Ubiquitous and Mobile Computing
CS 528

Smartphone Energy Drain in the Wild: Analysis and Implications

Jiang Miao, Fu Zhan
Computer Science Dept.
Worcester Polytechnic Institute (WPI)

Introduction

Why we need to focus on the battery life of modern smartphones?

- **1. The user experience has been severely limited by the phone battery life**

E.G.. A survey in May 2014 by research company GMI of 1000 Britons shows 89% rated long battery life as an “important” factor when buying a new smartphone , which means long battery life rated higher than all the other features.

- **2. Find a available method to improve battery performance is critical**

  Understanding the energy drain model:

  a. The energy drain of wireless interfaces such as WiFi or cellular

  b. Different users set very different device configurations

  c. Different users spend differing amounts of time each day on the phone.

  d. Different users install and play with different apps on their devices

  e. Different users can have very different usage patterns on a same App

Power Modeling for Phones

What a power model we need to do in this experiment:

- 1. It can be used to measure phones in the wild, not in lab
- 2. It can measure the energy drain of individual apps and services concurrently running on the phone
- 3. It can only be collected by modifying the Android framework or the kernel of the phones and should not rely on packet-level trace.
- 4. It can incorporate different uses’ behaviors such as WiFi beaconing, cellular paging and SOC suspicion.(SOC is short for System on a Chip, an integrated circuit that combines all the primary components of a mobilephone into a single chip)

Modeling Overview

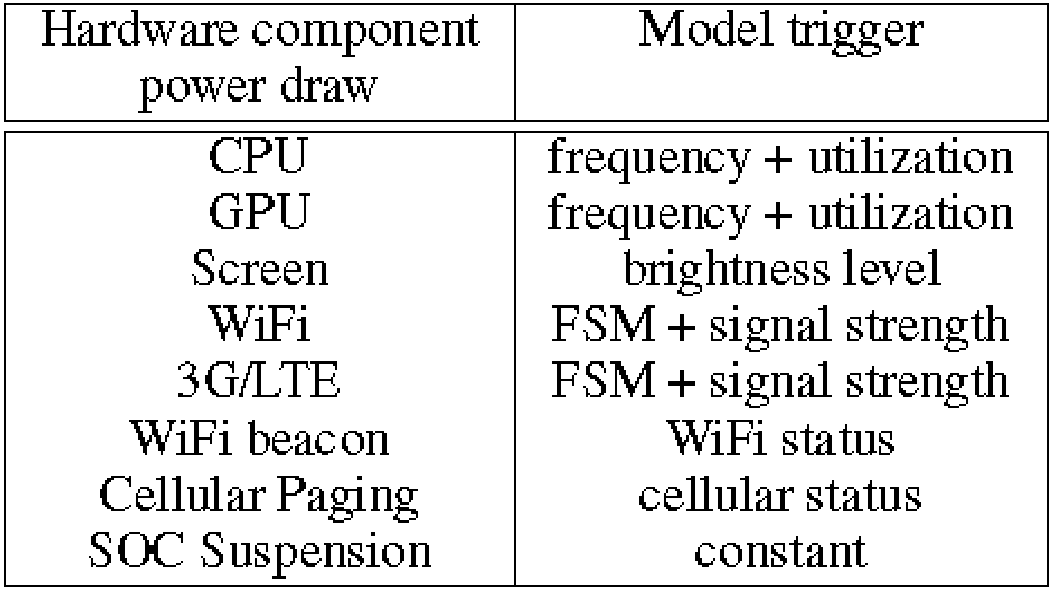

Hardware component power draw	Model trigger
CPU	frequency + utilization
GPU	frequency + utilization
Screen	brightness level
WiFi	FSM + signal strength
3G/LTE	FSM + signal strength
WiFi beacon	WiFi status
Cellular Paging	cellular status
SOC Suspension	constant

- The set of components showing significant power draw on Galaxy S3 and S4
- Our model described which assumes different components are independent

Modeling Details

- CPU: Using CPU microbenchmarks to achieve the relationship between the CPU power draw and CPU operating frequency

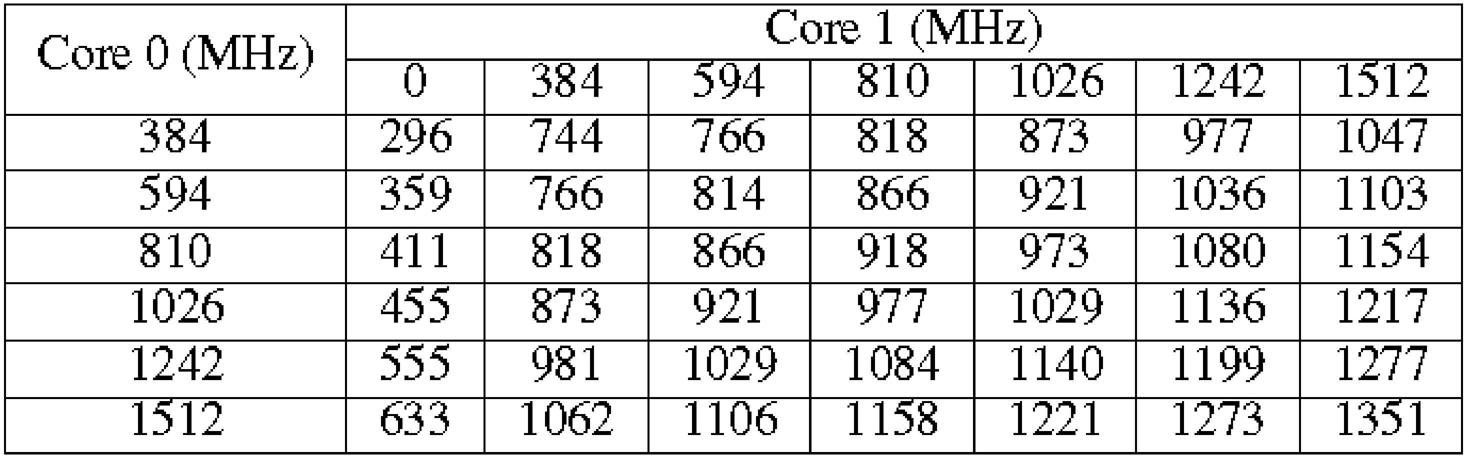

Core 0 (MHz)	Core 1 (MHz)						
	0	384	594	810	1026	1242	1512
384	296	744	766	818	873	977	1047
594	359	766	814	866	921	1036	1103
810	411	818	866	918	973	1080	1154
1026	455	873	921	977	1029	1136	1217
1242	555	981	1029	1084	1140	1199	1277
1512	633	1062	1106	1158	1221	1273	1351

This table shows the CPU power draw at 100% CPU utilization for Galaxy S3 under a range of frequencies.

