

Characterizing and Modeling the Impact of Wireless Signal Strength on Smartphone Battery Drain

Ning Ding
Purdue University

Abhinav Pathak
Purdue University

Daniel Wagner
University of Cambridge

Y. Charlie Hu
Purdue University

Xiaomeng Chen
Purdue University

Andrew Rice
University of Cambridge

ABSTRACT

Despite the tremendous market penetration of smartphones, their utility has been and will remain severely limited by their battery life. A major source of smartphone battery drain is accessing the Internet over cellular or WiFi connection when running various apps and services. Despite much anecdotal evidence of smartphone users experiencing quicker battery drain in poor signal strength, there has been limited understanding of how often smartphone users experience poor signal strength and the quantitative impact of poor signal strength on the phone battery drain. The answers to such questions are essential for diagnosing and improving cellular network services and smartphone battery life and help to build more accurate online power models for smartphones, which are building blocks for energy profiling and optimization of smartphone apps.

In this paper, we conduct the first measurement and modeling study of the impact of wireless signal strength on smartphone energy consumption. Our study makes four contributions. First, through analyzing traces collected on 3785 smartphones for at least one month, we show that poor signal strength of both 3G and WiFi is routinely experienced by smartphone users, both spatially and temporally. Second, we quantify the extra energy consumption on data transfer induced by poor wireless signal strength. Third, we develop a new power model for WiFi and 3G that incorporates the signal strength factor and significantly improves the modeling accuracy over the previous state of the art. Finally, we perform what-if analysis to quantify the potential energy savings from opportunistically delaying network traffic by exploring the dynamics of signal strength experienced by users.

Categories and Subject Descriptors

C.4 [Computer System Organization]: Performance of Systems—*Modeling techniques*; C.2.3 [Computer System Organization]: Computer Communication Networks—*Network Operations*

General Terms

Experimentation, Measurement

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGMETRICS'13, June 17–21, 2013, Pittsburgh, PA, USA.
Copyright 2013 ACM 978-1-4503-1900-3/13/06 ...\$15.00.

Keywords

Signal strength, smartphone, energy, battery drain, power model

1. INTRODUCTION

The smartphone market has been growing at a phenomenal rate. A recent study [3] finds that out of the world's 4 billion mobile phones, over one billion are smartphones, and projects that by the year 2014, there will be more smartphone users surfing the Web than desktop users. Despite such an incredible adoption rate of smartphones, the user experience has been, and will remain, severely limited by the phone battery life. As such, understanding and optimizing the energy consumption of apps running on mobile devices is of significant importance.

A major source of smartphone energy consumption is accessing the Internet via 3G or WiFi [4, 20, 18] when running various interactive apps and background services. The energy consumed by the wireless interfaces is perceived as a necessity – after all, a smartphone provides the user with a ubiquitous portal to the Internet. Ideally, accessing the Internet should consume an amount of energy commensurate with the amount of traffic being transported and the (peak) throughput supported by the wireless technology used.

In practice, the wireless channel can be noisy and a fundamental law governs wireless networking performance: wireless channel capacity is upper-bounded by the signal-to-noise ratio (SNR), which measures the ratio of the level of a desired signal to the level of background noise, as dictated by the Shannon-Hartley theorem. In other words, poor signal strength can significantly affect the achievable network performance. How to achieve the upper-bound channel capacity under a given SNR is at the very center of wireless communication research.

In this paper, we contend that poor wireless signal strength not only affects network performance, but also – in the context of energy-constrained mobile devices perhaps more importantly – can significantly inflate the actual energy consumption by the wireless interface to be much higher than under good signal strength, while transferring the same amount of network traffic. The more obvious impact on energy drain is that reduced signal strength triggers rate adaptation at the PHY layer to lower the data rate which elongates packet transmission and hence increases power consumption by the radio. There are also several additional, less obvious, impact factors. Weak signal strength can result in retransmissions, *e.g.*, in the link layer in 3G and at the MAC layer in WiFi, and even at the transport layer, all of which lead to extra radio power consumption. Retransmissions at the transport layer (*e.g.*, TCP) can further increase the number of times and hence the total duration the wireless interface stays in the tail power state [4, 19, 20], wasting “tail” energy. Further, poor signal strength can cause disassociation and reassociation with the access point which incur extra energy drain.

Despite the above intuitive understanding of the energy impact of poor signal strength and much anecdotal evidence that smartphone users frequently have experienced quicker battery drain when the wireless signal strength is poor, we have a rather limited understanding of (1) how often smartphone users experience poor signal strength, and (2) the quantitative impact of poor signal strength on the phone battery drain. The answers to such questions are essential to assessing cellular network service qualities, and diagnosing and reducing their negative impact on smartphone battery life. Further, answering question (2) above enables us to develop an online power model for smartphones that is more accurate than previous work [19], by explicitly capturing the potentially significant impact of signal strength. Accurate online power models for smartphones form building blocks for energy profiling and optimization of smartphone apps [18], and the new signal-strength-aware power model enables what-if analysis of techniques that explore signal strength dynamics to reduce the energy drain of wireless interfaces, *e.g.*, by delaying and aggregating network traffic for latency insensitive apps and services.

In this paper, we conduct to our knowledge the first measurement and modeling study of the impact of wireless signal strength on smartphone energy consumption. Our study makes four specific contributions and findings.

- We collected a cellular and WiFi signal strength and traffic volume trace from 3785 smartphones, geographically distributed over 145 countries, each trace covering at least a one-month period with an average of 4.2 months. Our trace analysis shows that (1) WiFi and 3G are dominant wireless technology choices among users, (2) individual users are experiencing significant signal variations in daily life and during active phone usage, (3) on average 43% and 21% of their foreground data are transferred during poor 3G and WiFi signal strength, respectively, (4) 19% of 3G transfer and 4% of WiFi transfer can be classified as background data and thus can be potentially deferred to times with better signal strength, and (5) the signal variations are correlated with popular user locations to some extent which also manifests itself in a correlation with time of day.
- We then conducted controlled experiments to quantify the extra energy consumption for data transfers that is induced by poor wireless signal strength and its breakdown to different factors including rate adaptation, power control, link layer retransmissions and TCP retransmissions. Our measurements refined packet-driven power models for WiFi and 3G to be signal-strength-aware and hence much more accurate than before. Our energy impact analysis shows that for WiFi the lower bit rate and MAC layer retransmission when the Received Signal Strength Indicator (RSSI) drops from -50dBm to -90dBm cost 810.5% more energy for a typical mobile download of 100KB with 30ms server RTT; for 3G the increased energy on data transfer and RRC state demotion when the RSSI drops from -85dBm to -105dBm dominates the extra energy consumption, resulting in 52.0% more energy for the same 100KB download with 30ms server RTT.
- We further derived a new system-call-driven power model for WiFi and 3G that improves the state of the art [19] by incorporating the impact of signal strength and RTT. Our evaluation shows that the new model drastically improves the modeling accuracy over the prior art. Specifically, under poor signal strength, the new model reduces the energy estimation error from 61.0% to 5.4% for WiFi, and from 52.1% to 7.2% for 3G.
- We present a case study showing how our new system-call-

driven power model can be used to perform what-if analysis of the effectiveness of energy optimization techniques that explore signal strength dynamics. Specifically, our analysis shows that for three selected subsets of users who experienced predominantly good, fair, and poor signal strength, opportunistically delaying background network traffic till the signal strength becomes better, can reduce the total energy consumption of data communication by up to 23.7% and 21.5% under WiFi and 3G, respectively.

The rest of the paper is organized as follows. §2 briefly reviews the impact of signal strength on the network layers in 3G and WiFi. §3 presents the trace analysis. §4 gives the context for the remaining impact analysis, power modeling, and what-if analysis. §5 presents our controlled measurement study of the energy impact of signal strength on WiFi and 3G. §6 presents our new power model which incorporates the signal strength and RTT factors, and §7 presents a case study of using the power model to perform what-if analysis. We discuss related work in §8 and conclude in §9.

