any concerns about sharing the data with the scientific community. In order to attend to this request from our participants, the research team compiled an hourly consumption report (in kWh) for each participant. This was later made available in the form of a spreadsheet with pre-defined pivot tables. Additionally, once all the three deployments were completed, all the collected data was compiled and released to the research community for free (Pereira et al., 2014).[9]

4.4. Financial considerations

In order to provide an overview of how much studies of this nature can cost, we now report on the costs involved in our deployments. More concretely, we explore the costs associated with acquiring the monitoring hardware, and the energy required to run the energy monitors.

[9] SustData Dataset: http://aveiro.m-iti.org/data/sustdata.

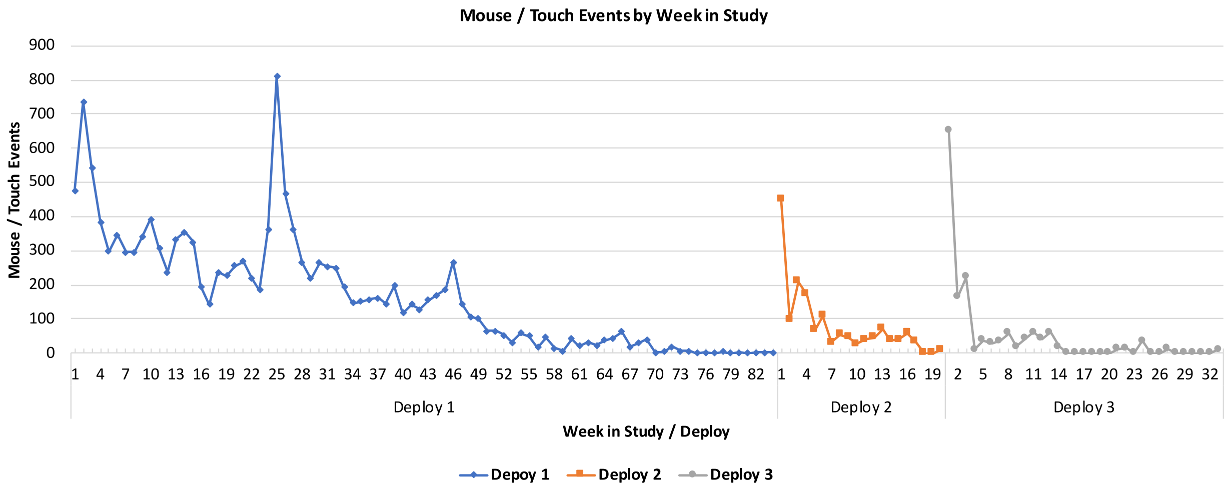

Fig. 8. The total number of user interactions with the eco-feedback per week into the deployment, for each of the three deployments.

