A Versatile High Frequency Electricity Monitoring Framework for Our Future Connected Home

Benjamin Völker(✉), Marc Pfeifer, Philipp M. Scholl, and Bernd Becker

Chair for Computer Architecture, University of Freiburg, Freiburg im Breisgau, Germany
{voelkerb,pfeiferm,pscholl,becker}@informatik.uni-freiburg.de

Abstract. In our homes a lot of devices are powered by electricity without us knowing the specific amount. As electricity production has a large, negative environmental impact, we should be more aware about how devices consume power and how we can adapt our daily routine to decrease our electricity requirements. Methods such as *Non-Intrusive Load Monitoring* (NILM) can provide the user with precise device level electricity data by measuring at a single point in a houses' electricity network. However, the time resolution of most off-the-shelf power meters is not sufficient for NILM or the meters are locked down for security reasons. Therefore, we have developed our own versatile energy metering framework which consists of a high frequency electricity metering device, a versatile backend for data processing and a webapp for data visualization. The developed hardware is capable of sampling up to 32 kHz, while the software framework allows to extract other power related metrics such as harmonic content. The system's application ranges from providing transparent electricity usage to the user up to generating load forecasts with fine granularity.

Keywords: Load monitoring · NILM · NIALM · Electricity feedback

1 Introduction

Electricity is omnipresent without us really thinking about it. We use it while we brew our morning coffee, while on our way to work, while at work and while spending free time with our family. In 2014, the yearly global average electricity consumption per capita was 3353 kWh [5] worldwide which is topped by an average person in the US at 12 305 kWh. This is equivalent to approximately 6.49 t of CO_2 per capita per year in the US according to [6]. Besides CO_2, electricity production is the largest contributor to global greenhouse gas emission with approximately 29% [1]. As a way to reduce our electricity consumption without giving up our morning coffee D [8] stated that an "appliance specific breakdown" should be provided instead of just the monthly whole house energy

J. L. Afonso et al. (Eds.): SESC 2019, LNICST 315, pp. 221–231, 2020.
https://doi.org/10.1007/978-3-030-45694-8_17

consumption. This would make end users more aware about their electricity consumption. If a device level electricity breakdown is provided, energy hungry devices e.g. the 30 year old freezer can be identified easily. Further, gamification aspects could be applied to our daily electricity consumption. As an example, a contextual positive energy feedback - possibly a simple reward - could be provided if the TV has been used less often during the last weak. Such feedback can be applied by integrated device level electricity data with standard human activity recognition systems.

However, as long as devices do not expose their electricity consumption, retrofittable solutions need to be applied. Device level metering can be added using off-the-shelf power measurement units like *Kill-A-Watt* [9]. This approach is known as intrusive load monitoring (ILM). It requires the user to attach measurement units to all or a subset of appliances across the residence. This is a laborious process and for devices directly connected to the mains like the stove impractical for non-electricians. Furthermore, if a large set of these measurement units is deployed, their total self-consumption might not be negligible. An alternative approach to ILM is Non-intrusive Load Monitoring (NILM). This approach uses a single measurement unit installed at the houses fuse box which measures the composite load of all appliances in the home. Disaggregating this composite load into the load of each individual appliance is challenging. Hence, disaggregation algorithms typically require to analyze different macroscopic as well as microscopic features in both time and frequency domain. Therefore, data of high sampling-rates is required. For example the algorithms used in [4] and [16] use power harmonics up to the 20th which requires a sampling-rate of at least 2 kHz according to the *Nyquist Theorem*.

Such data can not be provided by existing infrastructure in our homes, since typical smart meters - which may internally sample with such high frequencies - only provide the electricity data with a temporal resolution of <1 Hz. Furthermore, due to security concerns, interfacing with these smart meters is often restricted to the grid operator.

In this paper we discuss the requirements for a versatile, retrofittable electricity monitoring system. We further propose a hardware and software framework that meets these requirements. The resulting embedded system is comprised of a smart meter like device installed in the residences' fuse box and a software backend which gathers the sampled raw voltage and current measurements and distributes electricity information via network or unix domain sockets to arbitrary clients. An example client is a disaggregation module splitting the total electricity consumption into the individual devices consuming power. We have installed the system in the office kitchen of our chair to analyze its long term performance. It provides real time electricity feedback and is used to record datasets which can be used to evaluate load disaggregation algorithms.

