Sandra Helps You Learn: the More You Walk, the More Battery Your Phone Drains

Chulhong Min[1], Chungkuk Yoo[1], Inseok Hwang[2,*], Seungwoo Kang[3], Youngki Lee[4], Seungchul Lee[1], Pillsoon Park[5], Changhun Lee[1], Seungpyo Choi[1], Junehwa Song[1]

[1]School of Computing, KAIST, [2]IBM Research – Austin, [3]Computer Science and Engineering, KOREATECH, [4]School of Information Systems, Singapore Management University, [5]Division of Web Science and Technology, KAIST

[1,5]{chulhong, ckyoo, seungchul, pillsoon.park, changhun, spchoi, junesong}@nclab.kaist.ac.kr, [2]ihwang@us.ibm.com, [3]swkang@koreatech.ac.kr, [4]youngkilee@smu.edu.sg

ABSTRACT
Emerging continuous sensing apps introduce new major factors governing phones' overall battery consumption behaviors: (1) added nontrivial persistent battery drain, and more importantly (2) different battery drain rate depending on the user's different mobility condition. In this paper, we address the new battery impacting factors significant enough to outdate users' existing battery model in real life. We explore an initial approach to help users understand the cause and effect between their physical activity and phones' battery life. To this end, we present *Sandra*, a novel *mobility-aware* smartphone battery information advisor, and study its potential to help users redevelop their battery model. We perform an extensive explorative study and deployment for 30 days with 24 users. Our findings reveal what they essentially learned, and in which situations they found Sandra very helpful. We share the lessons learned to help in the design of future mobility-aware battery advisors.

Author Keywords
Continuous sensing; battery; smartphone; user perception

ACM Classification Keywords
H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION
Smartphones' battery life is users' major interest in their everyday lives. When battery level is low, they often ask themselves: "*How long will my phone last from now?*" and "*What should I do to keep my phone alive until I get home?*" Throughout years of experiences, many users have acquired empirical practices to make an educated guess about the lifetime of their phones at a given battery level. Also, they know well-known factors draining their phone battery faster, e.g., frequent app use, long calls, GPS use, brighter screen, and weaker cellular signals. We believe that these practices help users develop an implicit battery model. Research prototypes and even commercial apps have been presented to supplement such users' own models, by providing an accurate estimate of remaining battery [44], identifying major power-draining features [26], and deactivating noncritical features to extend battery life [32].

We advocate that it is time to update such battery models in users' mind. The new players changing the rule of the game are rapidly prevailing continuous sensing apps (CSAs hereinafter) such as Google Fit, Apple Health, S-Health, and Moves. CSAs introduce new major factors governing phones' battery consumption: (1) nontrivial persistent battery drain is added, and more importantly, (2) the user's different mobility conditions result in different battery drain rates. In other words, the more a user walks or runs, the more background power his phone consumes. We believe these factors are unexpected and counterintuitive to many users. The following scenario highlights CSAs interfering with the existing battery model in users' minds.

Becca, a young New Yorker, is walking to Downtown for her dinner appointment with Eric. It is going to be her two-hour-long exercise walk; this week she decided to do more workouts and installed a new fitness tracker app. She left home with 25% battery in her phone, but she is certain that she shall have enough battery to call him when she arrives. She knows the phone battery drops only a few percent an hour as long as she does not use it. She finished her walk keeping her phone untouched. However, she gets shocked as her phone is turned off, screwing up her dinner plan. She does not understand why this erratic battery drop happened.

Why did such an erratic battery drop happen that Becca did not understand? CSAs run in the background and monitor the user's physical activity. For battery saving, their internal logics often adopt conditional sensing pipelines, e.g., a location-tracking app triggers power-hungry GPS only when energy-efficient accelerometers detect a noticeable motion potentially indicating walking or biking [21, 35]. As

* The corresponding author

UbiComp '15, September 7-11, 2015, Osaka, Japan.
Copyright 2015 © ACM. 978-1-4503-3574-4/15/09…$15.00.
http://dx.doi.org/10.1145/2750858.2807553

such, CSAs not only consume battery in the background continuously, but also their battery uses vary depending on users' mobility. Without understanding this, users may perceive growing disparities between their estimation of the near-future battery consumption and the actual outcomes.

In this paper, we address the new battery-impacting factors of CSAs which are prominent enough to outdate users' existing battery model in real life, and explore an initial approach to help users count the new cause and effect between their physical activity and their phones' battery lives. We first expose the macroscopic impact that today's commercial CSAs have already been exhibiting on our smartphone battery experiences. Next, we present *Sandra*, a novel *mobility-aware* smartphone battery information advisor and study its potential to help users redevelop the battery models in their mind with the new factors of CSAs.

Our explorative study exposes the significance of the problem. We measure the significant difference in the phones' battery life, in terms of both average standby consumption and its variance upon the changes of user mobility conditions. We also find the conflicts between the existing battery models in users' minds and real experiences with unexpected or counter-intuitive battery behaviors potentially attributable to CSAs. Despite the growing prevalence of commercial CSAs, we identify that the users are mostly unaware of the CSAs' operation, and more importantly, the causality that their own physical activities leads to largely different battery consumptions rates.

