

On the Impact of User Feedback on Energy Consumption

Christian Bunse¹

Abstract

The energy consumption of mobile devices is of increasing importance for users, developers and vendors. Frequent recharging cycles have an impact onto the mobility of a device and thus, on user experience. While hardware and software are known to have an impact on energy consumption, user behaviour (when and how a device is used) is often not in the focus of attention. Research in on “smart-metering” has shown that user behaviour is a non-negligible impact factor. This paper reports about the results of an ongoing study on the impact of user feedback onto the energy consumption of mobile devices (i.e. Smartphones). The hypothesis being, that the frequent provision of energy-related feedback on user actions will significantly reduce energy utilization. First results support this hypothesis and indicate that feedback frequency and feedback clearly linked to specific user activities are key in this regard.

1. Introduction

Energy is one of the most limiting factors for the growth of information and communication technologies, especially when it comes to mobile devices. Usually such devices are battery powered and thus, have limited energy resources. Beneath hardware, the software running on a device utilizes the device hardware and therefore has an impact onto energy consumption too [6]. An often neglected, but nonetheless important, third impact factor is user behaviour or more precisely the actual usage of a device. The user decides on which application or service to start, or which sensors and communication means are going to be used. Changing this behaviour might therefore reduce energy needs. Unfortunately, user behaviour often deviates from a standard rational choice model. From a user perspective comfort and performance are regarded higher than energy consumption. Single or isolated energy saving instructions are often without effect.

This paper presents a study that examines the effects of frequent, action- and energy related user feedback onto the energy consumption of mobile devices. This is achieved by comparing the energy needs of different Smartphone user-groups: Users that are frequently informed about the energy costs of their actions and users who are not. It is assumed that members of the first group will, in the long run, use significantly less energy than users of the other group. The hypothesis being that frequent and prompt feedback will positively alter user behaviour. It is important to note that the results presented in this paper are preliminary and that the actual study is ongoing. Thus, the examination of long-run learning effects is not within the scope of this paper.

The remainder of the paper is structured as follows: Section two discusses related work. Section three provides details about the underlying research hypothesis, study design and study execution. Section four presents and analyses collected data and discusses possible threats to validity. Finally, section five concludes the paper and provides an outlook on future research.

2. Related Work

Several research projects have been conducted regarding the energy consumption/needs of mobile systems. Often these address hardware or software related approaches [7]. Research that belongs to

¹ University of Applied Sciences Stralsund, 18435 Stralsund, Germany, Christian.Bunse@fh-stralsund.de, Department of Electrical Engineering and Computing Science

the hardware category, usually attempts to optimize the energy consumption by optimizing hardware usage, such as [3], or by introducing energy-efficient hardware devices and techniques, such as [12], [13]. Research in the second category attempts to understand how software and its development affects energy consumption. Research in this category can be further classified according to factors such as networking/communication, application nature, memory management, and algorithms. Concerning networking and communication projects such as [4] provide new routing techniques that are aware of energy consumption. Other efforts of this category focus on providing energy-aware protocols for transmitting data in wireless networks [11] and ad-hoc networks [5]. Memory consumption is another important factor concerning a system's energy consumption. In this regard work such as [8] provides energy-aware memory management techniques. In battery powered systems, it is not sufficient to analyze algorithms based only on time and space complexity. Energy-aware algorithms such as [7] supporting randomness, or [9] focusing on cryptographic networks were published.

Beneath technical approaches towards optimizing energy needs, users and their behaviour are additionally identified as an important factor. [16] defines microbenchmarks for emailing and web browsing, and evaluates applications from these domains. [15] proposes an approach for profiling the power consumption of mobile applications and thus a means for comparing the consumption of similar services (Energy Labeling). [2] compares the energy consumption characteristics of mobile networking technologies and how this influences user interaction. [14] presents a framework that automatically recognizes user daily activities and energy needs in real time using sensors on a smart phone. Still missing are (empirical) studies that actually monitor user behaviour over a longer period of time to better understand the relationship of behaviour and energy consumption.

3. Study Design

3.1. Subjects

The experimental subjects used in the study are students (bachelor and master) across all departments and cohorts of the University of Applied Sciences Stralsund. Volunteers owning a suitable device were asked for participation and were informed about the nature of the study. Subjects were motivated by making it clear that they would gain valuable experience from participating and that their data will help in improving battery life of future devices. As a result, 30 subjects volunteered to take part. The subjects knew that data would be, anonymously, collected and that an analysis would be performed on the data, but were unaware of the experimental hypotheses that were being tested. While the call was open to all students the majority of subjects are studying computer science. Before attending, subjects were asked to answer a small questionnaire on their background, used devices, typical usage scenarios, and expectations/experience regarding the energy consumption of mobile (smart-) phones. It appeared that most subjects (>90%) use high-end devices and (~75%) had the subjective impression that the energy consumption of their device is too high. Interestingly only a minority (10%) of them were aware that their behaviour might have an impact while the majority believed in in-efficient hardware (40%), energy mismanagement (OS) (29%), or badly programmed apps (18%).