Modeling Details

- Screen: Design a power model based on screen brightness and ignored screen content to reduce the overhead.

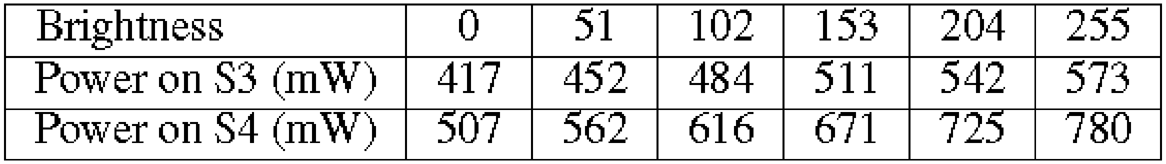

Brightness	0	51	102	153	204	255
Power on S3 (mW)	417	452	484	511	542	573
Power on S4 (mW)	507	562	616	671	725	780

Galaxy S3 and S4 screen power for 6 sample brightness levels

- GPU: Recording the duration of each GPU frequency and state combination every 1 second to predict the GPU power draw

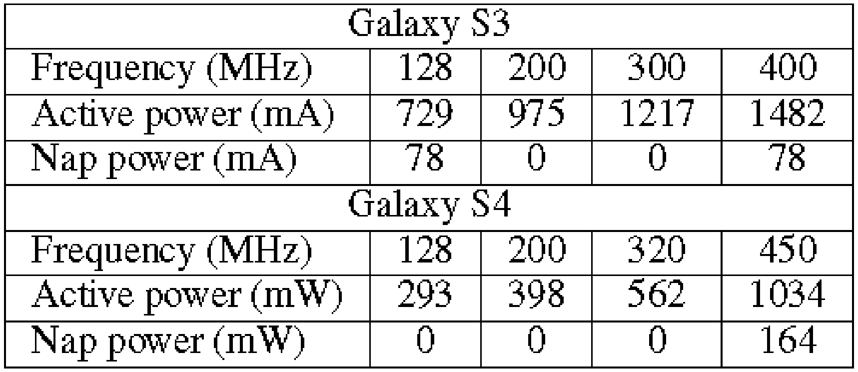

Galaxy S3				
Frequency (MHz)	128	200	300	400
Active power (mA)	729	975	1217	1482
Nap power (mA)	78	0	0	78
Galaxy S4				
Frequency (MHz)	128	200	320	450
Active power (mW)	293	398	562	1034
Nap power (mW)	0	0	0	164

We run GPU microbenchmarks to generate workload and in the meanwhile measure the power draw using the power meter.

Modeling Details

- WiFi/3G/LTE State: They have multiple power states and the power draw and duration at the Active state which is affected by the wireless signal strength

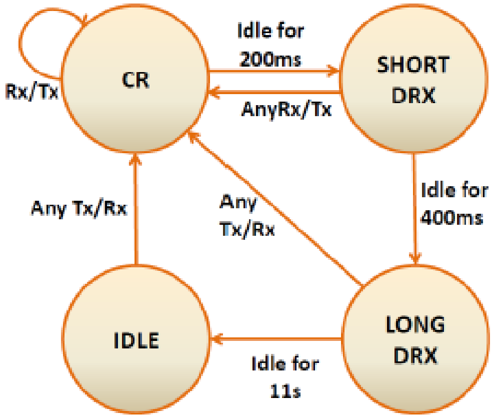

(a) LTE

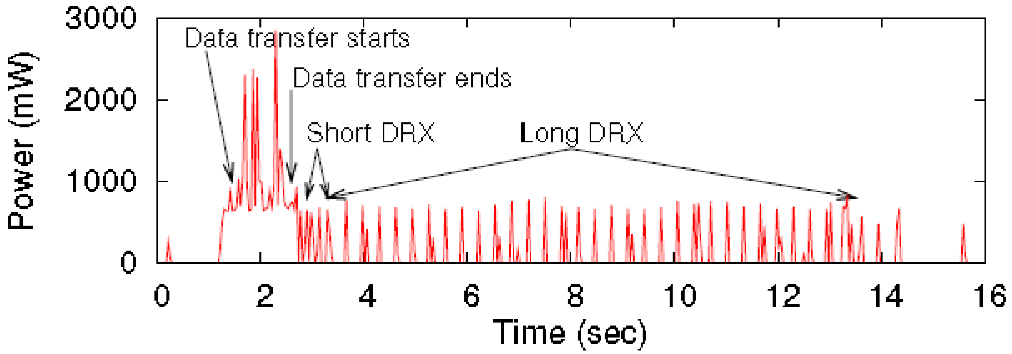

Figure 18: LTE Power states on Galaxy S3

IDLE(The interface is in idle states): when the User Equipment (UE) does not send or receive any data.
CR(Continuous Reception): When the UE sends or receives any data, the interface enters the CR state and consumes high power.
Short DRX(Discontinuous Reception): After the UE finishes data transfer and becomes idle for 200ms, the interface consumes little power but wakes up frequently to check for incoming traffic
Long DRX: The interface enters the Long DRX state after staying in Short DRX for 400ms without receiving any data.

Modeling Details

- WiFi Beacon: When the WiFi radio is associated with an APP and in power saving mode, the WiFi radio wakes up at fixed intervals to receive beacons from the APP. 1.1 mA for screen off and 3.3 mA for screen on for both Galaxy S3 and S4.
- Cellular paging: a celluar network, the base station periodically broadcasts a message during the 3G/LTE Idle state to signal incoming downlink data. The values are 8.3 mA on S3 and 2.3 mA on S4.
- SOC during suspension: We turn the screen and WiFi off, set the phone in airplane mode; soon after the SOC is put in suspension by the power manager, and we measure the SOC base power draw in this state. The constant power draw of the SOC suspended state are 3.8 mA and 5.1 mA for Galaxy S3 and S4