2. BACKGROUND

In this section, we first review the basics about the multiple power states of 3G and WiFi. We then examine the impact of signal strength on the extra energy consumption at various network layers.

2.1 Power States and Transitions

A wireless device like a 3G or WiFi radio can be in several operating modes, known as *power states* for that device, each draining a different amount of power. Each device has an idle base state which is the power state where that particular device consumes least power. A device can further have one or more levels of productive power states, depending on the workload. Finally, both 3G and WiFi radios exhibit a tail power phenomenon where the radio stays in a high power state after active usage and continues to consume energy in anticipation of more communication [4, 20, 19], before eventually returning to the base power state.

In a 3G network, the power state of a user equipment (UE) is determined by the Radio Network Controller (RNC) via the Radio Resource Control (RRC) protocol. The power states of the 3G device we measured on an HTC Nexus One phone are shown in Figure 1(a): (1) *IDLE*: A UE is in the IDLE state when it does not send or receive any data; the 3G radio draws nearly zero power. (2) *FACH*: At low transfer rate, the UE establishes the connection and enters the FACH state which does not have a dedicated channel. The 3G radio consumes moderate power in the FACH state. (3) *DCH*: At high transfer rate, the UE enters the DCH state, which usually has its dedicated data channel but also consumes high power. (4) The transitions from IDLE to FACH and from FACH to DCH – known as promotions – take a certain amount of time and consume a certain amount of power. (5) Similarly, the transitions from DCH to FACH and from FACH to IDLE have specific timeouts and incur tail power costs. In particular, in T-Mobile network 3G stays in DCH and FACH after active usage for 3.6s and 3.3s, respectively, before switching to the next lower power state.

The power states and transitions for WiFi are simpler, as shown in Figure 1(b). In WiFi, the tail power state is a result of 802.11 Dynamic Power Saving Mode (PSM) [5]; after transmitting/receiving, the device will stay in the high power state for a pre-defined duration (called PSM timeout) before going into PSM.

The above power state behavior of 3G and WiFi suggests that the total energy drain of 3G/WiFi in carrying out a fixed amount of traffic equals the sum of the energy consumed while in productive power states and while in the tail state, which is determined by the




Figure 1: 3G and WiFi power state machine for HTC Nexus One phone. All parameters are for under good signal strength.

total duration they stay in these states, $T_{productive}$, T_{tail} :

$$E_{total} = P_{productive} * T_{productive} + P_{tail} * T_{tail}$$

The total time spent in a productive state is affected by the data rate and retransmissions. The total time spent in the tail state depends on the total number of times the device enters the tail state and the duration of each episode.

2.2 Energy Impact of Weak Signal Strength

Weak wireless signal strength can result in smartphone apps consuming significantly more energy than under good signal strength. In particular, weak signal strength can increase the total time spent in productive and tail power states, $T_{productive}$, T_{tail} , as well as transmit power, $P_{productive}$, when sending and receiving data by affecting the behavior of almost all layers of the network stack.

Rate adaption. Rate adaptation is a dynamic, continuous link adaptation process where the modulation, coding and other signal and protocol parameters of the wireless communication are adapted to match the dynamically changing channel conditions, in order to reduce bit error rates and improve the effective rate of transmission. In general, weak signal strength leads to reduced data rate, which effectively elongates per-packet transmission time and thereby the total time the device stays in the productive state, $T_{productive}$.

WiFi switches between multiple modulation types and data rates in the PHY layer. A large number of autarate schemes have been proposed, including link-layer approaches [13, 14, 25] and physical-layer approaches [10, 22]. 3G blurs the notion of PHY and MAC and rate adaptation forms part of the tightly integrated link adaptation. In the 3G link layer, poor signal strength therefore not only leads to reduced data rate, which directly elongates the transmission time of a packet, but also increases the retransmissions of transport block sets of information bits within a packet, which indirectly elongate the transmission time.

Power control. Another form of link adaptation that could be triggered by signal strength variation is transmission power control, which performs intelligent selection of transmitting power to achieve good link data rate, network capacity, and/or geographic coverage. Increased transmission power can increase the SNR and hence reduce the bit error rates at the receiver, or allow transmission at a higher data rate using rate adaptation. However, it can also increase the power consumption of the transmitting radio, *e.g.*, of a mobile device, and increase the interference to other users in the same frequency band.

WiFi MAC Retransmission. Poor signal strength can lead to dropped packets and retransmissions at the WiFi MAC layer. Packet loss is discovered when a sender does not receive the ACK frame within a fixed time interval and the MAC retransmits the same packet up

Table 1: Trace statistics.

Devices > 1 month trace (3/6/9/12 months)	3785 (1728/835/463/232)
Aggregate trace duration	1334 years
Median trace duration	82 days
Countries of origin	145
Mobile operators	896
Unique phone types	487
Rate of mobile RSSI reading	when signal changes, effective: 1/min
Rate of network usage reading	every 5 minutes

to a fixed number of retries. Clearly retransmitting the same packet multiple times increases the energy consumption per packet.

TCP Retransmission. When the signal strength is extremely weak, retransmissions at the link layer may not guarantee eventual packet delivery. When this happens, it triggers packet retransmission and potentially reduces the sending rate from congestion control at the TCP layer which again leads to extra energy consumption.

Re-Association and Handoffs. Weak wireless signal strength due to the mobility of smartphone users, *e.g.*, during driving while connected to 3G, or walking on a campus covered with WiFi, can result in disassociations and reassociation with the wireless network (3G tower or WiFi AP), as well as handoffs between 3G and 2G [17], which consume significant amount of extra energy. We note that this can happen even if the user is not actively interacting with the smartphone and no app requires network access.

Tail Energy. Conceptually, the tail power state phenomenon is independent of the signal strength. However, for TCP flows, since the radio is in the tail state in between consecutive TCP windows, retransmissions due to packet loss can increase the number of times the radio enters the tail state and hence the total tail state duration.

Wait Energy. All of the above categories of extra energy consumption due to weak signal strength incur on the wireless device itself. Weak signal strength can further have a secondary effect that wastes energy on other components of a smartphone. In particular, the delay induced by the slower transfers from lower data rate or retransmissions can cause other phone components such as CPU and screen to stay awake longer while waiting for the network transfer to complete. This secondary effect is app usage dependent and we leave quantifying it as future work.

3. TRACE ANALYSIS

In this section, we present a trace analysis of the signal strength experienced by 3785 smartphone users from 145 countries.

3.1 Trace Collection

We developed a free Android app that users can install on their smartphone to contribute data anonymously. We collected information from 3785 volunteers worldwide. Each user trace ranges from 1 month to 19 months in length, with an average of 4.2 months (median 84 days). The collected data include cellular and WiFi signal strength, operator name, bytes transferred, coarse (network-based) location, screen on/off state and battery level. The detailed characteristics of the trace are shown in Table 1.

The very first observation we make is that the vast majority of the users spend most of the time using 3G cellular networks and WiFi. In particular, we find that over 80% (60%) of the 3785 users spent at least 44% (63%) of the time using 3G (UMTS or HSPA) networks or WiFi. We therefore focus on these two wireless technologies in the rest of the paper. For clarity, we only plot the distribution of various statistics for a stratified sample of 100 users drawn from




Figure 2: Wireless technologies over time.




Figure 3: Overall 3G signal strength (left) and while screen was on (right) as observed by 100 sampled devices.

the whole dataset of 3785 users—incorporating the data from all users does not alter the visual characteristics. Figure 2 shows the fraction of time (*i.e.*, non-cumulative, as in the other figures of this type) spent in different wireless technologies for the stratified sample of 100 users. Each horizontal row (*y*-axis) corresponds to a single user, and each unique color segment shows the fraction of time (along the *x*-axis) for which the primary data communication mechanism was a particular technology. For example, the 20th row contains a dark green segment of width 0.1, indicating this particular user made use of some WiFi network 10% of the time.

3.2 3G Signal Strength

We make three key observations from analyzing the 3G signal strength experienced by 3785 users.