Sandra, our novel smartphone battery life advisor, is designed to highlight the different impacts from the user's different mobility conditions. Sandra features (1) a forecaster providing multiple standby battery life estimates under different future mobility conditions and (2) an archive providing a retrospective summary of past battery drain rates categorized by different mobility conditions. For user-perceived simplicity, Sandra represents the mobility conditions into different combinations of high-level movement conditions (either *stationary*, *walking,* or *transport*) and major location conditions (either *indoors* or *outdoors*). Sandra has gone through an extensive deployment study for 30 days with 24 users. We report our vivid findings from the users about what they essentially could learn, and in which situations they found Sandra particularly helpful. We also share our lessons to help in the design of future mobility-aware battery advisors.

We note that Sandra is neither a reconfiguration tool extending the phone's battery life nor an omniscient battery life predictor considering every possible variable. The major goals of Sandra are to enlighten users about the new causal factors of their own mobility changes impacting their phones' standby life, and to help them learn such new factors' impact with their existing battery models in mind.

Our main contributions are threefold: (1) we bring forward to the UbiComp community the new user-centric issues of battery-impacting causal factors amplified by proliferating CSAs, that users' physical activities lead to shorter smartphone standby life; (2) we articulate those impacts with real devices, by real users, and under real-life situations; (3) we premier a working prototype smartphone service to help users be aware of the new causal factors and develop newer practices to expect their smartphones' battery life along with their own mobility conditions.

RELATED WORK

A rich body of literature reported how people use batteries of their smartphones. We classified them into works on (1) *battery interface*, (2) *battery management*, and (3) *battery diagnosis*, and elaborate each category in more detail.

Battery interface: Researchers have attempted to find what types of battery information are useful for mobile users and when such information should be provided. Rahmati et al. investigated the influences of a battery interface on users' battery management strategies [38, 39]. They addresses the limitations of conventional battery interfaces and suggests that a new battery interface can guide users to use the limited battery more efficiently. Truong et al. proposed a task-centered battery interface that provides the expected lifetime of a mobile phone when it executes a specific set of applications [42]. Ferreira et al. proposed an interactive battery interface to enhance energy efficiency [9]. Jung et al. proposed an active battery interface that actively provides energy-related information via toast messages [16].

We focus on the *background-running* CSAs and their large battery drains whereas the existing works focus on those of typical foreground apps. To our knowledge, we are the first to explore battery interfaces for CSAs, and report the importance of user mobility factors in battery management.

Battery management: Banerjee et al. examined users' battery charging behaviors and found that recharges mostly occur based on time and location, not remaining battery levels [3]. Ravi et al. proposed a new battery management strategy by predicting the next recharging opportunity. If the battery is likely to be exhausted before the next opportunity, it warns users [40]. Ferreira et al. investigated the recharging habits of users and discussed design chances for battery management [8]. However, these works did not take CSAs into considerations. This limitation leads us to study enhanced battery strategies for CSAs, especially considering continuous and nonlinear power consumptions.

There have also been research efforts to systematically manage CSAs' power consumption. As a common platform for CSAs, they reduce energy use of CSAs [18, 22, 33], coordinate resource use conflicts among CSAs [15, 19], and balance energy use and accuracy of CSAs considering the tradeoff [4, 20]. Unlike these, we aim to help CSA users manage their battery for themselves. We believe our work can complement such system-driven battery management.

Battery diagnosis: Prior works studied abnormal battery drain of smartphones [1, 26, 30, 34]. They systematically

Context	Application/System	Energy optimization
Location	SensLoc [21]: Identify semantically meaningful places	GPS duty cycling, Movement-based WiFi scan
	SurroundSense [2]: Recognize logical places via ambience fingerprinting	Hierarchical sensing with sequential filtering
Physical activity	UbiFit Garden [5]: Encourage physical activity based on inferred user activity status	* Duty-cycling/ movement-based activity monitoring
	Calorie Monitor [25]: Estimate daily caloric expenditure based on activity status	* Duty cycling for movement detection * Movement-based GPS trigger
Conversation	SocioPhone [24]: Monitor conversational turns	VAD-based processing trigger
	SocialWeaver [29]: Perform conversation clustering and build conversation networks	Encounter-based monitoring trigger, VAD-based processing
Emotion	EmotionSense [37]: Recognize emotion, proximity, speaker of conversation	Non silence-based filtering, encounter-based filtering
Sound-related events	SoundSense [28]: Recognize human voice and classify ambient sounds	Hierarchical processing with pre-filtering
Sleep quality	iSleep [12]: Monitor an individual's sleep quality	Sound-based processing trigger
Indoor/ outdoor	IODetector [45]: Detect if a user is indoor or outdoor	Duty cycling

Table 1 Energy optimization of CSAs; *the asterisk mark (*) means that the mentioned optimization strategies can be applied although the proposed systems do not adopt them in the papers*

detected abnormal battery drain and its causes, e.g., energy bugs and misconfigured apps. Based on the information, they give users a guide for battery life improvement. Unlike these works, we focus on making users aware of the impact of users' mobility conditions on the phone's battery life.

