

EnerSim: An Energy Consumption Model for Large-Scale Overlay Simulators

Christian Gross¹, Fabian Kaup, Dominik Stingl¹, Björn Richerzhagen,
David Hausheer¹, Ralf Steinmetz

Multimedia Communications Lab, Peer-to-Peer Systems Engineering,
Technische Universität Darmstadt

Email: {chrgross, fkaup, stingl, bjoern.richerzhagen,
hausheer, steinmetz}@kom.tu-darmstadt.de

Abstract—Determining the energy consumption in large-scale overlay simulations is still an open issue as most existing simulation frameworks are agnostic to that aspect. Especially simulations including mobile devices, such as smartphones or tablet PCs, can benefit from having a energy consumption model in simulations such that newly developed large-scale overlay mechanisms can be evaluated with respect to their energy consumption on mobile devices. Therefore, this paper proposes a component-based energy consumption model, which is based on measurements of existing smartphones. The model causes little computational overhead, thus, being suitable for large-scale simulations. A brief evaluation shows that using our model, the energy consumption can be estimated with a mean error of $\pm 4.7\%$. Furthermore, the measurements conducted to derive the model show that WiFi and Bluetooth communication are one, respectively two, orders of magnitude more energy efficient than cellular communication.

Index Terms—Energy Efficiency, Energy Consumption Model, Overlay Simulations, Measurements

I. INTRODUCTION

Energy efficiency of today's communication is becoming more and more important. According to various studies the energy consumption of today's information and telecommunication (ICT) varies between 2 to 10% of the world-wide energy consumption [7], [10]. Furthermore, it is said that mobile traffic is growing at a rate of 92% per year, and will constitute 7.5% of the world's global IP traffic by 2015 [8]. Hence, the development of energy efficient mobile communication is of major importance. However, to enable an energy efficient communication, newly developed mechanisms and protocols have to be evaluated with respect to their energy consumption already at the design stage using testbeds or simulations. While the measurement of energy consumption in testbeds can easily be done, existing simulation frameworks for overlay networks [2], [6], [13] are agnostic to that aspect.

To overcome this problem, this paper presents the following contributions: (i) The power consumption of the Bluetooth, WiFi, and 3G communication of two widely deployed smartphones has been measured. The

measurement results show that WiFi and Bluetooth communication are one, respectively two, orders of magnitude more energy-efficient than cellular communication with respect to the energy spent per transmitted byte. (ii) Based on our measurement results, a component-based energy consumption model is derived, which can be applied in discrete event-based simulations, thereby enabling the simulative evaluation of mobile network communication. Although focusing on mobile communication, the proposed energy consumption model is easily extensible, thus, it can also be used to model the energy consumption of other types of communication, e.g., wired communication.

II. MEASUREMENT METHODOLOGY

For measuring the power consumption of smartphones a sense resistor without a pre-amplifier has been used as shown in Figure 1, which is introduced between the battery and the device, causing a voltage drop proportional to the current to be measured. According to




Figure 1. Schematic of the measurement circuit.

Zhang et al. [16] the maximum voltage drop should be chosen as 30 mV in order to ensure a high measurement accuracy. For measuring the voltage drop, a 12 bit A/D conversion card (Meilhaus 1208LS) was used. The card allows analog measurements on up to four channels in differential mode. As the maximum resolution of the card in the $\pm 5V$ range is $\frac{10V}{2^{12}} \approx 3$ mV, the voltage drop must stay in the range between $\Delta U_{\min} = 6$ mV and $\Delta U_{\max} = 30$ mV. Today's smartphones consume currents in the range of 1 mA up to 350 mA. Therefore, at least two different resistors are needed to measure the power consumption in idle mode and under heavy load accurately. Each resistor is connected to the battery via a jumper in order to enable an online switching

¹This work has been [co-]funded by the German Research Foundation (DFG) in the Collaborative Research Center (SFB) 1053 "MAKI – Multi-Mechanism-Adaptation for the Future Internet.

between the different resistors without disconnecting the smartphone from the battery. Furthermore, a direct connection to the battery without resistor was added to the circuit to enable the boot process. Booting the smartphone while being connected to a resistor with higher resistance causes a high voltage drop and, thus, the boot process would fail as the voltage at the device is too small.