2 Related Work

The authors of the The Reference Energy Disaggregation Data Set (REDD) [11] used a custom build recording interface to recorded the whole house electricity consumption of six different homes for several weeks. To measure the aggregated current and voltage signals at the mains, they used *NI-9239* analog to digital converters. To scale down the mains voltage onto the measurement range of the ADC, they used an 1:100 oscilloscope probe (*Pico TA041*). For current sensing, they used split core current transformers (*SCT-013*). The sampling rate was 16 kHz with an ADC resolution of 24 bit. A laptop was connected to log the raw data and to send low frequency versions to an external server.

The authors of the UK Domestic Appliance-Level Electricity dataset (UK-DALE) [10] recorded 655 days of data with a sampling rate of 16 kHz of a single home in the UK. They used an off-the-shelf USB sound card with a stereo line in interface connected to a PC. The voltage was transformed to line level using AC-AC transformer and a voltage divider while current is transformed using split core current transformer. However, only a single phase was recorded. Recording the whole house supply would require to use three sound cards (in Europe). The authors of [7] presented a hardware and software system to provide the device level electricity consumption for a building. Their hardware system is comprised of a dedicated power monitoring chip (ADE7880), a microcontroller to interface with the chip and an ethernet interface. Unfortunately, they do not provide any information about achievable sampling rates or the cost of the proposed system. The software consist of an applet that is able to identify on- and off-events of devices which have an active power consumption that is "significantly different of each other".

Kriechbaumer et al. [13] have developed a measurement system embedded into an off-the-shelf power strip. They used six hall effect sensors (*ACS712*) to measure the current consumption at each plug and one AC-AC converter to measure the overall plug's voltage. The analog output of these sensors is converted into the digital domain using seven 12 bit unipolar ADCs (*MCP3201*). A combination of a microcontroller and a single board PC handles the digital data acquisition. Therewith, they can measure up to six different devices with an adjustable sampling frequency of up to 50 kHz. Their setup allows a "mobile" data collection of all devices with power plugs. But they cannot measure devices directly connected to the mains like the lighting or the stove. More recently, the same authors developed a measurement systems which can be installed at a residences fuse box to measure the aggregated power consumption. They used hall effect current transformers (*HAL 50-S*) for current sensing and 6 V AC-AC transformers for voltage sensing. Using a 16 bit ADC (*AD7656A*) and an FPGA, they are capable of sampling raw current and voltage waveforms up to 250 kHz. However, they do not reveal any information about the costs of their system. Since the used components (Latte Panda, 3x Hall Effect Sensors, Lattice FPGA, ADC, etc.) already cost 400 €, a relatively large system price (≈500 €) is assumed.

Kriechbaumer et al. have proposed a set of requirements for an electricity data acquisition hardware (DAQ) which is summarized in Table 1. We have added the requirement "**R11**: safe external interfaces" since we are dealing with high voltages, we have to make sure that interfacing with the system is safe.

Table 1. Requirements for electricity data acquisition hardware, adapted from [13]

R1 high sampling rates	**R6** synchronized world clock
R2 long term recordings	**R7** precise time-stamping
R3 common file format	**R8** high resolution and accuracy
R4 data compression	**R9** persistent data storage
R5 low price (per appliance)	**R10** resources for data processing
	R11 safe external interfaces

Our contribution is three fold:

1. We provide a versatile framework comprised of both hardware and software to provide high frequency electricity data in near real-time. The time delay between data recording and provision is determined by the transmission channel alone. This transmission can be performed over a wired or wireless channel. The data can be further distributed to other client modules using unix domain (IPC) or network socket connections.
2. The provided raw measurements are of high horizontal (up to 32 kHz) and vertical (24 bit) resolution.
3. Instead of being restricting to a specific task, the framework can adapt to the users need and either serve as just a recording interface for long term electricity datasets (see **R2**), provide real-time electricity data to further data analytic modules or directly display historic and real-time consumption to the end user.

