Smarter Buildings for the Smart Grid? Let Them Forecast Their Power Consumption!

(Demo Abstract)

Andreas Reinhardt*, Delphine Christin†, Kai Li*, Salil S. Kanhere*, Jin Zhang*
*School of Computer Science and Engineering, The University of New South Wales, Sydney, Australia
andreasr@cse.unsw.edu.au, kail@cse.unsw.edu.au, salilk@cse.unsw.edu.au, jinz@cse.unsw.edu.au
†Secure Mobile Networking Lab, Technische Universität Darmstadt, Germany
delphine.christin@seemoo.tu-darmstadt.de

***Abstract*—An increasing number of buildings are equipped with embedded sensing systems in order to capture what is happening within. These *smart buildings* process collected sensor data to increase user comfort and safety, cater for ambient assisted living, or help the residents save energy. However, saving energy is not always beneficial to the power grid, especially when renewable sources are present. More specifically, the volatile nature of their primary energy carriers (e.g., fluctuating wind speeds) may lead to situations where significant surplus energy is being generated, which must be consumed in order to keep the power grid stable. Likewise, when unexpected drops in the generation occur, utilities must react and possibly even disconnect loads. At present, grid operators only react to the observed power consumptions, and the efficacy of the measures taken to maintain grid stability is moderate. We demonstrate how the sensor infrastructure present in smart buildings can be leveraged to accurately *predict* future power consumptions. Our system is based on commercially available device-level measurement units that transmit consumption data to a central building server. The server extracts characteristic fingerprints from historical power consumption data and uses time series pattern matching in order to detect similarities. Our demo shows that long-term predictions of an appliance's power consumption can be made, even when an appliance has been in operation for less than a minute.**

I. Introduction

Power grids that comprise renewable sources often experience high generation dynamics. The lack of energy storage components in today's power grids, however, necessitates that surplus energy provided during periods of high winds or intense sunshine must be consumed at the time of its generation. Renewable sources may also yield unexpected low output power due to the volatility of their primary energy carrier. The prevalent solution to maintain grid stability in such situations is to keep conventional power plants ready as reserve plants and even to intentionally disconnect renewable plants [1], yet neither of them represents a sustainable solution.

Accurate estimations of supply and demand help operators to alleviate this situation, yet only a few countries have accurate short-term forecasts of their renewable generation [2]. Predictions for electricity demand currently provide even less detail, as they are almost exclusively based on stochastic user models (load profiles) [3]. Although smart meters are widely deployed in numerous countries, the collected data is also rarely used to extrapolate future consumption characteristics in practice.

Instead of addressing the challenge of load prediction from the power grid perspective, we propose the use of sensor data that is available from smart building infrastructure. More precisely, our system utilizes distributed power consumption metering units, sometimes referred to as "smart plugs". Although more and more of these devices are deployed in households in order to assess appliance energy consumption, the collected power data is rarely used for the benefit of the power grid. We demonstrate how accurate power consumption forecasts can be extracted from distributed power consumption data. These predictions can be forwarded to the grid operator to facilitate the matching between supply and demand, and thus ultimately bolster the efficiency of the smart grid.

II. Demonstration System

The overall architecture of our demonstration system is depicted in Fig. 1. In essence, it is based on distributed metering units, which are connected between appliances and wall outlets. The metering units use a low-power wireless communication link to forward their readings to a gateway, which in turn relays them to a server via a wired connection. The server maintains a database of previously collected data in order to permit matching real-time consumption readings against the historical data. Furthermore, it comprises signal conditioning and preprocessing components to improve the reliability of the underlying data, and thus reduce the risk of erroneous predictions. Lastly, the server visualizes how accurate the projected consumptions match the real-time data.

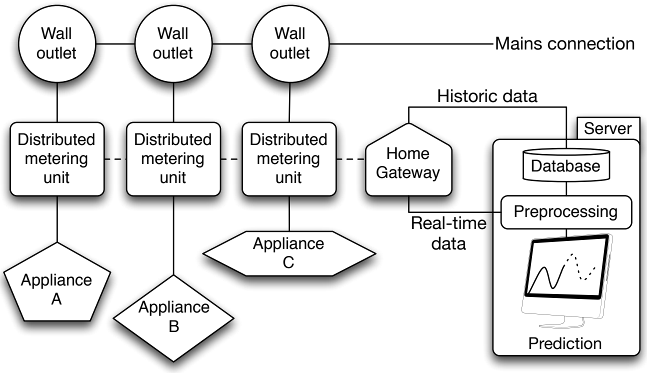

Fig. 1. System architecture of the proposed demo

A. Hardware Components

As indicated in Fig. 1, our demonstration system will comprise several different household appliances and the distributed metering units connecting them to the wall outlets. The data collection methodology follows the approach proposed in our previous work [4]. A gateway is employed to translate the data from the wireless sensors to an Ethernet connection, which is then interfaced to the server. Detailed descriptions of the individual hardware components are provided as follows.