Modeling Validation

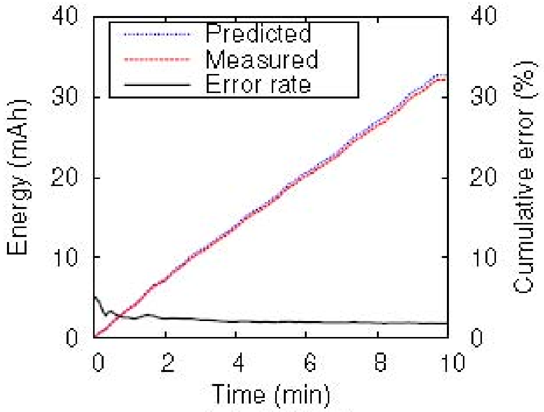

(a) CPU

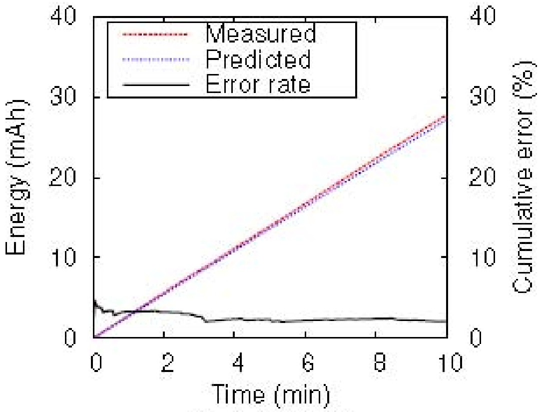

(b) WiFi

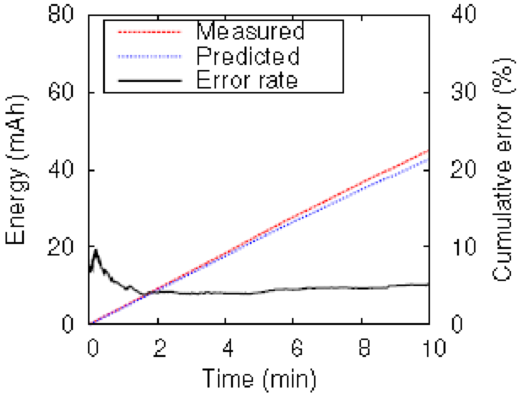

(c) LTE

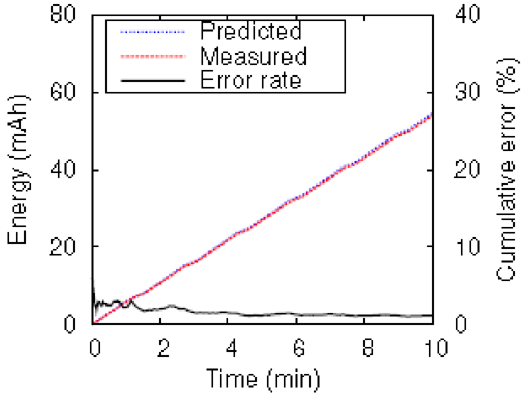

(d) GPU

a) CPU
b) CPU+NetWork
c) CPU+NetWork
d) CPU+GPU+ Screen

Modeling test

Precondition: Install 25 top apps on Google Play including 11 games, 7 online chat apps, 4 music apps and 3 news apps.

- Screen on: A normal user performed similar operations for the same type of apps, 2-3 minutes each, under WiFi and under LTE.
- Screen off: A user login to all these apps. Then we left the phone screen-off for 1 hour with either WiFi or LTE connectivity.

The cumulative estimated energy drain beyond 20min	**WiFi**	**LTE**
Screen on	10%	10.3%
Screen off	4.1%	5.0%

CPU Time Analysis

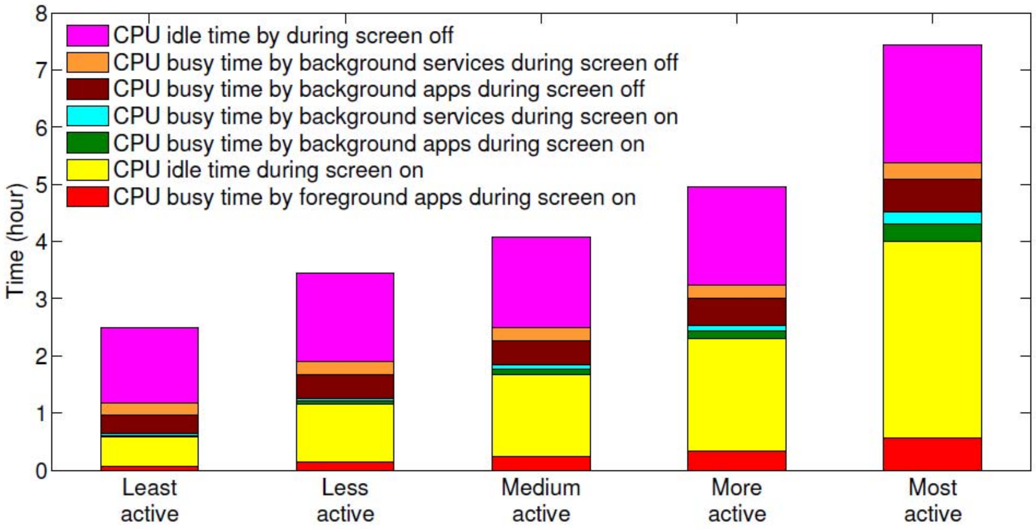

Average daily CPU time breakdown of 5 groups of the 1520

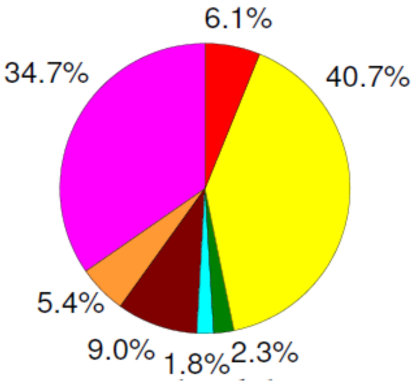

Daily CPU time percentage breakdown, average over all users

- **CPU idle:** 40.7% during screen-on and 34.7% during screen-off
- **Screen-on vs. screen-off:** total CPU busy time during screen-on and screen-off periods are 10.2% to 14.4%
- **Services vs. apps:** background services account for about 28.1% of total CPU busy time (7.2% in absolute)

