

Towards Energy Efficient Smart Phone Applications: Energy Models for Offloading Tasks into the Cloud

Michele Segata*[†], Bastian Bloessl*, Christoph Sommer*, Falko Dressler*

*Computer and Communication Systems, Institute of Computer Science, University of Innsbruck, Austria

[†]Systems and Networks, Dept. of Information Engineering and Computer Science, University of Trento, Italy

{segata,bloessl,sommer,dressler}@ccs-labs.org

Abstract—Many people use smart phones on a daily basis, yet, their energy consumption is pretty high and the battery power lasts typically only for a single day. In the scope of the EnAct project, we investigate potential energy savings on smart phones by offloading computationally expensive tasks into the cloud. Obviously, also the wireless communication for uploading tasks requires energy. For that reason, it is crucial to understand the trade-off between energy consumption for wireless communication and local computation in order to assert that the overall power consumption is decreased. In this paper, we investigate the communications part of that trade-off. We conducted an extensive set of measurement experiments using typical smart phones. This is the first step towards the development of accurate energy models allowing to predict the energy required for offloading a given task. Our measurements include WiFi, 2G, and 3G networks as well as a set of two different devices. According to our findings, WiFi consumes by far the least energy per time unit, yet, this advantage seems to be due to its higher throughput and the implied shorter download time and not due to lower power consumption over time.

I. INTRODUCTION

Smart phones are about to completely change our daily life. These phones are used for a multitude of applications that go far beyond the initial use as a mobile phone. With the incredible success of smart phones, this technology evolved at a fast pace and is now enabling even more advanced applications including games, video, time and contact management, email, and of course all kinds of social media services [1]. In order to make this possible, state-of-the-art smart phones come with powerful multi-core CPUs and typically support several communication technologies like Bluetooth, WiFi, and different cellular technologies like GSM, UMTS, or LTE-Advanced.

Despite the great technological improvements, one problem remains: the current generation of smart phones is limited by their high energy consumption [2]. Users experience this weakness by the fact that a modern smart phone has to be charged almost every day. Energy is especially critical when it comes to computationally expensive tasks such as to record, edit, and publish videos on the phone [3]. Since the processors of smart phones are already very energy efficient, it is worth investigating other possibilities to save energy.

The EnAct project¹ investigates exactly that. The main idea of the project is to offload computationally expensive tasks to the cloud instead of executing them locally. So instead

of spending energy for the local CPU to compute the task, potentially less energy is used to transfer a certain amount of data into the cloud and to download the result. Technically speaking, this is a remote procedure call where the function arguments and the results have to be transmitted over a wireless connection.

In this context, we have to distinguish two possible optimization objectives: First, to minimize the power consumption of the mobile phone, i.e., to maximize battery life time, and secondly, to decrease the overall power consumption [4], [5]. For the second optimization objective, also the energy that servers in the cloud consumes when performing the task needs to be included. Even though first intuition might not consider this approach reasonable, it is indeed possible to save energy since the computers used in the cloud are more powerful and efficient compared to the CPUs of typical smart phones. Therefore, the servers in the cloud might allow computing a given task with less energy. This optimization objective saves energy and thus natural resources as well as decreasing the CO₂ emission [6].

Independent from the actual optimization goal, one obviously has to balance the trade-off between the energy required to compute a task locally and the energy that has to be spent for the wireless transmission to and from the cloud. We investigated the second part of this trade-off, i.e., the power consumption of wireless communication. The goal is to come up with a model that allows predicting the energy consumption given the currently available communication technologies like WiFi, 2G, or 3G, the connection quality like Signal to Noise Ratio (SNR), and the estimated size of data to be transmitted. Together with the prediction of the energy consumption for computing a certain task locally and remotely, this is the main building block for a new generation of energy efficient and, most importantly energy-aware applications.

Our main contributions can be summarized as follows:

- We investigate the energy consumption of smart phones for data transmission using WiFi, 2G, and 3G connections. For this, we developed a measurement methodology and performed an extensive set of experiments using different phones and network connections.
- Based on our findings, we developed an energy consumption model that can be used in the decision process whether to offload computational expensive tasks to the cloud or performing them locally.