1) Appliances: Several household appliances will be used in order to demonstrate the functional principle of our prediction system and to subsequently show its accuracy and real-world applicability. The set of appliances may include a toaster, a water kettle, a paper shredder, and further devices that allow us to showcase the concepts outlined in Sec. IV.

2) Power sensors: The distributed metering units connect between the appliances and the wall outlets and collect readings of each appliance's real power consumption at least once per second. For our demonstration system, we use the Plugwise Sting [5] devices to collect data due to their electric safety approval and the simple deployment. The Plugwise system communicates sensor readings wirelessly using low-power IEEE 802.15.4 radio transceivers.

3) Gateway: Only a single sink node is supported by the Plugwise network, yet multiple systems might possibly require access to the collected data. As a result, we have added a gateway device to the demonstration system. The gateway functionality is realized by a Raspberry Pi [6] system, on which a dedicated script continuously polls for power data and forwards them to all interested recipients. A display can be connected to the gateway to show a live view of the current device power consumptions.

4) Server: The server system will be realized by a computer system with sufficient computational power to perform time series pattern matching, which represents the basis for our load forecasting. Besides performing the computations, the server system also displays the resulting predictions on a screen.

B. Software Components

The methodology for processing both historical and real-time power consumption data is visualized in Fig. 2. In a first step, a repository of signatures is established from historical data that have been collected earlier and are stored in the database shown in Fig. 1. This process is shown on the left side of the figure. Subsequently, i.e., during the system's regular operation, the real-time incoming data (shown in the lower part) is matched against the signature repository to find the closest match and thus predict the future power consumption characteristics. We apply the following processing steps.

1) Interpolation: Packet losses may occur due to the wireless channel between the power sensors and the gateway. Likewise, more than one reading per second may be received under good conditions. In order to allow our system to correctly predict future consumptions, a unified data basis is however inevitable. In this step, the system thus interpolates and re-samples the input data to one sample per second.

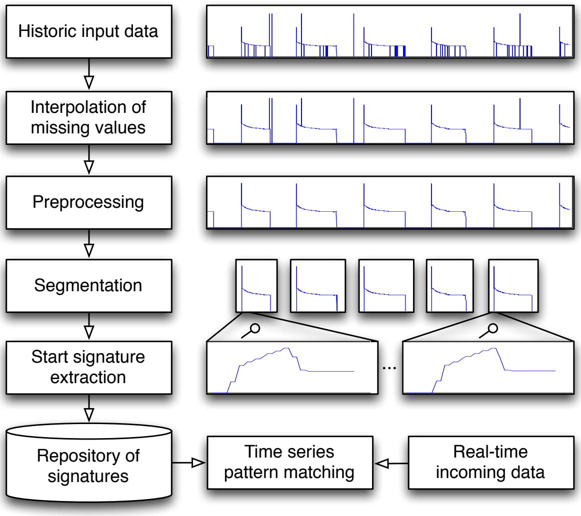

Fig. 2. Processing flow of the load predictor

2) Preprocessing: Across several installations of our power sensors, we have occasionally observed reports of erroneous power consumption readings in excess of 17,000 watts. We have hence added a preprocessing step to eliminate these values and further contribute to a coherent data representation.

3) Segmentation: After the preprocessing, the continuous time series of past consumption data is divided into individual activity segments. The resulting time series representations of the appliance's power consumption are then used as templates, against which real-time consumption data is matched during the system's regular operation.

4) Signature extraction: A large number of complete activity segments may be extracted if the database contains many historical traces. In order to cater to the system's scalability, we thus store a replica of the beginning of each activity segment in a dedicated signature repository. Each of the resulting signatures contains a reference to its underlying segment, such that the segment can be loaded when a signature match has been detected.

C. System Operation

Once all segments and the corresponding signatures have been extracted from the historical data stored in the database, the system is ready to predict appliance power consumptions. During regular operation, the server receives the readings from the gateway and applies the preprocessing and interpolation steps to the data. Subsequently, the input data are matched against the signature repository by means of time series correlation calculations. The segments linked to the traces with the highest correlation values are then returned as prospective candidates for the appliance's future power consumption pattern. This step is repeated for every incoming sample, i.e., every second, and the segment selection is refined accordingly. Once the system has detected a sufficiently close match between a signature and the real-time input data, it reports the expected remaining activity duration and the corresponding estimated power consumption.