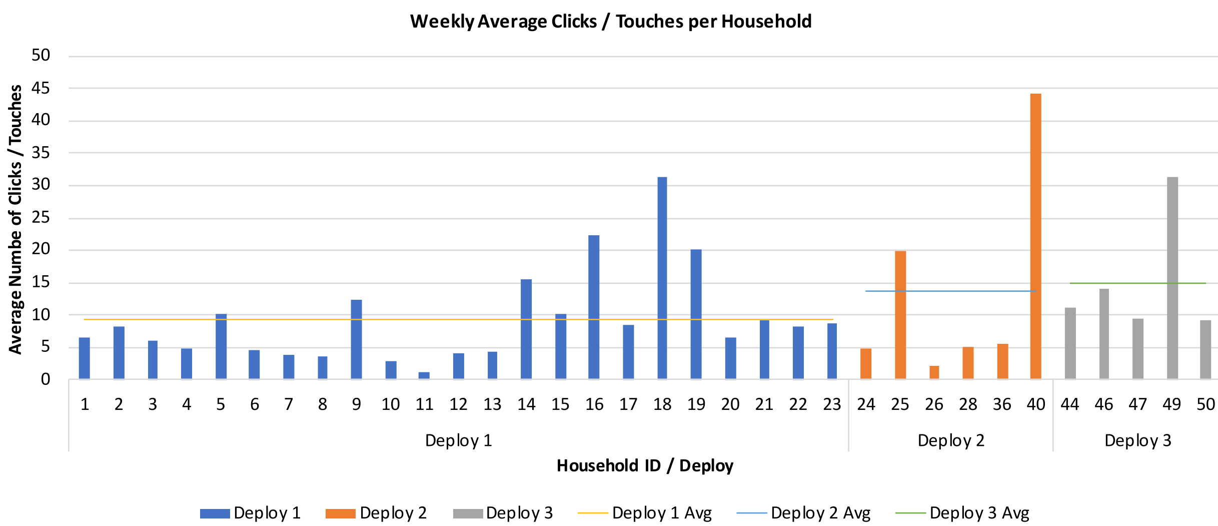

Fig. 9. Average user interactions per household for the whole duration of each deployment

Table 1
Baseline hardware costs of the single- and multi-house energy monitors (prices in Euro).

Item	Unit cost	Single-house		Multi-house	
		Qt.	Total	Qt.	Total
Netbook	230	1	230	1	230
Current sensor	10	1	10	10	100
Voltage sensor	15	1	15	3	45
Audio splitter	3	1	3	–	–
LabJack U6 DAQ	338	–	–	1	338
Tablet	99	–	–	10	999
			258		**1783**[a]

[a]Total per house: 170.3 (with tablet); 71.3 (without tablet).

4.4.1. Hardware acquisition

Regarding the hardware costs associated with the single- and multi-house energy monitors, the baseline costs were estimated based on the acquisition prices of the different components that comprise each solution. We also consider that the multi-house platform can monitor up to 10 houses.

The individual costs of each component are presented in Table 1, showing a comparison between the single-house and two different versions of the multi-house energy monitor (with and without tablet).

As alleged, monitoring multiple houses from one single location is significantly cheaper than installing hardware in every house (170.30€ vs. 258€) . Furthermore, it can also be observed that a substantial part of the costs is associated with the need to provide eco-feedback, hence indicating that the solutions will become much more cost-effective when eco-feedback is provided using channels that do not require additional hardware.

In order to better understand the costs associated with energy monitoring deployments, we also compare a multi-house solution with a single-house monitor based on a low-cost credit card-size embedded computer and one hypothetical multiple-sensor smart-meter.

With regards to the multi-house solution, we consider that the Raspberry Pi 3[10] embedded computer is used as a processing unit. As for the single-house solution, we consider that the system requires a dedicated processing unit and a dedicated DAQ board. More precisely, we assume that the single-house monitor is comprised of the BeagleBone Black rev. C[11] and the

10 Raspberry Pi, https://www.raspberrypi.org.

11 Beaglebone, https://beagleboard.org.

Table 2
Baseline hardware costs of the single- and Multi-house energy monitors (prices in Euro).

Item	Unit cost	Single-house		Multi-house		Multi-sensors	
		Qt.	Total	Qt.	Total	Qt.	Total
Raspberry Pi 3	40	–	–	1	40	–	–
Beaglebone Black	45	1	45	–	–	–	–
LabJack U6 DAQ	338	–	–	1	338	–	–
PRUDAQ	60	1	60	–	–	–	–
Whole house meter	75	–	–	–	–	1	75
Individual plug	31	–	–	–	–	10	310
Current sensor	10	1	10	10	100	–	–
Voltage sensor	15	1	15	3	45	–	–
			130		**523**[a]		**385**

[a]This is the price for 10 houses.

Table 3
Estimated energy costs of the components that compose the monitoring solutions.

Item	Watts	Day		Month		Year	
		kWh	Euro	kWh	Euro	kWh	Euro
Netbook	30	0.72	0.12	21.6	3.46	259.2	42.05
Embedded computer	5	0.12	0.02	3.6	0.58	43.2	6.91
Individual plug	1	0.002	0.004	0.72	1.18	8.64	1.40

PRUDAQ[12] high-speed ADC. Lastly, with respect to the multiple-sensor smart-meter, we consider that the system is able to monitor 10 different loads as well as the aggregate consumption by means of an additional whole house smart-meter. More precisely, we consider the CurrentCost[13] smart-meter, due to the fact that it is one of the cheapest solutions in the market. The baseline costs for the different solutions are presented in Table 2.