¹Energy Aware Computing, <http://www.en-act.eu/>

II. RELATED WORK

Fundamentally, the philosophy of the EnAct project is the same as in SmartDiet [7]: to exploit the cloud to reduce energy consumption. In particular, the idea is to offload computationally expensive operations to the cloud in order to save energy. One discriminating factor in the decision of whether a particular task should be offloaded or not is the energy consumed to transfer the data to be processed to the network, and to transfer the results back to the phone. Clearly, if the energy consumed for this transfer is higher than the energy required by the CPU to compute the result, offloading is not beneficial.

There have been several different approaches to measure the power consumption of different kinds of communication in smart phones. They differ not only in how the power consumption has been measured (e.g., battery power level vs. power meter measurements), but also in which different aspects of the communication have been considered (e.g., transfer time, data amount, traffic patterns, SNR, and network load).

In [8] energy measurements of WiFi and UMTS data transfers have been published which take different load situations as well as different SNR levels for both technologies into account. Besides the fact that WiFi is more energy efficient when transferring large amounts of data, the measurements also show a significant difference in power consumption under different channel load and SNR levels for UMTS.

The energy measurements in [5] concentrate on the analysis of Round Trip Time (RTT) related effects that are experienced in WiFi and 3G networks. For small amounts of data (10 kB and 100 kB) the consumed energy is much higher using 3G. This is explained by the much higher RTT of around 220 ms for 3G compared to 25 ms for WiFi. The authors also mention that they experienced unexpected energy consumption of WiFi when using the power-save mode. This can happen in low latency environments, where the RTT is in the order of the power saving interval. In this case, the variance induced by power-saving modes can cause TCP retransmissions, increasing the overall energy consumption.

In [9] the authors perform a detailed measurement campaign on the Openmoko Neo Freerunner, the HTC Dream, and the Google Nexus One. The authors propose a model to estimate power consumption for the Neo Freerunner for each usage scenario as a function of time. For our aims, this is not enough, as we are interested in determining the power consumption as a function of the number of data bytes to transfer, either via WiFi, or via a cellular network.

In [3], the authors focus on the impact of network conditions on the energy consumption of a mobile device while performing video delivery. In particular, they consider different communication scenarios, i.e., with and without background traffic in the network, close to and far from the access point, to obtain different signal strengths. Moreover, they also consider different video encodings and different transport layer protocols (i.e., UDP and TCP). The results presented in the paper show that the load in the wireless network can have a huge impact

on the overall consumption, as the competition for the channel increases. Furthermore, it is shown that being far from an access point (i.e., having a low signal strength), increases the power consumption and that, in general, TCP is less energy consuming than UDP. The paper does not develop an actual model, but highlights factors that have a huge impact on the energy consumption and thus, should be taken into account when estimating the energy consumption.

In [10] a measurement study of GSM, 3G, and WiFi is performed in order to develop an energy consumption model for the three technologies. This model is then used to develop a protocol which aims to achieve highly energy efficient data transfers. The findings of the authors about cellular technologies show that the estimation of energy consumed by a data transfer has to consider the state machine of the communication protocol. Both 3G and GSM can be characterized by three states, a high power state, a low power state, and a idle state. The energy required by the data transfer is not only the one consumed by the transfer itself, but also by the time the phone remains in the high-power and low-power states before switching to idle. Even though such states do not exist in WiFi, the authors argue that if the phone is not associated to any network, the energy consumption model should also take into account the scan and the association procedure. In this paper, however, we assume that in order to consider a technology as a candidate for data transfer, connection must be already established. We thus derive a model which estimates only the amount of energy required by the transfer itself.

Similarly, in [11] a comparison of the three wireless technologies is presented, deriving per-bit energy consumption for different packet sizes, and for different situations (i.e., sending and receiving). The paper gives some fundamental insights and highlights the constant overhead of data transmissions, leading to lower energy consumption of bursty traffic, but it does not develop an energy consumption model.

