

TUCap: A Sensing System to Capture and Process Appliance Power Consumption in Smart Spaces

Andreas Reinhardt
Technische Universität Clausthal
Clausthal-Zellerfeld, Germany
reinhardt@ieee.org

Abstract—The collection of power consumption data from electrical appliances is a key enabling element for grid-related services such as load forecasting or anomalous consumption pattern detection. Device-level sensors (*smart plugs*) have found widespread use to collect such data. However, they prevalently report an electrical appliance’s power consumption at a rate of one reading per second. With mains voltage frequencies of 50/60 Hz, undersampling and the consequent loss of spectral information result from the use of such sensor devices. Moreover, most smart plugs only report readings of an appliance’s real power consumption. Important supplementary features like the phase shift between voltage and current or the magnitude of reactive power are not available for retrieval from commercially available devices. In this work, we overcome these limitations of smart plugs by presenting TUCap, an embedded sensing system to capture appliances’ electrical load signatures in minute detail. Our design caters to the provision of a high information content by capturing voltage and current waveforms at a sampling rate of 36 ksps. Thus, spectral components are implicitly included in collected traces. Moreover, TUCap uses local data processing routines to detect and eliminate redundancies. Thus, a high data fidelity is maintained while achieving significant reductions of network traffic. All functionalities are implemented in a proof-of-concept system design and evaluated in practice.

I. INTRODUCTION

Since the publication of George Hart’s seminal work on non-intrusive load monitoring (NILM) in 1985 [1], research activities to extract and utilize the information content in electricity consumption data have seen an almost exponential increase. A key insight of *energy analytics* is that power consumption patterns are suitable indicators to infer information on both appliance and user activities in residential, commercial, and industrial settings. Many data processing algorithms have consequently been presented to analyze energy consumption data, e.g., to characterize household types [2, 3], infer building occupancy [4], or predict future electricity consumption [5, 6]. Their designs often also highlight one of the principal limitations of NILM [7, 8]: It relies on the usage of a single measurement point (usually a smart electricity meter). Thus, differentiating between appliances of the same type operated in different locations within a dwelling is complicated. In order to overcome this limitation, deployment strategies for sensors to capture electricity consumption in a more fine-grained fashion have emerged, e.g., by adding sensors for each electrical circuit in a home [9, 10] or even inserting appliance-level sensors (*smart plugs*) into the power cords of all relevant devices [11].

Supplementary to the option of increasing energy analytics accuracy through deploying more sensing points in a dwelling, the data processing capabilities are also strongly dependent on the sampling rate at which readings have been collected [12]. Unfortunately, the majority of today’s commercially available plug-level sensors suffer from shortcomings in this regard. Smart plugs often report values only once per second or even less often, thus transients and spikes of shorter duration commonly remain undetected. Furthermore, appliances with inductive or capacitive components incur a reactive power consumption, whose detection requires the synchronous sampling of voltage and current. As most smart plugs are not fitted out with voltage sensors, they are unable to differentiate between real and reactive power consumption. At last, mains voltage and current consumption waveforms rarely resemble perfect sinusoids. In contrast, slight distortions are ubiquitous due to the wide presence of non-linear loads, e.g., switch-mode power supplies. To fully capture these characteristics, sampling frequencies have to satisfy the Nyquist sampling theorem [13], i.e., be greater than twice the highest distorting frequency. This is not the case for smart plugs that report readings at 1 Hz.

The limitations of commercially available smart plugs have motivated many researchers to design platforms to measure electrical energy consumption, such as Plug [14], ACme [15], SEM [16], SmartMeter.KOM [17], WCSN [18], or YoMo [19]. While these devices are capable of sampling appliance power consumptions at higher resolutions than their commercially available counterparts, their data processing functionalities are often limited. In fact, research platforms either forward data to a collection device without any prior processing (i.e., they report data at the native sampling rate), or apply lossy data processing algorithms to return characteristic values (e.g., RMS current, crest factor, etc) at the 1 Hz interval prevailing among commercial platforms. The former approach, however, results in an enormous bandwidth requirement, whereas the latter solution disallows for the detection of short-term fluctuations. Such features are, however, vital for energy analytics.

The unavailability of a solution to capture high-resolution data while minimizing the bandwidth demand during periods of constant power consumption has motivated us to create TUCap. We present an overview of its architecture and the hard- and software design decisions taken in Sec. II. Subsequently, we discuss considerations for the system’s practical use in Sec. III. At last, we conclude this paper in Sec. IV.

