# A Fine-grained Event-based Modem Power Model for Enabling In-depth Modem Energy Drain Analysis

XIAOMENG CHEN, Purdue University, USA
JIAYI MENG, Purdue University, USA
Y. CHARLIE HU, Purdue University, USA
MARUTI GUPTA, Intel Corp., USA
RALPH HASHOLZNER, Intel Deutschland GmbH, Germany
VENKATESAN NALLAMPATTI EKAMBARAM, Intel Corp., USA
ASHISH SINGH, Intel Deutschland GmbH, Germany
SRIKATHYAYANI SRIKANTESWARA, Intel Corp., USA

Cellular modems enable ubiquitous Internet connectivities to modern smartphones, but in doing so they become a major contributor to the smartphone energy drain. Understanding modem energy drain requires a detailed power model. The prior art, an RRC-state based power model, was developed primarily to model the modem energy drain of application data transfer. As such, it serves well its original purpose, but is insufficient to study detailed modem behavior, *e.g.,* activities in the control plane.

In this paper, we propose a new methodology of modeling modem power draw behavior at the event-granularity, and develop to our knowledge the first fine-grained modem power model that captures the power draw of all LTE modem radio-on events in different RRC modes. Second, we quantitatively demonstrate the advantages of the new model over the state-based power model under a wide variety of context via controlled experiments. Finally, using our fine-grained modem power model, we perform the first detailed modem energy drain in-the-wild study involving 12 Nexus 6 phones under normal usage by 12 volunteers spanning a total of 348 days. Our study provides the first quantitative analysis of energy drain due to modem control activities in the wild and exposes their correlation with context such as location and user mobility.

CCS Concepts: • **Hardware → Power estimation and optimization**; • **Networks →** *Network protocols*;

Additional Key Words and Phrases: Power modeling, Cellular modems, Carrier networks, LTE protocols, Performance analysis

**ACM Reference Format:**
Xiaomeng Chen, Jiayi Meng, Y. Charlie Hu, Maruti Gupta, Ralph Hasholzner, Venkatesan Nallampatti Ekambaram, Ashish Singh, and Srikathyayani Srikanteswara. 2017. A Fine-grained Event-based Modem Power

45

Authors' addresses: Xiaomeng Chen, Purdue University, 465 Northwestern Avenue, West Lafayette, IN, 47907, USA, chen1058@purdue.edu; Jiayi Meng, Purdue University, 465 Northwestern Avenue, West Lafayette, IN, 47907, USA, meng72@purdue.edu; Y. Charlie Hu, Purdue University, 465 Northwestern Avenue, West Lafayette, IN, 47907, USA, ychu@purdue.edu; Maruti Gupta, Intel Corp. 2111, NE 25th Ave, Hillsboro, OR, 97124, USA, maruti.gupta@intel.com; Ralph Hasholzner, Intel Deutschland GmbH, Am Campeon 10-12, Neubiberg, 85579, Germany, ralph.hasholzner@intel.com; Venkatesan Nallampatti Ekambaram, Intel Corp. 2111, NE 25th Ave, Hillsboro, OR, 97124, USA, venkatesan.nallampatti.ekambaram@intel.com; Ashish Singh, Intel Deutschland GmbH, Am Campeon 10-12, Neubiberg, 85579, Germany, ashish.singh@intel.com; Srikathyayani Srikanteswara, Intel Corp. 2111, NE 25th Ave, Hillsboro, OR, 97124, USA, srikathyayani.srikanteswara@intel.com.


2476-1249/2017/12-ART45 $15.00
https://doi.org/10.1145/3154504

Model for Enabling In-depth Modem Energy Drain Analysis. *Proc. ACM Meas. Anal. Comput. Syst.* 1, 2, Article 45 (December 2017), 28 pages. https://doi.org/10.1145/3154504

## 1 INTRODUCTION

Cellular modems enable ubiquitous Internet connectivities to modern smartphones which are fundamental to their rapid market penetration and advancing the society to the mobile era in the past decade. In doing so, they also become a major contributing factor to smartphone battery drain. As the limited battery life has remained a major hindrance to the smartphone users' mobile experience, understanding and optimizing modem energy drain is of greater importance.

Understanding modem energy drain requires measurement studies of modem behavior under normal usage, *i.e.,* subject to the real context experienced by the billion users in the wild (*e.g.,* [9]). Measuring the power draw of the modem in such studies can not be done using a powermeter which cannot be easily deployed and further can only output power readings for the whole phone. Instead, such studies require a detailed power model that can accurately estimate the modem power draw due to all of its activities under normal usage, including not only application data transmissions, but also all the control-plane activities such as interactions with the carrier networks which are needed to maintain the ubiquitous, persistent connectivity to the Internet.

Existing LTE modem power models [9, 11] were developed primarily to model the modem energy drain of application data transfer under the hypothesis that application data transfer dominates the modem energy drain. Such power models are *RRC-state based* (or simply state-based for short) finite-state machines that essentially inherit the RRC (Radio Resource Control) states (and DRX (Discontinuous Reception) modes) in the modem RRC State Machine, and annotate each RRC state with a single average power to approximate the power draw while the modem stays in that state. These power models are therefore *coarse-grained*; they operate at the RRC-state granularity and do not model the detailed modem activities happening while staying inside any one of the RRC states.

While the prior state-based modem power model has served reasonably well its original purpose, *e.g.,* a large-scale in-the-wild study [9] has shown modem data transfer accounts for over half of the total modem energy, it has three significant limitations: (1) It fundamentally cannot capture any detailed modem activities in the control plane such as multiple radio switching events between modem radio-on and radio-off in one RRC state, or multiple modem procedures or timer intervals [4] in one radio-on period. This limitation in turn has two consequences. (2) It can only assign macroscopic parameters to each RRC state such as average duration and power, and as a result can suffer high energy drain estimation error in all RRC states, including Continuous Reception, Tail, and RRC Idle. (3) It cannot be used to study the energy drain of modem control-plane activities. Besides application data transmission, the modem in each UE performs many complex interactions with the cellular infrastructure in order to support connectivity in the presence of UE mobility, network dynamics and varying tower coverage. Such activities cannot be captured by the state-based power model, yet can potentially provide many new insights into the actual operations of cellular protocols under normal usage by the billions of UEs and the corresponding energy drain of the modems that implement the cellular protocols. Such insights are of significant interests to mobile operators, modem vendors, and the 3GPP standard committees.

To overcome these limitations of the prior art, in this paper, we develop to our knowledge the first fine-grained empirical modem power model that captures the power draw of all LTE modem radio-on episodes in different RRC modes. Such a fine-grained power model enables new, microscopic analysis of the cellular network control plane that exposes energy issues in LTE under different network conditions and guides potential improvement of LTE protocol design and hardware chip design under such conditions. The new model will directly benefit LTE protocol desiners, modem designers, as well as Operating System designers for UEs. Although detailed power modeling tools

for modem designers exist, those tools are typically based on complex circuit-level modeling (*e.g.,* [7, 22]), they are proprietary, and they are limited to measuring energy consumption in controlled experiments in the lab. In contrast, our new model provides an effective way to measure and perform energy analysis of modems running in the wild.

Developing such a fine-grained power model faces two major challenges. First, to capture detailed modem control-plane activities, *i.e.,* interactions with the cellular infrastructure, we need to be able to log all such modem activities. To overcome the challenge, we exploit the observation that chipsets by a major modem vendor which reside in major Android smartphones expose modem messages via `/dev/diag`, and leverage a recently developed tool called MobileInsight [14] that logs the exposed modem messages and further decodes raw cellular logs in binary streams to readable LTE protocol messages.

Second, we need to find the right granularity for the fine-grained power model. Intuitively, modem activities boils down to many individual over-the-air packets exchanged with eNodeB (E-UTRAN Node B) and internal events such as LTE timers and measurement, which however are simply too fine-grained to be modeled individually. The key insights in our approach are that (1) the model just needs to capture all radio-on episodes or events, (2) each modem event follows well-defined LTE protocol procedures, each consisting of a sequence of modem messages, and (3) there are only a small, finite number of such radio-on events.

Motivated by these insights, we define a *modem event* as a set of modem operations that happen in a single radio-on period, and develop an *event-based power model* that models modem power draw at the granularity of events. We show such an event-based power model is *practical* – it can be derived using a commodity powermeter *fine-grained* – it captures the precise duration and accurate power draw of each event, and *complete* – it captures both control and data events.

We summarize our contributions and highlight of findings of our measurement study as follows.

- We propose a new methodology of modeling modem power draw behavior at the *event-granularity*, which overcomes the challenge of modeling individual modem activities (*e.g.,* messages or internal activities) which are too fine-grained to be distinguishable by the power-meter, yet can accurately capture the energy drain of all modem radio-on episodes.
- Using the methodology, we develop to our knowledge the first fine-grained modem power model that captures the power draw of all modem radio-on events of LTE. As LTE evolves over time, we believe our methodology can be applied to derive fine-grained power models for them.
- We show the fine-grained power model achieves far better modeling accuracy than the state-based model via controlled experiments using a Nexus 6 phone under a diverse set of UE usage and signal strength conditions. Comparing 5-second-interval average power estimation by the two models with the ground-truth power readings by the powermeter shows that, under good, medium and poor signal condition, respectively, (1) During continuous reception of data, the new model has less than 20mA absolute error in 60%, 59% and 70% of intervals, while the old model under-estimates modem power by more than 20mA in 82%, 61% and 94% of the intervals; (2) During RRC tail (Short and Long DRX cycles), the new model has less than 5mA absolute error in 80%, 85% and 81% of the intervals, while the old model has less than 5mA absolute error in 28%, 33% and 23% of intervals; (3) During RRC Idle, the new model has less than 0.2mA absolute error in 80%, 80% and 41% of the intervals, while the old model has less than 0.2mA absolute error in 78%, 27% and 48% of the intervals.
- We present three case studies from a modem energy drain in-the-wild study involving 12 Nexus 6 phones under normal usage by 12 volunteers over a period of 2 months, to showcase the new types of analysis enabled by the event-based modem power model: (1) The new event-based model refines the LTE tail energy in RRC Connected. In particular, it exposes that durations of full tails can vary for different network carriers and within a network carrier, and detects

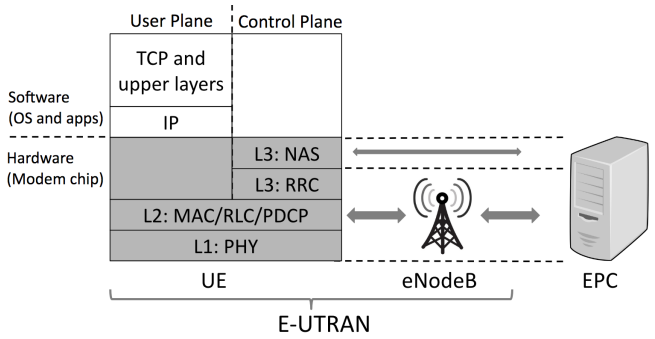


Fig. 1. LTE Protocol stack on the left. LTE network architecture on the right.

occurrences of significantly high BSR (Buffer Status Report) energy drain in the middle of tails. (2) The new model captures the energy drain of LTE NAS (Non-Access Stratum) procedures, in particular, EMM (EPS Mobility Management) specific procedures. The trace shows in 20% user-days, there are more than 3 attach procedures, 4 detach procedures, and 24 tracking area update (TAU) procedures per day, and there is positive correlation between TAU energy increase and RRC connection count caused by user-plane data. (3) The new model captures RRC-Idle procedures such as PLMN (Public Land Mobile Network) search and cell selection/re-selection. It also exposes a correlation between PLMN search energy and signal strength and user location, and between cell selection/re-selection and user mobility status.