In [12] an accurate consumption model is derived which is able to predict the energy consumption for different components of the smart phone, e.g., CPU, display, WiFi card. Such model was then implemented into a freely available application, called PowerTutor. The consumption model has been derived from two HTC phones (Dream and Magic), however, it is not clear whether it can also suit other phones. A more general investigation is needed.

The work presented in [13] is most closely related to ours. Here, a data transmission power model for IEEE 802.11g is developed, i.e., the paper focuses only on WiFi. The model takes into account average data rate, internet flow characteristics, different WiFi states, and power saving modes. The goal of the model is to predict the average energy consumption for different data rates (in terms of TCP throughput). Due to this it is more oriented to typical internet traffic flows, while we are interested on modeling consumption needed to send a specific amount of data from the phone to the access point, and vice versa.

III. MEASUREMENT ENVIRONMENT

The first step towards energy measurements in the lab is to setup and validate a measurement environment. In this section we give an overview over the methodology, the utilized devices, the validation of our measurement setup and present an Android application to conduct reproducible measurements. The application can automatically initiate file transmissions to avoid human interactions with the phone while performing the measurements.

Regarding measurement methodology, due to the lack of functionality of the Android OS, the only possible way is by connecting a physical measurement device to the smart phone. The most common way is by connecting a volt-meter and an ampere-meter between the battery and the phone [3], [5], [9]–[12], [14], [15]. When Android APIs are used, instead, the only metric that can be collected is the percentage of battery consumed [8]. In some cases it is possible to get information about each single hardware component [9], or to get information from an application developed to this specific purpose [10]. This, however, depends on the software and the hardware of the phone. For example, in [9] the analysis of each single hardware component was possible because the circuit schematics are openly available. In [10] one of the phones under analysis was a Nokia N95, for which an energy profiling application is available.

A. Methodology

The relationship between the physical quantities of interest is

$$E(t) = \int_0^t P(t)dt = \int_0^t U(t) \cdot I(t)dt,$$

where $E(t)$, $P(t)$, $U(t)$, and $I(t)$ represent energy, power, voltage, and current at time t respectively. The measured quantities from which the energy is derived are I and U . A photograph of the assembly with a Samsung Galaxy S2 is shown in Figure 1. The straightforward schematic of our measurement circuit is depicted in Figure 2. As shown in the figure, we perform the measurements by intercepting the current flow between the battery and the mobile phone.

B. Measurement Device

In theory, we would record $U(t)$ and $I(t)$ and integrate their product over time to get the accumulated energy consumption $E(t)$ at time t . In practice, we measure a time-discrete, sampled version of the process. For our measurements we use the Voltech PM1000+. This device can be connected to a computer to record the measurements. The accuracy of the samples is very high, however, the PM1000+ supports only sampling frequencies of up to 1 Hz.

Given the coarse time resolution of the samples we checked how the power samples are actually recorded by the meter. Basically, two sampling strategies can be employed. The first one is sample-and-hold, where an instantaneous value is measured, whereas with the second strategy an average over the whole sampling period is recorded. The sample-and-hold method is hardly usable in our context since the underlying

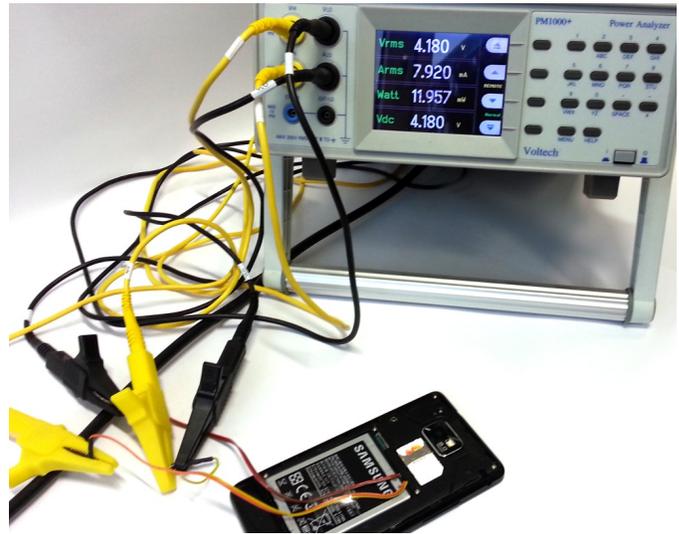


Figure 1. Photograph of the measurement setup: we connect a Voltech PM1000+ to directly intercept current from (and measure voltage at) the battery of the phones.