As it can be easily observed, the multi-house solution is much more cost-effective than the other solutions (e.g., it costs 60% less than the single-house option — 52.30€ vs. 130€). On the contrary, in a multiple-sensor solution, the information comes at much higher costs. For example, even if we consider only two individual plugs, the cost per house would still be higher than the single-house version (137€ vs. 130€) (see Fig. 10).

4.4.2. Electric energy consumption

In this work, we also look at the energy that is needed to run the energy monitoring devices, as this will impact the overall conclusions regarding the savings produced by the eco-feedback interventions. We considered the instantaneous power usage of each solution and projected the total consumption in kWh and EUR/kWh for different periods of time. To calculate the monetary cost we assume a baseline value of 16 cents per kWh, which is the current rate of the local provider in Portugal. The obtained results are presented in Table 3.

Then, given these estimates, we projected the costs in energy after one year of providing eco-feedback with a number of different energy monitoring scenarios. More concretely, we considered the following scenarios:

1. Single-house with a netbook (used in the first deployment),
2. Multi-house with a netbook (used in the second and third deployments),
3. Single-house with an embedded computer,
4. Multi-house with an embedded computer,
5. Multiple sensors, considering ten plugs, and
6. Ten individual plugs and a single-house meter.

12 PRUDAQ, https://github.com/google/prudaq.
13 CurrentCost, http://www.currentcost.com.

The obtained estimates for each case are presented in Fig. 11. As it can be observed, even though the notebook provides all the components needed to conduct eco-feedback research studies, the amount of energy that is consumed by that device is much higher than all the other solutions. This is particularly evident in the single-house solution where each monitored house represents a monthly energy cost of 3.5 Euros.

Another relevant observation is the fact that the ability to constantly monitor the energy consumption of 10 individual appliances will cost 1.7 Euros per month. The higher costs associated with multiple sensor solutions become even more evident when compared with those associated with the NILM solutions based on embedded computers. For example, the single-house solution has a monthly cost of 57 cents and the multiple house solution costs only about 5 cents a month.

5. Research implications and way forward

We now summarize the research implications of this work, discuss its limitations and highlight some possibilities of future work.

5.1. Research implications

We now draw some implications of this work to future research in this particular or similar fields that involve monitoring and proving eco-feedback.

On the technical side, we have found that connectivity plays a key role in such research studies and that unfortunately, it is not always readily available, which in the end can represent a considerable increase to the maintenance and update costs. Furthermore, we have also learned that on the contrary to what many early vendors of smart-meters claim, the installation is not necessarily straightforward, and will most likely require the work of professional electricians. Additionally, we have observed that conducting studies where variables are constantly monitored will inevitably result in an explosion of data and that as a consequence of this the choice of the right database technology is of crucial importance particularly in terms of the user experience (i.e., time necessary to query the data).

On the social side, we have found from our deployments that security and intrusiveness are a major concern of the participating families, thus letting us conclude that the best way to have participants engage in similar research studies is by providing a combination of transparent energy monitoring and ubiquitous eco-feedback. In other words, people would prefer not to see any monitoring hardware, and the eco-feedback should be provided using different communication channels. Furthermore, we have learned that conducting and validating real-world experiments can be very challenging due to caveats that result mostly from the busy agendas of the participants. With respect to privacy, our studies have reported on very little concerns from our users.