Understanding on contextual power consumption: Recently, research communities actively proposed a variety of CSAs to monitor location, physical activity, conversation, emotion, heartrate, etc. [2, 5, 7, 12, 13, 14, 17, 21, 23, 24, 25, 28, 29, 35, 37, 45]. They often employ human context-dependent energy optimization strategies (see Table 1). They adopt conditional sensing pipelines, leveraging low-power sensors in earlier processing stages to trigger high-power sensors only when a user is in a relevant context. For example, a location tracker can detect a user's motion by power-efficient accelerometers, and the power-hungry GPS is triggered only when significant motions are detected [7, 21, 35]. Likewise, the encounters detected by a relatively cheaper Bluetooth scan can trigger expensive sound-based speaker detection for conversation monitoring [24, 29]. Due to such context-dependent logic, the actual power consumption largely depends on how frequently and how long the user's physical behaviors or environmental conditions match the triggering conditions in the logic.

	Context		No-CSA	Google Fit	Moves	Dieter	Accupedo
Nexus 5(5.0)	Indoor	Stationary	45.8	48.4	86.2	69.9	100.3
		Walking	42.5	91.5	198.9	88.4	322.7
	Outdoor	Stationary	54.2	78.1	104.2	87.7	94.3
		Walking	47.9	53.3	208.0	88.5	343.1
Nexus S(4.1.2)	Indoor	Stationary	102.1	128.9	337.4	120.8	120.8
		Walking	96.1	148.9	354.1	380.1	514.8
	Outdoor	Stationary	103.2	118.1	322.9	127.7	163.5
		Walking	98.6	155.1	314.3	362.9	541.7

Table 2. Standby power consumption (mW)

EXPLORATIVE STUDY

To expose the problem's significance, we demonstrate the real-world impact attributable to CSAs' battery behaviors. We focus on two major dimensions: (1) the quantitative battery impact under various user mobility conditions and real-life scenarios, and (2) the user-perceived impact due to CSAs' counterintuitive battery behaviors straying from the users' expectation of their phones' battery consumption.

In this paper, we focus on the mobility conditions as main battery-impacting factors; our main targets are mobility-based CSAs. While diverse CSAs have been proposed in a research domain, widely-used commercial CSAs are mostly mobility-based ones, e.g., pedometer, physical activity tracker, and path tracker. We discuss the battery impact on CSAs with other types of contexts in the Discussion section.

Quantitative Impact: Nonlinear Battery Drains of CSAs

We report our experimental results showing to what extent CSAs *amplify* the impact from user mobility changes to the phone's battery life. We acknowledge that users' location- or movement-changes may change many factors potentially influencing the phones' standby battery drains, e.g., cellular radio strength, Bluetooth devices nearby, etc. However, it is unknown how much additive battery impact CSAs introduce under the given conditions. To explore this question, we conducted two experiments, each with and without a CSA: measuring (1) the phone's battery drain under different mobility conditions and (2) the phone's day-long cumulative battery drain during the course of a real user's diurnal life with natural mobility variations.

For all measurements, we used Nexus 5 and Nexus S and kept the following variables identical: OS version (Android 5.0 for Nexus 5, Android 4.1.2 for Nexus S), cellular carrier, battery age (brand new), and factory-default apps except a CSA for the "with CSA" phones. We used four commercial CSAs, Google Fit[1], Moves[2], Dieter[3], and Accupedo[4]. To ensure identical mobility and environment conditions both

[1] https://play.google.com/store/apps/details?id=com.google.android.apps.fitness
[2] https://play.google.com/store/apps/details?id=com.protogeo.moves
[3] https://play.google.com/store/apps/details?id=net.ultracaption.dieter
[4] https://play.google.com/store/apps/details?id=com.corusen.accupedo.te

with and without a CSA, we had our subject carry multiple phones at the same time in the same way. As we were interested in the phones' standby power impacts, we kept the displays off and had him not interact with the phones.

Standby powers under different mobility conditions

Table 2 lists the standby power consumptions for different CSAs and mobility conditions, measured by a Monsoon power monitor which the user carried on a cart while he moves. Each mobility condition continued for 1 hour and the average powers are shown. Table 2 convinces us that running a CSA generally drains extra standby power, and these extras indeed vary depending on the user's mobility condition. The average increment is 171% compared to No-CSA under the same conditions. Many cases confirm that the Walking condition drains a significantly larger amount of standby power than the Stationary condition.

Besides, we also observe that the increment highly depends on each CSA. On Nexus S, when stationary, the standby powers with Moves are outstandingly larger than with the other CSAs. The increment when walking compared to stationary is also largely different; Accupedo results in 3-4 times higher power drain when walking than stationary, but Moves incurs a relatively small increment when walking.