3 Hardware Setup

The measurement system is a self build prototype board encapsulated in a DIN housing to allow a rail mount inside a houses' fuse box. The voltages of all three main legs are measured using voltage dividers with a ratio of ≈1:1000. The currents are measured using split core burden-less current transformers (*YHDC SCT-013*) with a ratio of 1:2000. Such electromagnetic current transformers are inexpensive solutions to measure high currents without changing the electrical wiring. Compared to other contact-less current sensing techniques, they also show a high linearity at a high measurement range. The used split core variants are also fairly easy to install. The induced smaller current is transformed into a voltage using two 7.87 Ω burden resistors per channel. The analog signals are fed into a dedicated electricity monitoring chip (*ADE9000* [2]). The *ADE9000* has a seven input ADC with a sampling rate of up to 32 kHz (see **R1**), a resolution of 24 bit, and a very high signal to noise ratio of 96 dB which meets requirement **R8**. This allows to measure up to 3 voltage and 4 current channels covering the three phase input of typical homes in Europe (L_1, L_2, L_3 and N). As noted, the *ADE9000* is dedicated for electricity monitoring and, therefore, has additional internal hardware to calculate electricity related metrics such as active, reactive and apparent power, phase shift etc. These measures and the raw voltage and current samples can be retrieved via a galvanically isolated SPI interface (see **R11**) as shown in Fig. 1. The main controller inside the measurement system is a WiFi and Bluetooth enabled microcontroller (*ESP32*). Since this microcontroller is very powerful with its 240 MHz dual core processor, it can simply relay the sampled data e.g. over a TCP connection or execute/incorporate several algorithms for energy analysis (see **R10**).

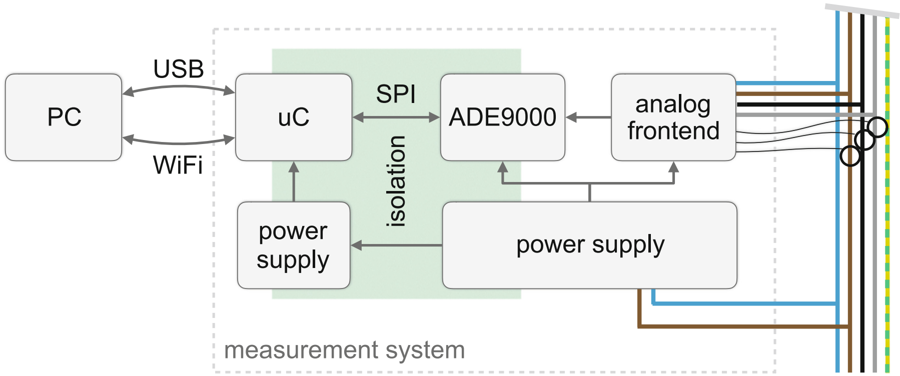

Fig. 1. Overview of the measurement system.

In addition it can also serve as a lightweight web-server in order to provide a simple user feedback without the requirement of any external server. If we take a closer look at the systems' typical environment such as a shielded fuse box located in the basement, we expect WiFi connection to be poor. Therefore, we added the ability to connect an external antenna over a *SMA* connector. A *USB-serial* interface (*FTDI-FT2232H*) allows to change the systems firmware or to gather the electricity data over a reliable wired channel instead of WiFi.

Measuring the voltage directly using voltage divider has a very high accuracy (see **R8**) compared to indirect methods but leads to safety requirements. We focused on meeting these requirements not only to meet **R11** but to meet standards according to IEC 61010-1, IEC 60950, DIN EN 60664-1 (clearance distances) and IEC 61000-4-5.

The price of the measurement system is roughly 150 € depending on the production volume, making it an affordable solution to retrofit smart metering to an existing home (see **R5**).

4 Software Backend

The software backend is split into multiple individual modules. These modules communicate over TCP sockets. This allows to distribute or centralize all calculations.