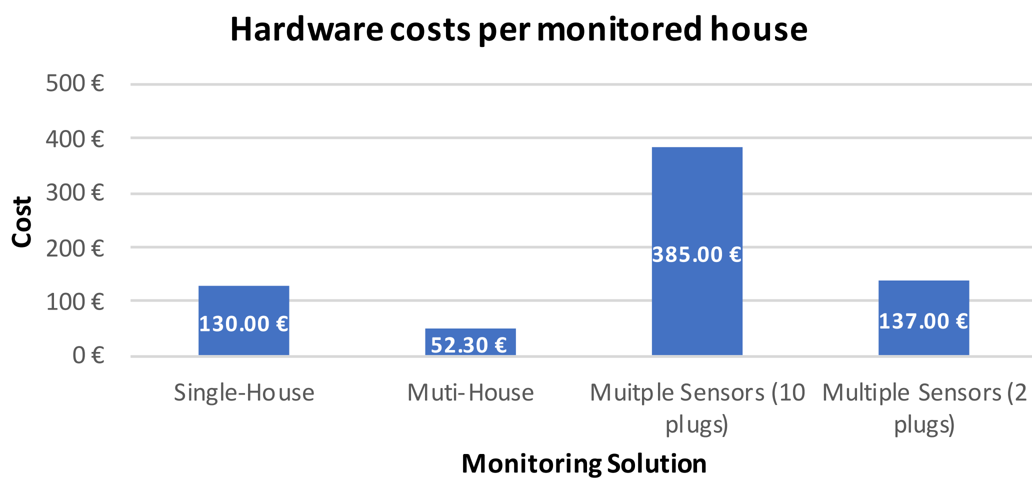

Fig. 10. Single- vs. Multi-House vs. Multiple Sensors: Hardware costs associated with monitoring one house.

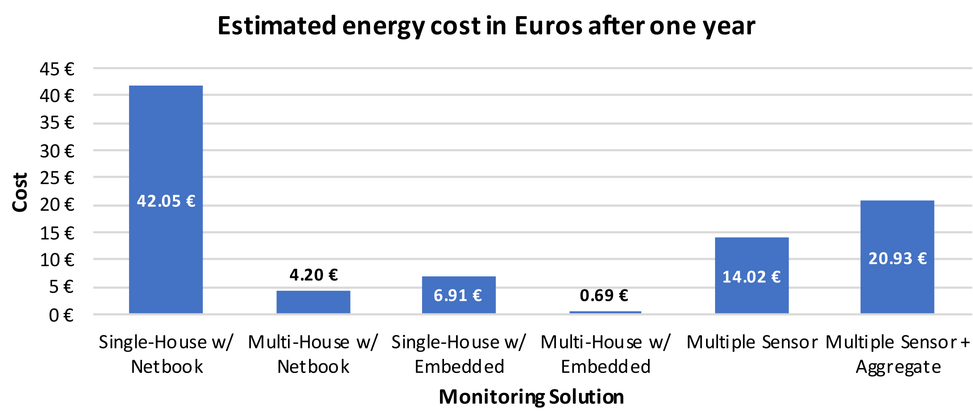

Fig. 11. Estimated energy costs of different energy monitoring solutions after one year.

Still, we should stress that this research was conducted before the recent scandals of abusive personal data use in the context of social networks. As such, privacy and personal data protection are two topics that should always be addressed when dealing with real-world deployments, independently of the technology and the main research objectives.

Lastly, on the financial side, we have learned that NILM solutions tend to be significantly less expensive than multiple-sensor technologies, not only in the initial acquisition costs but also in terms of the energy needed to run the systems. For example, we have seen that a multiple-sensor solution with only two individual plugs will still have a slightly higher acquisition cost and will consume 30% more energy than a single-house NILM system. Likewise, we have also learned that despite being limited to apartment buildings, the multi-house platform presents several advantages when compared to the single-house platform, in particular, the fact that more houses can be monitored from a single location and also the fact that it has the potential to be further expanded to monitor two and three-phase electric systems without extra hardware. Ultimately, considering all the costs associated with the different solutions, we believe that it is safe to say that in the long-term, even the more conservative NILM solutions (i.e., with a very limited number of disaggregated appliances) tend to have higher potential as tools to save energy than the multiple sensor solutions.

5.2. Limitations and way forward

Although we have reached the goals that were initially set in this work, there are some limitations that we would like to acknowledge and point out possible solutions.

First, the fact that the first deployment was conducted using a very unconventional smart-meter may have a direct influence on how the device was received by the householders. Furthermore, the second and third deployments were conducted without the need to install any hardware inside the households, which of course did not pose any issues related to the installation and security of the devices. As such, in future work, we should seek to further validate the social issues related to the physical location, security, and intrusiveness of conventional smart-meters.

A second limitation of this work is the fact that our deployments were only targeted a very specific segment of consumers living in a modern city, which may also have implications on how these technologies are received and perceived by the participants. Consequently, in future live deployments of NILM and eco-feedback technologies we should target different consumer segments as these may have different needs and perceptions regarding smart-meters and eco-feedback technology. These new segments include, for example, consumers from rural areas and most of all, consumers with micro-production installations, e.g., solar PV systems.