Energy Analysis

Energy breakdown by activities

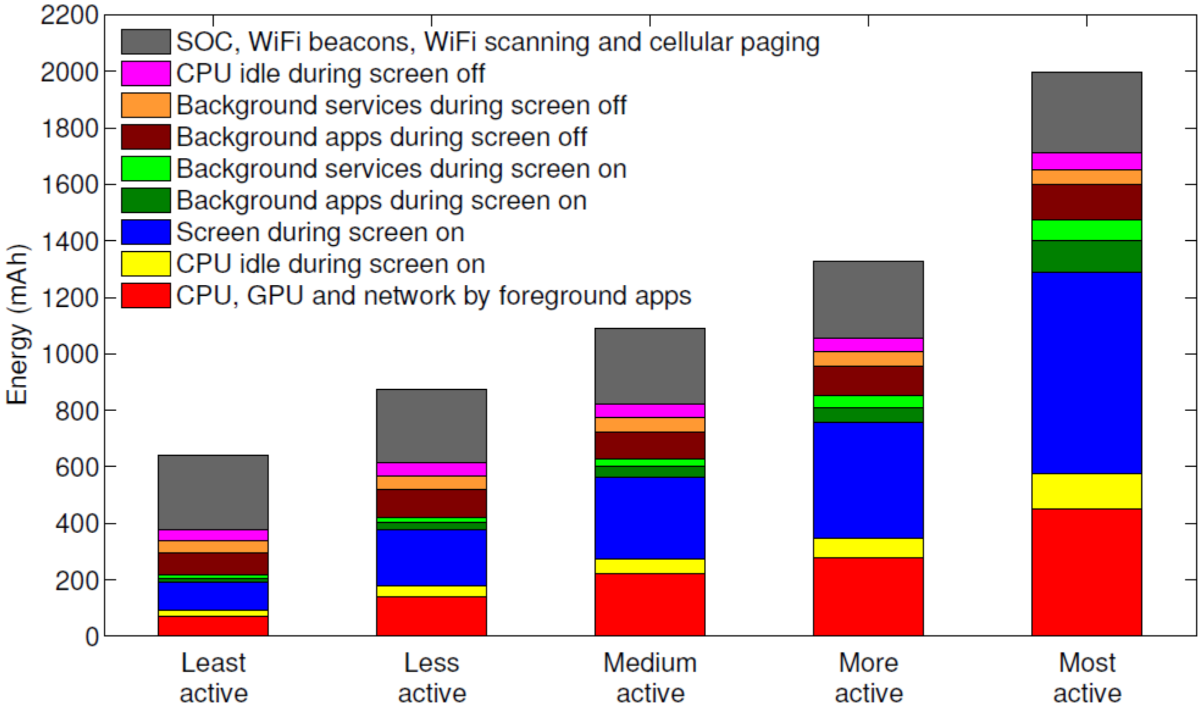

Average daily energy drain breakdown of 5 groups of 1520 users

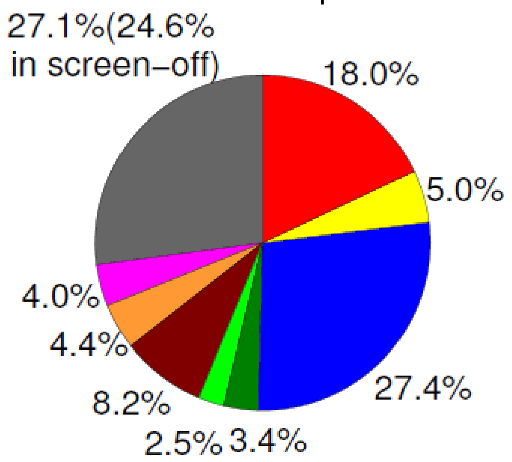

Daily energy percentage breakdown, averaged over all users

- **Suspended state energy :** SOC, WiFi beacon, WiFi scanning and cellular paging activities, account for 24.6%
- **Screen energy:** 58.8% energy incurred during screen-on periods
- **Useful energy in Screen-on vs. screen-off:** apps and services during screen-off 12.6% , during screen-on 23.9%
- **CPU idle energy:** only drains on average 9.0% of the total energy.(but CPU spends 75.4% of the total CPU time in idle)

Energy Analysis

Energy breakdown by components

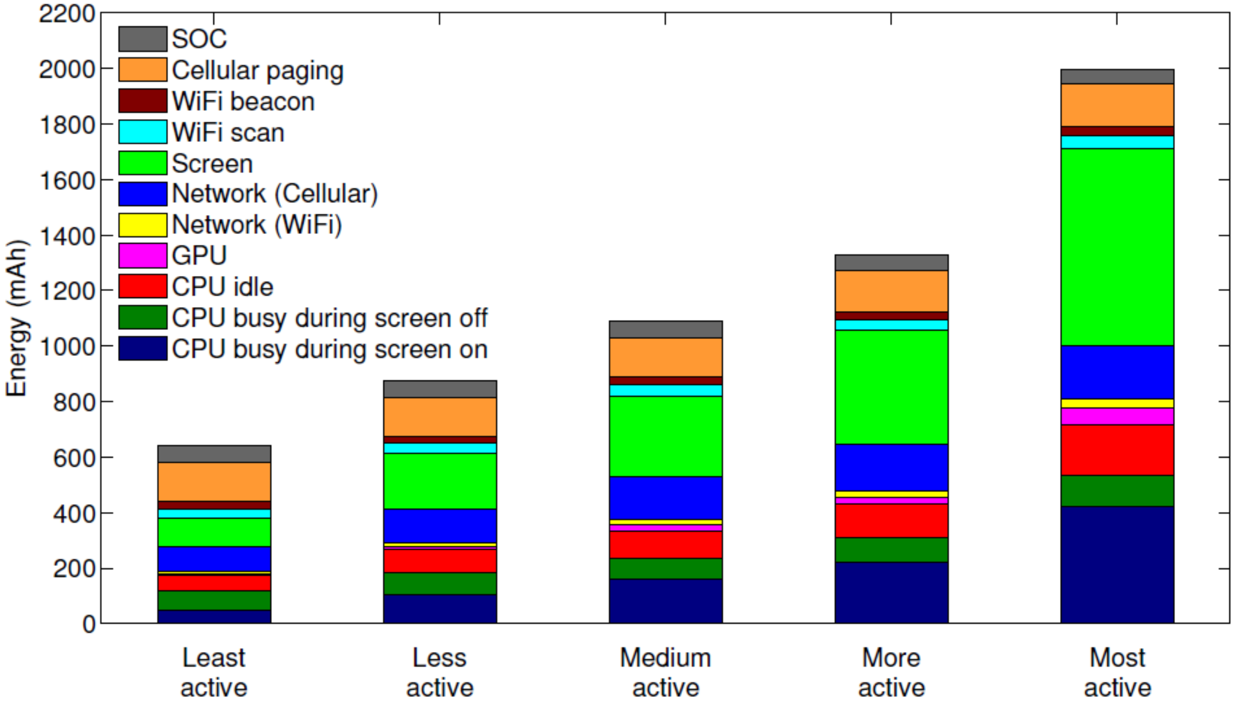

Average daily energy drain breakdown of 5 groups of 1520 users

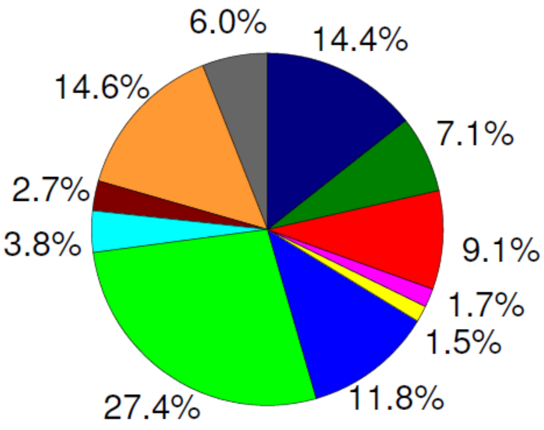

Daily energy percentage breakdown, averaged over all users

- **Cellular paging vs. WiFi beacon:** cellular paging 14.6% and WiFi beacon 2.7%,cellular paging is a significant energy hogger.
- **Cellular vs. WiFi :** cellular 11.8% and WiFi 1.5%, cellular drains significantly more energy than WiFi, tail energy
- **CPU:** The busy CPU energy during screen-on is twice that during screen-off
- **GPU:** 1.7%, used by foreground apps during screen-on periods