The main module that interfaces directly with the measurement system is the *Data Collector* module (see Fig. 2). It is connected to it over a TCP or USB serial connection. The measurement system can be configured using the data collector to only provide raw voltage and current signals, active ($P(n)$), reactive ($Q(n)$) and apparent ($S(n)$) power or the signals harmonics. However, to simplify the data collection process and to restrict data throughput, the standard configuration of the measurement system is to provide raw voltage and current waveforms, since other measures can be directly calculated from it but not vice versa. The calculation of active, reactive and apparent power is therefore performed at the data collector module on the basis of the mains frequency f_l (e.g. 50 Hz in Asia and Europe and 60 Hz in North America). The according formulas are shown in (1), (2) and (3). This stream is further referred to as the lower frequency stream.

$$P(n) = \frac{1}{N} \cdot \sum_{i=0}^{N-1} U(i) \cdot I(i) \tag{1}$$

$$S(n) = \frac{1}{N} \cdot \left(\sum_{i=0}^{N-1} U(i) \cdot \sum_{i=0}^{N-1} I(i) \right) \tag{2}$$

$$Q(n) = \sqrt{S(n)^2 - P(n)^2} \tag{3}$$

P, Q and S are calculated for each non-overlapping window n. The length of the window is $N = \frac{f_s}{f_l}$ with f_s being the sampling frequency of the raw voltage and current signals. The sampling frequency f_s can be set based on the system's requirement up to 32 kHz. To keep the amount of data reasonable while still allowing to retrieve information about higher frequency harmonics, a sampling rate of 8 kHz is typically chosen. This value is also based on the findings of [3] that, "there may be little additional benefit between 15–40 kHz because of the noise in that range in real buildings".

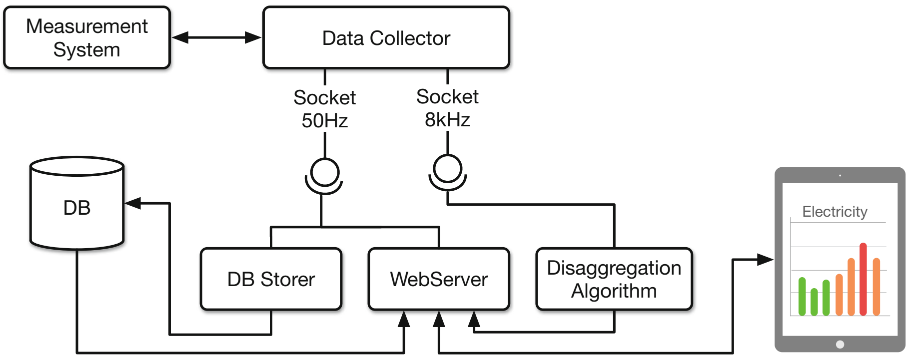

Fig. 2. Overview of the proposed system.

Other modules can subscribe for specific data streams i.e. the calculated power signals or the raw voltage and current signals.

One client type is the *DB Storer* module. It subscribes to the stream which outputs the calculate active, reactive and apparent power (at frequency f_L) and stores it into a *NoSQL* database. Since NoSQL does not require a fixed predefined data layout, it allows to change storage schemes later on.

Another type of client is the *Disaggregation* module. Its main purpose is to disaggregate the composite load into the load of each individual appliance connected. As a typical disaggregation algorithm also analyzes higher harmonic features such as frequency components in the signal up to 4 kHz (e.g. [16]), it requires input data of higher frequency. Therefore, the module 'subscribes' to the raw voltage and current stream. The disaggregation result can further be subscribed by other modules such as a web-server. The web-server is a another client module. It connects to the lower frequency stream for a real-time electricity feedback and also to the database for historic electricity data. The data is transported via web-sockets to the client. The client can be an arbitrary browser running on a PC or mobile device.

5 Software Frontend

The frontend is an interactive web application. Therewith, the user has the flexibility to use any tablet, phone or PC to interface with the system. An overview of the web application is shown in Fig. 3. A dashboard like modular layout allows to display just the information the user wants to be shown. We have seen that mounting a tablet displaying the current real-time electricity consumption e.g. at a house's kitchen makes the user aware of sudden changes of the power consumption. The user can than relate these changes directly to their actions (e.g. turning on the stove).