In order to determine the values for the different resistors R_1 , R_2 , and R_3 , the range of the current drawn by the smartphone has to be identified. The minimum current draw I_{min} and, thus, the maximum battery lifetime T_{max} can be observed during the standby phases of the smartphone. The maximum current draw leads to the minimum lifetime T_{min} , which can be observed under heavy usage, assuming all components running in high power state. Knowing the battery capacity W_{batt} , the nominal battery voltage U_{nom} , and the minimum and maximum battery lifetime T_{min} and T_{max} , the average current I_{avg} drawn in the two scenarios can be computed based on Equation 1.

$$I_{avg} = \frac{W_{batt}}{U_{nom} \cdot T_{max}} \quad (1)$$

Based on the average current I_{avg} , and the desired range for the voltage drop $[\Delta U_{min}, \Delta U_{max}]$, the resistor range can be computed following Equation 2,

$$[R_{min}, R_{max}] = \frac{\Delta U}{I_{avg}} \quad (2)$$

with $\Delta U \in [\Delta U_{min}, \Delta U_{max}]$ and $I_{avg} \in [I_{min}, I_{max}]$. The parameters to determine the values of R_1 , R_2 , and R_3 in the experiments are shown in Table I. From the

Parameter	Symbol	Value
Battery Capacity	W_{batt}	22,680 Ws
Nom. Batt. Voltage	U_{nom}	4.1 V
Min. Batt. Lifetime	T_{min}	14,400 s
Max. Batt. Lifetime	T_{max}	777,600 s
Current Range	$[I_{min}, I_{max}]$	[6.5 mA, 350 mA]
Resistor Range	$[R_{min}, R_{max}]$	[0.085 Ω , 1 Ω]

Table I
PARAMETER LISTING

calculated resistors range the following high precision resistors have been chosen: $R_1 = 0.05 \Omega$ ($\epsilon = 0.5\%$), $R_2 = 0.1 \Omega$ ($\epsilon = 0.5\%$), $R_3 = 1.0 \Omega$ ($\epsilon = 1\%$).

A. Measuring the Power Consumption

In order to determine the power consumption of the different components, the power consumption of the particular component has been isolated as suggested by Rice et al. [12]. First, the baseline power consumption of the device is measured right before enabling a particular component. Subsequently, the measured baseline power consumption is subtracted from the values obtained while having the component under study enabled. Each networking element has a set of dedicated power consumption states representing its level of utilization, e.g.,

idle, low power, high power. Thus, an artificial workload has to be generated, which utilizes the component under study accordingly. While this can be done easily for components such as the display by adjusting the level of brightness, generating a dedicated workload for the network devices requires additional overhead. Therefore, an Android app is developed that transmits and receives a predefined amount of data, either via the cellular network, via Bluetooth, or via WiFi in infrastructure or ad hoc mode.

To obtain statistically significant results, each measurement is repeated ten times. Out of those values the 5th and 95th confidence intervals are computed. For sampling the power consumption a sampling rate of $300 \frac{1}{s}$ is used forming the set of measurement samples S . Out of those samples the energy consumption can be calculated following Equation 3.

$$P_{avg} = \frac{1}{|S|} \sum_{s \in S} P(s) = V_1(s) * \frac{V_2(s)}{R} \quad (3)$$

The A/D measurement card has an accuracy of 0.2% in the 5 V range or ± 10 mV. In addition, the resistors used have a limited accuracy of $\pm 1\%$. Hence, due to error propagation the error in the measured power consumption ΔP deviates and can be calculated as

$$\Delta P = \left| \frac{\partial P}{\partial V_1} \right| \Delta V_1 + \left| \frac{\partial P}{\partial V_2} \right| \Delta V_2 + \left| \frac{\partial P}{\partial R} \right| \Delta R. \quad (4)$$

with ΔV_1 , ΔV_2 , and ΔR denoting the maximum error of the particular parameter. For the measurement with $V_1 = 4.1$ V, $V_2 = 20$ mV this results in a deviation $\Delta P = 73$ mW using R_1 , $\Delta P = 37$ mW using R_2 , and $\Delta P = 4$ mW using R_3 . Depending on the measured power consumption (idle, low, or high power consumption) the resistor that causes the smallest measurement error is selected. All measured values are averaged over a large number of 30.000 samples to obtain statistically significant values.