## 2 BACKGROUND AND MOTIVATION

### 2.1 LTE network protocol stack

Figure 1 shows the LTE protocol stack implemented on the UE and the network architecture including UE, eNodeB and EPC (Evolved Packet Core). The left side of the figure shows the LTE protocol stack of a UE. The E-UTRAN (LTE Radio Access Network) protocol stack includes the bottom three layers (shaded) implemented inside the modem chipset. (1) L1 is PHY (physical layer), which carries all the information from MAC transport channels over the air interface, and takes care of link adaption, power control, cell search and other measurement for the RRC layer. (2) L2, including MAC, RLC and PDCH, together with L1 enables radio access between a UE and eNodeB. L1 and L2 serve both the user plane and the control plane. (3) L3 belongs to the control plane, including RRC and NAS. RRC is responsible for broadcast of system information, paging, establishment, maintenance and release of an RRC connection between a UE and eNodeB. NAS supports UE mobility and session management procedures to establish and maintain IP connectivity. On top of the E-UTRAN stack are the protocols implemented in software, *e.g.*, IP, TCP, HTTP, all belonging to the user plane.

### 2.2 RRC State Machine and prior modem power model

Figure 2 shows the commonly known RRC Finite State Machine (FSM) of LTE [11]. There are two basic states, RRC Connected on the left and RRC Idle on the right. In addition, DRX is a method used in mobile communication between a UE and the base station to conserve the battery of mobile devices [20]. Using DRX, a UE can be in one of three modes in RRC Connected: Continuous Reception, Short DRX, and Long DRX, but only in the DRX mode while in RRC Idle. In the rest of the paper, we refer to RRC states and modes interchangeably.

Figure 3 illustrates the typical DRX behavior (transitions among its four modes) of the LTE modem on a UE [2, 3], which boils down to many DRX cycles. A DRX cycle consists of an "opportunity"

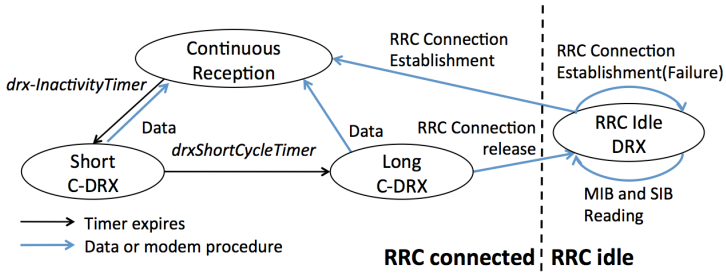


Fig. 2. LTE RRC state machine.

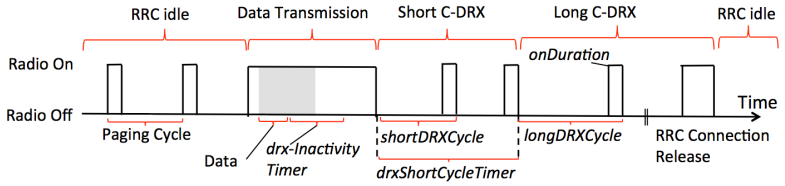


Fig. 3. Illustration of LTE DRX during RRC Connected and RRC Idle states.

period during which the modem radio is off and the UE stops monitoring PDCCH (Physical Downlink Control Channel), and an "on duration" period during which the modem radio is on and the UE monitors PDCCH. During RRC Connected, there are two types of DRX cycles, Short DRX cycles and Long DRX cycles, with the periodicity of shortDRXCycle and longDRXCycle, respectively. Long DRX cycles are beneficial for extending the battery life and short DRX cycles have faster response to data transfer from the network. During RRC Idle, a DRX cycle is a paging cycle, which is 1.28 second long under the default setting.

**RRC-state based modem power model.** Existing modem power models [11, 16, 17] are RRC-state based finite-state machines that (1) inherit the RRC states in the RRC Finite State Machine, (2) use data send/receive events that can be captured at the user/kernel level as the transition triggers among the states, and (3) annotate each RRC state with a single average power to approximate the power draw while the modem stays in that state. Such power models are therefore *coarse-grained*; they operate at the RRC-state granularity and do not model the detailed modem activities happening while staying inside any one of the RRC states.

To derive the average power assigned to each RRC state, controlled experiments are performed to drive the modem into the RRC Connected state and the average power draw at each state is measured using a powermeter. Such power models [11, 16] use either system calls or packets in the TCPdump trace as the triggers of transition from the RRC Idle state to the Continuous Reception state, and use estimated timeout durations as the triggers of the transitions among the remaining states.

### 2.3 Limitations of prior power model

The above RRC-state based power models were developed to model the modem energy drain of application data transfer. As such, they served well their original purpose (*e.g.,* [9, 11]), but they are insufficient to study detailed modem behavior, *i.e.,* activities in the *control plane*, for example, for detecting potential anomalies of a particular modem implementation of the LTE standard and opportunities for optimizing the LTE control protocol design.

Specifically, as shown in Figure 3, one RRC state may consist of multiple switching events between modem radio-on and radio-off, and one radio-on period may contain one or more modem procedures or timer intervals, where *modem procedures* are defined in the LTE protocol, which include RRC-Connection-Request, RRC-Connection-Release, data transmit, data receive, *etc.*, and *Timer intervals* are the intervals during which the modem waits until the expiration of LTE timers, such as the drx-InactivityTimer and onDuration timers. Such modem procedures and timer activities can only be logged at the PHY to RRC layers.

## 3 EVENT-BASED POWER MODEL

In this section, we introduce a fine-grained LTE modem power model that can capture all modem control activities, which enables researchers to gain insights into modem energy drain in the control plane, *e.g.,* to identify abnormal energy behavior and potential optimization of the LTE control protocol.

### 3.1 Challenges and approach

Developing such a fine-grained model power model faces two major challenges. First, we need to log all modem activities including air interface RRC messages exchanged between a UE and the network/eNodeB and PHY activities such as UE local inter/intra frequency idle mode measurements. To overcome the challenge, we exploit the observation that Qualcomm chipsets which reside in major Android smartphones expose modem messages via /dev/diag, and leverage a recently developed tool called MobileInsight [14] which logs the exposed modem messages and further decodes raw cellular logs in binary streams to readable LTE protocol messages.

A second challenge is how to derive a fine-grained power model that accurately models the modem power draw due to all modem control and data activities. As discussed in §2.3, one radio-on period contains one or more modem procedures, which can be captured by the modem logs. One obvious option is to model the modem power draw at the modem procedure level. However, the procedures in the same radio-on period may overlap with each other, in which case the powermeter can not isolate the energy drain due to individual procedures. Additionally, some procedures, such as security activation, happen within several milliseconds, which is of comparable or smaller time scale compared to the powermeter resolution and logging delay, making it impractical to accurately align the powermeter reading among those small procedures.

To avoid the limitations of modeling at the procedure level, we propose to model modem power behavior at a slightly higher level, the "event" level. In the following, we first introduce the notion of modem events. We then present the methodology for building the event-based power model, followed by the derived fine-grained power model.

### 3.2 Modem events

We define a *modem event* as a set of modem operations that happen in a single radio-on period. The modem events thus defined have the following properties:

- The modem radio keeps on during an event.
- No two events share the same radio-on period.
- Procedures in one event belong to the same RRC state.
- Procedures in one event serve the same purpose according to the LTE protocol.
- Each type of events have the same power behavior, and consequently can be easily modeled using the same, simple function in the new power model.

We denote a modem power model at the granularity of events thus defined as an *event-based power model.*

Modeling the power of each modem event has the following advantages: (1) **Practical:** It can model individual modem events using commodity power measurement tools, *e.g.*, the 5Khz Monsoon powermeter [1]. (2) **Fine-grained:** It can capture the start and end time of each event and the procedures inside the event, whileas the state-based power model can only capture the triggers of transitions from RRC Idle to Continuous Reception and has to estimate the rest of model parameters. (3) **Complete:** It can capture both control and data events, whileas RRC state-based power models could only model data transmissions. In short, such an event-based power model can capture microscopic modem energy behavior, which is needed to detect potential modem abnormal energy drain and optimization opportunities.

### 3.3 Step 1: Classifying Modem Events

In principle, the sequence of procedures and timers that happen in each radio-on period, which form a modem event, and each radio-off period, can be found in the LTE protocol specification [4]. For example, the first procedure or timer in the sequence indicates entering a radio-on or radio-off period, and the last one indicates exiting the period. Therefore, we can use a fine-grained FSM similar to the RRC State Machine to model and capture the transitions between radio-on periods and radio-off periods. Afterwards, we just need to classify each radio-on period as one of the modem events based on the actual transitions of that radio-on period.

From the RRC State Machine, we observe one modem event corresponds to (1) one RRC state, (2) one distinct in-bound inter-state or intra-state transition, and (3) one distinct out-bound inter-state or intra-state transition. Since the intra-state transitions are hidden in the RRC State Machine, we uncover them in the fine-grained FSM and classify the modem events as follows.

First, we map each state in the RRC State Machine with a distinct in-bound transition and a distinct out-bound transition as shown in Figure 2 into a sub-FSM (shown in Figure 19 in Appendix), where a node in the sub-FSM is either a radio-on period or a radio-off period, and a transition is either a modem procedure or an LTE timer. In total, there are seven sub-FSMs corresponding to 7 distinct modem events.

Second, we connect the sub-FSMs expanded from each RRC state using the inter-state transitions from the original RRC state machine, to form a single fine-grained FSM, shown in Figure 4 - this is the new event-based FSM modem power model. In this fine-grained FSM, each radio-on state along with an in-bound transition and a distinct out-bound transition maps to a modem event. For example, the "Short C-DRX Radio on" state with an in-bound transition of the shortDRXcycle timer and an out-bound transition of the onDuration timer maps to event E3 (C-DRX wakeup without data). Given a modem activity trace, we can use the FSM to extract occurrences of modem events by driving the FSM using the modem messages in the trace.

The detailed behavior of the 7 modem events are as follows.

**E1: Continuous Reception after connection establishment.** This event is initiated by the UE during the RRC Idle mode, by sending RRC-Connection-Request. After the UE receives RRC-Connection-Setup from E-UTRAN, the connection establishment is successful and it enters the RRC Connected mode. Afterwards, continuous data reception proceeds and drx-InactivityTimer is reset whenever a data block is sent or received. The time when drx-InactivityTimer expires after the last data transfer marks the end of the event.