Day-long battery drains under real life mobility variations

Figure 1 illustrates the diurnal battery drains of two phones carried by a user at the same time for his usual day. Both phones' HW and SW conditions were identical except that a CSA was running on only one phone. To incorporate usual background app operations, we ran Gmail, Hangout, and Google Calendar with the same Google account. Our custom logger ran on both phones and periodically recorded the battery levels. On the CSA-phone, our logger records the mobility conditions as well, but there is no need to be concerned about its power cost since our logger retrieved the mobility pattern by using the application programming interfaces (APIs) of the CSA after the collection session.

The results of Figure 1 (a) are astonishing; the phone with Google Fit depletes its battery roughly three times as fast. We also find clearly steeper drains for a moving user. Figure 1 (b) also shows similar patterns. These observations convince us that the cumulative battery impact of CSAs and the mobility-dependent variance would be at a sufficiently large scale that ordinary users may perceive the differences in their daily life. Still, they may hardly discover this new causality as they only see the cumulative battery impact from many factors including their own foreground phone usages. Still we believe that such a macroscopic impact of CSAs may have given memorable, unexpected experiences to users regarding their phones' battery behaviors. We explore this hypothesis through the user study below.

User Perceptions with CSAs' Battery Behaviors

We recruited 24 participants through our university bulletins, who were actually using CSAs (6 females, mean

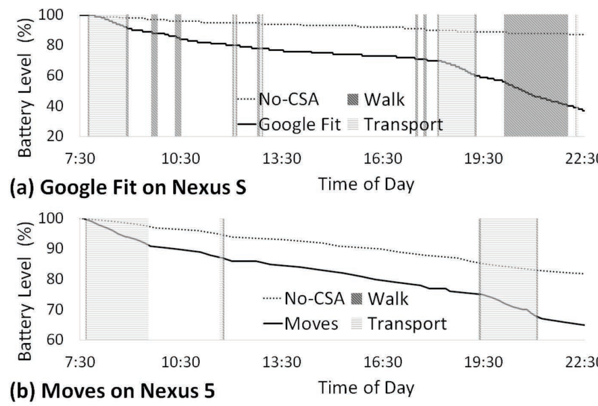

Figure 1. Diurnal battery drains of CSAs

age: 23.7). We used the methods of semi-structured one-on-one interviews 0.5–1 hour long; the interviews were audio-recorded and transcribed. Two researchers analyzed the transcripts individually and discussed together to reach their consensus of high-level themes [41]. Each participant was compensated an equivalent amount of USD 9.

Our participants were using Google Fit (14 users), S-health[5] (2 users; fitness tracker preinstalled in Samsung Galaxy series), Dieter (1 user; a major fitness tracker in Korea), or Moves (7 users). Google Fit and S-health were preinstalled; Dieter and Moves were installed by the users. While these apps' features vary, they basically provide pedometry and basic mobility classifications. They stay in the background, classify the users' mode of locomotion (e.g., walk, run) and estimate the step count or the distance moved. Google Fit and Moves provide APIs so that third-party apps can retrieve the users' mobility conditions nearly free of power.

Our major high-level study questions included: "*What is your common understanding of factors influencing your phone's battery life?*", "*Do you usually feel concerned about your phone's battery level?*", "*What is your common understanding of how CSAs actually operate?*", "*Have you suspected that your phone's battery life might become shorter after installing a CSA?*", "*Have you suspected that your phone's battery life might become more random beyond your expected variance after installing a CSA?*" Upon positive responses, we proceeded with open-ended interviews to collect detailed episodes they experienced.

Stereotypes that standby battery drains should be minimal: We found most of our participants already knew well-known causes of battery drains, e.g., display, GPS, network, and voice calls. All participants stated that their phones are supposed to drain minimal power when they are not using the phones. When they have not used the phone for a few hours but find noticeable battery drain, they think something is wrong or a suspicious task is running and sometimes attempt to hunt it down by task-killer apps.

[5] http://shealth.samsung.com/

Prevalent sensitivity to the remaining battery life: When our participants were asked "*Do you usually feel concerned about your phone's battery level?*", they responded: "*In general, not much.*" However, following interviews derived their practices to ensure their phones not to deplete the batteries in the middle of a day. Many of them had multiple chargers in their routine locations, such as the office, car, and bedside. 8 carry an additional battery. 15 mentioned they often charge their phone at a restaurant or cafe.

We also found that, deviating from their usual routine often elevates their battery concerns. P16 uses his phone for GPS navigation when driving to a new destination. In such cases, he tries to predict the remaining battery life to make sure the GPS continues to work until he arrives and also he can keep using the phone for the rest of the day. In addition, 19 participants stated they get more sensitive to the battery life when their phones' battery levels drop below a certain point; (P22): "*I feel uncomfortable when I see my phone has less than 40% battery remaining.*" In such cases, they usually make an educated guess to estimate the battery life based on their years of smartphone experiences. If a risk of depletion is expected sooner or later without a chance of recharge, they consciously refrain from making a call or surfing the Web, trying to keep their battery drains minimal.