II. SYSTEM OVERVIEW

In order to enable energy analytics, both mains voltage and current flowing into an electrical appliance must be sampled at a high temporal and amplitudinal resolution. Therefore, three design decisions need to be made: Firstly, suitable transducers to translate voltage and current into analog signals need to be carefully selected. Secondly, an analog-to-digital conversion (ADC) step with high resolution must be part of the system in order to discretize samples and thereby enable their processing. Thirdly, a networked embedded system with the capabilities for local data processing is required in order to accomplish our objective of filtering out redundant readings while maintaining full data fidelity for when changes occur. We discuss our design choices for each of these aspects as follows.

A. Transducer Selection

The *load signature* of an electrical appliance, i.e., its characteristic power consumption during the course of its operation, is defined by the current flowing through the device as well as the voltage across its terminals.

TUCap samples current flows via a non-invasive AC current transformer that can be clipped on any conductor. The selected current transformer (model SCT-013) has a 2000:1 ratio, thus it transforms a primary current of 2 A into a secondary current of 1 mA. We burden the current transformer with a series of load resistors configured as a voltage divider that delivers an output voltage of ± 1 V for a ± 16 A primary current. The ± 1 V range has been chosen in order to be compliant with the input specifications of most ADC devices (cf. Sec. II-B). The current transformer has a specified frequency response up to 150 kHz, thus harmonics will be reflected in its output. As recommended per the data sheet [20], a passive first-order low-pass filter has been added on-board TUCap to avoid aliasing.

To capture the mains voltage signal, a 230 V/6 V transformer is present on TUCap. The transformer's output is burdened with a resistor network to proportionally scale voltages up to ± 433 V down to the ± 1 V input range of the ADC. Like on the current channel, a first-order anti-aliasing filter has been added to the board. The key intention behind using a transformer is to ensure galvanic decoupling, and thus render the use of a signal isolation device for the power and data connections unnecessary. A downside of using a transformer, however, is its non-linearity: Small transformers often introduce a phase shift that needs to be compensated for. We have hence measured the transformer's voltage transfer characteristics in Fig. 2, and determined the output voltage to lead (i.e., advance) the input voltage by $500\mu\text{s}$. A compensation for this phase shift thus has to be implemented in software.

The connection points for the chosen voltage and current transducers in an appliance's power cord are shown in Fig. 1.




Fig. 1. Transducer connection terminals in an appliance's power cord.

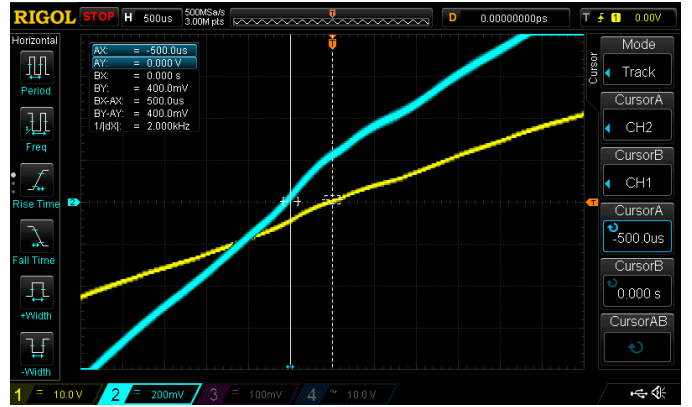


Fig. 2. Input (channel 1; yellow) and output (channel 2; cyan) voltages when supplying the transformer with $16\text{V}\sim$ at 50 Hz and using a 330Ω burden resistor. The cursors show a discrepancy of $500\mu\text{s}$, i.e., a phase shift of 9° .

B. Analog-to-Digital Conversion

While many microcontroller-based embedded sensing solutions feature integrated ADCs, their resolution is commonly limited to a narrow value range. For example, the ADC of the well-known TelosB mote provides an amplitude resolution of 12 bits. In contrast, integrated circuits for power metering often feature sample resolutions of 20 bits (effective) or even more. As our objective was to capture voltage and current at high temporal and amplitude resolutions, we have decided in favor of a dedicated energy metering integrated circuit. The selected Microchip MCP3910 [20] device serves as the analog front end and converts values from the voltage and current transducers into digital signals. It synchronously samples voltage and current and allows for the retrieval of the readings over a three-wire serial interface. In addition to sampling the transducers' outputs, the MCP3910 features an integrated amplifier with configurable amplification (up to 32x) which can be used to monitor small currents if needed. When operated with a crystal at 18.432 MHz, the device reliably achieves a sampling rate of 36 kSps at 23 bits resolution, thus allowing for the in-depth spectral analysis of the captured signal. A photography of our TUCap board is shown in Fig. 3.




