

EnergyBox: A Trace-Driven Tool for Data Transmission Energy Consumption Studies

Ekhiotz Jon Vergara^(✉) and Simin Nadjm-Tehrani

Department of Computer and Information Science, Linköping University,
Linköping, Sweden
{ekhioetz.vergara, simin.nadjm-tehrani}@liu.se

Abstract. Although evolving mobile technologies bring millions of users closer to the vision of information anywhere-anytime, device battery depletions hamper the quality of experience to a great extent. We argue that the design of energy-efficient solutions starts by energy-awareness and propose EnergyBox, a tool that provides accurate and repeatable energy consumption studies for 3G and WiFi transmissions at the user end. We recognize that the energy consumption of data transmission is highly dependable on the traffic pattern, and provide the means for trace-based iterative packet-driven simulation to derive the operation states of wireless interfaces. The strength of EnergyBox is that it allows to modularly set the 3G network parameters specified at operator level, the adaptive power save mode mechanism for a WiFi device, and the different power levels of the operation states for different handheld devices. EnergyBox enables efficient energy consumption studies using real data, which complements the device-dependent laborious physical power measurements. Using real application transmission traces, we have validated EnergyBox showing an accuracy range of 94–99% for 3G and 93–99% for WiFi compared to the real measured energy consumption by a 3G modem and a smartphone with WiFi.

Keywords: Transmission energy · Trace-based simulation · Energy consumption studies · 3G · 802.11 · UMTS · WLAN

1 Introduction

Wireless interfaces account for a great energy cost on mobile devices [7, 13], but unfortunately, there seems to be some misconceptions about how energy consumption is influenced by data transmissions. While some studies employ only bulk data transfer to measure energy consumption [4], they do not measure the impact of the applications' real data pattern on energy consumption. The data transmission pattern depends on the real application operation (influenced by application developers) and the user-device interaction [6, 14], and drastically influences the energy per bit [3, 17].

In this paper we argue that reducing the energy consumption of wireless transmissions begins by being aware of the energy consumption characteristics

of different technologies. Since physical energy measurement studies are laborious and time-consuming, solutions are usually tested with little variation of factors such as network settings, hardware dependence, and synthetic data traces.

There are different approaches available to perform energy studies. While the professional testing equipment [1] or the physical power measurement equipment [19–21] provide high accuracy, they limit the studies to a fixed environment and incur high cost. Smart battery interfaces providing current values are available in some devices, and the software using them (e.g., CurrentWidget for Android or Nokia Energy Profiler) measures the aggregated current draw from all the device components. The error of the instant battery interface reading is usually high [5]. These and vendor specific development platforms [2] enable only device-dependent studies of aggregated consumption. We believe that a modular approach to carry out flexible and efficient energy studies isolating the transmission energy to complement physical energy measurements is essential.




Fig. 1. Overview of EnergyBox function.

Our work proposes the design of EnergyBox: a tool that enables accurate studies of data transmission energy consumption at the user end, using real traffic traces as input. We focus our efforts on the most widespread wireless technologies (3G and WiFi) and capture application data transmission energy footprint at the user device. EnergyBox captures the underlying states of operation of the wireless interfaces. The hardware dependence is overcome by using parametrised device specific power levels. For a given data trace, EnergyBox automatically outputs the operation states over time, so that when combined with device specific power levels it enables the computation of energy consumption. The general idea of our EnergyBox is shown in Fig. 1.

The data traces can be directly captured on different devices or synthetically created in order to study the impact of different transmission mechanisms under different configurations for these wireless technologies. EnergyBox assists researchers and application developers to immediately estimate energy consumption of data transmissions for a diverse range of test cases. The scenario is set up using different data traces, device power levels and network settings.

The tool accepts the statically configured parameters at the radio layer such as inactivity timers, the data buffer thresholds and a low activity mechanism used by the operator for state transition decisions in 3G communication. For

WiFi, it incorporates the adaptive power save mode commonly used as the power saving mechanism in the latest generation devices. We validate our EnergyBox against physical measurements and show an accuracy range of 94–99% for 3G and 93–99% for WiFi compared to physically measured energy consumption. The EnergyBox has already been used in a research environment to facilitate energy consumption studies complementing physical measurements [22, 23].

The rest of the paper is organised as follows: Sect. 2 provides the energy consumption background for 3G and WiFi. The design of EnergyBox is presented in Sect. 3 and this is evaluated against physical measurements in Sect. 4. Section 5 presents the related works and Sect. 6 covers concluding remarks and future work.