Limited understanding about CSAs' operation and battery drains: Our participants had divided ideas on how CSAs operate. Some participants did not know that a CSA keeps running even after they close its foreground activity. The others had understanding of background processes. In terms of battery drains, they had divided estimates about the battery drains by CSAs. A few participants had no estimate at all; (P14): "*It's counting my steps so should consume power anyway. I have no idea how much.*" They were asked to make a guess how their CSAs keep counting their steps. Their response varied; (P8): "*Maybe use GPS to measure the distance traveled, and then divide it by stride?*" and (P4): "*Tapping on the sensors and find out some patterns?*" However, we could not find an implication that CSAs' battery drain may depend on user mobility conditions.

Implications of erratic battery drains with use of CSAs: Importantly, we found strong implications that CSAs might be responsible for erratic battery drains largely deviating from the user's empirical expectation. 10 participants were suspicious of their phones' shortened battery life after using CSAs, but they were not confident if and how much CSAs would be responsible. (P7): *"I have a strong feeling that [the CSA] changed something, but [Android's battery consumption ranking] always shows the display at the top.*"

Specifically, we collected multiple episodes with large day-to-day battery life variations. P21 was puzzled: "*(... after installing a CSA) I am not sure. Some days it works as fine as before but some other days I find the battery level shockingly low in the middle of the day.*" Furthermore, P3's episode was more specific: "*My phone died unexpectedly and I was suspicious of [my newly installed Dieter app]. The day after I examined [Android's battery consumption ranking], but Dieter was not even close to the major ones. (...) Some days I do find Dieter highly ranked but it is really random day by day. (...) I don't trust the ranking.*"

Interestingly, we found a few embarrassing episodes with unexpected battery depletion soon after installing CSAs. P4 recalled: "*When I go to bed, I don't bother plugging my phone as long as a half or more battery remains. My charger is at the other side of my bedroom, and I play with my phone until the last moment I fall asleep. (…) I really don't like getting out of the bed again to plug my phone. (…) After installing [CSA], I found my phone powered off in the middle of night and the alarm didn't go off in the morning! I was late to work and really upset.*" Similarly, P2 recalled: "*I was traveling by [Korean bullet train] to [a city approximately 300 km away]. (…) I tried to call my mom to find her at the station but my phone was dead! (…) I made sure my phone was fully charged before leaving my home and I really don't understand how it could happen.*"

What makes the most utility of their CSAs: Unlike the participants' common practice of suppressing phone usage or killing bogus apps in case of temporary battery shortage, we found that even a short deactivation of a CSA may largely harm its overall utility. Many participants wanted to avoid data discontinuity in their CSAs. The primary reason for using CSAs is for life-logging to pursue a healthy life; if the data log discontinues in the middle, it would render the whole day's data less credible or even useless. (P10): "*One day I found my phone was powered off for 2,3 hours. I was unhappy that it missed my steps. That prevented me from achieving my 10,000-step daily goal for 30 days in a row.*" A discontinuation degrades their satisfaction for an even longer period. (P3): "*I updated my [CSA], but for some reason I was supposed to reactivate it manually. I forgot that for 4, 5 days. I got crazy when I found the empty days in my log! It corrupted my diet plan for the entire month.*"

Key Takeaways

We summarize our key findings that lead us to design Sandra, a novel mobility-aware battery information advisor:

- Users are mostly sensitive to battery drain of their phones. They manage remaining battery carefully based on their daily usage and recharging patterns.
- Users who used a CSA experienced unexpected battery depletion caused by background operation of the CSA.
- Based on the experiments, we showed that commercial CSAs increase the standby battery drain by 38~367%.
- We also showed that CSA's battery drain varies depending on user contexts.

DESIGN AND IMPLEMENTATION

Sandra Interface Overview

Based on the insights from the explorative study, we design Sandra featuring mobility-aware battery drain information.