3.2. Hypotheses

The relationship between user behaviour and energy consumption is in the focus of this study. The assumption being that keeping user informed about the implications of their actions will significantly improve the energy consumption of the devices they are using. More precisely the null hypothesis can be stated as: The energy consumption of devices whose users are frequently

informed about the energy related effects of their activities does not significantly differ from the energy consumption of devices whose users are not actively informed.

3.3 Experimental Design

In order to test our hypothesis subjects we selected a between subject design and randomly assigned subjects to one of three groups (10 subjects per group). The dependent variable being the information frequency and the independent variable being energy consumption. When performing experiments regarding the behaviour of humans, the Hawthorne effect, i.e. subjects change their behaviour in response to the fact of being observed, bears the danger of falsifying results. While this effect cannot be fully excluded it also correlates with the experience of a user. Experienced subjects know which factors do have a significant impact onto energy consumption. Behavioural changes might therefore result in larger effects. To control this random group assignment was performed. As the number of subjects was known before running the experiment, it was a simple to create groups of equivalent size, which is important to prevent the independent variables from becoming non-orthogonal. Subjects within the two treatment groups receive real-time feedback regarding the impact of their activities on energy usage. The difference between both groups is the frequency of feedback. Subjects within the control group are not provided with such feedback.

3.4 Infrastructure

The hypothesis of this study requires users to be informed in real-time about the impact of their activities onto the energy consumption of the device. This requires energy measurement at the level of user activities. Energy consumption could either be measured by using onboard software tools or by measuring voltage drops at battery level. Following [6] this study uses onboard which are sufficient for trend analysis. We developed an app that runs as a system service using the AIDL-Interface. The service measures app-usage in terms of time and energy [1]. Data is stored locally and is send daily to a central server. Additional sensor data (e.g. GPS) is not monitored, although correlations between energy usage and Geo-position might be an interesting future topic.

The service also creates system notifications to inform users about their energy related behaviour (treatment group). Such notifications provide information about the app being used, the time of usage and energy related information. The latter is not provided as a Joule value but in form of a “traffic light” symbol. Joules are provided on user request only. In addition, a summary report is created on (user definable) times. These reports inform users about the actual trend. On the server side incoming data is stored with a timestamp and provided as a CSV-file for further analysis.

3.5 Procedure

The week before the experiment started, subjects were asked to fill a questionnaire in order to learn about their background, experience, etc. Then subjects were given an overview of the study and its procedure. Subjects were assisted for preparing their devices and informed the type and nature of collected data, data security, etc. Subjects were told that there is no external access to their device, that personal data will not leave the device, and that data will be anonymized prior to sending. Finally, groups were formed and the study was started. Subjects of the treatment groups were asked to pay attention towards the energy related feedback, while subjects of the control group were not told anything in this regard but were free to use system energy information. Subjects of the first treatment group received feedback on a time basis by pushing a system notification every 5-10 minutes. The second treatment group was notified based on their actions (switching apps, making phone calls, etc.).

4. Results

Once the study was started, data was collected by the background service running on subject devices. Since the study is ongoing no final statistical analysis has taken place yet. Unfortunately, the dataflow of the first treatment group was sparse. During a feedback session it appeared that subjects quickly got annoyed by the high frequent feedback (5-10 minutes) that was not directly correlated with their actions. As a result, subjects soon stopped data collection. Thus, valid data is only available for the second treatment as well as the control group.

Figure 1 visualizes the accumulated energy consumption (KWh) of both groups (Control and Treatment) over a period of eighty days. As time advances, the gap between both groups continuously grows. This supports the assumption that real-time feedback enables continuous energy savings. At a more detailed level it can be concluded that the treatment group adapted their behaviour (energy wise) within the first days) while the control group did not.

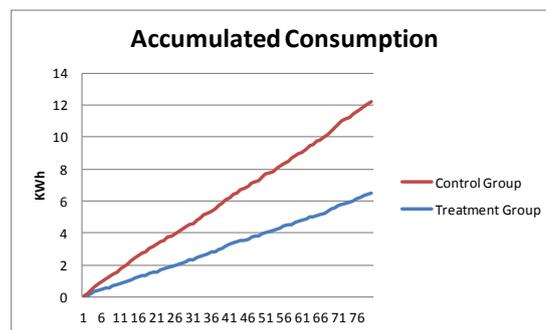


Figure 1: Accumulated Energy Consumption of both groups

This is supported by Figure 2, which compares the energy consumption of both groups at an early stage of the project (day 3) with those later in the experiment (day 40). While the behaviour of the control group has not significantly changed over time (from 0.0079 KWh to 0.0081 KWh) the treatment group does show a change (from 0.0079 KWh to 0.0059 KWh). Although a saving of 0.002 KWh per day is not a dramatic change, data supports the hypothesis that behaviour and thus energy consumption can be changed by keeping the user informed.