2 Background

We provide a background for communication energy footprint through illustrative measurements that show the main factors affecting the data transmission energy of 3G and WiFi at the device end. Figure 2 shows an overview of our measurement setup, which is also used for validation purposes in Sect. 4.

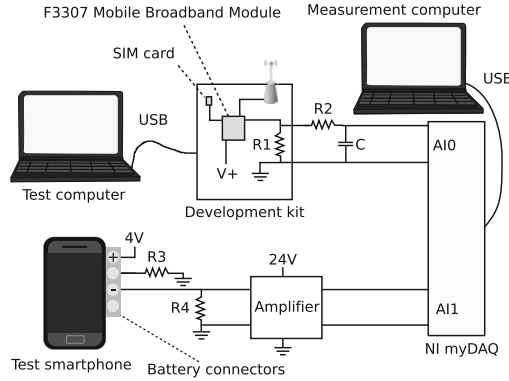


Fig. 2. Measurement setup for 3G and WiFi.

The 3G measurements were performed on a power-efficient mobile broadband module provided by Ericsson AB (KRY 901214/01, marked as Development Kit in 2). Since this exposes interfaces to measure the power consumption of the 3G modem easily isolating it from the rest of system (e.g., CPU or screen), we employ this platform instead of the 3G module in a smartphone. We measure the voltage drop over a shunt resistor ($R1=0.1\Omega$) and use a low-pass filter (approximately 16 Hz) to avoid any anti-aliasing effects [21].

The WiFi measurements were performed removing the battery of the smartphone under test, and adding a low-side sensing circuit ($R4 = 0.1\Omega$) with an

isolating amplifier (maximum transmission error of 0.4%). We isolate the transmission energy from the rest of the system as in earlier works [19]. The power consumption is sampled at 1 kHz in both cases.

2.1 Energy Consumption of 3G

The energy consumption of the user equipment (UE) in 3G is mostly influenced by the radio resource management performed at the network operator side by the Radio Network Controller (RNC). The RNC uses the Radio Resource Control (RRC) and the Radio Link Control (RLC) protocols. According to the RRC, the UE implements a state machine where the states imply different performance (response time and data rate) and power consumption. The UE states are Dedicated Channel (DCH), Forward Access Channel (FACH), and Paging Channel (PCH), sorted from highest to lowest power drain. The transitions between states are controlled using inactivity timers and RLC data buffer thresholds.




Fig. 3. Example power profile for 3G and RLC buffer thresholds.

Figure 3 shows the power consumption levels of the states experienced at one location using the operator TeliaSonera in Sweden when sending a packet burst. In PCH, the UE can be paged with the lowest energy drain but no data can be sent. Some signalling is required to establish a connection and move from PCH to DCH before sending any data (1-2s in Fig. 3).

The UE reports to the RNC the observed traffic volume, that controls state transitions to higher performance states via the RLC protocol. The RNC employs fixed thresholds over the RLC buffer data occupancy and triggers the transitions when the thresholds are exceeded. The table in Fig. 3 shows the measured threshold values for the different state transitions.

Inactivity timers are used to move the UE to lower performance states. The UE moves to FACH after T_1 with small or no data transmission (7s in Fig. 3). T_2 controls the transitions from FACH to PCH. Inactivity timers create energy

overheads known as *energy tails* due to the UE remaining in high consuming state while not transmitting anything [4, 21].

To reduce these energy overheads, the Fast dormancy mechanism allows the UE to signal the RNC the desire to switch to the lowest power state by sending a Signalling Connection Release Indication before the inactivity timeout. Some networks implement a low activity mechanism in DCH to release the transport channel and move to FACH when there is low traffic [9]. Thus, we see that the energy tails and the above mechanisms make the energy consumption dependent on the traffic pattern and operator settings in a complex way.

2.2 Energy Consumption of WiFi

The transmission energy consumption for WiFi is mostly influenced at the driver level in the WiFi station (the client handset). The station is in the Constantly Awake Mode (CAM) when it has the power-saving features disabled experiencing the best performance.




Fig. 4. Power levels of adaptive PSM for Xperia Arc.