Finally, the third limitation of this work is directly related to the "in the wild" nature of our deployments, which prevented us from conducting more controlled experiments targeting specifically the NILM problem. Consequently, we believe that in order to conduct more accurate studies of NILM technology when deployed in the wild, future studies should be conducted in more controlled environments. For example, in order to correctly assess the value proposition of NILM as a tool that helps users save energy, one needs to conduct A/B testing studies (Kohavi, 2015). In such studies, two randomized groups of households are provided with energy eco-feedback through the same communication channel, but only one of the groups has access to disaggregated consumption information.

6. Conclusion

In this paper, we have reported on the different technical, social and financial challenges behind the building and deploying energy monitoring and eco-feedback systems in real-world scenarios.

To conclude, we describe how the insights presented in this paper were put into practice in another real-world deployment of energy monitoring and eco-feedback that is being conducted in the scope of the Smart Island Energy Systems project (SMILE[14]), founded by the European Union's Horizon 2020 research and innovation programme (ACIF-CCIM, Prsma, EEM, M-ITI, Route Monkey, 2017; Prsma, M-ITI, EEM, ACIF-CCIM, 2018; Hashmi et al., 2019).

Energy monitoring hardware: In this deployment we are using conventional smart-meters that can be installed directly in the breaker-box, thus making the system much more secure. For computing we are using the Raspberry pi 3, that is powered using a power-socket on the main breaker box. Using the raspberry pi we reduce the energy consumption of the monitoring system, and also achieve a loosely coupled monitoring solution.

Communications: For communications we are currently using the built-in Wi-Fi and Ethernet, or 3/4G through a USB dongle. Nevertheless, this solution can be easily upgraded to use communication technologies like LoRa[15] and NB-IoT.[16]

Remote access: We are using the DWService,[17] a free and open-source software, that allows access to remote systems using a standard web browser. Besides being fully free and open-source, DWService is also a light-weight alternative to TeamViewer.

Proactive maintenance: Our monitoring hardware is configured to upload data to a cloud-based server every minute. If no data is uploaded after a pre-defined time period or if errors occur while communicating with the smart-meter, the system automatically logs the errors, and notifies the system administrator via e-mail. Likewise, on the server-side, if a client does not send its data past a pre-defined time period, the server initiates a communication with the client and reports the results to the system administrator via e-mail.

Database technology: We are using MongoDB to store all the data generated by the energy monitors. Also, to further improve the query times we resorted to caching mechanisms using Redis.[18]

Sample selection: The sample selection was not restricted to the city of Funchal. Instead, participants were recruited from all over Madeira island, and the only restriction is that they own a solar PV installation.

Eco-feedback: The feedback is provided online using a fully responsible web-application. This allows end-users to access information from virtually anywhere, as long as a browser and an Internet connection is available.

14 SMILE H2020, https://www.h2020smile.eu/.

15 LoRa Wireless Technology, https://www.semtech.com/technology/lora.

16 Narrow Band IoT, http://www.huawei.com/minisite/iot/img/nb_iot_whitepaper_en.pdf.

17 DWService, https://www.dwservice.net/en/home.html.

18 Redis, https://redis.io/.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The authors would like to acknowledge the members of the SINAIS project that were involved in the work-packages that lead to this paper. We would also like to acknowledge all the participating families, as well as *Empresa de Electricidade da Madeira*, and *NOS Comunicações S.A.* for their technical support. Finally, the authors would like to thank professor Mario Bergés, for his insights during the preparation of this research paper.

This research was funded by the CMU-Portugal SINAIS project (CMU-PT/HuMach/0004/2008), the Portuguese Foundation for Science and Technology (FCT) individual doctoral grant (SFRH/DB/77856/2011), and FCT research grant (UID/EEA/50009/2013).

Appendix A. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.egyr.2019.11.025.

References

ACIF-CCIM, Prsma, EEM, M-ITI, Route Monkey, 2017. Madeira Pilot Case Study Specification and Assessment. Technical report 4.1, EUROPEAN COMMISSION, Funchal, Portugal.

Barreto, M., 2014. Towards the design of eco-feedback tecnologies for families (PhD). Universidade da Madeira.

Batra, N., Singh, A., Whitehouse, K., 2016. Exploring the value of energy disaggregation through actionable feedback. In: 3rd International NILM Workshop.

Carrie Armel, K., Gupta, A., Shrimali, G., Albert, A., 2013. Is disaggregation the holy grail of energy efficiency? The case of electricity. Energy Policy 52, 213–234.