**E2: Continuous Reception during C-DRX.** During RRC Connected, if the UE wants to send data, it can initiate transmission at any time. If the UE has data to receive, it is delayed until a C-DRX wakeup. A C-DRX wakeup means when the shortDRXCycle timer expires during short C-DRX or the longDRXCycle timer expires during long C-DRX, the radio is turned on, and the onDuration timer is started. Two cases can happen before the onDuration timer expires, with data

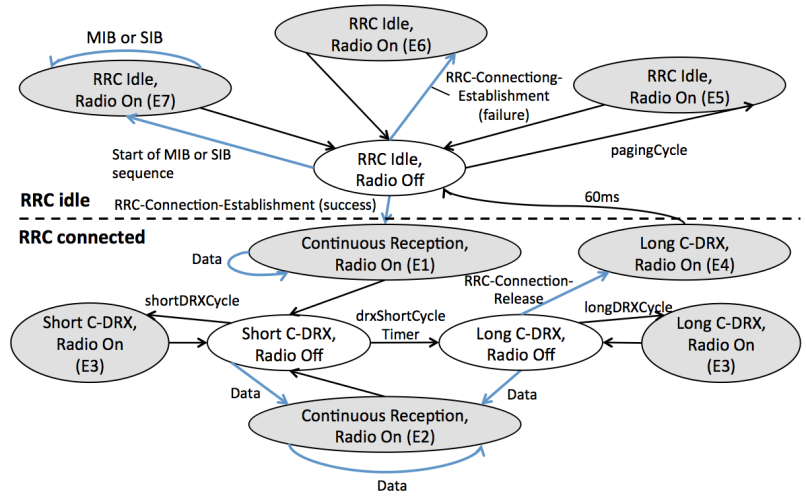


Fig. 4. The event-based FSM modem power model. Each shaded circle corresponds to one of the seven modem events.

or without data. Event E2 corresponds to the former case. The modem data reception behavior during data transmission is the same as in E1.

**E3: C-DRX wakeup without data.** This event corresponds to the case in which there is no data after C-DRX wakeup. In this case, the event ends when the onDuration timer expires. We do not distinguish Short and Long DRX cycles here because they share the same wakeup duration.

**E4: RRC connection release.** This event happens during the RRC Connected state, mostly initiated by E-UTRAN for common cases with reasons of loadBalancingTAURequired, cs-FallbackHighPriority, or others. It can also be initiated by the UE when the upper layers indicate barring of PCell. After a timeout of 60ms, the UE releases all radio resources and reset necessary states and enters the RRC Idle state.

**E5: DRX paging during RRC Idle state.** This event happens during the RRC Idle state, initiated by E-UTRAN. The UE modem wakes up to monitor one Paging Occasion (PO) per DRX paging cycle. The DRX paging cycle is decided by the defaultPagingCycle parameter or the UE-specified DRX cycle.

**E6: RRC connection establishment failure.** There are three main indicators of a failed RRC connection establishment: (1) Rejected connection by E-UTRAN indicated by receiving RRC-Connection-Reject, (2) Timer T300 expiration, and (3) Cell re-selection when timer T300 runs. After this event, the UE stays in the RRC Idle state.

**E7: MIB and SIB reading during RRC Idle state.** During RRC Idle, PLMN search and initial cell selection/re-selection require the UE to receive and decode system information of the cells in MIB (Master Information Block) and SIBs (System Information Blocks), in order to evaluate the quality of the camping cells and neighbor cells. The timing of reading MIB and SIBs are deciphered by the UE after knowing the SFN (System Frame Number) of the cell.

### 3.4 Step 2: Modeling event energy drain

With the seven modem events classified, we next need to measure the energy consumption of each event, via controlled experiments. The straight-forward approach is to reproduce each event on the UE, and measure its energy consumption using a powermeter connected to the UE. However, we

can only run micro-benchmarks or subject the UE to certain external conditions (*e.g.,* mobility), but not directly manipulate inside the modem and play one event at a time. To overcome this challenge, we adopt the following 4-step methodology: (1) Manually play a set of usage scenarios; (2) Log modem messages and measure power draw using the powermeter; (3) Parse events from message log using the FSM in Figure 4; (4) Correlate energy and events to derive the energy drain per event.

**Usage scenarios.** To increase the possibility of encountering each type of events, we design a set of diverse UE usage scenarios. First, we use 2 types of network traffic: (1) UE receives TCP packets every 1 minute and (2) UE sends TCP packets every 1 minute. The data sizes are 4K, 40K or 400K, interleaved in each experiment. Second, we exercise different signal conditions in two ways: (1) statistically find 3 types of signal conditions: good (RSRP $\geq -95dBm$), medium ($-95dBm >$ RSRP $\geq -110dBm$), and poor (RSRP $< -110dBm$); (2) drive the device in a car which exposes the device to any possible signal conditions. We then perform experiments under the 8 combinations or scenarios of network traffic types and signal condition types. The purposes of designing the above usage scenarios are as follows: (1) Network send/receive operations of different packet sizes are to simulate different types of user traffic. (2) Different signal strength are to simulate different network conditions. (3) Staying in one place and driving the phone in a vehicle are to simulate different user mobility. For each scenario, we subject the UE for 20 minutes and repeat for 4 times, of which two are used for power modeling, and the remaining two are used for validating the model (§4).

The UE used in the control experiments is a Nexus 6 phone, which runs Qualcomm Snapdragon 805 processor, with Gobi$^{TM}$ 4G LTE World Mode, LTE category 6 and RRC release number 11.

**Logging modem messages and measuring power draw.** We use MobileInsight (MI) [14] to log modem messages. MI records modem raw logs exposed by Qualcomm chips through a side channel on the UE (recorder) and decodes the raw logs to modem messages based on LTE protocol standard headers (decoder). We run the recorder on the testing UEs and the decoder offline on a laptop to obtain the modem messages.

We use a Monsoon powermeter [1] to measure the instantaneous power draw of the phone by attaching measurement pins to the phone's battery. The powermeter has a resolution of 5KHz, and outputs the instantaneous power of the whole device, which includes both CPU power and modem power under our testing setups. We estimate the CPU power using a utilization-based CPU model derived similarly as in [9], and remove the CPU power from the powermeter output to calculate the modem power.

**Extracting events from modem message log.** First, we identify modem procedures and LTE timers from the logged modem messages. Since the MAC layer is the lowest layer containing data packets logging, we treat MAC-Transport-Block as a data procedure. The detailed timer values are obtained from SIB2 and RRC connection reconfiguration messages.

Second, we drive the FSM in Figure 4 with the identified procedures to parse events. When an event is parsed, we record (1) the state entering time as event start time, (2) the state exiting time (not including re-enter transition) as event end time, and (3) extra field from procedures as appropriate, *e.g.,* signal strength and the number of sent/received bytes.

Because of the delay of MI logger logging and the delay between the MAC layer and PDCCH decoding, it is difficult to capture the exact time of radio-on based on LTE timers alone. We observe that whenever the radio is on, there are PHY-Serv-Sell-Measurement or PHY-Connected-Mode-Intra-Freq-Meas messages. We use these two messages and relaxed DRX cycle timers (by 10ms) to detect DRX wakeups as follows: whenever the DRX cycle timer expires at $t_1$ during the parsing, we look for the first PHY-Serv-Sell-Measurement and PHY-Connected-Mode-Intra-Freq-Meas appeared within $[t_1 - 10ms, t_1 + 10ms]$, and treat the time of that first message as the DRX wakeup time.

**Correlating energy with events.** The final step is to find the corresponding energy consumption of each event. This is accomplished by processing the powermeter reading to remove the CPU

power as discussed above, and integrating the processed powermeter reading between the start and end time of each event into the event energy.

### 3.5 Derived event-based power model

Using the methodology discussed above, we constructed the event-based power model for the Nexus 6 modem. We note that same as with the prior state-based power model, we need to train the new power model for each individual mobile device, because the power behavior of each device may be different due to different modem implementation or OS driver behavior. Once trained for a device, the model should be able to capture all control plane modem energy drain behavior in-the-wild, because we have carefully designed the user scenarios in training the new model to cover all typical daily usage scenarios, as discussed in Section 3.4.

**Model overview.** The power model takes as input a modem event along with its parameters including the type of the event, duration and extra field as appropriate, *e.g.,* signal strength, number of sent bytes and number of received bytes, and outputs the estimated energy of the event.

An example of modem events E1–E5 appearing in a continuous time interval from RRC Idle to RRC Connected and back to RRC Idle and the corresponding powermeter reading are shown in Figure 5(a) Detailed figures of the 5 events are shown in Figure 5(b)–(f). Similar examples of events E6 and E7 are shown in Figures 5(g)(h).

Next, we summarize the derived power model for each modem event, along with detailed explanation. The average value and standard deviation (in brackets) of the derived parameters are shown in Table 1.

**Continuous reception after connection establishment and during C-DRX (E1, E2).** As shown in Figures 5(a), the power draw of these two events contain three main parts.

**Start-up energy** is the energy consumed before continuous reception. For continuous reception after connection establishment (E1), it is consumed by modem ramp-up and the initiation of RRC connection before continuous reception. For continuous reception during C-DRX (E2), if the data is initiated by the network, *e.g.,* downlink data, it can only happen when a C-DRX cycle expires (*i.e.,* it is time for a short/long DRX radio-on event), and the start-up energy contains both ramp-up energy and the energy consumed during onDuration timer when it waits for the first received data packet; if the data is initiated by UE, *e.g.,* uplink data, it can happen any time before C-DRX cycle expires, and the start-up energy contains only the ramp-up energy. We use the start time of the event to differentiate the two cases: if the event starts when C-DRX cycle expires (*i.e.,* C-DRX wakeup), it is downlink data initiated by the network; otherwise, it is uplink data. We separately model the start-up energy of the above three cases as three constants.

**Base power** is the minimum power when the radio is on. The base power happens after data occurs, including the continuous reception duration and drx-InactivityTimer duration. The duration of continuous reception is decided by the start time and end time of MAC layer data blocks, and the duration of drxInactivityTimer can be explicitly read from the RRC-Connection-Reconfiguration message.

**Transmission power** on top of the base power is consumed by data transmission during continuous reception. Following the known correlation between data transmission power (and data rate) and signal strength [5, 9, 10], we model the transmission energy per byte as a function of signal strength.