Observation 1: *3G signal strength over the entire trace varies considerably across users and over time.* Figure 3 (left) shows the fraction of the total trace collection time of different signal levels experienced by the 100 sampled users. The figure for all 3785 users has the same visual characteristics. Research by the UK communications regulator Ofcom considers a good signal to have a power of -91.7 dBm or higher (shown in colors from orange to green in Figure 3).¹ By this definition, the 3785 users in our dataset saw a poor signal on average 47% of the time, and over 80% (60%) of them experienced signal strength weaker than -91.7 dBm for over 13% (31%) of the time.

Observation 2: *Signal strength also varies considerably during active usage across users and over time.* Clearly, periods of poor signal strength are only relevant to users if they notice the effects of




Figure 4: Data transferred over 3G with a given signal strength. Devices without mobile data traffic were excluded.

bad signal strength. These effects will be most pronounced when the user is actively using the phone. We define periods of active usage as periods when the device’s screen is turned on. Figure 3 (right) shows the distribution of signal levels experienced during active usage remains almost identical to that for the entire tracing period in Figure 3 (left). In particular, across our dataset of 3785 users we observe poor signal, *i.e.*, below -91.7 dBm, during 47% of active usage time, and over 80% (60%) of users experienced poor signal strength for over 15% (32%) of their active usage time.

Observation 3: *Users transfer significant amounts of data during poor signal strength periods.* Periods of poor signal strength matter most to users if significant amounts of data are transferred during those periods. We therefore investigate the amount of data transferred over mobile networks by the individual devices and further distinguish between foreground and background network traffic. Since there is no explicit indication of traffic class from the OS, we approximately partition data into foreground and background classes as follows. We assume that data transmitted while the screen is off are solely non-interactive background data, whereas data transmitted while the screen is on are primarily foreground data. This is an underestimate of the amount of background data which could happen while the screen is on. For the 3785 users, this approach labels on average 19% of the total bytes transferred on 3G as background data.

Signal strength is collected upon change through notifications from the OS. As it may change in the middle of a 5-minute network usage polling interval, we assume that the network usage is uniform across each interval. This allows us to calculate the number of bytes transmitted by each wireless device at different signal levels. We find that by this definition, the 3785 users performed on average 43% of their foreground data transfers during poor signal strength of below -91.7 dBm, with over 80% (60%) of the 3785 users seeing over 11% (26%) of their foreground data transfer during such poor signal strength periods.

The figures for background data transfers are identical to the above numbers for foreground data transfers. Background data transfers during poor signal strength in the user’s daily cycle can potentially be deferred until a strong signal is available, while we assume foreground transfers to be time-critical. In §7, we show how to use our new WiFi and 3G power model to estimate the energy reduction from deferring background data.

Finally, Figure 4 shows that the fraction of data transferred for the 100 sampled devices during different signal conditions are largely similar during screen on and screen off, suggesting the network behavior of users is in general not affected by poor signal conditions.

3.3 WiFi Signal Strength

We next analyze the WiFi signal conditions experienced by our 3785 users. For WiFi, we consider -80 dBm and below *poor signal*

¹<http://stakeholders.ofcom.org.uk/market-data-research/other/telecoms-research/mobile-not-spots/Appendix1>

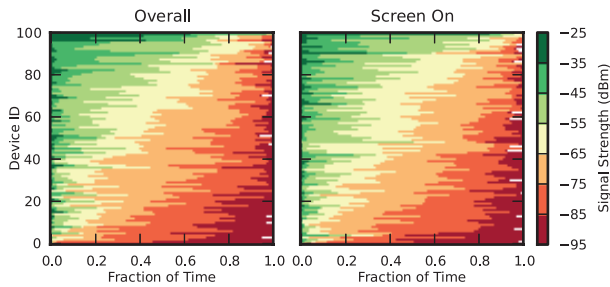


Figure 5: Overall WiFi signal strength (left) and while screen was on (right) as observed by 100 sampled devices.




Figure 6: Data transferred over WiFi with a given signal strength. Devices without WiFi data traffic were excluded.

strength, which can significantly affect data transfer time and energy drain, as shown in §5. We find the same three key observations made for 3G above hold true for WiFi as well. In particular, (1) the 3785 users on average experienced poor WiFi signal, *i.e.*, below -80 dBm, 25% of the time, and over 80% (60%) of them experienced poor WiFi signal strength over 5% (13%) of the time; (2) during active device usage on WiFi, the 3785 users saw poor signal on average 23% of the time, and over 80% (60%) of them experienced poor signal condition over 5% (13%) of the time; (3) foreground data traffic on WiFi is much more prevalent, with the 3785 users transferring on average 96% of their total WiFi data volume during active device usage. On average 21% of these data transfers occurred during poor WiFi signal strength, with over 80% (60%) of the 3785 users transferring over 2% (8%) of their foreground data during poor signal strength.

Figure 5 and Figure 6 plot the WiFi signal strength experienced and fraction of data transfers during different WiFi signal conditions, respectively, for the 100 sampled devices. We observe they look largely similar to the previous same figures for 3G.

We also find that significantly more data usage occurs over WiFi than over 3G. In fact, the ratio of data downloaded to uploaded changes from 6:1 when using cellular networks to around 20:1 on WiFi. One explanation for this could be that users consume different types of content on WiFi, such as streaming audio and video or downloading larger files. Indeed many market applications which perform background data transfer offer an option to limit this activity to WiFi networks only. However, this is probably motivated more by cost implications than energy usage.

3.4 Insights for Poor Signal Strength

The most likely explanation for the prevalence of poor signal strength experienced by a significant fraction of users is the compound effect of geographic variation in cellular network coverage and the fact that a user principally stays in a few locations throughout a day. Figure 7 gives supporting evidence by analyzing three selected users from the dataset in detail under 3G. The figure for

WiFi looks similar and is omitted due to page limit. We see that the distribution of signal strength at the top three most popular locations for each user has a considerably tighter distribution than their overall distribution of signal strength. As a consequence, we expect that a user’s daily routine will also give rise to a correlation between signal strength and time of day. This is confirmed by the circular plots in Figure 7. These daily cycles in signal strength suggest that all three users can benefit from a system that delays background data usage to exploit periods of improved signal strength.

4. OVERVIEW OF IMPACT ANALYSIS AND POWER MODELING

The motivation for quantifying the impact of wireless signal strength on device energy drain is to develop more accurate power models which will enable accurate energy profiling of mobile apps and ultimately help to optimize the energy efficiency of mobile apps.

Power models for mobile devices in general and wireless components such as WiFi, 3G and 4G radios have gone through three generations. The first generation of power models on smartphones (*e.g.*, [24, 26]) are based on the fundamental yet intuitive assumption that the (actual) utilization of a hardware component (*e.g.*, NIC) corresponds to a certain power state and the change of utilization is what triggers the power state change of that component. Consequently, these designs all use the notion of utilization of a hardware component as the “trigger” in modeling power states and state transitions. Such models do not take into account non-utilization-based power behavior of modern wireless components such as the promotion and tail power behavior of 3G and 4G, and thus can incur high modeling error.

The second generation of power models capture the non-utilization-based power behavior, in particular, non-utilization power states and transitions, using power state machines. In essence, such power state machines reverse-engineer the built-in state machine of the wireless radio, *e.g.*, the RRC state machine in 3G, and annotate each state or transition with power draw and duration. These include [4, 15, 20, 16] for WiFi and 3G and most recently [11] for a commercial LTE network. However, such models suffer two drawbacks. First, they use packet-level trace, *e.g.*, collected using TCP-Dump, as the triggers to drive the power state machine (*e.g.*, [20]), which are fairly heavy weight. Second, they cannot map the power activities of the wireless component back to the program entities in the app source such as subroutines and threads, which is essential for energy profiling.

In [19], the authors propose a third-generation power model which overcomes the above two drawbacks. The new power model uses system calls issued from apps as triggers in a system-call-driven finite state machine which captures both utilization and non-utilization based power behavior of smartphone components including wireless devices. The new model thus does not require packet-level traces, and allows direct mapping of power activities to program entities, making fine-grained energy profiling of apps a straightforward task [18].