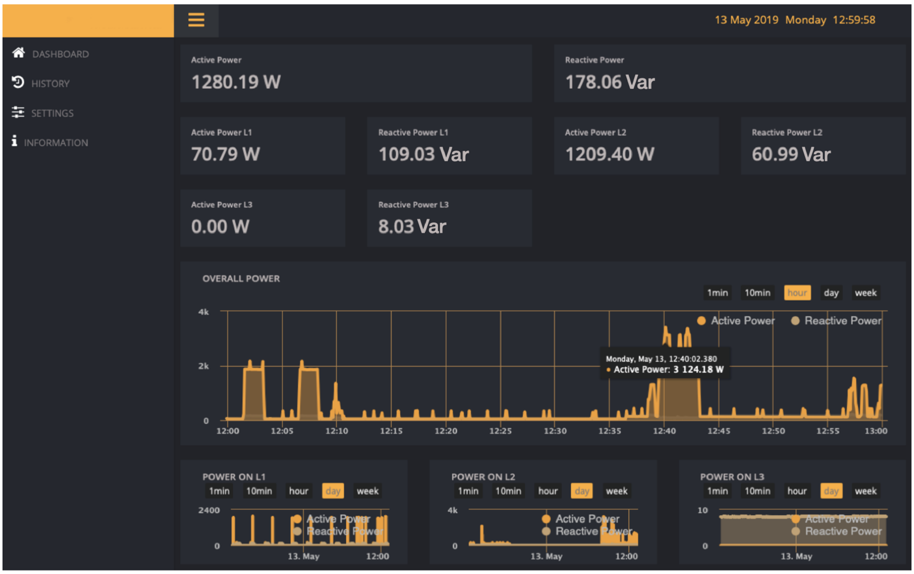

Fig. 3. Software frontend. Website showing the real time power consumption of our chair's office kitchen.

6 Flexibility

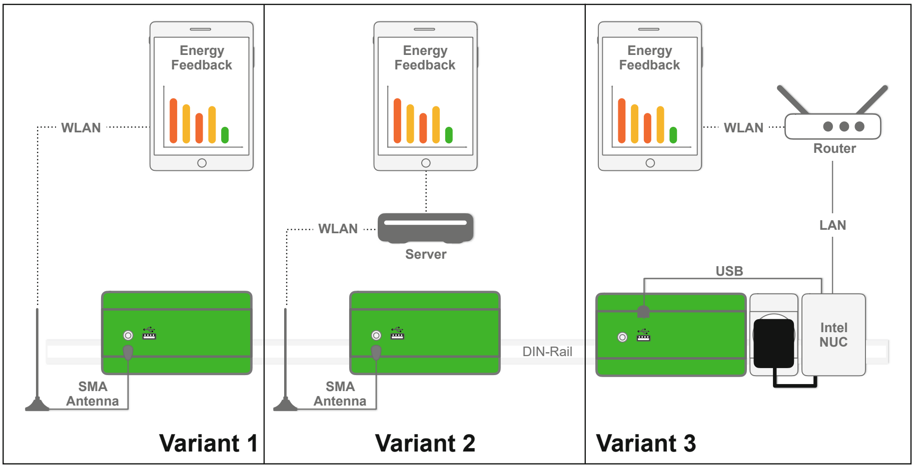

Fig. 4. Overview of different variants of the hardware setup. All variants involve the measurement system (ms) in green and a tablet. Variant 1: The ms opens a WiFi access point and performs lightweight data processing. Variant 2: An external server performs data processing and hosting a website for the tablet. Variant 3: An embedded PC is installed inside the fuse box for further data processing like electricity disaggregation; it is connected to the ms over USB and can further host the web server. (Color figure online)

Both hardware and software of the proposed framework allows for a large amount of user flexibility. Figure 4 shows the setup of different hardware variants. The basic and mandatory setup is variant 1.

Variant 1 is comprised of the measurement system explained in Sect. 3 and a tablet for visual feedback. A dedicated tablet app will connect over TCP to the measurement system and configures it to stream the data at the required sampling-rate. This data is then shown and processed in the tablet application.

Variant 2 adds an external server inside or outside the local network. This server is connected over TCP to the measurement system. It handles the additional data processing e.g. a disaggregation algorithm and hosts a web server as explained in Sect. 4. Therewith, instead of a dedicated tablet app, a web-app is displayed on the tablet. The web application is explained in Sect. 5. This further allows to exchange the tablet with any other device equipped with a web browser.

Variant 3 requires the most hardware but also shows the highest versatility. Instead of an external server, an embedded PC is installed inside the fuse box. It is connected over a USB serial connection to the measurement system. Therewith, the data is acquired over a robust wired connection and processed on a full scale processor. The PC might further perform data analysis and host the web application.

The modular structure of the overall framework allows to use it for generating real time electricity feedback using the web application, record long term electricity datasets by storing the raw data to file or to test and compare different NILM algorithms in situ with a live system by simply interfacing with the raw data socket.