Component Analysis

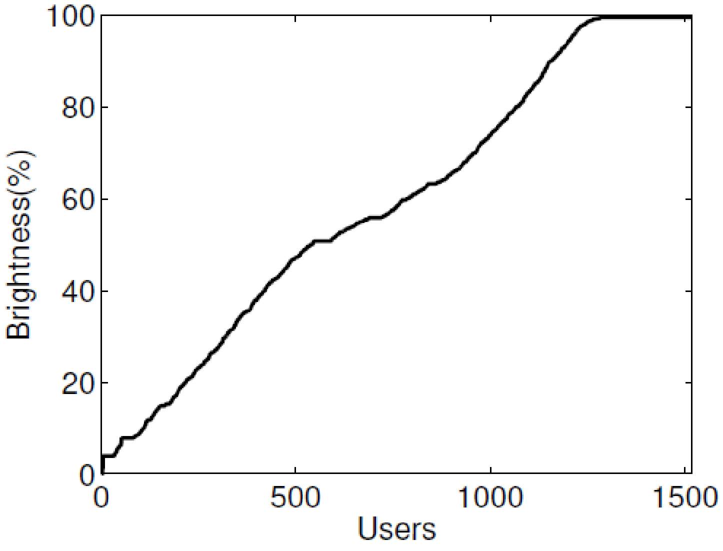

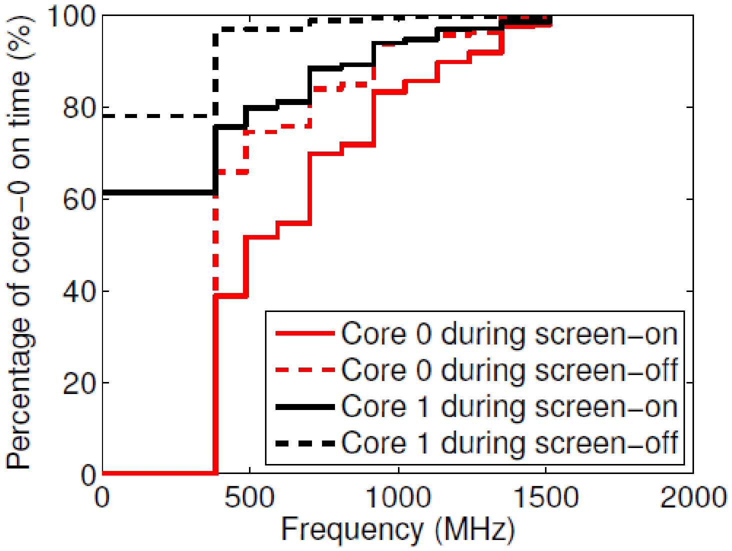

- **Screen** almost uniform distribution, with the average and median levels being 59.0% and 58.3%
- **CPU** (1)Core-0 is idle at the lowest 384 MHz for 38.8% of the time during screen-on but 65.7% of the time during screen-off. In screen-off background apps/services wake up and acquire some wakelocks and then wait for responses (2) Both core-0 and core-1 tend to be busy at higher frequencies during screen-on than during screen-off periods
- **GPU:** GPU drains only 1.7% while the screen drains 27.4%. Most non-game apps use little GPU.

Component Analysis

- **Networking**. compare the time spent and bytes transmitted over the two types of wireless technologies.

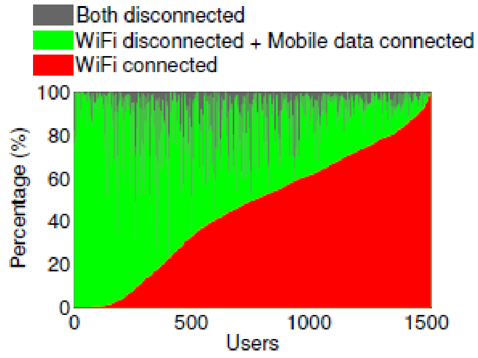

(a) Distribution of breakdown of time spent in WiFi and cellular states.

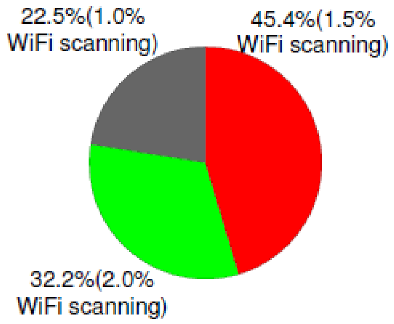

(b) Average percentage breakdown of time spent in WiFi and cellular.

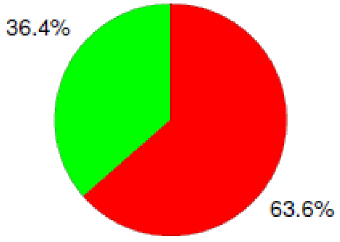

(c) Average percentage breakdown of bytes transmitted in WiFi and cellular.

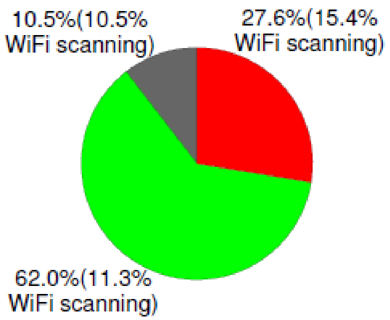

(d) Average percentage breakdown of energy drain by WiFi and cellular.

(a) almost uniform distribution in terms of the percentage time spent in WiFi between 0% to 100%.
(b) 45.5% of the time connected to WiFi, 32.2% to mobile data, and 22.5%from both. WiFi rarely perform scanning
(c) 63.6% of total traffic is transmitted over WiFi compared to 36.4% in cellular
(d) each device spent 1.4x more time in WiFi and transmitting 1.8x more bytes inWiFi, but drains 4.2x less energy in WiFi, excluding WiFi scanning energy

App Energy Analysis

38.1% of the daily energy drain of a device is by apps and services. The rest of the energy drain are largely fixed for a given hardware.