However, none of above models takes into account the signal strength factor, and hence will incur high modeling error in the presence of poor signal strength which we have shown to be very common. Concurrent to our work, [16] extends the 3G RRC power state machine in [20] to incorporate the impact of signal strength on DCH and FACH and their tails, but not on promotion transitions.

Given the above context, we proceed with our measurement and modeling study in three steps. First, we perform controlled experiments to quantify the impact of poor WiFi and 3G signal strength on the device energy drain, and in doing so effectively refine the

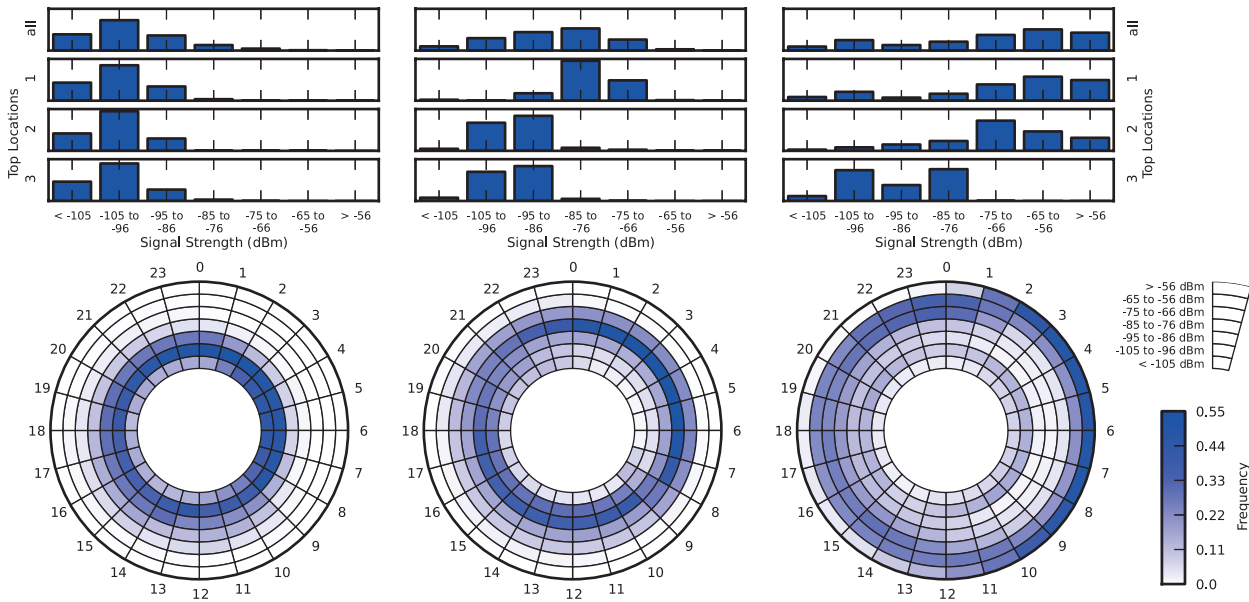


Figure 7: 3G RSSI per hour of the day for three devices experiencing predominantly bad, average and good signal (left to right).

Table 2: Mobile handsets used throughout the paper.

Handset	OS (kernel)	Cellular
HTC Nexus One	Android 2.3.4 (Linux 2.6.37)	T-Mobile 3G
Motorola Atrix 4G	Android 2.3.7 (Linux 2.6.32)	AT&T 3G
Sony Xperia S	Android 4.0.4 (Linux 3.0.8)	AT&T 3G

second-generation packet-driven power models to be signal-strength-aware and hence much more accurate (§5). For WiFi our analysis goes deeper than previous packet-level power modeling as we also quantify the impact on MAC retransmission and PHY rate adaptation. Second, we develop a refined third-generation, system-call-driven power model that takes into account signal strength (§6). Finally, we show how our new, signal-strength-aware system-call-driven power model enables what-if analysis of the effectiveness of energy optimization techniques that explore signal strength dynamics (§7).

5. CONTROLLED EXPERIMENTS

In this section, we present controlled experiments to quantify the impact of signal strength on the energy consumption of the WiFi and 3G devices and refine packet-driven power models.

5.1 Methodology

5.1.1 Devices and Tools

Table 2 lists the three smartphones used in our measurement study. The WiFi is provided by a Netgear WGR614v9 802.11 b/g wireless router. To measure the energy consumption, we use a Monsoon Power Monitor [2], which supplies power to the phone and samples the power draw every 0.2ms. During the measurement, we close all other apps and only run our benchmarks. Since our benchmark simply performs socket sending/receiving, the energy consumption of other phone components is close to 0 and the power draw measured by the powermeter is primarily the energy cost of WiFi or 3G NIC. Unless otherwise stated, we keep the phone’s screen on during the experiment, and subtract the display energy from the total energy consumption.

We focus on data download since our trace presented in §3.3 shows real life download volumes to far exceed upload volumes over both 3G and WiFi; results for upload were omitted due to page limit. The phone downloads data from a local server, which is connected to the wireless router via 100Mbps LAN. Since the RTT between the WiFi AP and the server (*server RTT* for short) is less than 1ms, we use the Linux traffic control API `tc` to emulate different RTTs when accessing the Internet. Similarly, for 3G, the server RTT is to emulate the Internet delay between GGSN and the server, *i.e.*, excluding the time for the packets to travel through the UMTS Terrestrial Radio Access Network (UTRAN), which is typically in the order of tens to hundreds of milliseconds [12].

The data download is implemented using a simple client/server program written in C and running on the phone and the server. The client opens a TCP socket, sends a request, receives data, and closes the socket. The data downloads are separated by 1s and 30s, for WiFi and 3G, respectively, to make sure that each transfer starts from the IDLE state of the interface. We conducted the measurement at late night, so that there is minimum interference in WiFi and the cellular network is lightly loaded.

5.1.2 Data Transfer

We set the base case data download size to 100KB, which is in the same order of typical transfer sizes in mobile web browsing. We set the default server RTT to 30ms, which is representative of typical Internet RTTs. We also vary the flow size and server RTT and study how they affect the impact of signal strength on energy consumption. For each configuration, we repeat the download 50 times; results presented below are averaged over the 50 trials.

5.1.3 WiFi Analysis

We control the WiFi signal strength by adjusting the distance between the phone and the AP, and record the signal strength by logging the RSSI in the `WifiStateTracker` and `WifiStateMachine` classes in the Android framework. To capture frame retransmissions, and to read the data rate of frames from the radiotap header, we set up two laptops right next to the phone and the AP in the monitor mode using the `Airmon-ng` tool in `Aircrack-ng` suite [1],

to eavesdrop WiFi link layer frames sent out by the phone and the AP, including DATA and ACK frames, control frames and beacons.

Energy breakdown. We break down the total energy of the WiFi NIC measured by the powermeter into six splits to quantify the impact on different network layers: (1) *Unique Tx*: Energy spent in the transmission of unique TCP ACK frames. (2) *Unique Rx*: Energy spent in the receiving of unique DATA frames. (3) *ReTx*: Energy spent in retransmitting TCP ACK frames. (4) *ReRx*: Energy spent in receiving retransmitted DATA frames. (5) *Idle*: Energy spent while waiting for frames from the AP, *i.e.*, in between frames or TCP windows. Its power equals the WiFi PSM tail power. (6) *PSM Tail*: Energy spent in the final WiFi tail power state, *i.e.*, after the flow completes while waiting for the PSM timeout to expire. The default PSM timeout on all three handsets is 210ms.

With the power profile from the powermeter and the eavesdropped link layer packet trace from the laptops, the above energy breakdown can be accomplished as follows.

First, we synchronize the time between the power profile and the monitor mode network traffic trace, by aligning the start of the first frame transmission with the start of the first WiFi power spike. Next, we categorize each frame sent or received by the phone WiFi NIC in the trace as either unique Tx/Rx or ReTx/ReRx, following their definitions above. Note ReTx/ReRx frames include both MAC and TCP retransmissions.