The IEEE 802.11 standard defines a Power Save Mode (PSM), which allows the stations to switch to low power mode during predefined periods of time when not transferring any data. The access point (AP) buffers downlink frames for the clients and the clients wake up periodically (at multiples of the beacon interval). The clients send a Power Save Poll (PS-Poll) message to the AP to receive each buffered frame. Recent smartphones implement a mechanism named Adaptive PSM to overcome the overhead and latency drawback of using this PS-Poll mechanism [10]. The client switches between the CAM and PSM modes based on heuristics (e.g., number of packets, traffic inactivity period or screen on/off). The client uses the power management field in null data frames to inform the AP about its current mode.

Figure 4(left) shows an adaptive PSM implementation in a Sony Ericsson Xperia Arc smartphone, where the station moves to CAM for sending the traffic, and switches back from CAM to PSM after a predefined inactivity timeout

(δ) without packet transmission. This δ timeout creates an energy tail in the similar way as in 3G. Repeated measurements have shown δ to be around 220 and 70 ms for the Samsung Galaxy SII and Xperia Arc respectively, much shorter than in 3G. Previous generation devices implement longer δ timeouts (e.g., 1.5 seconds for HTC Magic) [11, 16, 20]. Moreover, some drivers also implement packets per second thresholds (*Up* and *Down*) that trigger PSM-CAM and CAM-PSM transitions respectively [15].

While in the same state, the station consumes more power when the data rate increases. In order to show the impact of data rate on transmission power and δ , we create an uplink data stream varying inter-packet interval and packet size using the Xperia Arc. Figure 4(right) shows that for higher data rate (i.e., short inter-packet interval), the average power level increases substantially. When the inter-packet interval is increased to 70 ms, the station switches back to PSM since the inter-packet interval is greater than δ . This drops the average power level. To sum up, adaptive PSM leads to a high impact by the data pattern on the energy consumption of a WiFi station.

3 EnergyBox

EnergyBox captures the influence of the data transmission pattern on transmission energy consumption by performing trace-based iterative packet-driven simulation. The usage of real data traces means that the corresponding energy is realistic, and reflects the impact of the data pattern on real throughput and delay in the network. EnergyBox also accepts using synthetic data traces as input, creating repeatable tests for a given purpose. Given a packet trace and configuration parameters, the EnergyBox outputs the device states $S(t)$ over time. The total energy consumption is calculated by associating these states with power levels and integrating them over time. Device-specific power level values can be obtained through measurement platforms as the one described in Sect. 2 or in the literature [4, 16, 19, 21, 24]. These operation power levels abstract the hardware dependency allowing to quantify an application footprint on a given device. EnergyBox simulates the 3G network parameters specified at operator level and the adaptive PSM mechanism specified at the handset driver for WiFi. The design of the EnergyBox is rooted in measurements and a careful literature study. By providing our code to the research and development community we offer a simple to configure but still powerful approach to perform accurate data transmission energy measurement studies.

3.1 3G Model

The RRC state transitions are captured by a parametrised finite state machine that simulate the inactivity timers, the RLC buffers and a low activity mechanism in a packet-driven manner. Figure 5a shows the states and the state transitions we use in our 3G model.

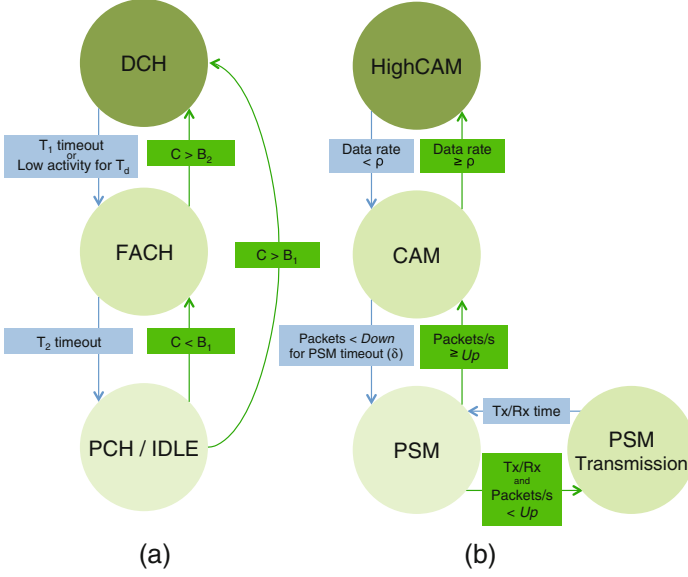


Fig. 5. Overview of the EnergyBox state machines for 3G (a) and WiFi (b).