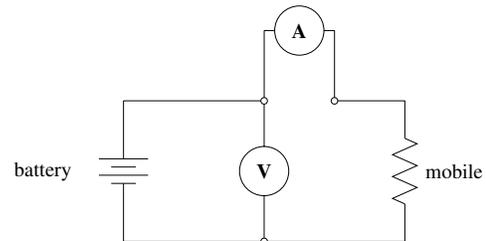


Figure 2. Schematic of the measurement setup. We intercept the current flow between the battery and the mobile phone.

technologies and protocols operate some orders of magnitude faster than the sampling interval. Therefore, a huge amount of repetitions would be needed in order to get significant results.

The second method, which calculates the average power over the sampling interval is much preferred, since we are mainly interested in the overall energy consumption, which is in turn completely reflected by the averaged samples.

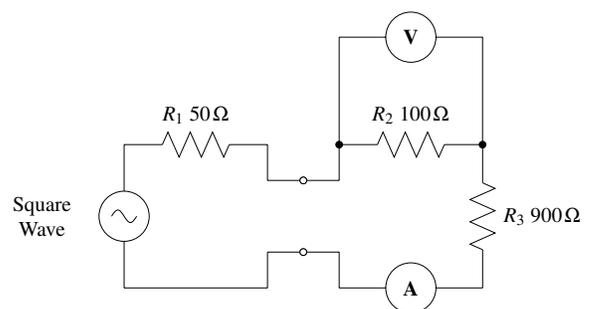


Figure 3. Schematic of the circuit that we used to determine the sampling strategy of the PM1000+.