**C-DRX wakeup without continuous reception and DRX paging during RRC Idle (E3, E5).** Since what the modem does during these two events is simply to wake up, measure and decode the paging channel, wait until the onDuraion timer expires, and go back to sleep, there is little variation during the events. For simplicity, we respectively model the energy drain of two events as constants. The derived parameters in Table 1 shows that the energy of DRX paging during RRC

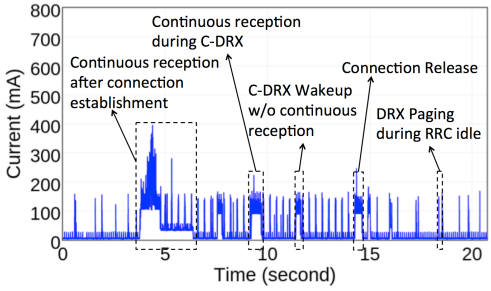


(a) Modem events E1–E5 appearing in a continuous interval from RRC Idle to RRC Connected and back to RRC Idle (good signal strength)

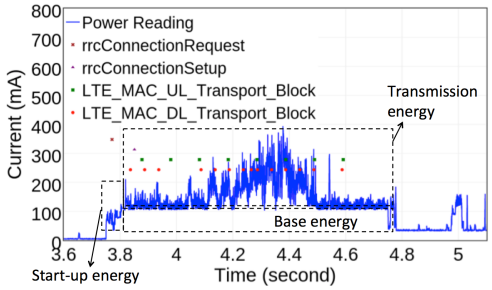


(b) E1: Continuous reception after connection establishment

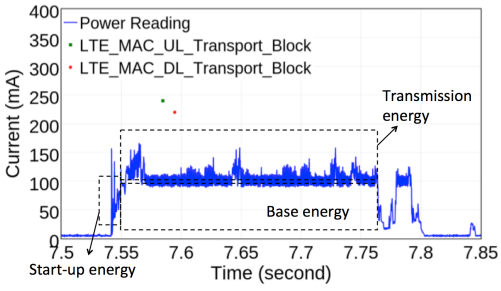


(c) E2: Continuous reception after C-DRX wakeup

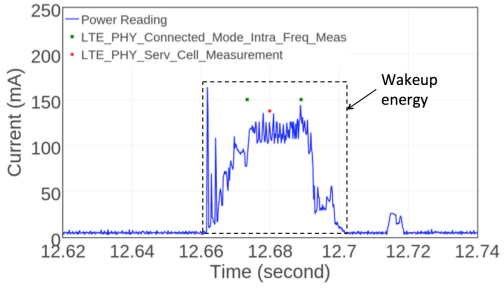


(d) E3: DRX wakeup without continuous reception

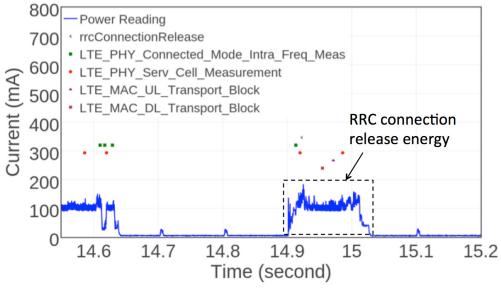


(e) E4: RRC connection release

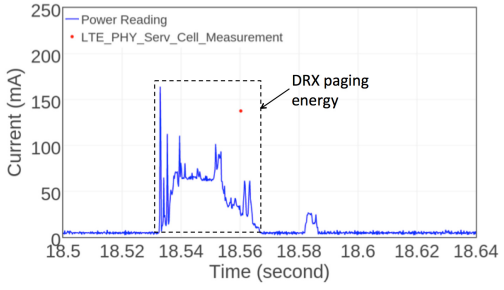


(f) E5: DRX paging during RRC Idle mode

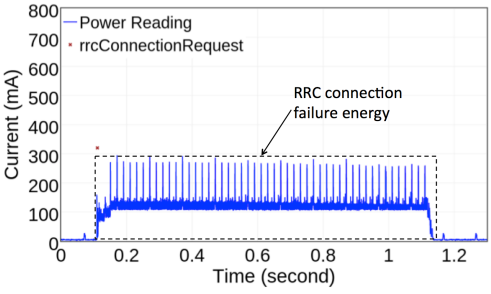


(g) E6: RRC connection establishment failure (poor signal strength)

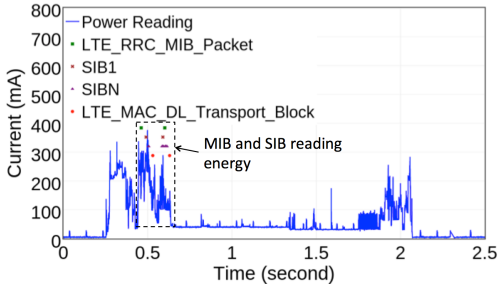


(h) E7: MIB and SIB reading during RRC idle (poor signal strength)

Fig. 5. Powermeter reading and modem messages of seven modem events.

Table 1. The average value and standard deviation (in parentheses) of the derived parameters of the event-based power model.

| | Energy / Power (Standard deviation) | | |
|---|---|---|---|
| | Good signal | Medium signal | Poor signal |
| **Continuous reception after connection establishment (E1) and Continuous reception during C-DRX (E2)** | | | |
| Start-up of E1 (mAs) | 7.26 (1.30) | | |
| Start-up of E2 initiated by network (mAs) | 10.00 (1.74) | | |
| Start-up of E2 initiated by UE (mAs) | 4.83 (0.25) | | |
| Tx Energy / KB (mAs/KB) | 0.42 (0.18) | 0.57 (0.08) | 1.26 (0.37) |
| Rx Energy / KB (mAs/KB) | 0.22 (0.05) | 0.36 (0.07) | 0.48 (0.12) |
| Base power (mA) | 95.65 (6.48) | | |
| **C-DRX wakeup without data (E3)** | | | |
| Total energy (mAs) | 3.04 (0.18) | | |
| **RRC connection release (E4)** | | | |
| Total energy (mAs) | 16.53 (1.68) | | |
| **DRX paging during RRC Idle mode (E5)** | | | |
| Total energy (mAs) | 1.02 (0.03) | 1.18 (0.19) | 1.31 (0.19) |
| **RRC connection establishment failure (E6)** | | | |
| Total energy (mAs) | 108.53 (9.14) | | |
| **MIB and SIB reading during RRC Idle mode (E7)** | | | |
| Total energy (mAs) | 16.82 (3.59) | | |

Idle significantly increases as the signal strength becomes worse, due to increasing data loss under poor signal strength.

**RRC connection release (E4).** As discussed in §3.3, after the UE receives RRC-Connection-Release, it waits for 60ms, releases all radio resources, resets necessary states, and enters the RRC Idle state. Thus it is safe to model the power draw as a constant because there is little variation in terms of modem radio behavior. The low standard deviation of the energy of the event occurrences shown in Table 1 validates the feasibility of the constant model.

**RRC connection establishment failure (E6).** Figure 5(b) shows the power drain of an RRC establishment failure event because of T300 timer expiration. We see after RRC-Connection-Request, the radio stays on for 1000ms, during which there are periodical spikes when the modem performs RACH attempts. We again model the event as a constant energy. The establishment failures with the other two reasons, rejected by E-UTRAN and cell re-selection during T300 timer, are rare cases and cannot be captured by our controlled experiments, and hence are not captured by our current model and we leave them as future work.

**MIB and SIB reading during RRC Idle mode (E7).** Figure 5(c) shows the power draw of an MIB and SIB reading event. For simplicity, we model the total energy of this event.

## 4 MODEL VALIDATION: ACCURACY AND OVERHEAD

In this section, we validate the event-based model by comparing its accuracy and overhead with the state-based model.

**Methodology.** We measure how the event-based power model improves the modeling accuracy over the state-based model using the two extra runs of each usage scenario from the same controlled

experiments in §3.4. We first build the state-based power model for Nexus 6 using the methodology in [9] (the details are in Table 2 in Appendix). We then compare the energy estimation using the event-based model and the state-based model with the powermeter reading subtracted by CPU power which serves as the ground truth modem energy drain. The CPU power is estimated using the methodology discussed in §3.4.

We quantify the estimation error of the model at 5-second intervals as follows. We integrate the power in every 5-second interval from model estimation and from the ground truth to get the per-5-second energy drain and hence average power draw, and then calculate the difference between estimated and ground truth average power of each 5-second interval.

In the following, we present the 5-second interval error of energy estimation for all Continuous Reception for user data, LTE tails, and RRC Idle periods in the controlled experiments (the two extra runs per scenario mentioned in §3.4) in turn.

### 4.1 5-second error during continuous reception for user data

Figure 6(a) shows the CDF of 5-second interval error of the two models during Continuous Reception for user data, under varying signal strength. We see that under good, medium and poor signal condition (defined in §3.4), respectively, the event-based model has less than 20mA absolute error in 60%, 59% and 70% of intervals, while the state-based power model has less than 20mA absolute error in 21%, 13% and 9% of the intervals, and under-estimates the modem energy by more than 20mA in 82%, 61% and 94% of the intervals.

To understand why the event-based model has much better accuracy, we plot in Figure 6(b) an example of how well the power draw timeline from the two power models fit with the ground truth power during one RRC Connected period. We calculate the estimated average power as the estimated energy of each event divided by the event's duration. Figure 6(c) further shows the comparison during Continuous Reception after RRC Connection Establishment for user data zoomed in from Figure 6(b). Compared with the ground truth power, the event-based model estimation matches the start and end of the power spikes better than the state-based model. This can be explained as follows: (1) Because the state-based model estimates the duration of continuous reception using the total user level data size and estimated modem data rate, the end of continuous reception is often not estimated accurately. (2) Since the state-based model does not capture Connection Establishment, it misses the beginning part of the power draw spike. In contrast, since the event-based model captures individual lower layer MAC layer data blocks, it accurately captures the ending of each continuous reception interval. Further, it can estimate the start time of each spike more accurately since it can capture connection establishment.

### 4.2 5-second error during tails

The continuous time interval when there are no modem activities involving transmission of packets originated from or destined to above the RRC layer during short C-DRX and long C-DRX was historically dubbed as the LTE *tail* [6], because the period follows active data transmission (Continuous Reception), does not involve data transmission, yet the modem is draining power. The fine-grained event-based model exposes four types of modem activities during an LTE tail: (1) wakeups after C-DRX cycle expiration and without continuous reception which corresponds to "DRX wakeup without continuous reception (E3)", (2) continuous reception for MAC layer control messages which corresponds to "Continuous reception during C-DRX (E2)", *e.g.,* BSR (Buffer Status Report), (3) continuous reception for RRC control messages which corresponds to "Continuous reception during C-DRX (E2)", *e.g.,* security-mode-command and UE-capacity, and (4) RRC connection release which corresponds to "RRC connection release (E4)". We denote a tail

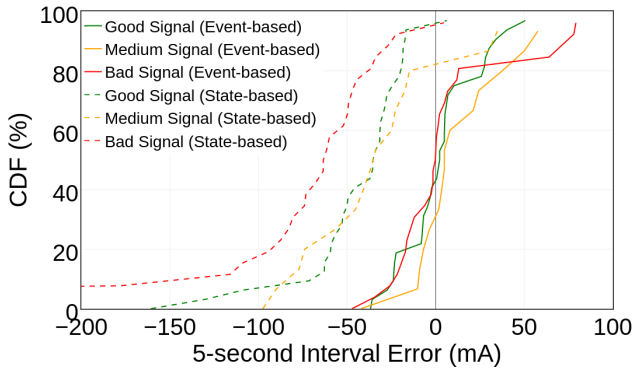


(a) CDF of 5-second interval error during continuous reception for user data

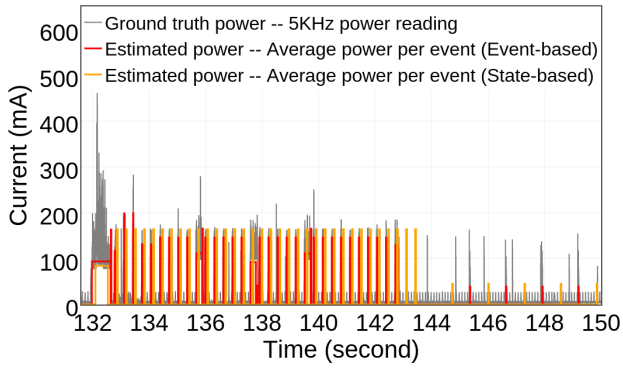


(b) Estimated and ground truth power during one RRC connected period.