We next calculate the duration of each frame transmission by dividing the frame size by its transfer rate, which is indicated in the 802.11 radiotap header. For every sent ACK or received DATA frames eavesdropped by the laptop, since the timestamp indicates the time when the frame is finished transmitting, we backtrack the start of the transmission by subtracting the duration from the timestamp of the frame, and account the energy during the transmission interval between the start and finish time as the Tx/Rx energy. The energy drain between frame transmissions is counted as idle energy.

Finally, the energy drains belonging to each category are summed together and give the six-way split of the total energy drain.

Refining packet-driven power models. Since the packet-driven models use the packet time in the packet trace to determine the state transition time, inferring the refined packet-driven power models to incorporate signal strength boils down to inferring the power draw of each state under different signal conditions. In particular, for WiFi, with the synchronized power profile and link layer packet trace, we derive the Tx/Rx power under different signal strength by calculating the average power draw of WiFi NIC during a burst of packet sending/receiving.

5.1.4 3G Analysis

We adjust the 3G signal strength by changing the location of the phone, and record the signal strength by logging the RSSI in the `GsmServiceStateTracker` class in the Android framework. We observe that during experiment the signal strength variation is within 2dBm at good and fair signal strength locations, and within 3dBm at poor signal strength locations.

Energy breakdown. Unlike WiFi, there is no easy way to eavesdrop link layer packet transmissions in 3G. We instead resort to TCPDump to capture the packet trace. As such, we cannot observe the block set retransmissions or changing data rate in each packet transmission. The compound effect of the two factors makes it difficult to infer either one based on the observed transmission time of a packet, and hence we cannot separate the link layer retransmission from unique Tx/Rx as in the WiFi case.

To break down the energy, we first synchronize the time between the packet trace, which is timestamped by the phone, and the power profile, by inserting controlled power events (*e.g.*, a CPU spike) and

recording its time on the phone. We then play the packet trace to infer the 3G state transitions by leveraging the prior knowledge of the 3G state machines, following the methodology studied in-depth in [21, 20]. The energy drain during the duration of each 3G state is counted as the energy for the state. The energy drain of DCH during data transfer is further divided into data transmission energy and idle energy, based on whether the phone is sending/receiving packets or waiting for packets from the server.

Following this methodology, we break down the energy into following categories: (1) *Promotion1*: Energy spent in the first promotion phase, *i.e.*, IDLE to FACH in T-Mobile network and IDLE to DCH in AT&T network. (2) *Promotion2*: Energy spent in the second promotion phase, *i.e.*, FACH to DCH in both T-Mobile and AT&T network. (3) *Data transmission*: Energy spent in receiving DATA and transmitting ACK packets. This entails the impact of link layer rate adaptation and retransmissions. (4) *Tail1*: Energy spent in the first 3G tail state, *i.e.*, tail of DCH state. (5) *Tail2*: Energy spent in the second 3G tail state, *i.e.*, tail of FACH state. (6) *Idle*: Energy spent while waiting in DCH for packets from the server, *i.e.*, in between TCP windows. The duration is primarily determined by the end-to-end RTT.

Refining packet-driven power models. With the above energy breakdown, we then calculate the average power draw for each 3G state under each signal strength to derive the signal-strength-aware power model for 3G. Note for both new WiFi and 3G models, we only present parameters at discrete RSSI values, and hence in using the models, we use linear interpolation to estimate the parameters for RSSIs that are between the discrete points.

5.2 WiFi Results

5.2.1 Data Transfer Energy Breakdown: Base Case

Due to page limit, we only show the measurement result for the Nexus One phone. Figure 8(a) plots the time and energy consumption to download 100KB data with 30ms server RTT.² While the flow time only changes mildly from -50dBm to -80dBm, we see a 142.2% increase at -85dBm, followed by a dramatic increase of 1345.5% at -90dBm, compared to -50dBm. Similarly, the energy cost of downloading increases by 113.3%, 810.5% at -85dBm and -90dBm, respectively.

To understand the reason for the flow time and energy inflation, we break down the energy consumption in Figure 8(b). We make the following observations.

First, from -50dBm to -80dBm, most energy drain is from idle energy and PSM tail energy, while at -85dBm and -90dBm the idle energy dominates the total energy consumption. The reason is that for transfers of 100KB at good or medium signal strength, *i.e.*, above -80dBm, the active NIC transmitting/receiving time is generally in the order of tens of milliseconds, much shorter than the aggregate idle time and PSM tail time, which are in the order of hundreds of milliseconds. When the signal strength further drops, the average idle interval between frames increases significantly, from 0.28ms at -50dBm to 0.84ms at -85dBm and 5.47ms at -90dBm, leading to drastically increased idle energy.

Second, both unique Tx/Rx energy and retransmission energy increase significantly under poor signal strength. The total energy

²For simplicity, we represent power draw using the current drawn by the phone in milliAmperes. The actual power consumed would be the current drawn multiplied by 3.7V, the standard battery voltage supply. Similarly, energy is reported in μ Ah (micro Ampere Hours), and the actual energy would be the μ Ah value multiplied by 3.7V. These metrics are used since smartphone batteries are rated using these metrics and hence is easy to correlate.

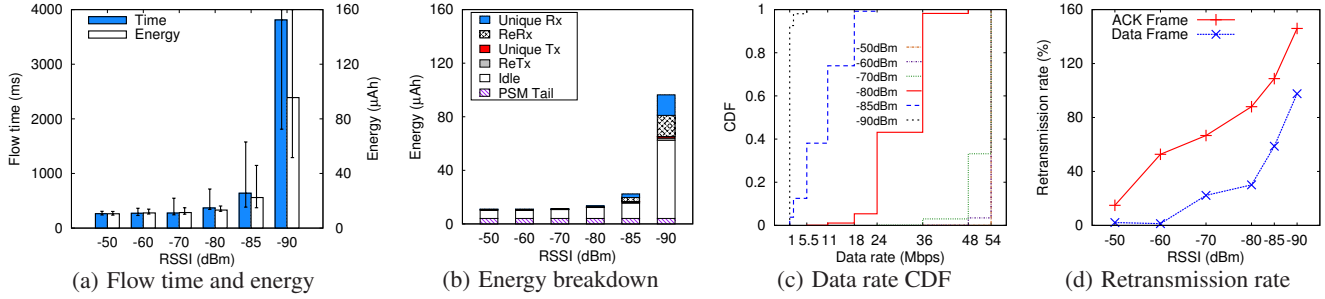


Figure 8: WiFi experiment: 100KB download with 30ms server RTT.




Figure 9: Time, energy and energy breakdown for downloading 10KB with 30ms server RTT under WiFi.

spent on data transfer, including unique Tx/Rx and ReTx/ReRx, increases from $0.60\mu\text{Ah}$ at -50dBm to $6.79\mu\text{Ah}$ at -85dBm , and further to $33.95\mu\text{Ah}$ at -90dBm . To understand the reason, we plot the frame data rate CDF in Figure 8(c). We observe that from -50dBm to -70dBm , most frames are transmitted at 54.0Mbps , which is the highest rate supported by 802.11g. When signal strength drops below -70dBm , frames are transmitted at much lower rates due to rate adaptation. For example, at -85dBm 74.0% of the frames are transmitted at rates equal to or lower than 11.0Mbps , while at -90dBm a majority of frames (92.6%) are transmitted at the lowest rate 1.0Mbps . Figure 8(d) plots the retransmission rate of DATA frames from the AP and of TCP ACK frames from the phone. The DATA frame retransmission rate has a sharp increase from 1.3% at -60dBm to 58.7% at -85dBm , and 97.7% at -90dBm , while the retransmission rate of ACK frames increases almost linearly from 14.9% at -50dBm to 146.0% at -90dBm .

Scanning. When the signal strength is below -80dBm , we observe the phone occasionally performs active scanning by exchanging Probe Request and Probe Response frames with the AP. Normally such a process takes from 40ms to 100ms. However, at poor signal conditions, both the Probe frames and their ACKs may get lost and it takes up to 1s to complete the scanning, with tens of retransmissions, consuming significant amount of energy on the phone.