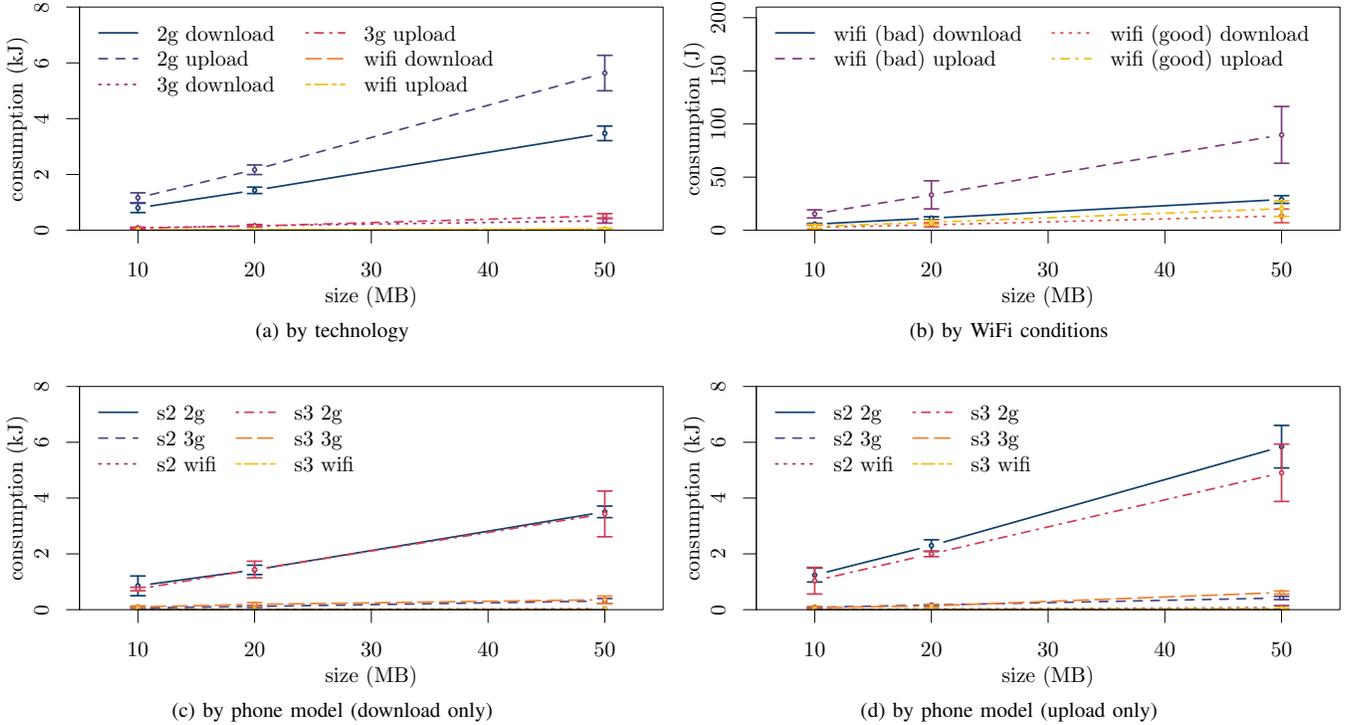


Figure 4. Energy consumed as function of transfer size.

C. Experimental Determination of the Sampling Strategy

In order to determine which sampling strategy is used by the meter, we set up another experiment and connected the Metrix MTX 3240 signal generator. A schematic of this experiment is depicted in Figure 3. The signal generator, which has an internal resistance of $50\ \Omega$ (R_1), has been connected to a $100\ \Omega$ resistor (R_2), where the voltage was measured. In addition, we had to add a $900\ \Omega$ resistor (R_3) due to the limited output power that the signal generator can provide.

To actually determine the employed sampling method, we generated a square wave with the signal generator and recorded the values with the PM1000+. The general idea is that if sample-and-hold is used, we would measure a random power value, while if the values are averaged over the sample interval we should record a constant value. By repeating the experiment with several different square wave frequencies, we discovered that the averaging method is used, fitting our purposes.

D. Measurement Application

In order to conduct reproducible measurements we automated the process so that no interaction with the phone is required during the course of the measurements. This is desirable since any interaction with the phone puts it out of power saving mode and switches the display on, polluting energy traces. For this reason we developed an Android application that allows us to define a measurement campaign in a simple textual format. The application requires a wake lock to avoid being interrupted when the phone switches to standby mode.

To define a measurement campaign we can specify a sequence of files that has to be downloaded or uploaded, and delays that we can insert between each transfer. These delays are important especially in measurements involving 3G networks, where the mobile might stay in a high power state even after data transfer is complete. This can be the case since the phone can maintain an active, even though idle, connection to its base station and remain in a power control loop that requires frequent, energy consuming signaling. This signaling is a technology induced energy overhead that has to be considered when taking the measurements.

IV. RESULTS

The measurement campaign has been conducted with Samsung Galaxy S2 and S3 smart phones, running Android v4.0.3 and v4.1.2 respectively. The phones were downloading and uploading files of different sizes from and to our web server, using different wireless technologies. The file sizes used were 10 MB, 20 MB and 50 MB, while the transmission technologies employed were 2G (EDGE), 3G (UMTS), and WiFi.

In order to understand whether different network conditions significantly affect energy consumption, we performed WiFi measurements on a dedicated access point (perfect link quality and small RTT) and on the university network. The university network is typically very crowded, in the sense that there are several other users associated that compete for channel access and cause cross traffic in the wireless and wired network. For cellular networks instead, we performed repeated measurements on different days and at different times, as we had no control