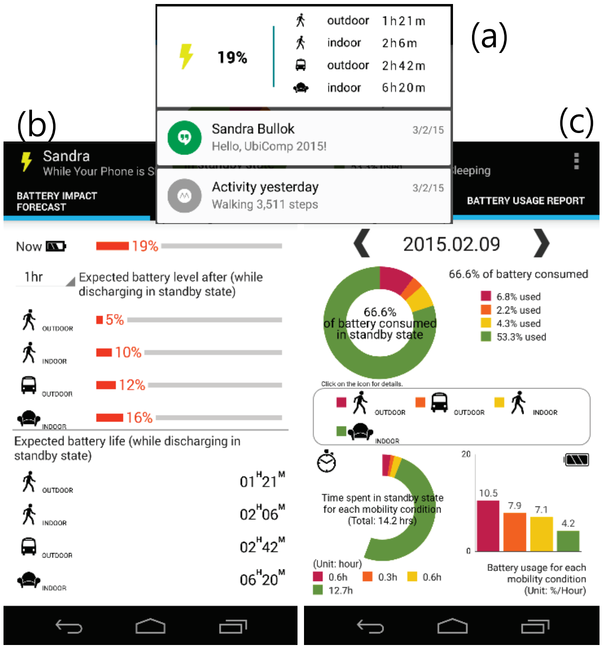

Figure 2. Screenshots of Sandra with Moves

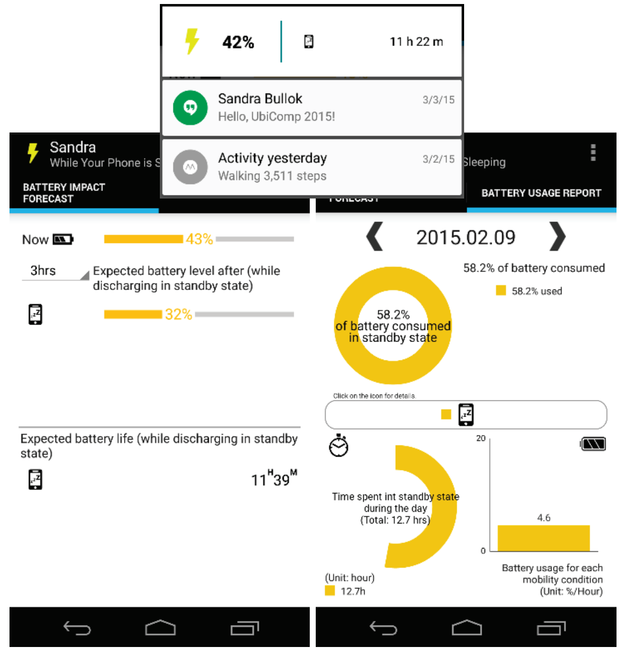

Figure 3. Screenshots of Sandra-lite with Moves

We highlight that its primary purpose is to help users build fundamental understanding that their mobility impacts their phones' battery behaviors. We anticipate that eliciting their responses and studying any actions they devise will lead us to an important scaffold for the next step of research, e.g., crafting actionable guidelines under mobility-dependent battery situations. In this light, the current design of Sandra refrains from providing users with any preset actions, and rather aims to study their unbiased responses. In the Discussion section, we discuss the potential directions toward tangible utility on top of the findings from Sandra.

Sandra presents mobility-aware battery drain information in two aspects. First, it provides expected standby times for a set of commonly occurring mobility conditions. We provide the battery information in the form of battery life (hr) since one of users' major concerns about their phones' batteries would be how long their phone will last. Also, such a form is more user friendly than a typical power metric (mW) or a relative battery usage (%). As in Figure 2(a), a Sandra interface shows expected standby time indicating that the phone will last 6.3 hours if a user is sitting at the office and 1.3 hours if he walks around outside. From this information, a user can better estimate how long the phone will last with respect to his potential activities in the next few hours, and take necessary actions such as recharging the phone or minimizing the use of apps. In Figure 2(b), Sandra displays the expected battery level after a specified amount of time. This helps users predict their future battery status when the battery level is relatively high, but a long-term activity is expected, e.g., 3-hour hiking for a trip.

Second, Sandra provides a retrospective battery use summary. Figure 2(c) shows how long a user spent under different mobility conditions, and how each condition contributed to the battery usage. For example, the battery drained 6.8% while a user walked outside on Feb. 9. This historical information is analogous to the per-app battery usage that Android provides by default, but importantly, Sandra provides mobility-aware battery usage. This interface intends to help users refer to the past days with similar distribution of mobility conditions when an unusual day is expected, e.g., a business trip or a weekend-date.

In designing Sandra, we allow users to access the information via a notification bar as well as within the app. The notification bar allows handy access to the information yet remains unobtrusive, helping users build awareness of CSAs' battery drain characteristics. Note that comparing multiple UI designs and finding the best one is not the main focus of this study. Many options are possible, e.g., on a lock screen [9], a toast message [16], or a widget, and it would be possible to extend Sandra to allow users to select one of them. Sandra shows a phone's expected standby time for the scenario that a user does not run any other apps. When a foreground app is running, we believe users can attribute the phone's battery drain to the foreground app.

Implementation

A key to provide mobility-aware battery use information is to estimate battery drain under different mobility conditions. To this end, Sandra first extracts each mobility condition and calculates battery drain rates for each one. Our implementation is based on the Android phones.

Mobility condition monitoring: Sandra leverages context information generated by CSAs, not having its own context inference logics. Recent CSAs such as Google Fit and Moves provide open APIs to let third-party apps utilize the contexts that the CSA monitors. Sandra taps onto this API to identify contexts. This approach has two key benefits: (1) it hardly adds a power overhead to provide mobility-aware battery information and (2) it naturally creates a suitable category of mobility conditions matching the running CSAs.

While current Sandra only supports CSAs which provide open APIs, Sandra can be extended for more commercial CSAs. Besides Google Fit and Moves, many fitness-related apps already provide open APIs to promote their ecosystem, including Apple Health[6] and Withings[7]. We expect that more CSAs will follow this trend in the future. Also, there is a growing effort to employ context recognition modules into mobile platforms, e.g., step detector in Android 4.4. This will allow Sandra to obtain context information requested by CSAs without additional cost.

Calculation of battery drains: Sandra calculates drain rates (%/hour) for detected mobility conditions. The rate is computed by using consecutive samples of battery levels provided by Android. While accurate power monitoring has been actively studied in research domains [36, 43, 44], we take a simple method of dividing the decreased battery levels by the duration [9, 16, 26] to avoid additional cost. Sandra logs the battery level upon its change, only when the phone is in a standby state; it skips when the phone runs a foreground app, it is being charged, or its screen is on.