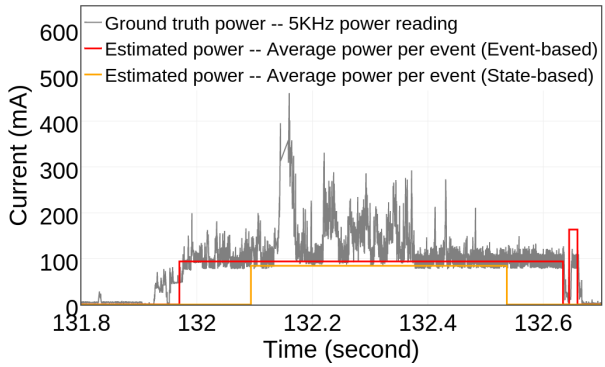


(c) Estimated and ground truth power during continuous reception after RRC connection establishment

Fig. 6. Accuracy of the two models during continuous reception of user data.

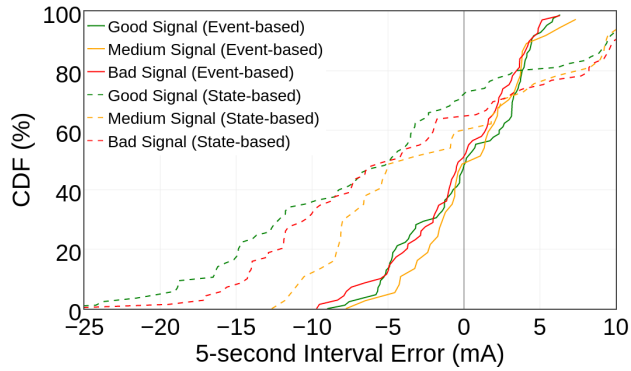


(a) CDF of 5-second interval error during tails

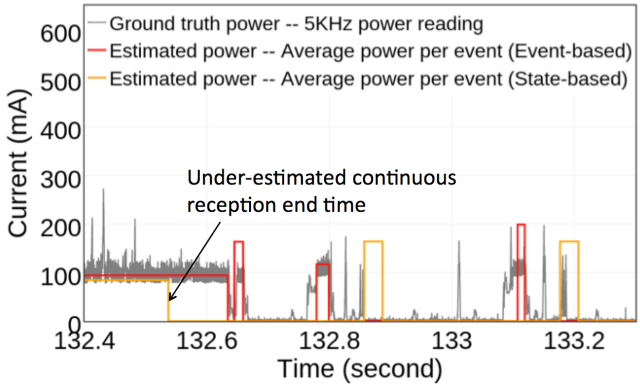


(b) Power draw comparison amid transition from continuous reception into a tail

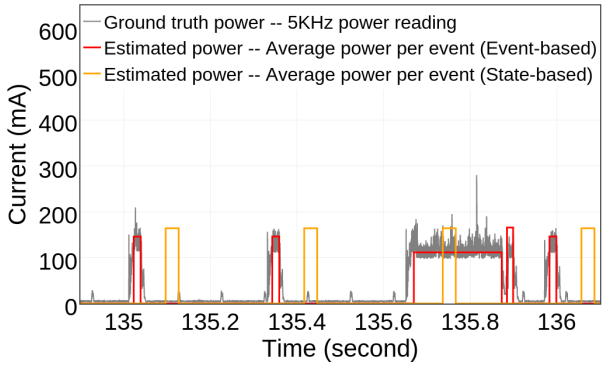


(c) Power draw comparison in the middle of a tail

Fig. 7. Accuracy of the two models during LTE tails.

that ends with RRC connection release as a *full tail*, and a tail that ends with the next continuous reception as a *partial tail*.

Figure 7(a) shows the CDF of 5-second interval error when the modem is in all tails, including full tails and partial tails. We see that the event-based power model has less than 5mA absolute error in 80%, 85% and 81% of the intervals, while the state-based power model has less than 5mA absolute error in 28%, 33% and 23% of intervals, and under-estimates energy by more than 5mA in 52%, 43% and 52% of intervals, under good, medium and poor signal condition, respectively.

We again plot estimated and ground-truth power lines to illustrate the two reasons for why the event-based model estimates tail energy more accurately. (1) *Accurate start time:* Figure 7(b) shows the power of the beginning part of the tail right after continuous reception, zoomed in from Figure 6(b). Since the state-based estimates each wakeup with some fixed duration and periodicity, it under-estimates the end time of continuous reception here, which is used as the start time of C-DRX wakeup cycles, *i.e.,* it incorrectly estimates the C-DRX wakeups with a time offset compared to the ground truth. In contrast, the event-based model is able to accurately estimates the end time of continuous reception using MAC layer data blocks, and thus the following C-DRX wakeup timing is estimated accurately. (2) *Accurate duration:* Figure 7(c) shows the power line in the middle of the tail, zoomed in from Figure 6(b). We see the third spike is continuous reception for control-related modem activities which lasts much longer than the DRX cycle onDuration interval; it can only be captured by the event-based power model but was incorrectly guessed using the default onDuration value by the state-based power model.

### 4.3 5-second error during RRC Idle

Figure 8(a) shows the CDF of 5-second interval error during RRC Idle intervals. We see the event-based power model has less than 0.2mA absolute error in 80%, 80% and 41% of the intervals, while state-based power model has less than 0.2mA absolute error in 78%, 27% and 48% of the intervals, and under-estimates the energy for more than 0.2mA in 17%, 71% and 50% of the intervals, under good, medium and poor signal condition, respectively. Even though under poor signal strength, the state-based approach has slightly higher percentile of absolute error less than 0.2mA (48% of the intervals) compared with event-based (41% of the intervals), overall the state-based approach underestimates the energy consumption. The reason the state-based power model significantly under-estimates modem energy for medium and poor signal conditions is that unlike the event-based model, it cannot capture cell selection / re-selection and PLMN search modem activities during RRC Idle, which happen more often during weak signal strength.

We again plot estimated and ground-truth power lines during RRC Idle to illustrate the two reasons for why the event-based model estimates tail energy more accurately. Figure 8(c) shows the power draw of the the estimated average power from the two power models along with the ground truth power reading during an RRC Idle interval. (1) *Accurate RRC release time and hence DRX paging start time:* The event-based model captures the exact RRC connection release time, while the state-based model estimates the time of RRC connection release as the end of the estimated constant LTE tail duration, which is not accurate, as shown in Figure 8(b). Since the start time of DRX paging during RRC Idle is RRC Connection Release, the state-based model incorrectly estimates the DRX paging wakeups with a time offset compared to the ground truth. (2) *Capturing other modem activities during RRX Idle:* The second energy spike in Figure 8(c) is for cell re-selection, which is captured by the event-based model but not estimated by the state-based model.

### 4.4 Logging overhead and deployment requirement

**Overhead.** Since the event-based power model performs extra logging of modem activities using MI Logger, we compare the logging overhead of the two models. Note for the state-based model,

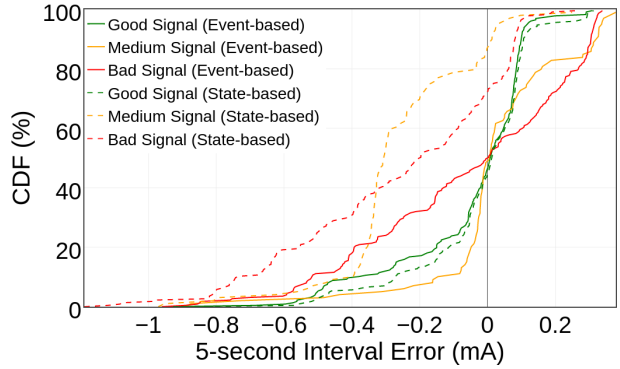


(a) CDF of 5-second interval error during RRC Idle

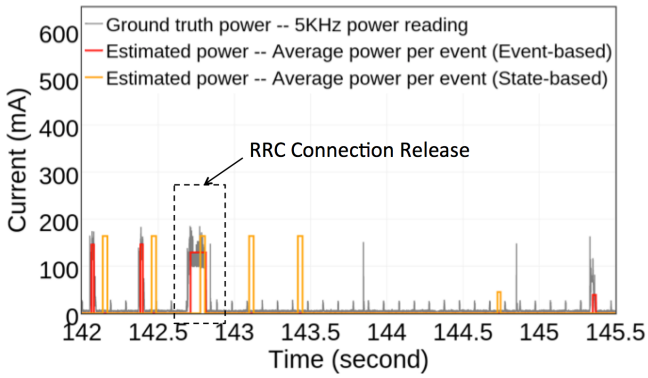


(b) Power draw comparison during RRC Connection Release.

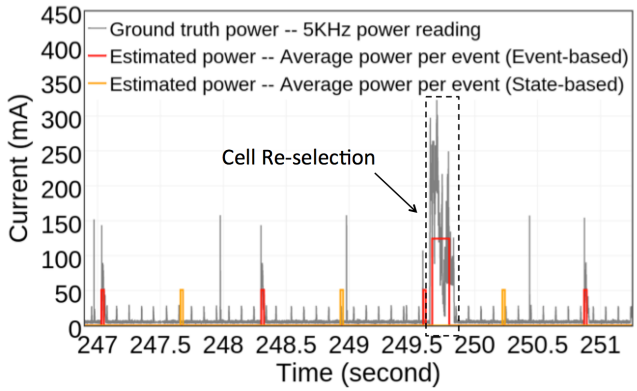


(c) Power draw comparison during RRC Idle

Fig. 8. Accuracy of the two models during RRC Idle.

we us the light-weight network-usage logger in [9], as opposed to the heavy-weight tcpdump used in [11].

Below we compare the CPU time overhead of two loggers in a few typical usage scenarios covered by the controlled experiments. We will further measure the MI logging overhead on UEs in the wild in §5.1.

In the first controlled experiment, a micro-benchmark running on the phone is receiving 40KB TCP packets every 1 minute under medium signal strength for 20 minutes. The average CPU overhead of MI Logger and the utilization-based logger are 0.047% and 0.72%, respectively. In the second controlled experiment, we measure the overhead of the two loggers when the phone screen is off and is either stationary or moving in a car. The average CPU overhead of the two loggers during are 0.020% and 0.48% when stationary, and 0.024% and 0.45% when moving, respectively. These results suggest logging overhead of MI Logger is insignificant compared to the packet logging needed for the state-based power model.

**Deployment requirement.** Deploying MI Logger on users' phones requires rooting the phone and setting the phone to the BP mode when starting up.

## 5 IN-THE-WILD STUDY

In this section, we report on an in-the-wild study of modem energy consumption of UEs under normal daily usage that makes use of the event-based fine-grained modem power model developed in §3. We present three detailed case studies to showcase the new types of analysis enabled by the event-based modem power model. We note that since it is impossible to hook user devices in the wild with a powermeter, we cannot provide ground truth data from a powermeter as we did in the controlled experiments in §4.

### 5.1 Trace collection

**Logging system.** To estimate the energy consumption of the user phones in the wild, in addition to MI Logger, we designed a light-weight My Logger to record upper layer device status and user external context, *e.g.,* user mobility, location, utilization of CPU and network, and screen status. The log files of both loggers are periodically compressed and uploaded to the server for offline energy drain analysis using the event-based power model. Uploads happen when the UE is connected to WiFi and the battery level is above 10%.

**Trace statistics.** We deployed 12 Nexus 6 phones with the two loggers preinstalled to 12 volunteers who used them for daily activities, with subscriptions to AT&T, Verizon, or CT-Mobile, over a 2-month period [1]. We collected a trace with 29 days per user on average.