III. MEASUREMENT RESULTS

The measurement methodology has been applied on two widely used smartphones: the Motorola Milestone 2 and the Google Nexus One. The results for the power consumption of the different components of both devices in idle mode are shown in Figure 2(a). Android devices distinguish between two different passive modes where the device consumes less power: the *suspend* and the *idle* mode [3]. In the suspend mode, all running applications are suspended to RAM, whereas in idle mode, background processing still takes place. Therefore, both devices consume twice as much power in the idle state compared to the suspended mode. Furthermore, when looking at the idle state power consumption of the networking devices, the components of the Milestone 2 device consume significantly more power than the Nexus One. The Bluetooth device on the Milestone 2 consumes almost twice as much power in the idle state than the

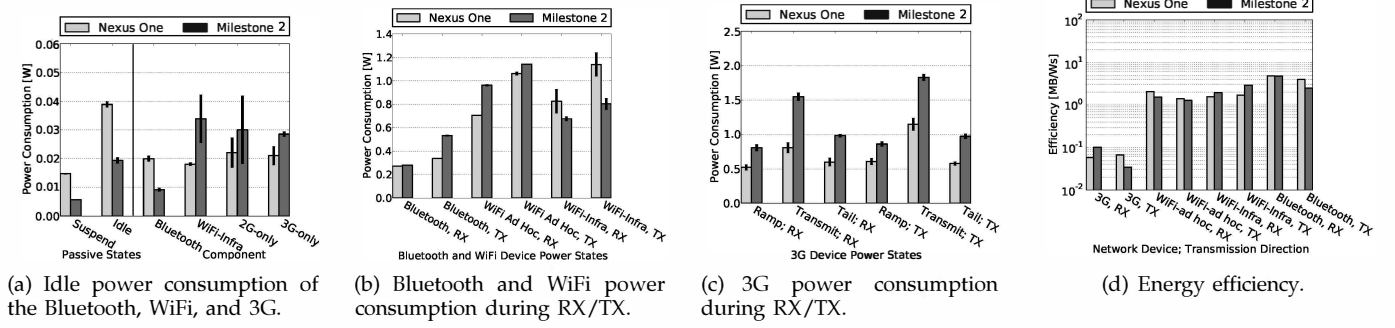


Figure 2. Power consumption of the network devices and resulting energy efficiency.

Nexus One and is as high as the Milestone 2 WiFi idle power consumption. The idle power consumption of the cellular network device in 2G mode is similar to the power consumption in 3G mode.

Figure 2(b) shows the power consumption of the Bluetooth device and WiFi device in ad hoc and infrastructure mode. Using the ad hoc mode, both devices consume slightly more power than in the infrastructure mode. Furthermore, sending of data causes a higher power consumption than receiving. Bluetooth, on the other hand, consumes less than half of the power of WiFi. Nevertheless, Bluetooth has a transmission range of about 10 m with a small bandwidth of 1 Mbit/s, whereas WiFi easily can bridge more than 100 m offering a bandwidth of up to 54 Mbit/s.

The 3G network device exhibits an additional set of power consumption states besides the high and low power state. When starting a transmission, the network device first switches to the *ramp* state for about 2 s followed by the *high* power state. After finishing a transmission, the device switches to the *tail* state, waiting for possible further transmissions to come in. After a period of 5 s without any further data being received or transmitted, the device again switches to the idle power state. The Milestone 2 consumes significantly more power in all of those states in comparison to the Nexus One as shown in Figure 2(c). The power consumption in the ramp and the tail state for each device is almost identical. Furthermore, similar to the WiFi device, sending data via the 3G device causes a higher power consumption than receiving.

Figure 2(d) shows the resulting energy efficiency defined as the energy spent per transmitted or received byte. Cellular communication is one magnitude less efficient than WiFi communication. Furthermore, WiFi infrastructure-based communication is slightly more efficient than WiFi ad hoc communication, due to the lack of proper power saving mechanisms for the ad hoc mode [4]. Finally, Bluetooth outperforms both WiFi and 3G communication and is up to two orders of magnitude more energy efficient.