Fig. 3. Key components on the TUCap sensor interface board.

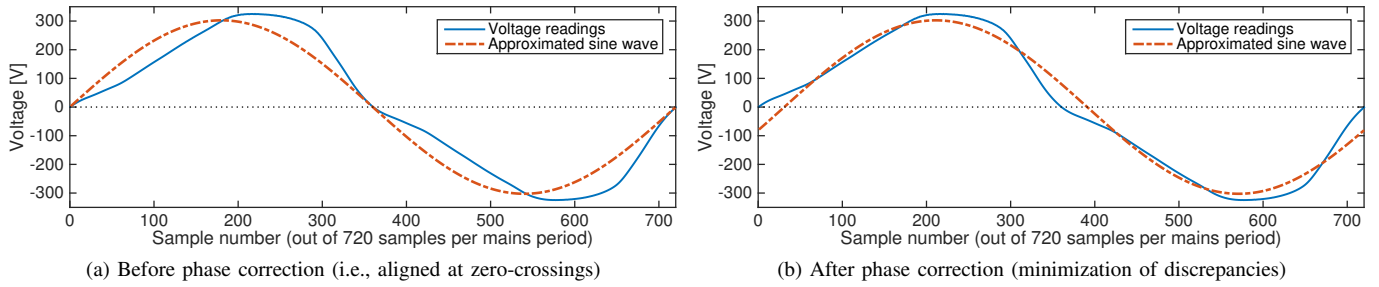


Fig. 4. Collected voltage waveforms and sine wave approximations before and after phase correction.

C. Data Retrieval and Preprocessing

In order to allow for the collection and transmission of power consumption readings and other values related to load signatures, we connect the MCP3910 analog frontend to a Teensy 3.0 system [21]. Whenever a falling edge is detected on the ADC's data ready signal (\overline{DR}), the Teensy retrieves the latest voltage and current values over the SPI interface and stores them in corresponding buffers. A zero-crossing detection routine has been implemented to identify the beginning of a mains voltage period. It triggers the execution of the data preprocessing operations when 720 samples have been collected with exactly one zero-crossing observed in the voltage waveform during this period.


As mentioned above, the MCP3910 is capable of sampling its voltage and current channels at a rate of 36 kSps and 23 bits of accuracy. Due to the word alignment of the Teensy's 32-bit ARM microcontroller, however, the measurements would need to be stored as 32-bit values, resulting in a data generation rate of more than 280 kB/s. The resulting data rate (2.25 MBps) is beyond the capacities of most communication links available to low-power embedded systems. To mitigate this issue, TUCap only retrieves 16-bit values from the ADC. The resulting amplitude resolution of 13 mV on the voltage channel and 0.5 mA for the current still allows for power measurements with a resolution better than 0.2 W.

Nonetheless, the need for local data preprocessing becomes apparent when considering the volume of data generated. TUCap thus executes two data preprocessing steps:


- 1) The sinusoidal nature of the mains voltage enables TUCap to approximate collected waveforms by a parametrizable sine wave. Instead of transferring 1,440 bytes of raw data samples, only two values are reported: The sine wave's amplitude and its phase shift. The amplitude of the parametrizable sine wave is computed by calculating the absolute value of the area under the curve of the voltage waveform and normalizing it to the absolute area under a sine wave ($\int_0^{2\pi} |\sin x| dx = 4$). Subsequently, the phase shift is determined by iteratively comparing the sampled voltage trace to parametrizable sine waves at different offsets until the observed discrepancy is minimal. This step is visualized in Fig. 4. Finally, the fixed offset introduced by the voltage transformer (shown in Fig. 2) is subtracted from the reported value.