**Logging overhead.** We extract the overhead of the two loggers in terms of CPU time and network utilization from the My Logger trace. On average, the daily CPU time overhead of the two loggers is 1.3% and the network traffic from uploading the trace is 3.2% of the daily WiFi usage.

### 5.2 Overall analysis

We first analyze the average daily time spent during the seven events for the 12 users, shown in Figure 9. We see the two largest portions, Continuous reception (E1 and E2) and paging during RRC Idle (E5), account for 54% and 24% of the total daily radio-on time, respectively.

We apply the event-based modem power model to the collected traces to break down the LTE modem energy among modem events, and draw the average daily energy of the 12 users in increasing order in Figure 10. The top 3 energy consuming events are continuous reception after

[1]Our experiments involving real users received exemption from the full requirement of 45 CFR 46 or 21 CFR 56 by the IRB of our organization.

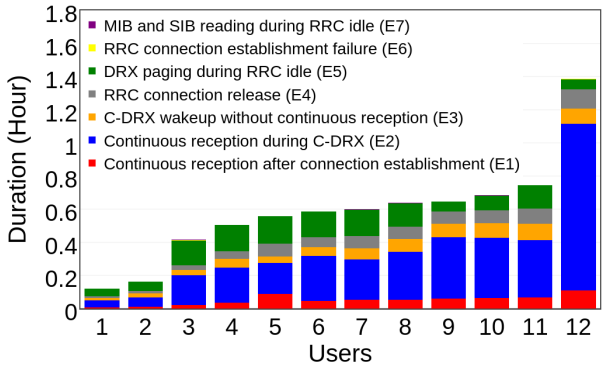


Fig. 9. Daily modem time breakdown across events.

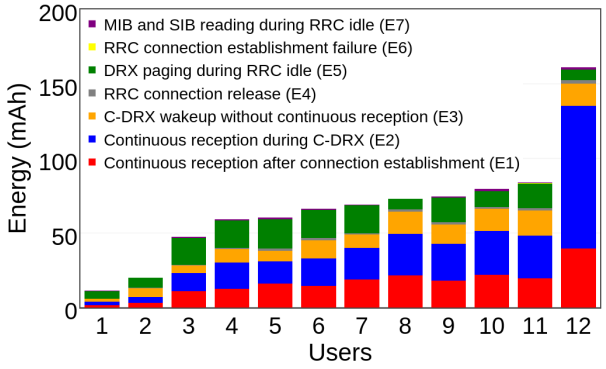


Fig. 10. LTE modem energy breakdown across events.

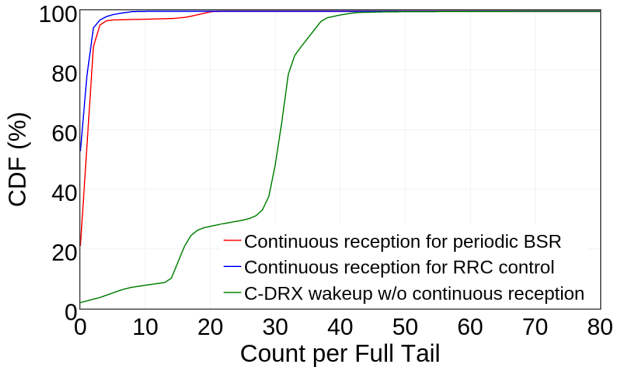


Fig. 11. CDF of count of C-DRX wakeups for different purposes in each LTE full tail.

connection establishment (E1), continuous reception during C-DRX (E2), and DRX paging during RRC Idle (E5). The medium daily energy of the three events across the users are 16.7mAh, 19.6mAh and 18.6mAh, respectively.

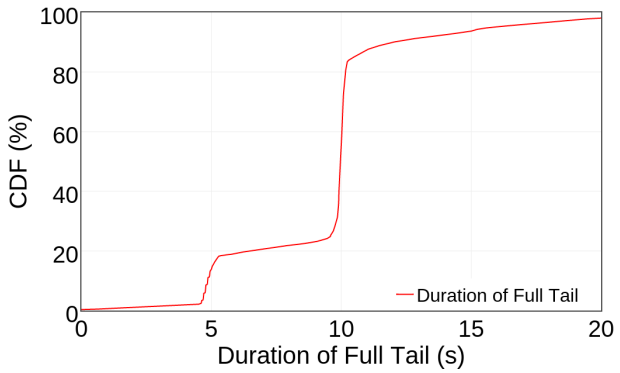


Fig. 12. CDF of duration full tails.

### 5.3 Case studies

*5.3.1 Refining LTE tail energy.* In the first case study, we show how the event-based power model enables in-depth analysis and diagnosis of LTE tail energy behavior on smartphones running in the wild.

**Exposing varying wakeup counts in tails.** Figure 11 shows the CDF of C-DRX wakeup counts caused by different modem activities in each LTE full tail. For 95% of the full tails, the number of C-DRX wakeups with continuous reception for periodic BSR and other control procedures during RRC Connected varies from 0 to 4. For 60% of the full tails, the number of C-DRX wakeups without continuous reception varies from 2 to 30. This result suggests that the counts of different types of C-DRX wakeups in different full tails can vary significantly and hence assigning the wakeup energy during tail as an average constant as in the state-based model is not accurate.

**Capturing precise tail duration.** Figure 12 shows the CDF of full tail durations. For 58% of the full tails, the tail time varies from 9.6s to 10.3s; these tails belong to the devices that connect to AT&T or Verizon network. For 16% of the full tails, the tail time varies from 4.6s to 5.3s; these tails belong to the devices which connect to CT-Mobile. For the remaining full tails, 2.4% of them uniformly range from 0s to 4.6s, 7.0% of them uniformly range from 5.3s to 9.6s, and 16.6% of full tails are longer than 10.3s. These findings indicate that the duration of full tails can vary significantly, and thus assigning the full tail duration with an average constant as in a generic state-based model is not accurate.

**More accurate tail energy estimation and breakdown.** We already discussed that the event-based model is more accurate in estimating LTE tail energy in §4.2. Here, we further show the new model can accurately break down tail energy among different modem control activities besides wakeups without continuous reception (E3).

We define one *RRC connection* as the time period between RRC connection establishment and RRC connection release. In this case study, we focus on RRC connection with user data. We apply the event-based power model to each RRC connection and break them down to continuous reception for user data and wakeup energy in tails, *i.e.,* tail energy, and further break down the latter among continuous reception for BSR, continuous reception for RRC control, C-DRX wakeup without continuous reception and RRC connection release. We randomly sample 40 RRC connections for each of the 12 users and draw the energy breakdown in Figure 13 in increasing order of the user data size in each RRC connection (right y-axis). We see that the energy of continuous reception for user data has an increasing trend as the data size increases, ranging from 0.01mAh to 0.51mAh. Further,

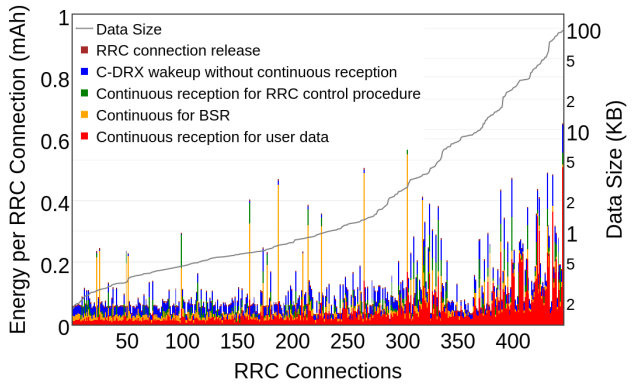


Fig. 13. Modem energy breakdown of sampled RRC connections.

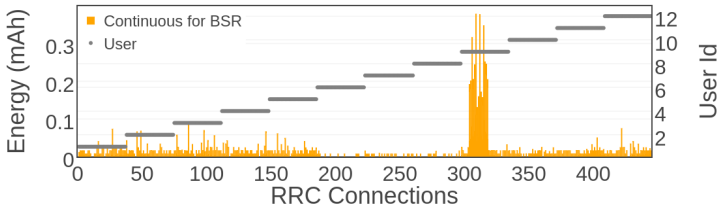


Fig. 14. Energy of continuous reception for BSR of sampled RRC connections, along with user Id.

among the 4 breakdowns of tail energy, for some RRC connections, the energy of continuous reception for BSR is significantly higher than other continuous reception for the other 3 causes.

**Diagnosing high BSR energy drain.** To analyze the reason for the high BSR energy consumption, we redraw the energy for continuous reception for BSR of the sampled RRC connections ordered by user Ids in Figure 14. We make several observations. First, the BSR energy for the samples of Users 6, 7 and 8 is lower than for the other users. It turns out they use the CT-Mobile network, while the other users use AT&T or Verizon networks. Second, all the samples with high BSR energy belong to User 9. We calculate the BSR energy of each RRC connection of User 9 and show them in Figure 15 along with the location information (using longitude since the latitude stayed the same) ordered in time. We see the BSR energy is significantly higher when the user stayed in the location in the middle part than in the other two locations. By looking at the location data, we confirmed that the location in the middle is a mountain area. Finally, we quantify the potential energy saving from removing the causes of high BSR energy drain. If we can lower the high BSR energy to normal BSR energy for User 9, 41.3% of her phone's total energy during RRC connections can be saved.

*5.3.2 Capturing energy drain of NAS procedures.* In the second case study, we show the new model can capture the energy drain of NAS procedures. We mainly focus on EMM specific procedures in the NAS layer, including attach, detach and TAU (Tracking Area Update).

The EMM specific procedures always happen during RRC Connected, either (1) in an RRC connection that was caused by this EMM specific procedure, or (2) after the first continuous reception in an RRC connection that was caused by some user-plane data. They are captured by the event-based model as either the "Continuous reception after connection establishment (E1)" or "Continuous reception during C-DRX (E2)" event.

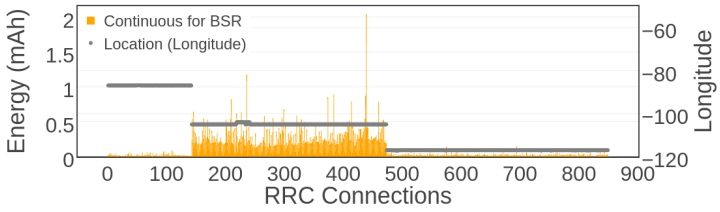


Fig. 15. Energy of continuous reception for BSR of all RRC connections of User 9, along with longitude.

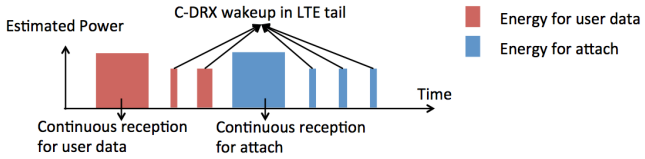


Fig. 16. Illustration of the last-trigger energy accounting policy.

Since continuous reception for both user-plane and NAS procedures could happen during the same RRC connection and thus share a same tail, a natural question is how to account the tail energy between them? We adopt the last-trigger policy [15] which accounts the tail energy to the last trigger that led to the tail. Specifically, the energy of a tail is always accounted to the continuous reception of user-plane data or NAS procedures that happened right before the tail. Figure 16 illustrates this accounting policy.