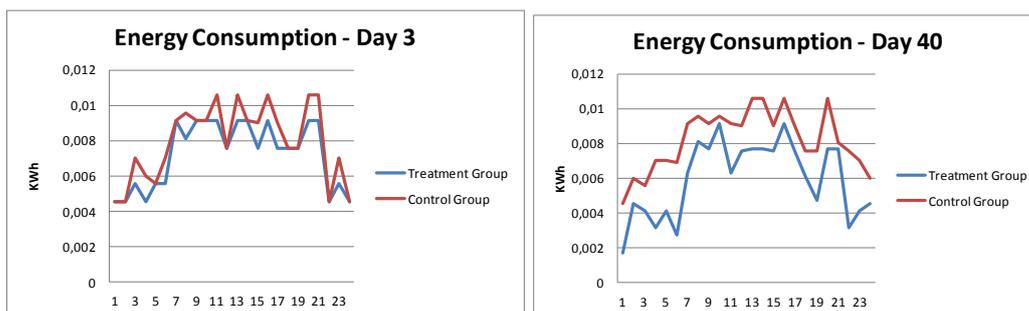


Figure 2: Energy Consumption – Day View

Another interesting fact is that energy consumption data reveals the daily activity (level) of users. It appears that it is possible to identify the currently executed app just by viewing at its energy profile. It can be clearly identified if a user is making phone calls, surfing the web or playing games. This confirms findings in the smart metering domain that energy profiling allows user identification [10] or even the movie currently playing. In this study we guaranteed anonymous data collection and therefore deleted all information that would allow data backtracing.

Although the study is ongoing we performed a first statistical analysis in order to test our hypotheses. Checking the data for normality revealed that it is substantially non-normal and that,

thus, a non-parametric test has to be used. Due to the early stage of the study we preliminary set the level of significance (i.e., α -level) to $\alpha = 0.05$. To test the hypothesis we decided to use the Kruskal-Wallis test since it is known to be robust against variations in group size. Results show that there is a significant difference between both groups (p -value < 0.03) and thus the null hypothesis can be rejected. Unfortunately, the test does not identify where the differences occur or how many differences actually occur. Thus, additional analysis is needed.

5. Threats to Validity

There are a number of threats to validity that can affect the dependent variables. While these threats limit generalization of this research, they do not prevent the results from being used in further studies.

5.1. Construct Validity

Construct validity is the degree to which the independent and dependent variables accurately measure the concepts they purport to measure. The energy consumption of a mobile device is a difficult concept to measure. Granularity, preciseness and timeliness are key in this regard. In the context of this study it is argued that using the onboard measurement facilities of Android are sufficient regarding the goals of the study. However, additional studies are required to investigate if other measurement approaches lead to different results.

5.2. Internal Validity

Internal validity is the degree to which conclusions can be drawn about the causal effect of independent variables on the dependent variable. The following possible threat was identified: • A maturation effect is caused by subjects learning as the experiment proceeds. Due to the long execution time information about the goals of the study and the expected results might reach subjects and alter their behaviour. However, since subjects were randomly assigned to their group, we assume that learning effects are equal to all groups and can therefore be neglected.

5.3. External Validity

External validity is the degree to which the results of the research can be generalized to the population under study and other research settings. The following possible threats were identified:

- The subjects were students and are therefore unlikely to be representative of Smartphone users in general. We assume that students are using their devices more intensively than the majority of users and thus, that only effect size will differ between students and other user groups.
- Using volunteers as subjects may affect validity (i.e., selection bias). Volunteers are almost certainly different from those subjects who do not volunteer. They are, by definition, motivated to participate and presumably expect to receive some benefit. These differences limit the ability to generalize the results of this study beyond the research sample and require future studies.

6. Summary and Conclusions

Energy resources are a limiting factor for most mobile devices and especially for smart-phones. The consumption of this resource is driven by the devices' hardware, running software, and the user. This paper reports about an (empirical) study that investigates the relationship between user behaviour and energy consumption. More specifically the study examines if keeping users informed about the energy "costs" of their activities lowers energy needs and thus, increases device uptime. Although the study is ongoing, data analysis supports our hypothesis that keeping users informed has indeed a significant impact (p -value below 0.03). Feedback related to concrete

user actions seems to be key in this regard. Time-triggered feedback seems to have a contrary effect as demonstrated by the drop-out of a complete group of subjects. However, there are a number of threats to validity that hinder generalization. Future studies have to address these threats and, more importantly, examine the effects of different forms of user feedback. Another study will investigate if subjects of the treatment group have permanently altered their behavior (learning effect).