One might challenge that smartphone's battery levels may not well represent the actual remaining energy due to nonlinear discharge of lithium-ion batteries [10] and errors in battery gauge circuits, resulting in slower battery drops at lower percentage values. For example, at a battery level shown as 10%, the true remaining usable time might be far longer than one tenth of the full battery life. To address this, we monitored the Android's native battery level values under consistent workloads run over time until the battery depletes. Figure 4 shows that the battery level decreases linearly to the elapsed time. This reasonably resolves the concern about possible errors at lower battery levels.

Sandra aggregates the battery logs for each mobility condition and calculates the drain rate. For the retrospective summary, Sandra uses the samples collected in a given day. For the forecast of impact, we take the samples collected in the most recent seven days to reflect the recent trend.

Sandra overhead: One might question that Sandra may incur additional battery overheads due to its continuous operation. However, our preliminary evaluation shows that Sandra incurs only a marginal cost. Using Monsoon power monitors, we measured the power consumption in four Nexus 5 devices at the same time under identical mobility conditions but running different CSAs: (1) Google Fit, (2) Google Fit and Sandra, (3) Moves, and (4) Moves and Sandra. For one hour, they consumed 68 mW, 71 mW, 75 mW, and 82 mW on average, respectively. It means that the

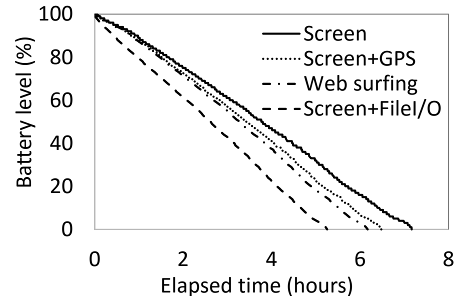

Figure 4. Battery levels under consistent workloads

net additional overheads by Sandra were only 3 mW and 7 mW in (2) and (4). The power consumed for Moves is higher because it retrieves context information via network upon a request. We believe that this is marginal because the average power of a smartphone is 659 mW assuming 15 hours of battery life with the fully charged battery, 2300 mAh for Nexus 5. Such insignificant overheads can be explained in that (a) the mobility condition retrieved by Sandra is only the by-product from readily running CSAs; (b) profiling per-mobility battery drain rates is done with *sparsely* sampled current mobility condition and drain rates, e.g., once an hour; and (c) the expected standby time is updated on the user's request, and its computation is very simple—dividing the current battery level by the pre-calculated drain rates for the current mobility condition.

EVALUATION

To study the user-perceived implications of Sandra, we conducted a 30-day deployment study with 24 users. The user pool was the same as those in the explorative study. When originally recruited, they opted in to our second invitation to use Sandra in their real life for up to 1 month. When re-invited for study with Sandra, they were given the option to actually participate in the study for a 30-day period and additional compensation equivalent to USD 54. Since Sandra is compatible with the APIs of Google Fit and Moves, we encouraged the participants who used S-health and Dieter to try either Google Fit or Moves for this study. Finally, we had 18 Google Fit users and 6 Moves users.

For the experiment, the participants installed Sandra on the phone they use. Their phones included Galaxy S3 (5), LG G2 (4), Galaxy S4 (3), Galaxy S5 (2), Nexus 5 (2), Galaxy Note 2 (2), Vega Iron 2 (1), Galaxy Grand (1), LG G3 (1), Vega LTE A (1), Galaxy Note 3 (1), and Galaxy S2 (1); the parenthesized number represents the number of phones. For the first 10 days, we set Sandra to work covertly without any information shown to the user; it periodically collected tuples of (*battery drain rate*, *mobility condition*) to evaluate each user's power coefficients under each mobility. For quantitative analysis on users' behavior changes, Sandra also silently collected a log of battery events (charging, discharging, level changes), smartphone usage, and Sandra usage. This collection continued for the whole 30 days.

For the remaining 20-day period, Sandra provided the user with battery life estimations and past history. Importantly,

[6] https://developer.apple.com/healthkit/
[7] http://oauth.withings.com/api

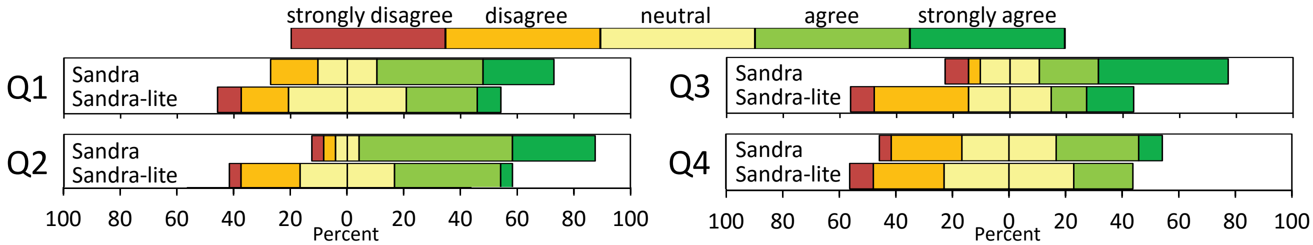

Figure 5. Participants' response distributions

Q1: "Did it bring changes to your existing understanding about your phone's stand-by battery drain?"
Q2: "Do you think the provided information is useful?"
Q3: "Did you find it helpful in managing your phone's battery?" *Q4: "Did you find it helpful in alleviating your battery concern?*