Figure 17 shows the count of EMM specific procedures per day per user in the trace. In 20% user-days, there are more than 3 attach procedures, 4 detach procedures, and 24 TAU procedures per day. Applying the event-based power model and the last-trigger accounting policy, we calculate the energy of the EMM specific procedures and draw them cumulatively for each user-day in Figure 18(a). We see that the detach energy is much lower compared to the other two procedures. This is because detach procedures are usually short, with 95.6% of them lasting less than 100ms. We also see that TAU energy varies significantly across user-days. To understand this, we analyze the correlation between other factors and TAU energy, by drawing TAU energy for each user-day, along with RRC connection count caused by user-plane data for that user-day, ordered by TAU energy, in Figure 18(b). We see that when TAU energy increases, the RRC connection count caused by user-plane data has an increasing trend. One possible reason is that release of RRC connections leads to load re-balancing TAU activities.

*5.3.3 Capturing energy of PLMN search and cell selection/re-selection.* We introduced PLMN search and cell selection/re-selection and how we detect them using in-modem trace in §3.3. Since they usually happen during the RRC Idle mode, in the following analysis, we focus on these two procedures during RRC Idle. The modem radio energy drain of these two procedures is mainly caused by MIB and SIB reception. The state-based power model cannot capture the occurrence of PLMN search and cell selection/re-selection, while the event-based power model captures them as "MIB and SIBs reading during RRC Idle (E7)" events.

Figure 18(c) shows the CDF of PLMN search and cell selection/re-selection count for each user-day in the trace. We see in 20% user-days, there are more than 10 PLMN search procedures and 151 cell selection/re-selection procedures. Figure 18(d) shows the energy of these two procedures for each user-day, ordered by the sum of their energy. We see that the rightmost 8 user-days have more than 3mAh of PLMN search and cell selection / re-selection energy per day, and for the remaining

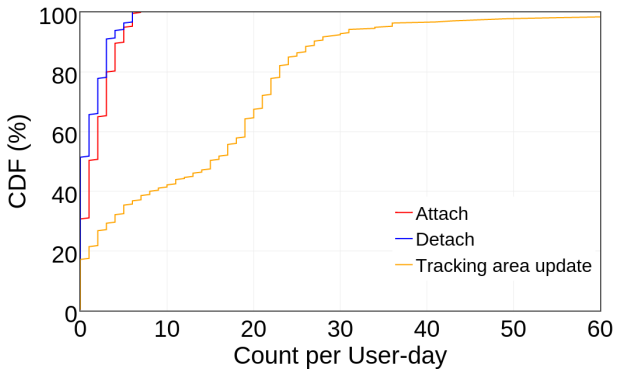


Fig. 17. Count of EMM specific procedures per user-day.

days, the combined energy is less than 3mAh and cell selection/re-selection energy is larger than PLMN search energy.

We further analyze the reasons for the above high PLMN search and cell selection/re-selection energy in the rightmost 8 days. Figure 18(e) shows cell selection/re-selection occurrence and user mobility in the user-day with highest cell selection/re-selection energy. We see cell selection/re-selection procedures happen more frequently when the user is moving. Figure 18(f) shows PLMN search occurrence and signal strength in the user-day with the highest PLMN search energy. We see that PLMN search procedures happen more frequently when the signal strength is low.

## 6 RELATED WORK

**Modem power analysis and modeling.** There have been a number of modem power analysis studies of smartphones. In [11], Huang *et al.* study the power characteristics and derive a power model for a commercial LTE network, and make comparison of power efficiency among WiFi, 3G and LTE. In [10], Ding *et al.* conduct a measurement study of the impact of poor signal strength on the smartphone energy consumption. Neither study captures the microscopic energy behavior of the modem. In [13], the authors perform controlled experiments using an eNodeB emulator and smartphones to measure the impact of Rx and Tx powers and Modulation and Coding Scheme on the phone's supply power draw to study the energy efficiency of different parameter configurations, following a slightly refined RRC-state-based power model. In [8], Catovic *et al.* analyze the impact of different scheduling mechanisms of SIB on the smartphone battery life in UMTS networks, but does not give a power model to estimate the energy drain under different scheduling schemes. In [12], Koc *et al.* develop an analytical model to estimate the power saving and delay through different DRX configurations, but do not perform any measurement on real user phones and applications.

**Smarpthone measurement study.** In [17], Qian *et al.* study the 3G network performance using 3GTest data. In [21], Xu *et al.* study the smartphone usage pattern via network measurement from cellular network provider. In [19], Sommers and Barford study the WiFi and cellular performance using the Speedtest.net data. None of these work however study the energy drain of real-world user behavior and apps in the wild. Chen *et al.* [9] perform a large-scale measurement study of energy drain of 1520 Galaxy S3 and S4 devices in the wild and present detailed analysis of where the CPU time and energy are spent across the devices and inside the various apps. Rosen *et al.* [18] examine a range of excessive energy consumption problems caused by app background network traffic through a two-year study of 20 smartphone users. Neither work studies the energy drain of the control plane in detail.

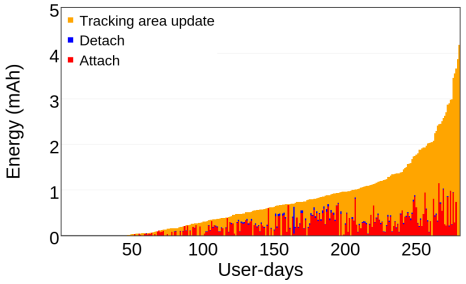


(a) Energy of EMM specific procedures per user-day

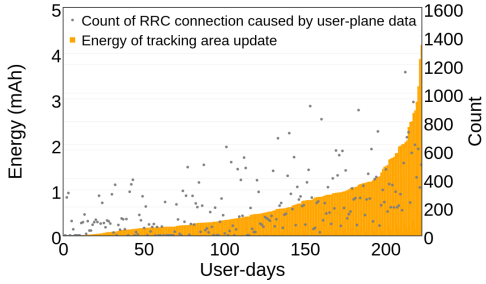


(b) Correlation between TAU energy and RRC connection count per user-day

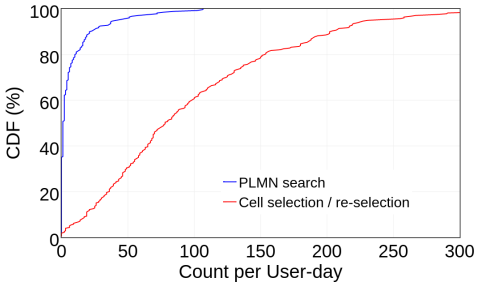


(c) CDF of count of PLMN search and cell selection/re-selection per user-day

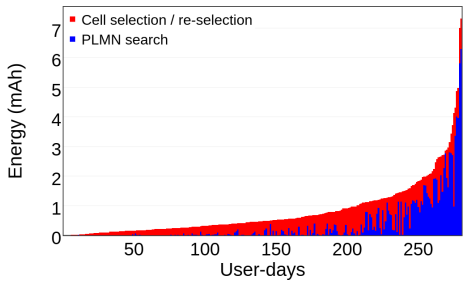


(d) Energy of cell selection/re-selection and PLMN search per user-day

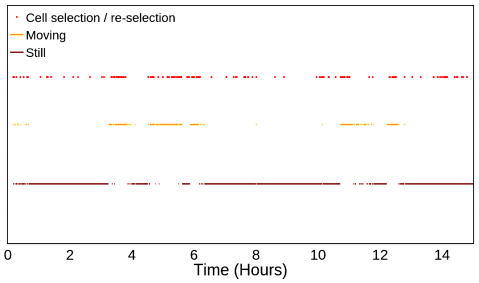


(e) Correlation between cell selection/re-selection frequency and user mobility

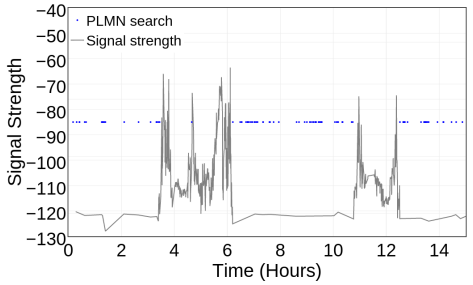


(f) Correlation between PLMN search frequency and signal strength

Fig. 18. Analysis of EMM specific procedures, cell selection/re-selection and PLMN search.

## 7 CONCLUSIONS

In this paper, we proposed the methodology of modeling modem power draw at the event-granularity, and developed to our knowledge the first fine-grained modem power model that captures the power draw of all LTE modem radio-on events in different RRC states. We showed that the fine-grained power model achieves far better modeling accuracy than the state-based model via controlled experiments using a Nexus 6 phone under a diverse set of UE usage and signal strength conditions. We further presented three case studies from a modem energy drain in-the-wild study involving 12 Nexus 6 phones under normal usage by 12 volunteers spanning a total of 348 days, to showcase the new types of analysis enabled by the event-based modem power model, including

Table 2. State-based power model of Nexus 6.

| Continuous reception parameters | | | | |
|---|---|---|---|---|
| RSRP(dBm) | Data rate (Mbps) | | Power (mA) | |
| | Tx | Rx | Tx | Rx |
| > −95 | 0.082 | 0.113 | 83 | 101 |
| −95 to −110 | 0.083 | 0.127 | 102 | 110 |
| < −110 | 0.083 | 0.116 | 114 | 104 |

| C-DRX parameters | | | | |
|---|---|---|---|---|
| | Power (mA) | On Duration (ms) | Total Duration (ms) | Periodicity (ms) |
| Short C-DRX | 100 | 20 | 20 | 20 |
| Long C-DRX | 320 | 20 | 11000 | 320 |
| DRX in RRC idle | 20 | 65 | N/A | 1280 |

analyzing and diagnosing tail energy, energy drain of NAS procedures, and energy drain of PLMN search and cell selection/re-selection.

## APPENDIX

### A1. Construction of event-based power model

In the first step of constructing the event-based power model, we map each state in the RRC State Machine with a distinct in-bound transition and a distinct out-bound transition as shown in Figure 2 into a sub-FSM in Figure 19, where a node in the sub-FSM is either a radio-on period or a radio-off period, and a transition is either a modem procedure or an LTE timer. In total, there are 7 sub-FSMs corresponding to 7 distinct modem events. in Figure 19. The states shown in double line circles are the states involved in the mapping in the original RRC State Machine, in solid circles with gray background are the states involved in the mapping in the expanded sub-FSM with radio-on periods, in solid circles with white background are the states involved in the mapping in the expanded sub-FSM with radio-off periods, and in dashed circles are the states not involved, respectively.

### A2. State-based power model for Nexus 6

Table 2 lists the parameters for the state-based power model for Nexus 6 used in the experiments.

### A3. Modem energy breakdown

The event-based modem power model allows us to break down the total modem activities according to the purposes of modem activites. Table 3 shows the detailed purposes and the mapping between them and the 7 modem events. We already discussed the first four and the last purposes in the three case studies. Each of the remaining two purposes, DRX paging during RRC Idle and RRC connection failure, directly map to the events E5 and E6.