- 2) Current consumption waveforms of electrical appliances often remain highly similar for extended periods of time. They mostly only alternate when the operation mode of the appliance changes. To demonstrate this in practice, we have used TUCap to record 100 periods of the waveform of a computer monitor's power consumption when playing back a video. The superposition of all recorded waveforms is shown in Fig. 5a and demonstrates this coherence. To avoid transmitting redundant data, TUCap thus captures a single waveform of the current and stores it in memory (similar to the notion of an *I-Frame* in MPEG video coding [22]). All signals captured subsequently undergo an element-wise comparison to this previously stored waveform. Unless elements show a difference greater than a user-configurable threshold value level ρ , TUCap does not report any update.

A supplementary analysis of practically determined deviations between an initial waveform and later traces is shown in Fig. 5b. It is visible that slight differences always occur due to the presence of noise, thus it is not meaningful to set $\rho < 20$. Other than that, the choice of ρ defines the frequency of reports and thus the accuracy of reported waveforms.



(a) Overlaid visualization of 100 current flow recordings taken sequentially. The black line shows the course of the waveform initially recorded.



(b) Maximum absolute differences observed between each of the 100 recordings and the initially stored current waveform.

Fig. 5. A highly similar current consumption pattern can be observed when comparing 101 current consumption recordings of a 20" TFT monitor.

III. PRACTICAL CONSIDERATIONS

During the design of TUCap, a set of practical insights into the operation of an energy monitoring platform were gained.

1) *Choice of Components*: Manufacturing tolerances of the used components make it complicated to lay both input channels out in a symmetric fashion. In particular, load resistors and voltage dividers must have closely matching values in order to attain symmetric voltage swings that feature identical amplitudes in both the positive and the negative direction. Despite the choice of resistors with 1% tolerance, we found the pre-selection of resistors through measurements using a high-accuracy LCR meter to be indispensable.

2) *Calibration*: An initial factory calibration step must precede the operation of TUCap in practice. To encounter component tolerances, the MCP3910 features calibration registers for both voltage and current. A reference AC power supply is inevitable to supply input voltages and currents of known amplitude and frequency to the device in order to determine and compensate any offsets in hard- or software.

3) *Shielding*: During some of our practical experiments, TUCap has infrequently been unable to retrieve data over the SPI connection. Insufficient shielding of the data connectors and their cable lengths could be identified as the primary source of error. An integration of the ADC and the microcontroller on a single circuit board and/or better cable shielding is thus recommended.

IV. CONCLUSIONS

Energy informatics research has received a lot of attention in the last decade, primarily due to the increasing availability of data collected by smart meters and smart plugs. A strong limitation of most commercially available platforms, however, is their limited sampling rate of 1 Hz or even less. In order to foster energy informatics research, we have thus presented TUCap. The synergistic combination of a high-resolution ADC to capture voltage and current waveforms with an embedded system to process and transfer collected data has emerged as an ideal solution for our use case. By choosing small values for ρ , highly accurate consumption traces can be recorded; in turn, selecting a larger value for the tolerance parameter reduces generated traffic and may thus enable TUCap's use in scenarios where bandwidth is scarce. As the next step in TUCap's development, we plan to implement more algorithms for local data analytics and compression on the embedded system, such as the ones surveyed in [23, 24].