- **Energy Drain, Screen-on vs. Screen-off**

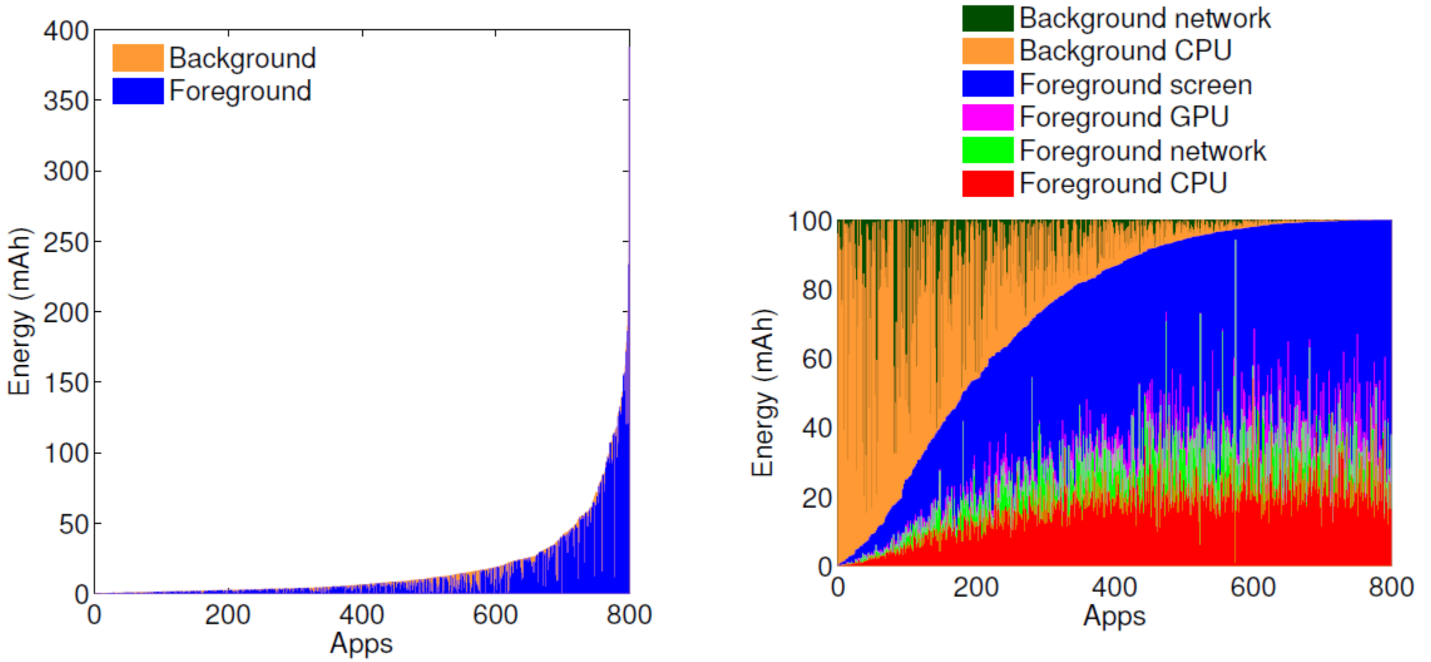

(a) Absolute breakdown (b) Relative breakdown

- background energy can be significant for many apps
- Within the foreground app energy, screen energy is the largest portion
- CPU and GPU energy dominates networking energy, ratios are 2.7x and 2.8x for foreground and background energy

App Energy Analysis

- **App Energy Drain Rate**

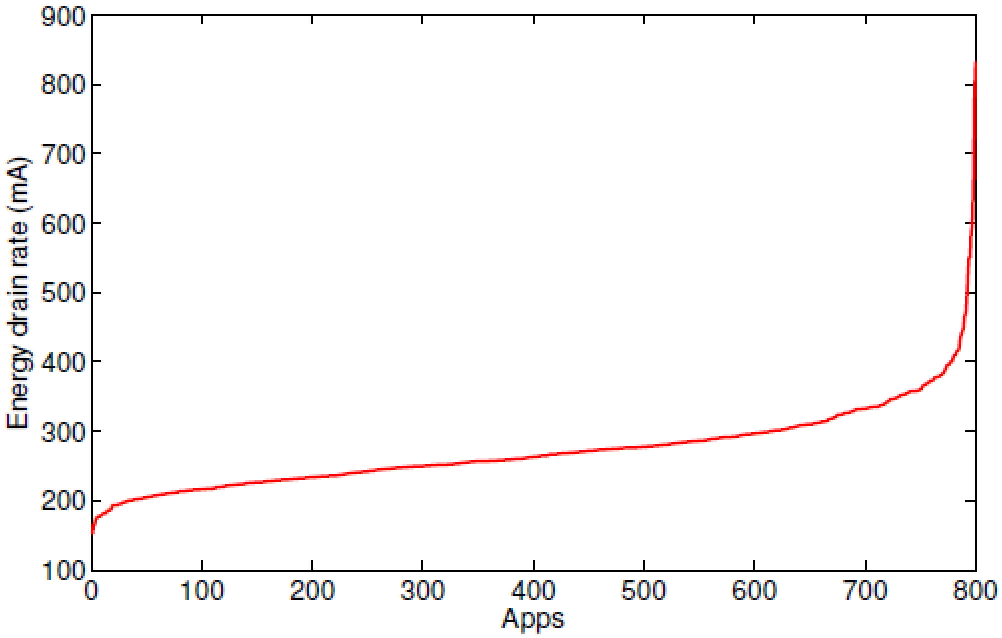

- drain rate (EDR): total foreground energy drain of an app divided by the total foreground time
- 92.6% of the apps have an average power draw between 200-400 mA