Darby, S., 2006. The Effectiveness of Feedback on Energy Consumption: a Review for DEFRA of the Literature on Metering, Billing and Direct Displays. Technical Report, Oxford, UK.

Ehrhardt-Martinez, K., Donnelly, K.A., "Skip" Laitner, J.A., 2010. Advanced Metering Initiatives and Residential Feedback Programs: A Meta-Review for Household Electricity-Saving Opportunities. Research Report E105, American Council for an Energy-Efficient Economy.

Esa, N.F., Abdullah, M.P., Hassan, M.Y., 2016. A review disaggregation method in Non-intrusive Appliance Load Monitoring. Renew. Sustain. Energy Rev. 66, 163–173.

Fischer, C., 2008. Feedback on household electricity consumption: a tool for saving energy?. Energy Effic. 1 (1), 79–104.

Froehlich, J., Findlater, L., Landay, J., 2010. The design of eco-feedback technology. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. In: CHI '10, ACM, New York, NY, USA, pp. 1999–2008.

Gonçalves, R., 2016. OpenDataHub: An Open Dataset Management System (Master). University of Madeira, Funchal.

Hansen, T.R., Bardram, J.E., Soegaard, M., 2006. Moving out of the lab: Deploying pervasive technologies in a hospital. IEEE Pervasive Comput. 5 (3), 24–31.

Hart, G., 1985. Prototype Nonintrusive Appliance Load Monitor. Technical Report, MIT Energy Laboratory Technical Report, and Electric Power Research Institute Technical Report.

Hashmi, M.U., Pereira, L., Bušić, A., 2019. Energy storage in madeira, Portugal: co-optimizing for arbitrage, self-sufficiency, peak shaving and energy backup. In: 2019 IEEE Milan PowerTech. pp. 1–6.

Huang, E.M., Mynatt, E.D., Trimble, J.P., 2007. When design just isn't enough: the unanticipated challenges of the real world for large collaborative displays. Pers. Ubiquitous Comput. 11 (7), 537–547.

International Energy Agency, 2014. 2013 Key World Energy Statistics. Technical Report, International Energy Agency, p. 82.

Jain, R.K., Taylor, J.E., Peschiera, G., 2012. Assessing eco-feedback interface usage and design to drive energy efficiency in buildings. Energy Build. 48, 8–17.

Karlin, B., Zinger, J.F., Ford, R., 2015. The effects of feedback on energy conservation: A meta-analysis. Psychol. Bull. 141 (6), 1205–1227.

Kelly, J., Knottenbelt, W., 2016. Does disaggregated electricity feedback reduce domestic electricity consumption? A systematic review of the literature. In: 3rd International NILM Workshop.

Kohavi, R., 2015. Online controlled experiments: Lessons from running a/b/n tests for 12 years. In: Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. In: KDD '15, ACM, New York, NY, USA, p. 1.

Konomi, S., Roussos, G., 2007. Ubiquitous computing in the real world: lessons learnt from large scale RFID deployments. Pers. Ubiquitous Comput. 11 (7), 507–521.

Kosonen, H., Kim, A., 2016. Quantifying plug load energy use in a LEED gold building—Lessons learned in the installation phase. In: Construction Research Congress 2016. American Society of Civil Engineers, pp. 1234–1243.

Langendoen, K., Baggio, A., Visser, O., Murphy loves potatoes: experiences from a pilot sensor network deployment in precision agriculture. In: Proceedings 20th IEEE International Parallel Distributed Processing Symposium, p. 8.

Mayhorn, E.T., Sullivan, G.P., Petersen, J.M., Butner, R.S., Johnson, E.M., 2016. Load Disaggregation Technologies: Real World and Laboratory Performance. Technical Report PNNL-SA-116560, Pacific Northwest National Laboratory (PNNL), Richland, WA (US).

Nalmpantis, C., Vrakas, D., 2018. Machine learning approaches for non-intrusive load monitoring: from qualitative to quantitative comparation. Artif. Intell. Rev. 1–27.

Nunes, N.J., Pereira, L., Quintal, F., Berges, M., Deploying and evaluating the effectiveness of energy eco-feedback through a low-cost NILM solution. In: Proceedings of the 6th International Conference on Persuasive Technology, Columbus, OH, USA.