5.2.2 Impact of Size

We next study the relative impact of signal strength on different transfer sizes. Figure 9 shows the flow time, energy consumption and energy breakdown for downloading 10KB using WiFi with 30ms server RTT. Similarly as in the 100KB download, the transfer time and energy consumption remain roughly the same as signal strength varies from -50dBm to -80dBm , but increase significantly by 70.7% and 40.8% at -85dBm , and by 586.1% and 263.4% at -90dBm . Compared to 100KB transfer, the transfer time and energy inflation due to poor signal strength is much less since the effect of lower data rate and higher retransmission rate are more pronounced in larger transfers.




Figure 10: Time and energy for downloading 100KB with different server RTT under WiFi.

Table 3: Signal-strength-aware WiFi power model. The unit of RSSI is dBm and the unit of Tx/Rx/Tail current is mA.

RSSI	Nexus		Atrix		Xperia	
	Rx	Tx	Rx	Tx	Rx	Tx
-50	117	165	115	175	120	251
-60	120	191	115	178	133	270
-70	125	251	139	197	133	308
-80	116	203	120	245	120	267
-85	104	207	118	242	118	256
-90	89	207	83	237	115	253
PSM Tail	Nexus		Atrix		Xperia	
	70		68		50	

5.2.3 Impact of RTT

We next study the relative impact of signal strength under different server RTTs. Figure 10 shows the time and energy consumption to download 100KB with different server RTTs. We observe that as the RTT increases from 30ms to 100ms, the flow time increases by 287.1% and 48.1% under -50dBm and -90dBm , respectively, and the energy drain increases by 144.8% and 26.6% under -50dBm and -90dBm , respectively. This suggests that RTT has a significant impact on the flow time and energy in WiFi downloading, at both good and bad signal strength. A large RTT leads to longer idle time between TCP windows, and hence more idle energy. The smaller increase ratio at -90dBm compared to -50dBm is because under -90dBm the idle interval between frames, due to heavy packet loss, dominates the flow time and energy consumption.

5.2.4 New Power Model

Table 3 shows the new signal-strength-aware packet-driven WiFi power model for the three handsets. The PSM tail power remains constant under different signal strength. We observe that for all three handsets, the average Tx and Rx power reach their peak at -70dBm (except the Tx power for Atrix). This is due to the compound effect of power control, rate adaptation and frame loss: as the RSSI drops, while the radio transmitting power tends to increase, the data rate decreases and more frame loss occurs, which leads to longer idle interval between frames, hence lower average power.

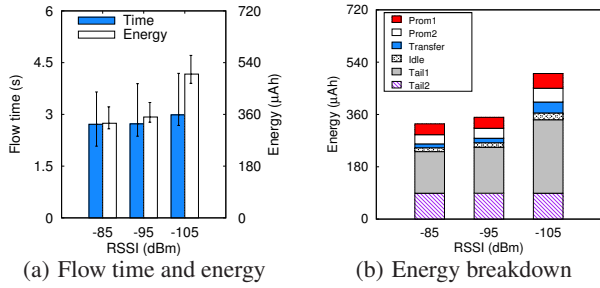


Figure 11: Time, energy and energy breakdown for downloading 100KB with 30ms server RTT under 3G.




Figure 12: 3G power snapshots for downloading 100KB with 30ms server RTT under different signal strength.

These three factors interplay and lead to the non-monotonic correlation between the signal strength and the average Tx/Rx power.

Another interesting observation is that, while in general Nexus One is the most power efficient in Tx/Rx among the three handsets, it has the highest PSM tail power; and Xperia S is the opposite.

5.3 3G Results

5.3.1 Data Transfer Energy Breakdown: Base Case

As with the WiFi experiment, we use 100KB download with 30ms server RTT as the base case and only show the result for the Nexus One phone due to page limit. Figure 11(a) plots the flow time and energy consumption for the base case. We see that from -85dBm to -95dBm, the flow time remains similar and the energy consumption increases mildly by 6.6%; while at -105dBm, the flow time increases by 7.3% and the energy consumption dramatically increases by 52.0%, compared to -85dBm.

To gain insight into the increase, we again break down the total energy, as shown in Figure 11(b). An immediate observation is that, most of the energy increase comes from increased data transfer and tail1 (DCH tail) energy. The powermeter output snapshots in Figure 12 confirm the significant power draw increase of the above two states.

5.3.2 Impact of Size

Figure 13 plots the flow time, energy consumption and energy breakdown for 10KB downloads under different signal strength. Compared to 100KB downloads, 10KB downloads on average only consume 6.3%, 5.8%, 8.0% less energy under good, medium and bad signal strength, respectively, as the promotion and tail energy dominate the total energy consumption under all signal strength.

5.3.3 Impact of RTT

Figure 14 plots the time and energy consumption to download 100KB with different server RTTs, under different signal conditions. We observe that the impact of server RTT on flow time and energy consumption is less significant in 3G compared to that in




Figure 13: Time, energy and energy breakdown for downloading 10KB with 30ms server RTT under 3G.




Figure 14: Time and energy for downloading 100KB with different server RTTs under 3G.

WiFi. Under -105dBm, the flow time and energy only increase by 25.9% and 9.4%, respectively, when the RTT varies from 10ms to 100ms, in contrast to 56.0% and 44.2% under WiFi at -90dBm. There are several reasons that contribute to the difference: (1) the 3G network used has a relatively large internal delay (about 100ms) which diminishes the server delay difference; (2) the promotion time in 3G, which is in the order of seconds and hence much larger than the server RTT, dominates the flow time for small and medium size downloads; and (3) the tail energy is also much higher in 3G than in WiFi, dominating the total flow energy consumption.

5.3.4 New Power Model

Table 4 shows the inferred new signal-strength-aware 3G power model for the three handsets. We see that while tail2 power stays constant across different signal strength, the power draw of all other states increases as the RSSI decreases, especially when the signal strength drops below -95dBm. In particular, compared to -85dBm, the tail1 power at -105dBm increases by 73.3%, 172.2% and 253.3%, for Nexus One, Atrix 4G and Xperia S, respectively.

5.4 Implications on App Design

We further draw implications from our new understanding of the energy impact of signal strength on how developers can optimize the energy drain of their apps. (1) App developers need to be conscious about the behavior of their apps in case of weak signal strength during the app development. For example, delaying elastic data transfers till good signal strength can potentially reduce the apps' energy consumption. A notification scheme supported in the OS or the framework that informs apps of when signal strength is good can help delay-tolerant apps (*e.g.*, peer-to-peer file sharing) to save energy by avoiding data transfers during poor signal strength moments. (2) Tail energy must be curtailed in case of small transfers since they consume a significant fraction of the total energy, especially under poor signal strength. (3) Aggregating data transfers in an app whenever possible can significantly reduce the total energy drain of the app, and the energy savings from doing so are even more pronounced when taking signal strength into account. For example, aggregating 10 separate 10KB transfers under -90dBm into

Table 4: Signal-strength-aware 3G power model. The unit of RSSI is dBm and unit of current of states is mA.

Nexus						
RSSI	Prom1	Prom2	Rx	Tx	Tail1	Tail2
-85	117, 0.6s	143, 0.9s	180	198	150, 3.6s	97, 3.3s
-95	123, 0.6s	150, 0.9s	192	281	160, 3.6s	97, 3.3s
-105	140, 0.9s	195, 1.0s	333	390	260, 3.6s	97, 3.3s
Atrix						
RSSI	Prom1	Prom2	Rx	Tx	Tail1	Tail2
-85	171, 1.8s	178, 0.9s	217	417	151, 4.0s	114, 10.0s
-95	197, 1.8s	203, 0.9s	287	497	188, 4.0s	114, 10.0s
-105	320, 2.0s	320, 1.0s	520	530	411, 4.0s	114, 10.0s
Xperia						
RSSI	Prom1	Prom2	Rx	Tx	Tail1	Tail2
-85	120, 1.8s	120, 0.9s	155	303	120, 4.0s	80, 10.0s
-95	127, 1.8s	127, 0.9s	215	425	190, 4.0s	80, 10.0s
-105	306, 2.0s	306, 1.0s	488	512	424, 4.0s	80, 10.0s

a single 100KB transfer under -70dBm can reduce the total WiFi energy consumed by a factor of 22.