IV. ENERGY CONSUMPTION MODEL

Based on the measurements presented in Section III, an energy consumption model is derived, which allows for the simulation of the energy consumption of mobile communication in overlay network simulators. The presented model is based on the following assumptions: (i) *Battery-powered devices* with a fixed battery capacity are assumed. (ii) As each device consists of a set of components, e.g., CPU, display, RAM, network devices, a *component-based model* is assumed. (iii) Each component of a smartphone has a finite amount of *power consumption states* P_i . (iv) The power consumption model is executed in a *discrete event-based simulation environment*, where the consumed energy is subtracted from the remaining capacity of the battery in discrete points of time and not continuously. (v) As the model is designed for overlay network simulators, message-oriented instead of packet-based communication is assumed.

To be applicable in large-scale simulations, a simulation model has to fulfill a variety of requirements: (i) Large-scale simulations usually require a large amount of computational resources. Hence, an energy simulation model must have a *low computational complexity* in order to be scalable. (ii) As the components of smartphones are constantly changing over time and new components are added to smartphones, the simulation model should be *expandable*. (iii) The simulation model should have a *modular design* allowing to toggle component state in an easily configurable way. (iv) The simulation model should approximate the power consumption of mobile communication of smartphones with a reasonable accuracy.

The energy consumption model consists of two main parts: The *battery model*, which reflects the remaining capacity of a device and the *power consumption model*, which describes how much energy is spent within a certain amount of time. The power consumption model considers a set of components C with each component $c_i \in C$ having a set of dedicated power consumption states $P_i \in \{P_{idle}, P_{low}, P_{ramp}, P_{high}, P_{tail}\}$. Those power states reflect the typical operation states of a component (off, idle, low, and high power). Whenever a component changes its power state, the amount of energy consumed

since the last state change is calculated and subtracted from the remaining capacity of the battery. The consumed energy W_i of a component $c_i \in C$ between two state changes is calculated following Equation 5,

$$W_i(t_{n-1}, t_n) = \int_{t_{n-1}}^{t_n} P_i(\tau) d\tau = P_i \cdot (t_n - t_{n-1}) \quad (5)$$

with $t_n - t_{n-1}$ denoting the time the component c_i spent in a certain power consumption state P_i . To deal with long times between two state changes, a threshold t_{max} is defined after which the consumed energy is subtracted from the remaining capacity of the battery.

The battery model describes the remaining energy of the battery. For the simulation model a simple linear battery model has been applied, which follows Equation 6 for computing the remaining energy $W(t)$ of the battery.

$$W(t_n) = W(t_{n-1}) - W_i(t_{n-1}, t_n) \quad (6)$$

Non-linear battery models [5], [9], [11] consider typical non-linear characteristics of Li-Ion batteries such as recovery and rate capacity effects, cycle aging effects, and temperature dependencies, but cause larger computational overhead. Since the focus of this paper is on the power consumption of network communication and other involved components of mobile devices rather than the exact remaining battery level, these approaches have not been considered for the presented simulation model.

V. MODEL VALIDATION

The power models are evaluated by calculating the power consumption of the device as given in Section IV, where several components are active at the same time and comparing this to the power consumption measured while the device is in the specific state. This is conducted for a selection of power states to verify the model. The power states evaluated and the respective errors are given in Table II. The power state is defined as: ad hoc enabled (AH), infrastructure (IS), duplex transmission (DX), and the screen states. The comparison of the model with the measured values are consistent. The average error of the measured states is 4.7% and hence, is in a range acceptable for the use in the simulator.

Device	State	Model	Measured	Error
N1	WiFi AH, Idle, Scr. On	1.246 W	1.213 W	2.7%
M2	WiFi AH, Idle, Scr. On	0.754 W	0.738 W	2.1%
N1	WiFi AH, DX, Scr. On	1.945 W	2.078 W	-6.4%
M2	WiFi AH, DX, Scr. On	1.482 W	1.528 W	-3%
N1	WiFi IS, Idle, Scr. Med.	0.838 W	0.762 W	10%
M2	WiFi IS, Idle, Scr. Med.	0.346 W	0.332 W	4.2%

Table II
COMPARISON OF THE MODEL AND THE MEASURED VALUES

VI. RELATED WORK

Related work has measured and optimized the power consumption of recent smartphones [1], [3], [4], [12]. The work presented in this paper goes beyond these approaches, since it proposes a model for simulating the energy consumption in discrete event-based simulations.