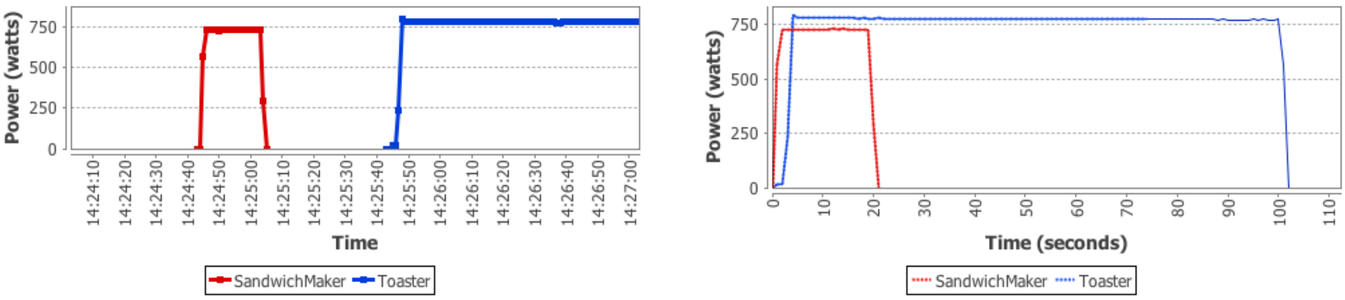

Fig. 3. Live trace view (left) and consumption predictions for both appliances (right)

III. Demonstration Setup and Requirements

The demand for space is largely dominated by the number of electric appliances that are being used. A table of 4x1 meters size with at least eight power outlets (230 volts) would be preferred, such that the following items can be accommodated:

- Four or more household appliances including the attached distributed metering units.
- The embedded gateway system and a computer monitor for a live visualization of the reported consumption data.
- One or two server computers with large monitors to show the prediction system in action and explain the underlying data processing steps.

Our Plugwise system uses IEEE 802.15.4 channel 15 and does not permit reconfiguration to another channel. While the system is resilient against incidentally received packets, a high concurrent utilization of this channel might lead to longer reporting intervals and thus to a degraded prediction accuracy. Other than that, the demonstration is self-contained, and no further infrastructure is required apart from the mains connections. The demonstration system can be set up and configured in less than one hour.

IV. Demonstrated System Functionalities

The following aspects of our system will be showcased:

1) Real-time data collection: The power sensors collect real-time consumption data of electric appliances and transmit these to the gateway. We will connect a display to the gateway and visualize the consumption readings at the time of their collection by the gateway.

2) Database retrieval: We will let the system retrieve individual traces from the server's database of historic data and plot them on the screen. Based on the graphical trace representation, we explain how the system performs automated trace preprocessing and segmentation and which parameters may be tuned to improve the quality of the automated process.

3) Visualization of the signature repository: An overview of the repository of extracted start signatures will be displayed on the screen in order to allow visitors to get an impression of their similarity.

4) Load prediction: We will allow visitors to choose an appliance of their choice and activate it while it is connected to one of the distributed power sensors. As the data are automatically forwarded to the forecasting system, a signature matching is automatically triggered. Once a match has been found, the expected power consumption until the appliance's deactivation will be forecast on the screen in real-time. The displayed user interface will look similar to Fig. 3, where a real-time view of the incoming power sensor data is shown on the left. Signature matches for both appliances have been detected, and projections of their power consumption are made in the window on the right.

5) Continuous prediction refinement: While the appliance is still connected to the sensor, the system continually updates its prediction to the most likely signature. Upon each match with a higher expected accuracy, a prediction estimate is emitted, against which the future consumption is then compared in the real-time visualization window.

V. Conclusion

In this demonstration, we will show the all necessary steps to forecast the energy consumption of electric appliances. This ranges from the real-time collection of consumption readings using smart plugs to the prediction of the appliance run-time of using time series pattern matching. The audience will be able to activate appliances of their choice and observe the forecasting process in real-time. Based on this individual forecasting, the future energy consumption of buildings can be predicted, allowing for a global optimization of the load balance in smart grids.

References

[1] The Telegraph, "Scottish wind farms paid £1 million to shut down one day," Online: http://bit.ly/10cWBjF, 05 May 2013.

[2] Elia Group, "Solar-PV Power Generation Data," Online: http://www.elia.be/en/grid-data/power-generation/Solar-power-generation-data, last access: June 2013.

[3] ELEXON Ltd, "Load Profiles and Their Use in Electricity Settlement," Online: http://www.elexon.co.uk/wp-content/uploads/2012/01/load_profiles.pdf, 2005.

[4] A. Reinhardt, P. Baumann, D. Burgstahler, M. Hollick, H. Chonov, M. Werner, and R. Steinmetz, "On the Accuracy of Appliance Identification Based on Distributed Load Metering Data," in *Proceedings of the 2nd IFIP Conference on Sustainable Internet and ICT for Sustainability (SustainIT)*, 2012, pp. 1–9.

[5] Plugwise B.V., "Plugwise - Smart Wireless Solutions for Energy Saving, Monitoring and Switching," Online: http://www.plugwise.com, last access: June 2013.

[6] Raspberry Pi Foundation, "Raspberry Pi – An ARM GNU/Linux box for $25. Take a byte!" Online: http://www.raspberrypi.org, last access: July 2013.