For each packet P_i in the trace and its timestamp $t(P_i)$, we calculate $\Delta_i = t(P_i) - t(P_{i-1})$ as the elapsed time between the two packets. Δ_i is used to simulate the inactivity timers T_1 and T_2 : if $\Delta_i > T_1$ or $\Delta_i > T_2$, we trigger the corresponding state transition. In order to account for the signalling time between states, every state transition has pre-defined a parametrised transition duration. This can be obtained from a simple power measurement (e.g., the delay observed in Fig. 3) or measuring the round-trip time for the different transitions.

There are four RLC buffer thresholds for PCH-FACH and FACH-DCH transitions (see Sect. 2): two uplink (B_1^u and B_2^u) and two downlink (B_1^d and B_2^d). The simulation of the RLC buffers is done as follows: we define Δ_i^u and Δ_i^d as the elapsed time from the last uplink or downlink packet (in the same direction). Note that since the size and direction (up/down) for each packet P_i is known, the current occupancy of each buffer in bytes can be computed. We denote this by C . The data is transmitted (i.e., the buffer is emptied setting $C = 0$) depending on the channel data rate. Given C and the channel data rate, the time to empty the RLC buffer can be calculated (T_e). T_e can be specified as a constant assuming constant data rate (e.g., $T_e = C/512$ kbps) or as function of buffer occupancy C based on data rate measurements in FACH for a network.

In order to simulate the low activity mechanism in DCH, we define T_d as the period of time over which the amount of data sent is monitored. We sum the size of packets during T_d and force the move to FACH when it is lower than a threshold. Fast dormancy is modelled by simply moving to PCH/Idle after a predefined time.

The current 3G model does not fully consider the impact of the received signal strength (RSS). Under poor radio-link conditions the transmission power is higher and the data rate is lower [21]. The data rate is captured by the recorded input data trace. However, in order to model the impact of RSS-power level we would have to feed the EnergyBox with a trace of received signal strength or similar indicator (e.g., Signal-to-Noise Ratio). This can be added to the implementation if the means of capturing such traces are available.

The strength of modelling 3G states in EnergyBox is that it is a generic way of capturing different state machines for different operators using a parametrised finite state machine: if an operator implements a RRC state machine where a single packet triggers a state transition from Idle/PCH to DCH, we only have to set the B_1^u threshold to 0 bytes.

3.2 WiFi Model

The EnergyBox WiFi captures the adaptive PSM mechanisms based on the inactivity timer and the number of packets per second. Handsets switch between two states (PSM and CAM) using adaptive PSM, but in order to model the high data rate behaviour of a handset, we define the state machine shown in Fig. 5b.

The station only wakes up for beacons in the *PSM* state. *PSM Transmission* state represents the sending or reception of packets in the PSM mode, where the station switches to a high power only during the transmission interval. The Tx/Rx time is defined as reconfigurable parameter, which is obtained by sending few packets and observing the power profile (similar to Fig. 4(left)).

In *PSM*, the transition to *CAM* is triggered whenever the number of packets per second count is higher than the Up threshold. When the number of packets per second is less than *Down* for a predefined timeout time (δ), the station switches back to *PSM*.

We have observed that the power drain of the station increases with the throughput as it was illustrated in Sect. 2. This behaviour is captured by computing the throughput over a short time window and comparing it to a ρ threshold. Once the time window has been set, the ρ threshold can be decided by sending a train of packets with different inter-packet interval (similar to Fig. 4(right)) and measuring the power drain. When the threshold is exceeded, the station switches to *HighCAM*, a state with a higher power level than *CAM*.

4 Evaluation

This section describes the evaluation of the EnergyBox accuracy against physical energy measurements using the setup described in Sect. 2. We define two metrics to quantify the accuracy of the EnergyBox: *energy accuracy* and *time accuracy*, which are described together with the methodology. The general methodology is similar for both WiFi and 3G and follows the following steps: (1) A set of applications are used to create different traffic patterns. We simultaneously collect packet traces and measure the *power trace* $P_d(t)$ of the device while creating

traffic from one application at a time. This is done using the toolset described in Fig. 2. The power trace $P_d(t)$ represents the ground truth. (2) From the measured power trace, we compute the expected *device states* $S_d(t)$ using power level thresholds. These thresholds are obtained earlier using measurements on a given handset. (3) The EnergyBox is fed with the packet traces and outputs the *inferred states* $S_i(t)$. (4) We compare the EnergyBox inferred states $S_i(t)$ against the device states $S_d(t)$ and compute the difference over time. T represents the total duration of a trace and $f(x, y)$ simply returns 1 if $x = y$, and 0 otherwise. Time accuracy represents the percentage of time that the inferred states and the measured states overlap (i.e. the higher the measure the better the accuracy over a time interval). It is defined as follows:

$$\text{Time Accuracy} = \frac{\int_0^T f(S_i(t), S_d(t)) dt}{T} \cdot 100 (\%) \quad (1)$$