App Energy Analysis

- **App Categories**

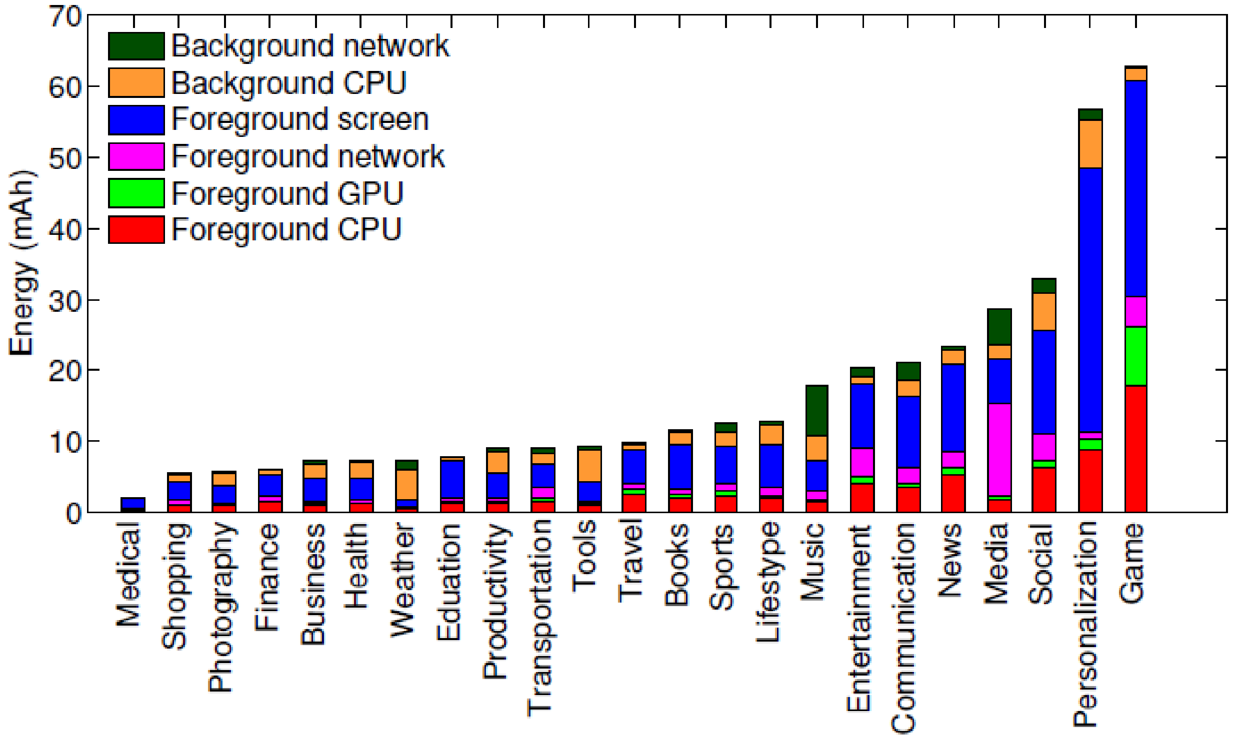

- **Total energy:** Games and Personalization drain most energy
- **Background energy:** varies significantly, little correlation with total energy drain.
- **Screen energy:** dominating chunk of the total energy.
- **GPU energy:** Game apps much higher than all other app categories
- **CPU energy:** The highest categories are Game, Travel, and Finance
- **Network energy:** Media and Music drain more energy

Evolution study

- **Device Evolution S3 vs. S4, Jellybean vs. KitKat**

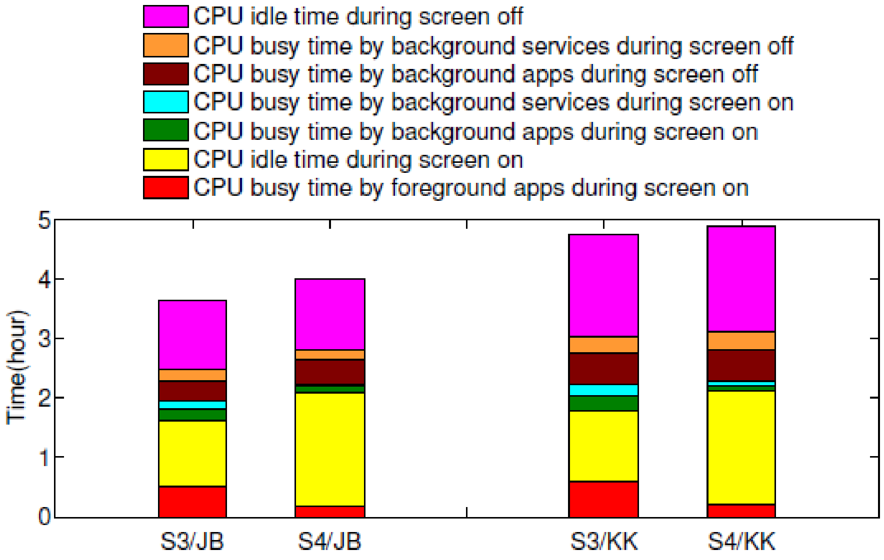

(a) Total CPU time breakdown

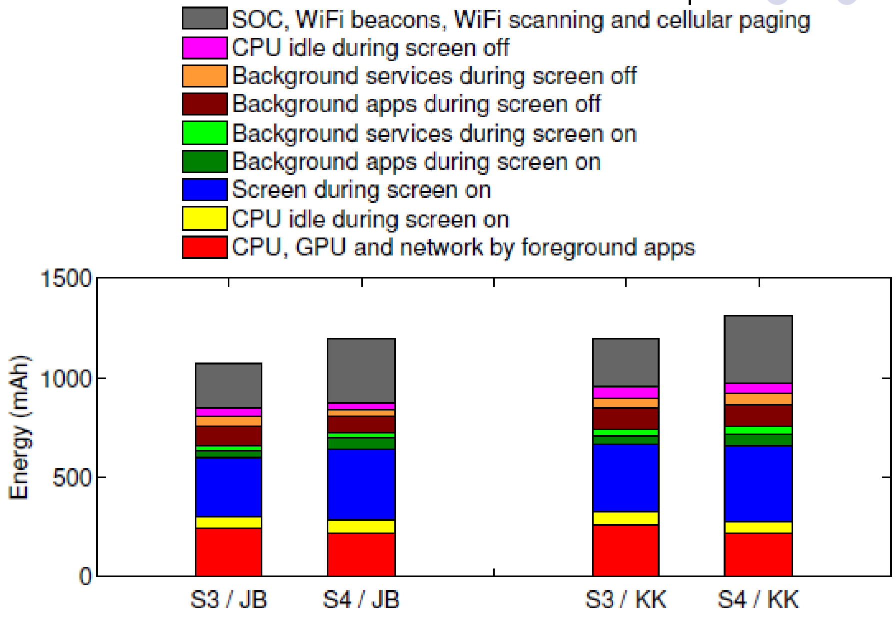

(b) Total energy breakdown

the average CPU time of S4/JB devices is 8.1% longer than S3/JB devices. mainly comes from increased screen-on time

8.1% increase in total CPU time of S4/JB over S3/JB devices translates into 11.3% energy increase

the average CPU time of S3/KK and S4/KK devices are 33.2% and 22.5% higher than the corresponding S3/JB and S4/JB. The increase appears to be mainly coming from increased background CPU time during screen-off, 37.6% and 38.4%.