Figure 20(a) shows the average daily modem energy breakdown across the purposes for each user in the 12-user in-the-wild study. The bars on the left and right in each bar pair draw the energy breakdown estimated using the event-based modem power model and the state-based modem power model, respectively. Figure 20(b) shows the average daily percentage energy breakdown for each purpose across the 12 users under the new event-based model. We make several observations. (1) The event-based model captures 22.2% more data reception energy and 26.0% tail energy on average for user data. (2) The new model captures the EMM procedure energy (1.55%) and cell selection/re-selection energy (1.12%) which were not captured by the state-based power model. (3) The majority of daily modem energy is split 3-ways: Continuous reception for user data (37.7%), tail for user data (34.2%), and DRX paging (25.4%).

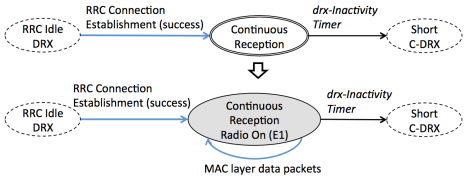


(a) Continuous reception to Event "Continuous reception after connection establishment" (E1)

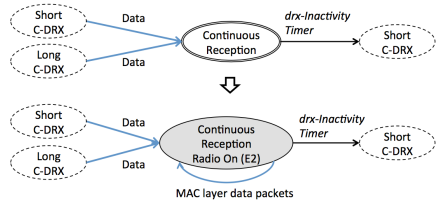


(b) Continuous reception to Event "Continuous reception after C-DRX wakeup" (E2)

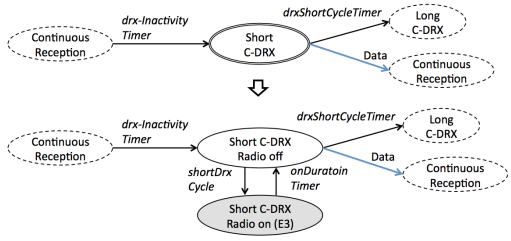


(c) Short C-DRX to Event "DRX wakeup without continuous reception" (E3)

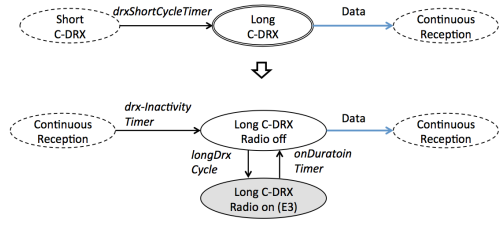


(d) Long C-DRX to Event "DRX wakeup without continuous reception" (E3)

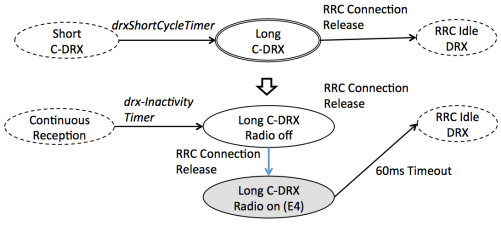


(e) Long C-DRX to Event "RRC connection release" (E4)

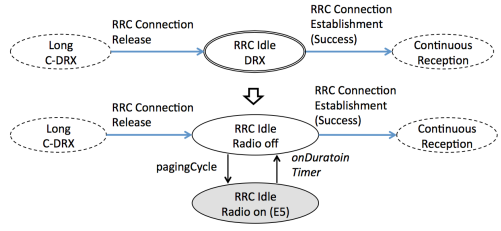


(f) RRC Idle DRX to Event "DRX paging during RRC Idle mode" (E5)

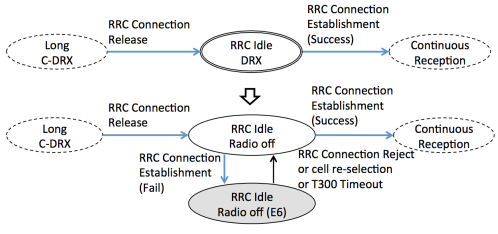


(g) RRC Idle DRX to Event "RRC connection establishment failure" (E6)

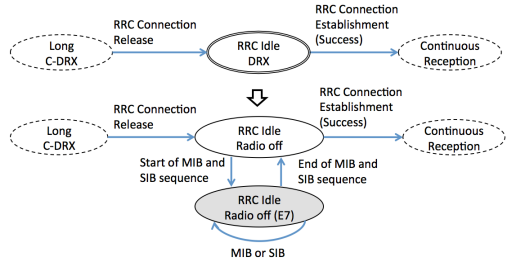


(h) RRC Idle DRX to Event "MIB and SIB reading during RRC Idle mode" (E7)

Fig. 19. Mappings from RRC states with a distinct in-bound transition and a distinct out-bound transition to sub-FSMs.

Table 3. Mapping between modem events and purposes.

| No. | Purposes | Modem events |
|---|---|---|
| 1 | Continuous reception for user data | E1, E2 |
| 2 | Tail for user data | E2, E3, E4 |
| 3 | Continuous reception for EMM procedures | E1, E2 |
| 4 | Tail for EMM procedures | E2, E3, E4 |
| 5 | DRX paging during RRC Idle | E5 |
| 6 | RRC connection failure | E6 |
| 7 | PLMN search and cell selection/re-selection | E7 |

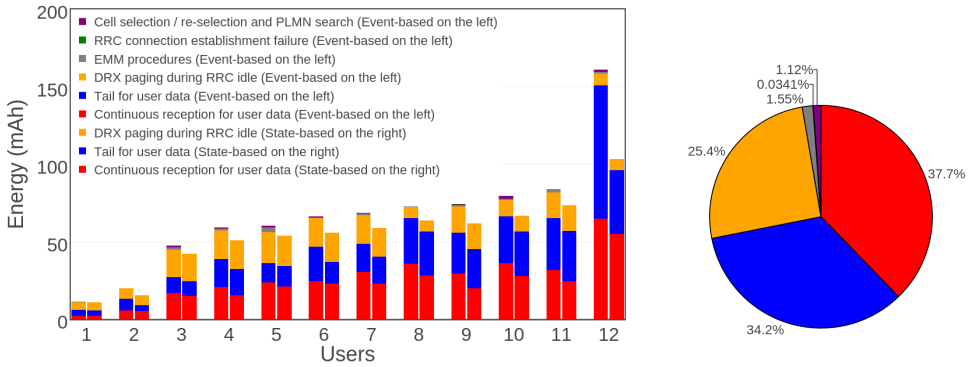


(a) Absolute energy. The two bars in each pair are for event-based and state-based models, respectively.

(b) Relative energy under the new model.

Fig. 20. Average daily modem energy breakdown across the purposes for each user.

## REFERENCES


[1] [n. d.]. Monsoon Power Monitor. ([n. d.]). www.msoon.com/LabEquipment/PowerMonitor/.

[2] 3GPP. 2011. *Medium Access Control (MAC) protocol specification.* 3GPP.

[3] 3GPP. 2011. *User Equipment (UE) procedures in idle mode.* 3GPP.

[4] 3GPP. 2012. *Radio Resource Control (RRC);.* 3GPP.

[5] Pilar Andres-Maldonado, Pablo Ameigeiras, Jonathan Prados-Garzon, Juan J. Ramos-Munoz, and Juan M. Lopez-Soler. 2017. Optimized LTE Data Transmission Procedures for IoT: Device Side Energy Consumption Analysis. In *Proc. of International workshop on application of green techniques to emerging communication and computing paradigms (GCC).*

[6] Niranjan Balasubramanian, Aruna Balasubramanian, and Arun Venkataramani. 2009. Energy consumption in mobile phones: a measurement study and implications for network applications. In *Proc of IMC.*

[7] D. Brooks, V. Tiwari, and M. Martonosi. 2000. Wattch: A framework for architectural-level power analysis and energy estimation. In *Proc. of ISCA.*

[8] A. Catovic, M. Narang, and A. Taha. 2007. Impact of SIB Scheduling on the Standby Battery Life of Mobile Devices in UMTS. In *2007 16th IST Mobile and Wireless Communications Summit.* 1–5. https://doi.org/10.1109/ISTMWC.2007.4299075

[9] Xiaomeng Chen, Ning Ding, Abiliash Jindal, Y. Charlie Hu, Maruti Gupta, and Rath Vannithamby. 2015. Smartphone Energy Drain in the Wild: Analysis and Implications. In *SIGMETRICS.*

[10] Ning Ding, Daniel Wagner, Xiaomeng Chen, Abhinav Pathak, Y. Charlie Hu, and Andrew Rice. 2013. Characterizing and Modeling the Impact of Wireless Signal Strength on Smartphone Battery Drain. In *SIGMETRICS.*

[11] Junxian Huang, Feng Qian, Alexandre Gerber, Z. Morley Mao, Subhabrata Sen, and Oliver Spatscheck. 2012. A Close Examination of Performance and Power Characteristics of 4G LTE Networks. In *Proc. of Mobisys.*

[12] Ali T Koc, Satish C Jha, Rath Vannithamby, and Murat Torlak. 2014. Device power saving and latency optimization in LTE-A networks through DRX configuration. *IEEE Transactions on wireless communications* 13, 5 (2014), 2614–2625.

[13] Mads Lauridsen, Laurent Noël, Troels Bundgaard Sørensen, and Preben Mogensen. 2014. An empirical LTE smartphone power model with a view to energy efficiency evolution. *Intel Technology Journal* 18, 1 (2014), 172–193.

[14] Yuanjie Li, Chunyi Peng, Zengwen Yuan, Jiayao Li, Haotian Deng, and Tao Wang. 2016. Mobileinsight: extracting and analyzing cellular network information on smartphones.. In *MobiCom*. 202–215.
[15] Abhinav Pathak, Y. Charlie Hu, and Ming Zhang. 2012. Where is the energy spent inside my app? Fine Grained Energy Accounting on Smartphones with Eprof. In *Proc. of EuroSys*.
[16] Abhinav Pathak, Y. Charlie Hu, Ming Zhang, Paramvir Bahl, and Yi-Min Wang. 2011. Fine-grained Power Modeling for Smartphones Using System-Call Tracing. In *Proc. of EuroSys*.
[17] Feng Qian, Zhaoguang Wang, Alex Gerber, Z. Morley Mao, Subhabrata Sen, and Oliver Spatscheck. 2010. Characterizing Radio Resource Allocation for 3G Networks. In *Proc. of IMC*.
[18] Sanae Rosen, Ashkan Nikravesh, Yihua Guo, Z. Morley Mao, Feng Qian, and Shubho Sen. 2015. Revisiting Network Energy Efficiency of Mobile Apps: Performance in the Wild. In *IMC*.
[19] Joel Sommers and Paul Barford. 2012. Cell vs. WiFi: On the Performance of Metro Area Mobile Connections. In *IMC*.
[20] Dario Vinella and Michele Polignano. 2009. Discontinuous reception and transmission (DRX/DTX) strategies in long term evolution (LTE) for Voice-Over-IP (VOIP) traffic under both full-dynamic and semi-persistent packet scheduling policies. *Project Group* 996 (2009).
[21] Qiang Xu, Jeffrey Erman, Alexandre Gerber, Zhuoqing Mao, Jeffrey Pang, and Shobha Venkataraman. 2011. Identifying diverse usage behaviors of smartphone apps. In *Proc. of IMC*.
[22] W. Ye, N. Vijaykrishnan, M. Kandemir, and M. J. Irwin. 2000. The Design and Use of SimplePower: A Cycle-Accrate Energy Estimation. In *Proc. of DAC*.

**Received August 2017; revised October 2017; accepted December 2017.**