(5) The EnergyBox assigns the measured device-specific handset power levels for each state to the inferred states obtaining the *inferred power trace* $P_i(t)$. The energy consumption is computed in Joules integrating $P_i(t)$ over the trace duration (T). Energy accuracy reflects the difference between the inferred energy consumption and the energy consumption of the measured power trace (the ground truth):

$$\text{Energy Accuracy} = \frac{\int_0^T P_i(t)}{\int_0^T P_d(t)} \cdot 100 (\%) \quad (2)$$

4.1 Methodology and Evaluation Settings

The packet traces are 5 minutes long, captured with *tcpdump* for WiFi and 3G. We demonstrate the reliability of the EnergyBox by covering a wide range of data patterns in terms of inter-packet interval, packet size and total amount of data transmitted created by commonly used mobile applications. Since the EnergyBox is deterministic (i.e., it creates the same output for a given input packet trace and the configuration parameters), each packet trace is fed only once in the EnergyBox. We employ 10 different packet traces coming from different applications for 3G and 9 for WiFi. Email has periodic small data transmissions, whereas Facebook and Web represent bursty downloads of bigger amounts of data. Spotify is a music streaming application that sends data in different bursts and Stream is a constant radio stream. Skype Chat represents instant messaging services which usually have smaller data transmissions. Skype Call and Video are audio and video conferences with some small chat messages. Finally, Youtube captures the user watching videos online. As opposed to doing several experiments per application, we chose to test several different applications with one trace each, similar to the approach in other current work [5, 13, 17].

For 3G, we set $T_1 = 4.1$ s and $T_2 = 5.6$ s and the RLC buffer thresholds shown in Fig. 3. The state transition times are 1.7, 0.43, 0.65 s for PCH-DCH, PCH-FACH and FACH-DCH respectively. We set the time to empty the RLC buffers, $T_e = 1.2 \cdot C + 10$ as a function of the buffer occupancy (C) based on repeated data

rate measurements in FACH following a previous methodology [17]. We enable the low activity mechanism (with $T_d = 4$ s and B_1^u as the threshold) and disable fast dormancy since the 3G module does not support it. The device-specific power levels used are the average power values measured on the 3G module for different states: DCH, FACH and PCH (1.3, 0.5 and 0.2 W). The evaluation was performed under similar and typical values of received signal strength.

For WiFi, we set $Up = 1$ and $Down = 1$, which are the settings of the stations used in our measurements. The Galaxy SII was used for the evaluation and δ was set to 220 ms based on observation (similar to Fig. 4(left)). We used 30, 250 and 500 mW for PSM, CAM and HighCAM respectively based on measured average power levels. The ρ data rate threshold was empirically set to 3 kB for a time window of 50 ms based on empirical experiments mentioned in Sect. 2.2. Since adaptive PSM in the Galaxy SII differentiates only two states (PSM and CAM), we perform the time accuracy evaluation of the EnergyBox WiFi using the two basic states of our WiFi state machine.

4.2 EnergyBox 3G

The average accuracy range of EnergyBox over the different traces for 3G is high, 95% and 98% for time and energy respectively. Figure 6 shows the time and energy accuracy for the different traces.




Fig. 6. EnergyBox 3G accuracy for different traces.

Regarding lower time accuracies (e.g., Skype Chat), the typical cause is that the inferred state is DCH state while the real state is FACH. The source of this deviation is the time to empty the buffers (T_e) of the RLC buffer simulation. Even though the modelling of this parameter is based on real measurements, the variations in number of active users and traffic at the operator network makes the real time to empty the buffers different from the measurement-based model

in some instances. The Skype Chat trace has many small packets. These stress the RLC data buffer simulation. However, the accuracy is still close to 90%.