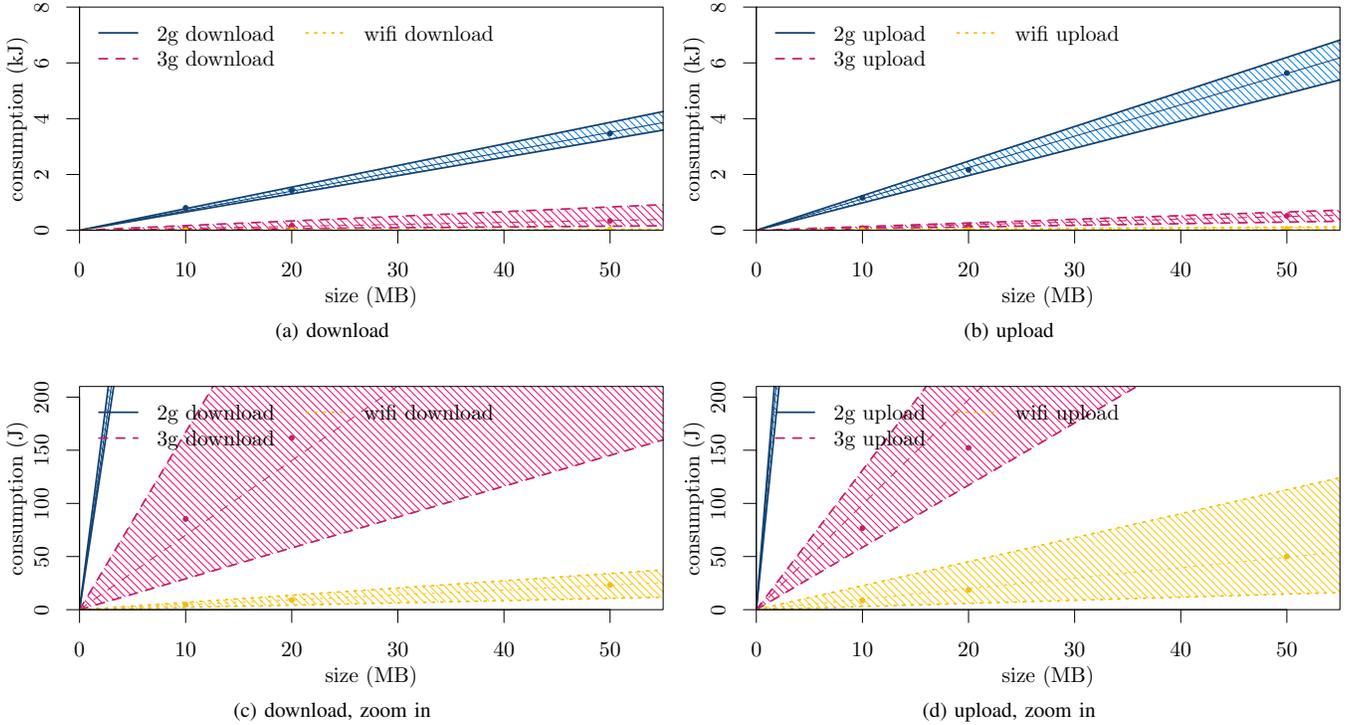


Figure 5. Energy consumption model for download and upload as function of transfer size. The shaded areas illustrate the span from β_{10} to β_{90} , the center lines correspond to the prediction by β_{avg} , the sample means for comparison are shown as dots.

on network quality.

As first result, we computed the energy required to transfer the data by multiplying the average power during the transfer by the transfer time. The average values with 95% confidence intervals are shown with different scalings in Figures 4a and 4b. Figure 4b shows a zoom-in of Figure 4a focusing only on WiFi, for better readability. In these plots we do not distinguish between different mobile phones.

First, it can be seen that 2G is by far the most consuming technology among the analyzed ones, followed by 3G, and then WiFi. This is in contrast to [10] where the authors report a better performance for 2G networks compared to 3G. This can be due to the fact that authors consider smaller file sizes (i.e., up to 500 kB) compared to what we are interested in. Even though cellular networks are hard to compare since their behavior heavily depends on the load, the link quality, and even on the disclosed configuration of the provider, we nevertheless consider the difference between 2G and 3G significant to state that 3G offers indeed better energy efficiency for considered transfer sizes.

Second, uploading is more expensive than downloading for cellular networks. This holds for WiFi too, as shown in Figure 4b, but the difference is not so pronounced. Moreover, WiFi link quality seems not to have a major impact on overall energy consumption, as the difference between good and bad WiFi links is smaller compared to the difference between WiFi and cellular networks. For this reason we keep good and bad link measurements merged together when deriving the

consumption model in Section V.