As NILM algorithms develop, they might require smaller sampling rates in the future without performance loss or different electricity related features such as the harmonic content. The proposed hardware is be able to adapt and supply just the features the algorithm requires, so that resampling or data conversion is outsourced to the measurement system itself which might even reduce the energy consumption of overall system itself.

7 Generating Datasets

Since the majority of existing NILM systems rely on supervised learning methods, one of the main remaining challenges is to obtain training data. We, therefore, added the ability to hook up a recording module onto the raw data streams. This recording module can store the raw voltage and current streams into file. We choose to store these streams as wavpack encoded floating point values inside a matroska multimedia container, as this allows to store multiple streams (e.g. the consumption of each main leg) in a single file meeting requirement **R3**. The chosen container format further allows to apply different audio encoders to the raw data. We have chosen wavpack since it features lossless compression with a high compression rate as stated in requirement **R4**. A discussion of different audio encoders is presented in [12] and [14]. The container format allows to store subtitle streams as well. Subtitles can be used to store ground truth labels together with the data in a single file. One remaining challenge is to generate these ground truth labels (e.g. ON-phases of a particular device) which has to be done manually after the recording. One attempt to automate this process is proposed in [15] by using additional intrusive sensors directly connected to the monitored appliances. If data with ground truth has been generated, it can later on be used to train supervised NILM algorithms.

We further used Variant 3 of the system and successfully recorded over one month of aggregated raw voltage and current waveforms sampled at 8 kHz. This field test shows that requirement **R2** is met. Data has been stored to disc at regular time intervals of 20 min (see **R9**) resulting in approximately 420 GB of overall data. Figure 5 shows the electricity consumption of our chair's office kitchen for one day. Timestamping was performed using the recording PC since it features an accurate world clock (see **R6**, **R7**).

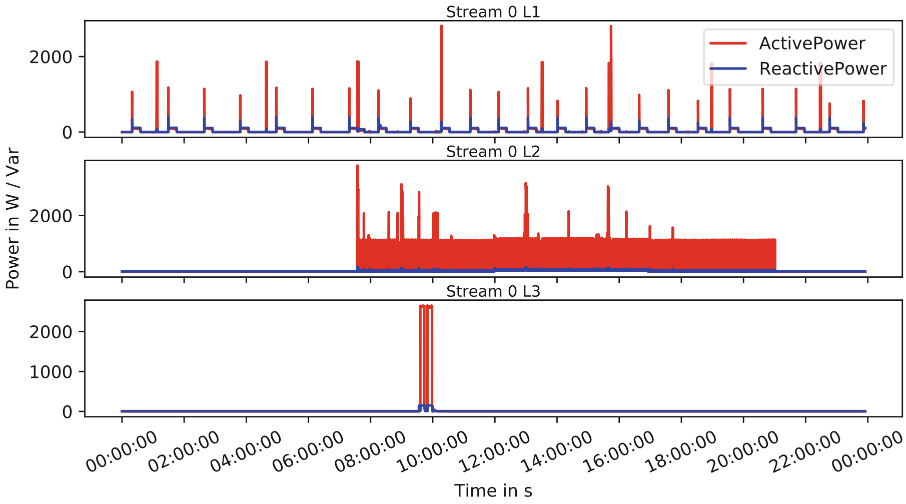

Fig. 5. One day electricity consumption of our chair's office kitchen. L_1 shows spikes of the refrigerator and heating cycles of the water therm. L_2 shows the coffee machine, the kettle and lighting and L_3 the consumption of the dishwasher which has been turned on at around 9 am.

8 Conclusions and Future Work

We have proposed a hardware and software framework to get the hands on real-time electricity data. The developed smart meter hardware allows to sample high resolution electricity data (raw voltage and current wave-forms) up to 32 kHz. This allows to further analyse high frequency components in the electricity signal. We discussed three hardware-setup variants which allow for different flexibility levels. Furthermore, we have introduced a software backend which can provide and store different electricity related metrics such as active and reactive power at various sampling-rate. We used a modular backend structure with socket communication between the modules to implement a publish and subscribe model. Different backend modules can subscribe to any data streams for further data processing. Moreover, we implemented a web server which hosts an interactive web application. The web app shows the real time electricity consumption as well as historic data.