The energy accuracy is high in all the traces. The lowest registered energy accuracy is for Email and Spotify traces, which is still higher than 94%. Figure 7 shows an example of the EnergyBox accuracy by comparing it to the real measurements for the Email trace.




Fig. 7. A fragment of a 3G measurement and EnergyBox 3G inferred output for the Email trace.

4.3 EnergyBox WiFi

The average accuracy of EnergyBox over the different traces for WiFi is really high with 99% and 98% for time and energy respectively. Figure 8 shows the accuracy for the different traces. Concerning time accuracy, the few discrepancies are originated from the value given to δ , which makes the switch from CAM to PSM somewhat earlier than in the real states.

Regarding the energy accuracy metric, the differences can be traced to traces with higher data rates (e.g., Skype Video and Stream). The simplicity of the fixed power level for *HighCAM* in our model does not completely cover the case of high data rate with a higher power level than the fixed one assigned to *HighCAM*. However, the discrepancy is still relatively small.

To summarise, the above comparisons provide evidence that EnergyBox provides high energy accuracy and device state information when using real traces.

5 Related Works

We categorise the main body of related work into two areas: energy consumption measurement studies and the characterisation of energy consumption of wireless networking at the user end.




Fig. 8. EnergyBox WiFi accuracy for different traces.

Energy measurement studies: an influential measurement study by Balasubramanian et al. [4] reports the categorisation of three different energy components in cellular networks: ramp, transfer and tail. The tail is the most consuming component and it is caused by the inactivity timers statically set by the network operators. Our previous works [3, 21] refine this study by performing physical measurements using a cellular modem that isolates the data transfer energy.

Qian et al. [18] perform a detailed study of the tail energy overhead using data traces retrieved from a network operator and later [17] point out how different applications inefficiently utilise the radio resources due to their data pattern.

The main advantage of our 3G measurement setup is that we are able to *isolate* the energy consumption for *data transfers* by validating the transmission-specific measurements in a 3G module, whereas earlier works do not isolate the energy for data transfers from the rest of the system when evaluating accuracy.

Wang et al. [24] present the energy consumption of data transfers for different packet sizes and transmission intervals over WLAN and cellular networks. Rice et al. [19] measure the WiFi energy consumption in a variety of smartphones. Several works focus in studying the energy consumption of WiFi based on bulk data transmissions [4, 7]. Studies using bulk data transmissions do not capture the impact of adaptive PSM on energy consumption. Various works [11, 16, 20] study the different adaptive PSM mechanisms in previous generation devices. A recent study by Pyles et al. [15] shows the *Up* threshold for different devices.

The main advantage of our validation base is the fine-grained measurements isolating data transfer energy. Moreover, most studies do not consider the impact of the data traffic pattern on energy consumption and focus on bulk transfers.

Wireless energy consumption characterisation: Balasubramanian et al. [4] model the energy consumption of 3G and WiFi using linear regression based on their bulk data download measurements. Their simple model does not capture the impact of the data pattern on the energy consumption. Oliveira

et al. [12] use a state based model for 3G modelling the RLC buffer occupancy as fixed data rate (kbps) without performing any validation of their model.

Some works [26, 27] analytically model different aspects of the energy consumption of 3G. Our approach is applicable to arbitrary data traces captured from any real application, whereas named works are limited to modelling web and streaming traffic.

The radio resource usage application profiler for 3G presented by Qian et al. [17] is the work closest to our EnergyBox for 3G. Their focus is on providing the application developer insights and hints about how to reduce the energy consumption due to 3G transmissions. We extend their 3G work by including in our 3G state-based model a low activity mechanism implemented by some operators, thereby improving the accuracy of the model. Moreover, our EnergyBox captures different RRC state machines within a single parametrised one.

The work by Harjula et al. [8] creates a device-specific power profile for different messaging intervals based on device measurements on a Nokia 95 for a single operator. Their model uses indicators such as average packet size and signalling frequency that need to be extracted from the data transmissions. Instead, EnergyBox directly works on data traces and derives the RRC states of different networks simulating the RLC buffer thresholds and state transitions. Our approach is more general and simplifies energy consumption studies.

Xiao et al. [25] present a detailed power level model for the 802.11g interface. The modelled data pattern based on traffic burstiness was only validated by simulating TCP download and upload traffic based on the amount of data transferred. In comparison, our adaptive PSM model is validated against a set of representative application traces. Dong et al. [5] propose to automatically self-generate a system energy model within a device (including WiFi and 3G) using the smart battery interface.