Furthermore, numerous work has been conducted to estimate the remaining capacity of Li-Ion batteries with a high accuracy using non-linear models [5], [9], [11], [14]. In network simulations, however, estimating the overall power consumption is more of interest, rather than estimating the remaining battery capacity. Therefore, in this paper only a simple linear battery model was chosen, which can be replaced with more sophisticated models in future versions of the simulation model. Finally, there are approaches for estimating the power consumption in packet-based simulations [15], which cannot directly be applied in the presented simulation model due to the use of a message-oriented design.

VII. SUMMARY

In this paper a state-based energy consumption model was presented, which allows for the assessment of the energy consumptions of large-scale overlay network simulations. A brief evaluation shows that using our model, the energy consumption can be estimated with a mean error of $\pm 4.7\%$. Subsequently, a state-based simulation model for estimating the power consumption of the Milestone 2 and the Nexus One has been derived.

REFERENCES

- [1] N. Balasubramanian, A. Balasubramanian, and A. Venkataramani, "Energy Consumption in Mobile Phones: a Measurement Study and Implications for Network Applications," in *ACM IMC*, 2009.
- [2] I. Baumgart, B. Heep, and S. Krause, "OverSim: A Flexible Overlay Network Simulation Framework," *IEEE Global Internet Symposium*, 2007.
- [3] A. Carroll and G. Heiser, "An Analysis of Power Consumption in a Smartphone," in *USENIX ATC*, 2010.
- [4] R. Friedman and A. Kogan, "On Power and Throughput Tradeoffs of WiFi and Bluetooth in Smartphones," *INFOCOM*, 2011.
- [5] M. Handy and D. Timmermann, "Simulation of Mobile Wireless Networks with Accurate Modelling of Non-Linear Battery Effects," in *Simulation*. Citeseer, 2003.
- [6] A. Montresor and M. Jelasity, "PeerSim: A Scalable P2P Simulator," *napa-wine.org*, no. 214412, 2009.
- [7] C. Panarello, A. Lombardo, G. Schembra, L. Chiaraviglio, and M. Mellia, "Energy Saving and Network Performance: a Trade-off Approach," in *ACM E-Energy*, 2010.
- [8] C. W. Paper, "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update 2010-2015," 2011.
- [9] M. Pedram, "An Analytical Model for Predicting the Remaining Battery Capacity of Lithium-Ion Batteries," *IEEE Transactions on Very Large Scale Integration Systems*, vol. 14, no. 5, 2006.
- [10] M. Pickavet, W. Vereecken, and S. Demeyer, "Worldwide Energy Needs for ICT: the Rise of Power-Aware Networking," in *IEEE ANTS*, 2008.
- [11] D. Rakhmatov and S. Vruthula, "An Analytical High-Level Battery Model for Use in Energy Management of Portable Electronic Systems," in *IEEE/ACM ICCAD*, 2001.
- [12] A. Rice and S. Hay, "Decomposing Power Measurements for Mobile Devices," in *IEEE PERCOM*, 2010.
- [13] D. Stingl, C. Gross, J. Rückert, L. Nobach, A. Kovacevic, and R. Steinmetz, "PeerfactSim.KOM: A Simulation Framework for Peer-to-Peer Systems," in *IEEE HPCS*, 2011.
- [14] S. Vruthula and D. Rakhmatov, "Battery Modeling for Energy-Aware System Design," *Computer*, vol. 36, no. 12, pp. 77–87, 2003.
- [15] H. Wu, S. Nabar, and R. Poovendran, "An Energy Framework for the Network Simulator 3 (ns-3)," *Symposium A Quarterly Journal In Modern Foreign Literatures*, vol. 3, 2011.
- [16] L. Zhang, B. Tiwana, and Z. Qian, "Accurate Online Power Estimation and Automatic Battery Behavior-Based Power Model Generation for Smartphones," in *ACM CODES+ISSS*, 2010.