Ojala, T., Kostakos, V., Kukka, H., Heikkinen, T., Linden, T., Jurmu, M., Hosio, S., Kruger, F., Zanni, D., 2012. Multipurpose interactive public displays in the wild: Three years later. Computer 45 (5), 42–49.

Pacala, Socolow, R., 2004. Stabilization wedges: solving the climate problem for the next 50 years with current technologies.. Science 305 (5686), 968–972.

Parker, D., Hoak, D., Meier, A., Brown, R., 2006. How much energy are we using? Potential of residential energy demand feedback services. In: 2006 Summer Study on Energy Efficiency in Buildings, Asilomar, CA - USA.

Pereira, L., Gonçalves, R., Quintal, F., Nunes, N., 2018. Data storage and maintenance challenges: The case of advanced metering infrastructure systems. In: Proceedings of the 5th International Conference on ICT for Sustainability. In: EPiC Series in Computing, EasyChair, Toronto, ON, Canada.

Pereira, L., Quintal, F., Barreto, M., Nunes, N.J., 2013. Understanding the limitations of eco-feedback: A one-year long-term study. In: Holzinger, A., Pasi, G. (Eds.), Human-Computer Interaction and Knowledge Discovery in Complex, Unstructured, Big Data. In: Lecture Notes in Computer Science, Springer Berlin Heidelberg, Maribor, Slovenia.

Pereira, L., Quintal, F., Gonçalves, R., Nunes, N.J., 2014. Sustdata: A public dataset for ICT4s electric energy research. In: Proceedings of ICT for Sustainability 2014. Atlantis Press, Stockholm, Sweeden.

Pereira, L., Quintal, F., Nunes, N., Bergés, M., 2012. The design of a hardware-software platform for long-term energy eco-feedback research. In: ACM SIGCHI Symposium on Engineering Interactive Computing Systems. In: EICS '12, ACM, Copenhagen, Denmark, pp. 221–230.

Peschiera, G., Taylor, J.E., Siegel, J.A., 2010. Response–relapse patterns of building occupant electricity consumption following exposure to personal, contextualized and occupant peer network utilization data. Energy Build. 42 (8), 1329–1336.

Prsma, M-ITI, EEM, ACIF-CCIM, 2018. Madeira Pilot Infrastructure Preparation and Kick-off. EUROPEAN COMMISSION, Funchal, Portugal.

Quintal, F., 2016. Exploring the Dimensions of Eco-Feedback in the Wild (PhD). Universidade da Madeira.

Quintal, F., Barreto, M., Nunes, N., Nisi, V., Pereira, L., 2013a. Wattsburning on my mailbox: A tangible art inspired eco-feedback visualization for sharing energy consumption. In: Human-Computer Interaction – INTERACT 2013. In: Lecture Notes in Computer Science, Springer, Berlin, Heidelberg, pp. 133–140.

Quintal, F., Pereira, L., Nunes, N., Nisi, V., Barreto, M., 2013b. Wattsburning: Design and evaluation of an innovative eco-feedback system. In: Human-Computer Interaction – INTERACT 2013. In: Lecture Notes in Computer Science, Springer, Berlin, Heidelberg, pp. 453–470.

Tateson, J., Roadknight, C., Gonzalez, A., Khan, T., Fitz, S., Henning, I., Boyd, N., Vincent, C.J., Marshall, I.W., 2005. Real world issues in deploying a wireless sensor network for oceanography. In: Refereed REALWSN 2005 Proceedings, Stockholm, Sweden.

US Energy Information Administration, 2016. International Energy Outlook 2016. Technical Report DOE/EIA-0484(2016), US Energy Information Administration, Washington DC, DC, USA.

Viana, A.R., 2011. Designing Services for Sustainability: The case of home energy monitoring (MSc). University of Porto, Porto, Portugal.

Vine, D., Buys, L., Morris, P., 2013. The effectiveness of energy feedback for conservation and peak demand : a literature review. Open J. Energy Effic. 2 (1), 7–15.

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

Zoha, A., Gluhak, A., Imran, M.A., Rajasegarar, S., 2012. Non-intrusive load monitoring approaches for disaggregated energy sensing: A survey. Sensors 12 (12), 16838–16866.