A recent study by Pathak et al. [13] present a system-call-based power modelling for smartphones, including device dependent 3G and WiFi models. Tracing system calls typically requires modifications to the operating system (e.g., root access), whereas we adopt a generic approach based on packet captures, which are generally available¹.

6 Conclusion and Future Work

Designing energy efficient solutions for wireless networks starts by energy awareness. To reduce energy waste, there is a need to provide visibility and understanding of the energy consumption that impacts battery lifetime at the user end as well as elsewhere in the network. Our contribution makes the energy footprint of the network interface explicit, and presents a tool that provides accurate energy consumption values given packet traces as input. We focus the EnergyBox on the most widespread wireless interfaces (3G and WiFi) and capture the parameters influencing the energy consumption most.

¹ Android phones provide traffic statistics that can be captured every millisecond and packet level capture is available without root permission from Android 4.0 on.

The design of EnergyBox is rooted in our physical energy measurements as well as a careful literature study. We build on the fact that the handset energy consumption in 3G is essentially driven by the inactivity timers and the RLC data buffer thresholds that control RRC state transitions at the network operator end. The energy consumption in WiFi is mostly influenced by the handset dependent adaptive PSM mechanism implemented at driver level in the handset. We capture these aspects within parametrised state machines in the EnergyBox and evaluate its accuracy against physical measurements. Our evaluation shows that the EnergyBox provides accurate energy consumption estimates.

Including the impact of the signal strength into EnergyBox and extending it towards other wireless interfaces such as the fourth generation cellular network (LTE) and Bluetooth would be interesting directions for future work. In its current state, however, the EnergyBox is found to be a valuable instrument in studying the energy consumption in different networking scenarios in our current research. It has substantially aided efficient energy-related studies by emulating different parameter setups and replacing energy measurements.

Acknowledgment. This work was supported by the Swedish national graduate school in computer science (CUGS). The authors wish to thank the support of Ericsson AB, and in particular B-O Hertz, Pär Emanuelsson and Claes Alströmer for providing the 3G developer kit and facilitating the measurement gathering phase.

References

1. T&m solution. rohde & schwarz. http://www.rohde-schwarz.com/en/applications/optimize-the-quality-of-experience-of-mobile-devices-application-card_56279-35727.html
2. Treppn profiler, qualcomm. <https://developer.qualcomm.com/mobile-development/development-devices/treppn-profiler>
3. Asplund, M., Thomasson, A., Vergara, E.J., Nadjm-Tehrani, S.: Software-related energy footprint of a wireless broadband module. In: Proceedings of the 9th ACM International Symposium on Mobility Management and Wireless Access, MobiWac '11. ACM (2011)
4. Balasubramanian, N., Balasubramanian, A., Venkataramani, A.: Energy consumption in mobile phones: a measurement study and implications for network applications. In: Proceedings of ACM Internet Measurement Conference IMC (2009)
5. Dong, M., Zhong, L.: Self-constructive high-rate system energy modeling for battery-powered mobile systems. In: MobiSys '11, pp. 335–348. ACM (2011)
6. Falaki, H., Mahajan, R., Kandula, S., Lymberopoulos, D., Govindan, R., Estrin, D.: Diversity in smartphone usage. In: Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services, MobiSys '10, pp. 179–194. ACM (2010)
7. Friedman, R., Kogan, A., Krivolapov, Y.: On power and throughput tradeoffs of wifi and bluetooth in smartphones. In: Proceedings of IEEE INFOCOM 2011, pp. 900–908 (2011)
8. Harjula, E., Kassinen, O., Ylianttila, M.: Energy consumption model for mobile devices in 3g and wlan networks. In: IEEE Consumer Communications and Networking Conference (CCNC 2012), pp. 532–537 (2012)