References

- [1] Bunse, C., Höpfner, H. (2012): OCEMES: Measuring Overall and Component-based Energy Demands of Mobile and Embedded Systems, GI-Jahrestagung 2012
- [2] Balasubramanian, N., Balasubramanian, A., Venkataramani, A. (2009): Energy Consumption in Mobile Phones: A Measurement Study and Implications for Network Applications, Proceedings of the 9th ACM SIGCOMM Conference on Internet Measurement Conference, Chicago, Illinois, USA
- [3] Chen, J.-J. & Thiele, L. (2008), Expected system energy consumption minimization in leakage-aware DVS systems, in 'ISLPED '08: Proceeding of the thirteenth international symposium on Low power electronics and design', ACM, New York, NY, USA, pp. 315–320.
- [4] Feeney, L. M. (2001), 'An Energy Consumption Model for Performance Analysis of Routing Protocols for Mobile Ad Hoc Networks', *Mobile Networks and Applications* 6(3), 239–249.
- [5] Gurun, S., Nagpurkar, P. & Zhao, B. Y. (2006), Energy consumption and conservation in mobile peer-to-peer systems, in 'MobiShare '06: Proceedings of the 1st international workshop on Decentralized resource sharing in mobile computing and networking', ACM, New York, NY, USA, pp. 18–23.
- [6] Höpfner, H., Schirmer, M., Bunse, C. (2012): On Measuring Smartphones' Software Energy Requirements, ICISOFT 2012 - Proceedings of the 7th International Conference on Software Paradigm Trends, Rome, Italy
- [7] Jain, R., Molnar, D. & Ramzan, Z. (2005): Towards understanding algorithmic factors affecting energy consumption: switching complexity, randomness, and preliminary experiments — Proceedings of the 2005 joint workshop on Foundations of mobile computing', ACM, New York, NY, USA, pp. 70–79.
- [8] Koc, H., Ozturk, O., Kandemir, M., Narayanan, S. H. K. & Ercanli, E. (2006), Minimizing energy consumption of banked memories using data recomputation, in 'ISLPED '06: Proceedings of the 2006 international symposium on Low power electronics and design', ACM, New York, NY, USA,
- [9] Potlapally, N. R., Ravi, S., Raghunathan, A. & Jha, N. K. (2006), 'A Study of the Energy Consumption Characteristics of Cryptographic Algorithms and Security Protocols', *IEEE Transactions on Mobile Computing* 5(2), 128–143.
- [10] Ruzzelli, Antonio G, Nicolas, C, Schoofs, Anthony, O'Hare, Gregory M. P. (2010): Real-Time Recognition and Profiling of Appliances through a Single Electricity Sensor, Proceedings of the Seventh Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks, SECON 2010, June 21-25, 2010, Boston, Massachusetts, USA
- [11] Seddik-Ghaleb, A., Ghamri-Doudane, Y. & Senouci, S.-M. (2006), A performance study of TCP variants in terms of energy consumption and average goodput within a static ad hoc environment, in 'IWCMC '06: Proceedings of the 2006 international conference on Wireless communications and mobile computing', ACM, New York, NY, USA, pp. 503–508.
- [12] Tuan, T., Kao, S., Rahman, A., Das, S. & Trimmerger, S. (2006), A 90nm low-power FPGA for battery-powered applications, in 'FPGA '06: Proceedings of the 2006 ACM/SIGDA 14th international symposium on Field programmable gate arrays', ACM, New York, NY, USA, pp. 3–11.
- [13] Wang, L., French, M., Davoodi, A. & Agarwal, D. (2006), 'FPGA dynamic power minimization through placement and routing constraints', *EURASIP Journal on Embedded Systems* 2006(1).
- [14] Wang, Y., Lin, J., Annavaram, M., Jacobson, Q.A., Hong, J., Krishnamachari, B., Sadeh, N. (2009): A Framework of Energy Efficient Mobile Sensing for Automatic User State Recognition, Proceedings of the 7th International Conference on Mobile Systems, Applications, and Services New York, NY, USA
- [15] Wilke, C., Richly, S., Püschel, G., Piechnick, C., Götz, S., Aßmann, U. (2012): Energy Labels for Mobile Applications, *Informatik 2012*, 42. Jahrestagung der Gesellschaft für Informatik e.V. (GI), 16.-21.09.2012, Braunschweig
- [16] Wilke, C. Richly, S., Götz, S., Assmann, U. (2013): Energy Profiling as a Service, *Informatik 2013*, 43. Jahrestagung der Gesellschaft für Informatik e.V. (GI), Informatik angepasst an Mensch, Organisation und Umwelt, 16.-20. September 2013,