6. NEW SYSTEM-CALL POWER MODEL

In this section, we develop a new system-call-driven power model for WiFi and 3G that improves the state of the art [19] by incorporating the impact of signal strength and RTT.

6.1 Current Power Model

The system-call-driven power model [19] for WiFi and 3G essentially looks similar to those in Figures 1(b) and 1(a), except the send/receive transitions are triggered by send/receive system calls (as opposed to packet transmissions).

The system-call-driven power model is derived in two steps, using a set of micro benchmarks and a powermeter [19]. First, the power state machines for individual system calls are derived, by capturing state transitions and power draw and duration at each productive state and tail state. The duration at an active power state is derived using linear-regression on the data transfer size. Second, the power state machines for individual system calls are integrated by observing the power profile for concurrent system calls.

As mentioned in §4, the current system-call-driven model for WiFi and 3G does not take into account the signal strength, which can affect various model parameters. It also ignores the effect of RTT, which can affect the power states visited during a system call.

6.2 WiFi Model

Observations. To develop the new model, we first make several observations about the effects of signal strength and RTT on the power profile of a data transfer system call. Figure 15 shows the power draw over time in a 100KB download with 30ms server RTT under different signal strength. A TCP download involves multiple TCP windows of packets, with packets in a window typically arriving in a cluster. We refer to the power spike corresponding to each TCP window as a window spike.

First, the current model ignores server RTT, and assumes the device stay in the productive state throughout data transfer due to the system call. This is a simplification. Figure 15(a) shows for a typical 30ms server RTT, the device effectively enters the tail power state in between TCP windows. The duration it stays in the idle tail state depends on the RTT value, and hence the power model needs to explicitly take RTT into account.

Second, when the signal strength drops, we observe several changes in the power profile as shown in Figure 15(b): (1) the frames are




Figure 15: WiFi power snapshots for downloading 100KB with 30ms server RTT under different signal strength.




Figure 16: Effective receiving rates for WiFi and 3G under different signal strength.

transmitted at lower bit rates from rate adaptation, and the window spikes become wider; (2) the retransmission rate increases, and more frames are transmitted or received within each window, making each window spike even wider; (3) the power draw of spikes changes, as summarized in Table 3. In summary, poor signal strength causes TCP window spikes to have longer duration and changed power draw.

Modeling a single system call. To extend the system-call-driven power model for WiFi and 3G to take into account wireless signal strength, we need to log the signal strength and end-to-end RTT, in addition to system calls. We record the WiFi and 3G signal strength by logging the RSSIs in the `WifiStateTracker`, `WifiStateMachine` and `GsmServiceStateTracker` classes in the Android framework. The RTT of TCP connections can be retrieved from `tcp_info` struct via the `Netlink` API in Linux.

We model the power behavior of a data transfer system call in 3 steps: (1) For a transfer of size s , we first estimate the number of TCP windows and the size of each window, following the AIMD behavior. This proves to be reasonably accurate for small transfers, *e.g.*, on the order of tens of KB, which account for the majority of the mobile traffic [8]. (2) For each TCP window w , we determine its duration t_w using linear regression on the window size, and power p_w according to Table 3, under each signal strength. (3) We then include the idle periods of staying in the PSM tail state in between the TCP windows, with the duration being the end-to-end RTT minus the window duration, as well as the final PSM tail after the last window.

The duration of a TCP window spike t_w is calculated by dividing the window size by the effective sending/receiving rate under the current signal strength. The effective rate is derived by calculating the average transfer rate of TCP windows during the training phase using micro benchmarks and the powermeter. Figure 16(a) plots the effective receiving rate for a window of 20KB measured under different signal strength. We see a sharp decrease of the effective rate from -70dBm, which captures the effect of lower bit rate and higher retransmission rate in the link layer, as shown in §5.2. The window spike power p_w is a function of signal strength and is directly measured during training phase. The values will be the same as in the packet-driven power model, *i.e.*, in Table 3.

Table 5: Website traffic used in the validation experiments.

Website	Flow #	Total size	RTT
Amazon	10	164.1KB	28ms
Gmail	6	852.2KB	37ms
Wikipedia	9	176.4KB	91ms
Youtube	13	341.2KB	37ms




Figure 17: Model validation result for Nexus One at poor signal strength locations.

Modeling multiple system calls. Once the power behavior for individual transfer is derived, we integrate them to model the power behavior for concurrent system calls. A second transfer may begin before the first transfer’s power trace finishes. There are two possible scenarios: (1) If the second transfer starts when the first transfer is in final PSM tail state, then the resulting power trace can be generated by taking the maximum power of the overlapped period and keeping other periods unchanged. (2) If the second transfer starts when the first transfer has not finished all of its windows, then the window spikes of the two transfers may overlap. Whenever two window spikes overlap, the spike that starts later will be push back, till after first spike finishes. All the subsequent window spikes of that transfer will also be pushed back by the same offset.

6.3 3G Model

Extending the system-call-driven model for 3G follows the same process as for WiFi, with the following slight complication. The 3G power state machine has two productive states, and accordingly two promotion and two tail transition states. Hence in modeling a data transfer system call, we need to keep track of the transitions among those states. We infer the conditions, *e.g.*, data transfer bytes, for different transitions following the methodology in [21]. We then apply these conditions in modeling the power behavior of the consecutive windows of a transfer. Figure 16(b) plots the effective receiving rate under different signal strength for T-Mobile and AT&T 3G network.

6.4 Model Validation

We have incorporated the new model for WiFi and 3G into eprof [18] with 1K lines of code.

To validate the model, we place the phone with the powermeter in locations with different signal strength, and run benchmarks to generate network traffic, while logging the system calls, signal strength and server RTT. We then feed the logs into eprof and compare the predicted energy consumption on wireless interfaces with the energy reading from the powermeter.

Since Android apps generally consume considerable amount of energy on display, CPU and other components (*e.g.*, GPS), which complicates the energy validation, we resort to trace driven experiment to generate traffic. We first load several popular mobile websites using the Android browser and record the size and RTT of each flow, as summarized in Table 5. We then host the same contents on our local servers, emulate the RTT using `tc` and write a simple Android app to perform the downloading.

Figure 17 shows the validation result for the old model (*i.e.*, without incorporating signal strength and RTT) and the new model for Nexus One at poor signal strength locations, where the WiFi signal is below -80dBm and 3G signal is below -100dBm. We see that for WiFi, while the error rate of old model ranges from 16.1% to 61.0%, the new model prediction error stays within 5.4%. Similarly, for 3G the new model has an error rate of less than 7.2%, in contrast to 19.3% to 52.1% of the old model.

7. WHAT-IF ANALYSIS

In this section, we show how our new, signal-strength-aware system-call-driven power model enables what-if analysis for the effectiveness of energy optimization techniques that explore signal strength dynamics. In particular, we study how much energy could be saved for a subset of the users from our 3785-user trace by exploring the simple technique of delaying background network traffic to a later time when the signal strength is good.

Trace. We selected three subsets, each containing 5 user traces over a 20-day period, who experienced primarily poor, moderate, and good 3G signal conditions. Similarly we selected three subsets of user traces for WiFi. Recall the user trace contains network usage in 5-minute bins, and the signal strength upon every change.

Setup. Since the system-call-driven power model requires network system call trace as input, we synthesize a flow trace from the 5-minute bin network usage trace as follows. We first divide the traffic in each 5-minute bin into flows with an exponential flow size distribution with a mean of 4344 bytes, and assign the start time for each flow following a uniform distribution within the 5 minutes period. Then we assign end-to-end RTTs to each flow following a uniform distribution between 5ms to 100ms for WiFi [7] and 100ms to 700ms for 3G [12].