9. Holma, H., Toskala, A.: WCDMA for UMTS: HSPA evolution and LTE. Wiley Online Library: Books, John Wiley & Sons (2010)
10. Krashinsky, R., Balakrishnan, H.: Minimizing energy for wireless web access with bounded slowdown. In: Proceedings of the 8th Annual International Conference on Mobile Computing and Networking, MobiCom '02, pp. 119–130. ACM (2002)
11. Manweiler, J., Roy Choudhury, R.: Avoiding the rush hours: Wifi energy management via traffic isolation. *IEEE Trans. Mob. Comput.* **11**(5), 739–752 (2012)
12. Oliveira, T., Ursini, E., Timoteo, V.: Simulation inspired model for energy consumption in 3g always-on mobiles. In: IEEE 2nd National Conference on Telecommunications (CONATEL 2011), pp. 1–7 (2011)
13. Pathak, A., Hu, Y.C., Zhang, M.: Where is the energy spent inside my app?: fine grained energy accounting on smartphones with eprof. In: Proceedings of the 7th ACM European Conference on Computer Systems, EuroSys '12, pp. 29–42 (2012)
14. Paul, U., Subramanian, A., Buddhikot, M., Das, S.: Understanding traffic dynamics in cellular data networks. In: Proceedings of IEEE INFOCOM 2011, pp. 882–890 (2011)
15. Pyles, A.J., Qi, X., Zhou, G., Keally, M., Liu, X.: Sapsm: smart adaptive 802.11 psm for smartphones. In: Proceedings of the 14th International Conference on Ubiquitous Computing, UbiComp '12. ACM (2012)
16. Pyles, A.J., Ren, Z., Zhou, G., Liu, X.: Sifi: exploiting voip silence for wifi energy savings in smart phones. In: Proceedings of the 13th International Conference on Ubiquitous Computing, UbiComp '11, pp. 325–334. ACM (2011)
17. Qian, F., Wang, Z., Gerber, A., Mao, Z., Sen, S., Spatscheck, O.: Profiling resource usage for mobile applications: a cross-layer approach. In: MobiSys '11, pp. 321–334. ACM (2011)
18. Qian, F., Wang, Z., Gerber, A., Mao, Z.M., Sen, S., Spatscheck, O.: Characterizing radio resource allocation for 3g networks. In: Proceedings of the 10th Annual Conference on Internet Measurement, IMC '10, pp. 137–150. ACM (2010)
19. Rice, A., Hay, S.: Measuring mobile phone energy consumption for 802.11 wireless networking. *Pervasive Mob. Comput.* **6**, 593–606 (2010)
20. Rozner, E., Navda, V., Ramjee, R., Rayanchu, S.: Napman: network-assisted power management for wifi devices. In: MobiSys '10, pp. 91–106. ACM (2010)
21. Vergara, E.J., Nadjm-Tehrani, S.: Energy-aware cross-layer burst buffering for wireless communication. In: Proceedings of the 3rd International Conference on Future Energy Systems: Where Energy, Computing and Communication Meet, e-Energy '12. ACM (2012)
22. Vergara, E.J., Prihodko, M., Nadjm-Tehrani, S.: Mobile location sharing: an energy consumption study (poster). In: Proceedings of the 4th International Conference on Future Energy Systems: Where Energy, Computing and Communication Meet, e-Energy '13. ACM (2013)
23. Vergara, E.J., Sanjuan, J., Nadjm-Tehrani, S.: Kernel level energy-efficient 3g background traffic shaper for android smartphones. In: Proceedings of the 9th International Wireless Communications and Mobile Computing Conference (IWCMC 2013) (2013)
24. Wang, L., Manner, J.: Energy consumption analysis of wlan, 2g and 3g interfaces. In: Proceedings of the 2010 IEEE/ACM International Conference on Green Computing and Communications & International Conference on Cyber, Physical and Social Computing, GREENCOM-CPSCOM '10, pp. 300–307. IEEE Computer Society (2010)

25. Xiao, Y., Savolainen, P., Karppanen, A., Siekkinen, M., Ylä-Jääski, A.: Practical power modeling of data transmission over 802.11g for wireless applications. In: Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking, e-Energy '10, pp. 75–84. ACM (2010)
26. Yang, S.R., Yan, S.Y., Hung, H.N.: Modeling umts power saving with bursty packet data traffic. *IEEE Trans. Mob. Comput.* **6**(12), 1398–1409 (2007)
27. Yeh, J.H., Chen, J.C., Lee, C.C.: Comparative analysis of energy-saving techniques in 3g pp and 3gpp 2 systems. *IEEE Trans. Veh. Technol.* **58**(1), 432–448 (2009)

Energy Efficiency in Large Scale Distributed Systems
COST IC0804 European Conference, EE-LSDS 2013,
Vienna, Austria, April 22-24, 2013, Revised Selected
Papers

Pierson, J.-M.; Da Costa, G.; Dittmann, L. (Eds.)

2013, XI, 312 p. 127 illus., Softcover

ISBN: 978-3-642-40516-7