REFERENCES

- [1] G. W. Hart, "Prototype Nonintrusive Appliance Load Monitor," MIT Energy Laboratory and Electric Power Research Institute, Tech. Rep., 1985.
- [2] M. Weiss, A. Helfenstein, F. Mattern, and T. Staake, "Leveraging Smart Meter Data to Recognize Home Appliances," in *Proceedings of the IEEE International Conference on Pervasive Computing and Communications (PerCom)*, 2012.
- [3] C. Beckel, L. Sadamori, T. Staake, and S. Santini, "Revealing Household Characteristics from Smart Meter Data," *Elsevier Energy*, vol. 78, 2014.
- [4] W. Kleiminger, C. Beckel, and S. Santini, "Household Occupancy Monitoring Using Electricity Meters," in *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*, 2015.
- [5] V. Bianco, O. Manca, and S. Nardini, "Electricity Consumption Forecasting in Italy using Linear Regression Models," *Energy*, vol. 34, no. 9, 2009.
- [6] A. Reinhardt, D. Christin, and S. S. Kanhere, "Can Smart Plugs Predict Electric Power Consumption? A Case Study," in *Proceedings of the 11th International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services (MobiQuitous)*, 2014.
- [7] M. Gulati, S. S. Ram, and A. Singh, "An in Depth Study into Using EMI Signatures for Appliance Identification," in *Proceedings of the ACM Conference on Embedded Systems for Energy-Efficient Buildings (BuildSys)*, 2014.
- [8] M. Kahl, A. Ul Haq, T. Kriebchaumer, and H.-A. Jacobsen, "WHITED – A Worldwide Household and Industry Transient Energy Data Set," in *Proceedings of the 3rd International Workshop on Non-Intrusive Load Monitoring (NILM)*, 2016.
- [9] A. Marchiori, D. Hakkarinen, Q. Han, and L. Earle, "Circuit-Level Load Monitoring for Household Energy Management," *IEEE Pervasive Computing*, vol. 10, no. 1, 2010.
- [10] D. Jung and A. Savvides, "Estimating Building Consumption Break-downs Using ON/OFF State Sensing and Incremental Sub-meter Deployment," in *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems (SenSys)*, 2010.
- [11] T. W. Hnat, V. Srinivasan, J. Lu, T. I. Sookoor, R. Dawson, J. Stankovic, and K. Whitehouse, "The Hitchhiker's Guide to Successful Residential Sensing Deployments," in *Proceedings of the 9th ACM Conference on Embedded Networked Sensor Systems (SenSys)*, 2011.
- [12] J. Liao, G. Elafoudi, L. Stankovic, and V. Stankovic, "Non-intrusive Appliance Load Monitoring using Low-resolution Smart Meter Data," in *Proceedings of the 5th IEEE International Conference on Smart Grid Communications (SmartGridComm)*, 2014.
- [13] H. Nyquist, "Certain Topics in Telegraph Transmission Theory," *Transactions of the American Institute of Electrical Engineers*, vol. 47, 1928.
- [14] J. Lifton, M. Feldmeier, Y. Ono, C. Lewis, and J. Paradiso, "A Platform for Ubiquitous Sensor Deployment in Occupational and Domestic Environments," in *Proceedings of the 6th International Symposium on Information Processing in Sensor Networks (IPSN)*, 2007.
- [15] X. Jiang, S. Dawson-Haggerty, P. Dutta, and D. E. Culler, "Design and Implementation of a High-Fidelity AC Metering Network," in *Proceedings of the 8th International Conference on Information Processing in Sensor Networks (IPSN)*, 2009.
- [16] T. Weng, B. Balaji, S. Dutta, R. Gupta, and Y. Agarwal, "Managing Plug-loads for Demand Response Within Buildings," in *Proceedings of the 3rd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings (BuildSys)*, 2011.
- [17] A. Reinhardt, D. Burkhardt, P. S. Mogre, M. Zaheer, and R. Steinmetz, "SmartMeter.KOM: A Low-cost Wireless Sensor for Distributed Power Metering," in *Proceedings of the 6th IEEE Workshop on Practical Issues in Building Sensor Network Applications (SenseApp)*, 2011.
- [18] D. Porcarelli, D. Balsamo, D. Brunelli, and G. Paci, "Perpetual and Low-cost Power Meter for Monitoring Residential and Industrial Appliances," in *Proceedings of the Design, Automation Test in Europe Conference Exhibition (DATE)*, 2013.
- [19] C. Klemenjak, D. Egarter, and W. Elmenreich, "YoMo: The Arduino-based Smart Metering Board," *Computer Science – Research and Development*, vol. 31, no. 1, 2016.
- [20] Microchip Technology Inc., "MCP3910 – 3V Two-Channel Analog Front End," Data Sheet available online at <http://ww1.microchip.com/downloads/en/DeviceDoc/20005116B.pdf>, 2014.
- [21] PJRC LLC., "Teensy USB Development Board," available online at <https://www.pjrc.com/teensy/index.html>, last access on 26 June 2017.
- [22] D. Le Gall, "MPEG: A Video Compression Standard for Multimedia Applications," *Communications of the ACM*, vol. 34, no. 4, 1991.
- [23] M. Kahl, A. Ul Haq, T. Kriebchaumer, and H.-A. Jacobsen, "A Comprehensive Feature Study for Appliance Recognition on High Frequency Energy Data," in *Proceedings of the Eighth International Conference on Future Energy Systems (e-Energy)*, 2017.
- [24] M. Ringwelski, C. Renner, A. Reinhardt, A. Weigel, and V. Turau, "The Hitchhiker's Guide to Choosing the Compression Algorithm for Your Smart Meter Data," in *Proceedings of the 2nd IEEE Conference and Exhibition / ICT for Energy Symposium (ENERGYCON)*, 2012.