As mentioned, generating ground truth data for better datasets remains challenging. Therefore, we are planning to instrument our chair's office kitchen with an RFID based access system to log the time when a specific device is used and the person who uses it. We hope that this will improve the quality of the generated datasets which might be also used for e.g. human activity recognition. According to the application rule *VDE-AR-N 4101* for German fuse boxes, a *RJ45* connector has to be installed in new meter cabinets allowing to e.g. connect new smart meters. Therewith, we plan to develop an updated version of our measurement system which will feature an Ethernet connector for a reliable wired connection and a desktop class processor to be able to centralize all calculations without loosing the flexibility to distribute them.

References

1. United States Environmental Protection Agency: Sources of greenhouse gas emissions (2014). https://www.epa.gov/ghgemissions/sources-greenhouse-gas-emissions. Accessed 24 Jan 2018
2. Analog Devices: High Performance, Multiphase Energy, and Power Quality Monitoring IC, April 2017. Rev. A
3. Armel, K.C., Gupta, A., Shrimali, G., Albert, A.: Is disaggregation the holy grail of energy efficiency? The case of electricity. Energy Policy **52**, 213–234 (2013)
4. Bouhouras, A.S., Chatzisavvas, K.C., Panagiotou, E., Poulakis, N., Parisses, C., Christoforidis, G.C.: Load signatures development via harmonic current vectors. In: 2017 52nd International Universities Power Engineering Conference (UPEC), pp. 1–6. IEEE (2017)
5. World Energy Council: Electricity use per household (2014). https://wec-indicators.enerdata.net/household-electricity-use.html. Accessed 24 Jan 2018
6. Emse, H.: Co2 calculator (2019). http://www.klimaneutral-handeln.de/php/kompens-berechnen.php#rechner. Accessed 23 Oct 2019
7. Felgueiras, M.C., Cruz, N., Martins, F., Martins, R.: Buildings sustainability-the non-intrusive load-identification system contribution. J. Clean Energy Technol. **4**(5), 367–370 (2016)
8. Fischer, C.: Feedback on household electricity consumption: a tool for saving energy? Energ. Effi. **1**(1), 79–104 (2008)
9. P3 International: Kill-a-watt (2017). http://www.p3international.com/products/p4400.html. Accessed 24 Jan 2019
10. Kelly, J., Knottenbelt, W.: The UK-DALE dataset, domestic appliance-level electricity demand and whole-house demand from five UK homes. Sci. Data **2**, 150007 (2015)
11. Kolter, J.Z., Johnson, M.J.: REDD: a public data set for energy disaggregation research. In: Workshop on Data Mining Applications in Sustainability (SIGKDD), San Diego, CA, vol. 25, pp. 59–62 (2011)
12. Kriechbaumer, T., Jorde, D., Jacobsen, H.A.: Waveform signal entropy and compression study of whole-building energy datasets. In: Proceedings of the Tenth ACM International Conference on Future Energy Systems, e-Energy 2019, pp. 58–67. ACM (2019)
13. Kriechbaumer, T., Ul Haq, A., Kahl, M., Jacobsen, H.A.: Medal: a cost-effective high-frequency energy data acquisition system for electrical appliances. In: Proceedings of the Eighth International Conference on Future Energy Systems, e-Energy 2017, pp. 216–221. ACM (2017)
14. Scholl, P.M., Völker, B., Becker, B., Laerhoven, K.V.: A multi-media exchange format for time-series dataset curation. In: Kawaguchi, N., Nishio, N., Roggen, D., Inoue, S., Pirttikangas, S., Van Laerhoven, K. (eds.) Human Activity Sensing. SSAE, pp. 111–119. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-13001-5_8
15. Völker, B., Scholl, P.M., Becker, B.: Semi-automatic generation and labeling of training data for non-intrusive load monitoring. In: Proceedings of the Tenth International Conference on Future Energy Systems, e-Energy 2019. ACM (2019)
16. Völker, B., Scholls, P.M., Schubert, T., Becker, B.: Towards the fusion of intrusive and non-intrusive load monitoring: a hybrid approach. In: Proceedings of the Ninth International Conference on Future Energy Systems, e-Energy 2018, pp. 436–438. ACM (2018)