Another interesting aspect is how different phone models compare in terms of energy consumption. This is fundamental, since if completely different results are obtained, it will not be possible to develop a device-independent consumption model. For this reason, we plot the energy consumption as function of transfer size, split by mobile phone model, for download (Figure 4c) and for upload (Figure 4d).

The figures shows that, both for download and upload, no significant difference can be seen. This result, however, might be biased due to the fact that the phones are developed by the same manufacturer (but representing two different generations of smart phones), so we plan to perform measurements with other devices in order to further confirm this hypothesis.

V. CONSUMPTION MODEL

Based on the measurement results we now aim to obtain an energy consumption model that could give us an indication of the expected energy consumption for uploading or downloading a given amount of data using one of the wireless communication technologies available. This metric then helps deciding whether offloading the task into the cloud is a reasonable alternative.

Since the energy consumption highly depends on current channel conditions, environmental factors, and even disclosed network configurations of the provider, the consumption suffers from a high variability, which should be captured by the model. Besides the mean energy, the model should therefore also provide an indication of the range where the energy

Table I
SLOPE PARAMETERS (IN J/MB) FOR THE ENERGY CONSUMPTION MODEL.

	β_{10}	β_{avg}	β_{90}
download WiFi	0.21	0.46	0.68
download 3G	2.90	7.01	16.69
download 2G	65.28	70.17	77.36
upload WiFi	0.29	0.99	2.26
upload 3G	5.86	9.88	13.16
upload 2G	97.99	112.35	123.93

consumption would likely fall into. Further, the model should be as simple as possible, and it should capture the fact that the device is always connected, i.e., no energy needs to be spent for establishing a channel.

Seeing that the measurement data exhibits a pronounced linear relationship between amount of data and energy consumption – and taking into account that no energy consumption should be forecast for not transferring anything – we opted for a simple linear model with no intercept

$$y[J] = \beta x[\text{MB}], \quad (1)$$

using both linear (to obtain β_{avg} , the average energy consumption per MB) and quantile regression (to obtain β_{10} and β_{90} , the 10th and 90th percentile) for fitting the parameters. The advantage of this quantile regression over just providing the standard deviation, for example, is that it has the straightforward interpretation that 80% of the outcomes fall in the given interval.

We give the resulting slope parameters in Table I and illustrate the fit between model and measurement results in Figure 5, plotting the span from fitted 10th to 90th percentile and the fitted mean, along with the sample mean of our measurement results.

VI. CONCLUSION AND FUTURE WORK

In this paper, we presented a measurement methodology for the energy consumption of smart phones. Using the energy model we derived from an extensive set of measurements, we are now able to determine whether offloading a computationally expensive task into the cloud eventually helps saving energy on the phone. In particular, we compared the energy required to download/upload a certain amount of data using different communication technologies, namely WiFi, 2G, and 3G. We have shown that WiFi is the most energy-efficient technology, followed by 3G and 2G, and that uploading is in general more demanding than downloading, especially for cellular networks. We have also shown that different mobile devices are comparable in terms of energy consumption. Based on our measurement results, we developed a linear model which estimates the energy required to download or upload a certain amount of data, given a particular transmission technology. This model could then be employed to perform decisions on whether to offload a computational task to the cloud or not, depending on how much energy such task requires to be performed on the phone.

The main advantage of the model is that it does not depend on any network parameters like link quality or network load. However, the drawback might be that it loses precision given that these factors are not taken into account. A direction for future research could therefore be to investigate the question if current network conditions should be considered and if so which network metrics have to be monitored by the mobile device.