Delayed flow scheduling. Our simple signal-strength-aware flow scheduling works as follows. For each flow, we determine if it belongs to background traffic following a probability of `BG_RATIO`. For each background flow, if the signal strength at its start time is below `RSSI_THRESHOLD` and the signal strength increases to above the threshold within the following `MAX_DELAY_TIME`, the flow will be delayed till when the signal condition crosses the threshold. We assume signal strength can be predicted, following techniques such as [23]. The `RSSI_THRESHOLD` is set to -80dBm and -100dBm for WiFi and 3G, respectively.

Results. We use the new power model to calculate each user’s energy consumption on data communications, with and without the delay scheduling. Figure 18 plots the energy reduction from the delay scheduling for 15 users over the 20-day period, with `MAX_DELAY_TIME` of 2 hours and 12 hours, respectively. We see almost linear increase in energy saving as the background traffic increases. We also observe a large variation in energy reduction across different users, ranging from 0.002% to 23.7% under WiFi and from 1.1% to 21.5% under 3G, with `BG_RATIO` of 40% and `MAX_DELAY_TIME` of 12 hours. To understand the reason, we plot the RSSI over time for three selected 3G users, who have little, moderate and a lot of energy saving (not shown due to page limit). We clearly observe that the more fluctuation the signal strength over time, the more energy saving from the delay scheduling.

8. RELATED WORK

We have already discussed related work on WiFi and 3G power modeling and how our new model differs from them in §4. To our best knowledge, our work is the first measurement study of 3G and WiFi signal strength experienced by a large number of users in




Figure 18: Energy reduction from delay scheduling. Each line represents a unique user.

daily life. In the following, we focus on previous work on the impact of poor signal strength on mobile device energy consumption.

Schulman *et al.* [23] observe that cellular signal varies by location, and strong signal reduces energy cost of communicating, and develop track-based signal strength prediction and energy-aware scheduling algorithms for two specific workloads by deferring or prefetching data during good signal conditions. In [9], the authors measure the power draw of WiFi-based phones to increase slightly under poor signal strength, when dynamic power control is enabled. In [6], the authors perform an in-depth study of power dissipation of smartphone components and find GSM dissipates 30% more energy when transferring at poor signal strength. Different from these work, we systematically quantify and break down the impact of poor WiFi and 3G signal strength on all relevant layers of the network stack, and for varying flow sizes and server RTTs.

9. CONCLUSION

In this paper, we have performed the first measurement study of 3G and WiFi signal strength experienced by a large number (3785) of smartphone users over 1 to 19 months of daily usage. Our analysis has shown that smartphone users are routinely experiencing significant 3G and WiFi signal variations during daily active phone usage. Further, the 3785 users we studied performed on average 43% and 21% of their foreground data transfers during poor 3G and WiFi signal strength, respectively. Our trace analysis motivates the need for diagnosing cellular network services to improve the user experience.

We further performed controlled experiments to quantify the energy impact of poor signal strength on data transfers, and drew implications on energy-efficient app design. We then developed a new system-call-driven power model that improves the accuracy over the previous state of the art by taking into account the impact of signal strength and RTT. Using our new model, we show that simply delaying of background traffic can reduce the total energy consumption of data communication by up to 23.7% and 21.5% under WiFi and 3G, respectively, assuming a maximum delay of 12 hours. Our new power model enables more accurate energy profiling for smartphone apps [18] and what-if analysis of current and future optimization techniques that explore signal strength dynamics. We plan to explore these avenues in our future work.

10. REFERENCES

- [1] Aircrack-ng. <http://www.aircrack-ng.org/>.
- [2] Monsoon power monitor. <http://www.monsoon.com/LabEquipment/PowerMonitor/>.
- [3] Smartphone statistics 2011. <http://www.digitalbuzzblog.com/2011-mobile-statistics-stats-facts-marketing-infographic/>.
- [4] N. Balasubramanian, A. Balasubramanian, and A. Venkataramani. Energy consumption in mobile phones: a measurement study and implications for network applications. In *Proc of IMC*, 2009.
- [5] I.-S. S. Board. Wireless lan medium access control (mac) and physical layer (phy) specification. *Electronics*, 1999(802.11), 1997.
- [6] A. Carroll and G. Heiser. An analysis of power consumption in a smartphone. In *Proc. of USENIX ATC*, 2010.
- [7] N. Ding, A. Pathak, D. Koutsonikolas, C. Shepard, Y. C. Hu, and L. Zhong. Realizing the full potential of psm using proxying. In *Proc. of IEEE INFOCOM*, 2012.
- [8] H. Falaki, D. Lymberopoulos, R. Mahajan, S. Kandula, and D. Estrin. A first look at traffic on smartphones. In *Proc. of IMC*, 2010.
- [9] A. Gupta and P. Mohapatra. Energy consumption and conservation in wifi based phones: A measurement-based study. In *SECON '07*, 2007.
- [10] G. Holland, N. Vaidya, and V. Bahl. A rate-adaptive mac protocol for multihop wireless networks. In *Proc. of ACM MOBICOM*, 2001.
- [11] J. Huang, F. Qian, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck. A close examination of performance and power characteristics of 4g lte networks. In *Proc. of Mobisys*, 2012.
- [12] J. Huang, Q. Xu, B. Tiwana, Z. M. Mao, M. Zhang, and P. Bahl. Anatomizing application performance difference on smartphones. In *Proc. of Mobisys*, 2010.
- [13] A. Kamerman and L. Monteban. WaveLAN ii: A high-performance wireless LAN for the unlicensed band. In *Bell Labs Technical Journal*, 1997.
- [14] M. Lacey, M. H. Manshaei, and T. Tuletti. IEEE 802.11 rate adaptation: A practical approach. In *Proc. of ACM MSWiM*, 2004.
- [15] C.-Y. Li, C. Peng, S. Lu, and X. Wang. Energy-based rate adaptation for 802.11n. In *Proc. of ACM MobiCom*, 2012.
- [16] R. Mittal, A. Kansal, and R. Chandra. Empowering developers to estimate app energy consumption. In *Proc. of ACM MobiCom*, 2012.
- [17] A. Pathak, Y. C. Hu, and M. Zhang. Bootstrapping energy debugging for smartphones: A first look at energy bugs in mobile devices. In *Proc. of Hotnets*, 2011.
- [18] A. Pathak, Y. C. Hu, and M. Zhang. Where is the energy spent inside my app? fine grained energy accounting on smartphones with eprof. In *Proc. of EuroSys*, 2012.
- [19] A. Pathak, Y. C. Hu, M. Zhang, P. Bahl, and Y.-M. Wang. Fine-grained power modeling for smartphones using system-call tracing. In *Proc. of EuroSys*, 2011.
- [20] F. Qian, Z. Wang, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck. Profiling resource usage for mobile applications: a cross-layer approach. In *Proc. of Mobisys*, 2011.
- [21] F. Qian, Z. Wang, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck. Characterizing radio resource allocation for 3g networks. In *Proc. of IMC*, 2010.
- [22] B. Sadeghi, V. Kanodia, and E. Knightly. Opportunistic media access for multirate ad hoc networks. In *Proc. of ACM MOBICOM*, 2002.
- [23] A. Schulman, V. Navda, R. Ramjee, N. Spring, P. Deshpande, C. Grunewald, K. Jain, and V. N. Padmanabhan. Bartendr: a practical approach to energy-aware cellular data scheduling. In *Proc. of ACM Mobicom*, 2010.
- [24] A. Shye, B. Scholbrock, and G. Memik. Into the wild: studying real user activity patterns to guide power optimizations for mobile architectures. In *MICRO*, 2009.
- [25] S. Wong, H. Yang, S. Lu, and V. Bharghavan. Robust rate adaptation for 802.11 wireless networks. In *Proc. of ACM MobiCom*, 2006.
- [26] L. Zhang, B. Tiwana, Z. Qian, Z. Wang, R. Dick, Z. Mao, and L. Yang. Accurate Online Power Estimation and Automatic Battery Behavior Based Power Model Generation for Smartphones. In *Proc. of CODES+ISSS*, 2010.