Evolution study

- **Cellular Evolution: 3G vs. LTE**

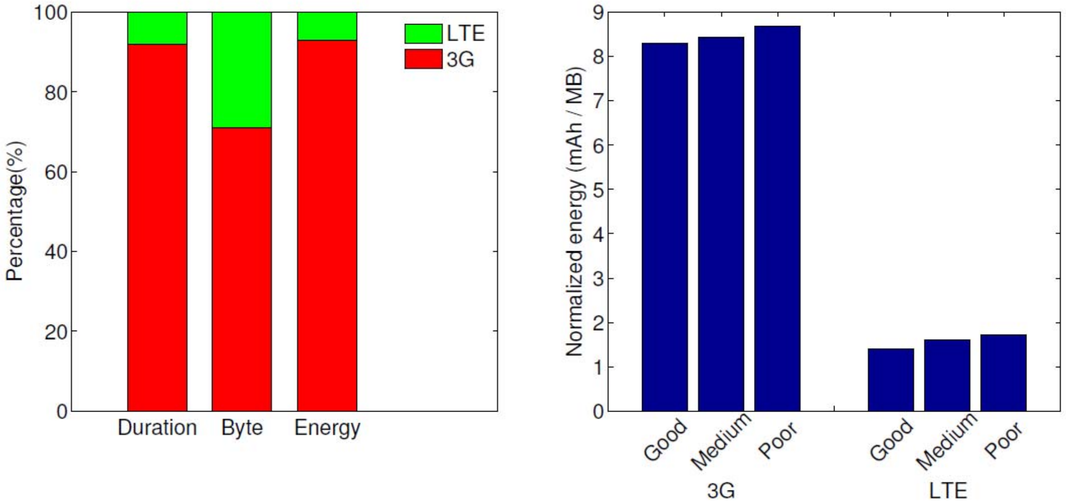

(a) Percentage energy, bytes, and duration (b) Normalized energy by bytes

(a) Although LTE accounts for 29.0% of total bytes are transmitted, it only accounts for 8.7% of the time connected to 3G/LTE and only consumes 6.9% of the total 3G/LTE network energy

(b) on average 3G drains 5.9x, 5.3x, and 5.1x more energy per MB transmitted under good medium and poor signal strength

Evolution study

- **App Evolution: App Updates**

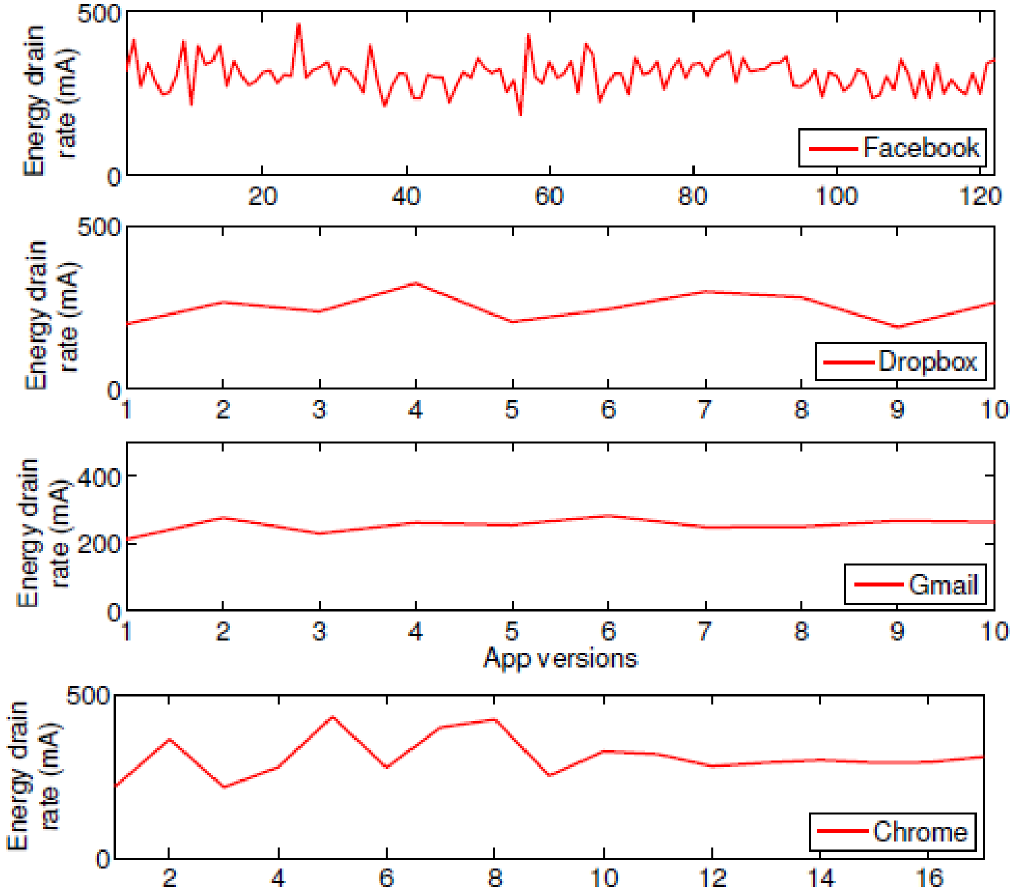

- Facebook has more than 120 versions Its fluctuating and high foreground power happens mainly due to its frequently updated new features.
- Dropbox and Gmail have less foreground power variation and a low average foreground power. (1)synchronize with servers in background, minimizing foreground network energy. (2)simple UI lower CPU and GPU energy.

Related Work

- **Power modeling of smartphones.** Already discussed various previous work on power modeling of smartphones in §2.1.
- **Measurement study.** study the energy drain of mobile apps in the wild. This paper collected trace from a much broader user base, developed a power model that captures both utilization-based and FSM-based components (for WiFI, 3G and LTE), and performed detailed activity and energy analysis across devices, components, apps, and technology and app evolutions.

Conclusion

- developed a hybrid utilization-based and FSM-based model
- much insight on energy drain across devices (users), device components, apps, and multiple technology and app evolutions
- draw implications to SOC vendors, cellular carriers, and app developers on better system, network, and app design to extend battery life

References

- Xiaomeng Chen. Ning Ding. Abhilash Jindal.Smartphone Energy Drain in the Wild: Analysis and Implications. Purdue University West Lafayette, IN 47907-1285
- theGuardian, your smartphone's best app? battery life, say 89% of britons. www.theguardian.com/technology/2014/may/21/ your-smartphones-best-app-battery-life-say-89.
- TechCrunch, mobile app usage increases in 2014, as mobile web surfing declines.
- H. Falaki, D. Lymberopoulos, R. Mahajan, S. Kandula, and D. Estrin. A first look at traffic on smartphones. In Proc. of IMC, 2010.
- A. Garcia-Saavedra, P. Serrano, A. Banchs, and G. Bianchi. Energy consumption anatomy of 802.11 devices and its implication on modeling and design. In CoNEXT, 2012.
- J. Huang, et al. A close examination of performance and power characteristics of 4g lte networks. In Proc. of Mobisys, 2012.

References

- Q. Xu, J. Erman, A. Gerber, Z. Mao, J. Pang, and S. Venkataraman. Identifying diverse usage behaviors of smartphone apps. In Proc. Of IMC, 2011.
- J. Huang, F. Qian, Z .M. Mao, S. Sen, and O. Spatscheck. Screen-off traffic characterization and optimization in 3g/4g networks. In IMC, 2012.
- 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.