REFERENCES

- [1] P. Zheng and L. Ni, *Smart Phone and Next Generation Mobile Computing*, Morgan Kaufmann, 2006.
- [2] G. P. Perrucci, F. H. P. Fitzek, and J. Widmer, “Survey on Energy Consumption Entities on the Smartphone Platform,” in *75th IEEE Vehicular Technology Conference Fall (VTC2011-Spring)*. Budapest, Hungary: IEEE, May 2011, pp. 1–6.
- [3] R. Trestian, A.-N. Moldovan, O. Ormond, and G.-M. Muntean, “Energy consumption analysis of video streaming to Android mobile devices,” in *18th IEEE/IFIP Network Operations & Management Symposium (NOMS 2012)*. Maui, HI: IEEE, April 2012, pp. 444–452.
- [4] A. Berl, E. Gelenbe, M. Di Girolamo, G. Giuliani, H. De Meer, M. Q. Dang, and K. Pentikousis, “Energy-Efficient Cloud Computing,” *The Computer Journal*, vol. 53, no. 7, pp. 1045–1051, 2010.
- [5] E. Cuervo, A. Balasubramanian, D.-k. Cho, A. Wolman, S. Saroiu, R. Chandra, and P. Bahl, “MAUI: Making Smartphones Last Longer with Code Offload,” in *8th International Conference on Mobile Systems, Applications and Services (MobiSys 2010)*. San Francisco, CA: ACM, June 2010, pp. 49–62.
- [6] G. Fettweis and E. Zimmermann, “ICT Energy Consumption - Trends and challenges,” in *11th International Symposium on Wireless Personal Multimedia Communications (WPMC’08)*, Lapland, Finland, September 2008.
- [7] A. Saarinen, M. Siekkinen, Y. Xiao, J. K. Nurminen, M. Kempainen, and P. Hui, “SmartDiet: offloading popular apps to save energy,” *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 42, no. 4, pp. 297–298, August 2012.
- [8] H. Petander, “Energy-Aware Network Selection Using Traffic Estimation,” in *1st ACM Workshop on Mobile Internet Through Cellular Networks (MICNET 2009)*. Beijing, China: ACM, September 2009, pp. 55–60.
- [9] A. Carroll and G. Heiser, “An Analysis of Power Consumption in a Smartphone,” in *USENIX Annual Technical Conference 2010 (USENIX ATC 2010)*. Boston, MA: USENIX Association, June 2010, pp. 21–21.
- [10] N. Balasubramanian, A. Balasubramanian, and A. Venkataramani, “Energy Consumption in Mobile Phones: A Measurement Study and Implications for Network Applications,” in *9th ACM SIGCOMM Conference on Internet Measurement (IMC 2009)*. Chicago, IL: ACM, November 2009, pp. 280–293.
- [11] L. Wang and J. Manner, “Energy Consumption Analysis of WLAN, 2G and 3G interfaces,” in *IEEE/ACM Int’l Conference on Green Computing and Communications & Int’l Conference on Cyber, Physical and Social Computing (GreenCom-CPSCOM 2010)*. Hangzhou, China: IEEE, December 2010, pp. 300–307.
- [12] L. Zhang, B. Tiwana, R. Dick, Z. Qian, Z. Mao, Z. Wang, and L. Yang, “Accurate Online Power Estimation and Automatic Battery Behavior Based Power Model Generation for Smartphones,” in *IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS 2010)*. Scottsdale, AZ: IEEE, October 2010, pp. 105–114.
- [13] Y. Xiao, P. Savolainen, A. Karppanen, M. Siekkinen, and A. Ylä-Jääski, “Practical Power Modeling of Data Transmission over 802.11g for Wireless Applications,” in *1st International Conference on Energy-Efficient Computing and Networking (e-Energy 2010)*. Passau, Germany: ACM, April 2010, pp. 75–84.
- [14] A. Rice and S. Hay, “Decomposing Power Measurements for Mobile Devices,” in *IEEE International Conference on Pervasive Computing and Communications (PerCom 2010)*. Mannheim, Germany: IEEE, April 2010, pp. 70–78.
- [15] A. Schulman, T. Schmid, P. Dutta, and N. Spring, “Phone Power Monitoring with BattOr,” in *17th ACM International Conference on Mobile Computing and Networking (MobiCom 2011), Poster Session*. Las Vegas, NV: IEEE, September 2011.