we tried to avoid an unwanted novelty effect from Sandra's look and feel. To rule this out, we created Sandra-lite, a purposely downgraded Sandra which has the same look and feel as Sandra but provides only a single standby life expectancy without *per-mobility* breakdown (see Figure 3). We had the participants use Sandra and Sandra-lite for 10 days each. Half of the group used Sandra first and Sandra-lite later; the other did in the opposite order. We set Google Fit and Moves users equally distributed to each subgroup.

Upon completion, each participant answered two identical written questionnaires for her experiences with Sandra and Sandra-lite, in 5-point Likert-scales (Strongly Disagree, Disagree, Neutral, Agree, Strongly Agree). Then, 1-hour semi-structured one-on-one interviews followed. All interviews were recorded, transcribed, and iteratively coded by two researchers [41]. Figure 5 illustrates the distribution of responses for select questions. Each distribution is centered at "Neutral (yellow)." The width of each color segment represents the frequency of each response.

Updating Users' Battery Models in Their Mind

Figure 5 (Q1) shows the responses about if Sandra or Sandra-lite changed their understanding about the phones' standby battery drain. The difference between Sandra and Sandra-lite (p-value = 0.023) is indicative that such change would be attributable to the per-mobility breakdown of battery usage. Subsequent interviews revealed that, for most of them, it was their first time realizing that their phones consume different power depending on their mobility, even without using the phones. P3 said, "*I have never thought my move affects my phone's battery use. When I was told so the first time I thought of it as nonsense. (...) One day I purposely walked to work to see if [Sandra] was saying the truth. I was shocked to see the battery drop was pretty close to what was predicted in the walking scenario.*" From those who positively responded with Sandra-lite, we found that it was because they had not experienced a battery information interface representing the remaining battery in unit of time, despite third-party apps available providing the feature.

We also found that Sandra helped them develop their own thoughts about the power efficiency of CSAs. 9 participants mentioned that seeing the per-mobility battery differences led them to think CSAs are developed in a power-efficient way. P4 said, "*Google Fit seems energy efficient. When I don't move, it is so wise doing a lot less computation and saving battery.*" In contrast, P15 said, "*(The battery usage report shows) way higher battery rate for walking; I didn't expect my Google Fit spends so much power.*"

Usefulness of Per-Mobility Battery Drain Information

We examined if and how much the participants perceived it useful to be aware of the per-mobility battery impact of CSAs. With per-mobility battery information, we observed a significant increase in the positive responses for overall user-perceived utility (p-value=0.005) (see Figure 5 (Q2); agree: 13, strongly agree: 7). Interestingly, they found Sandra useful in certain situations, e.g., when the battery level was low and recharging was not available soon. Below, we discuss notable cases where Sandra was useful.

In-situ arrangement: 12 participants reported episodes when Sandra helped them realize insufficient battery life for their imminent mobility conditions, and thereby they improvised necessary arrangements in-situ. P3 stated, "*I was traveling to meet my friend in our hometown several hours away. Soon after I got on board the train, [Sandra] let me know my phone would not make it until the arrival. So I told my friend early enough about it, and discussed precisely when and where to meet.*" P2 said, "*When I was taking the train to my parent's home, [Sandra] told me that my phone will last only for a few hours on the 'transport' condition. I tried to use my phone little to keep it alive.*"

Two disagreed that per-mobility battery information was informative to help their battery planning (Q3, strongly disagree: 1, disagree: 1). They complained about the lack of guidance upon imminent battery depletion; (P13): "*I want to know how longer my battery would last if I kill Google Fit.*" For Sandra-lite, 10 participants did not find it useful (Q3, strongly disagree: 2, disagree: 8). A common reason was doubts about its accuracy. (P8): "*[Sandra-lite] expects different standby life day by day. It's not consistent.*"

Acquiring new everyday practices: 9 participants mentioned that Sandra helped them develop new battery-related knowhow or strategies. (P1): "*I learned that the battery drains faster when I'm moving. So I consciously suspend Google Fit when I drive far. It is not a fitness for me anyway.*" (P16): "*I work at my office on weekdays but mostly hang out outside on weekends. Now I know battery* *runs out higher when moving, I try to make sure full battery before leaving.*" Some reported unpleasant experiences with Sandra-lite; (P22): "*I went to skiing with [Sandra-lite]. It showed more than 10 hours of battery life, but after 3-4 hours of skiing, less than an hour left… so embarrassing!*"

Feeling less nervous under limited battery: Before the study, we expected that Sandra would effectively relieve their concern when their phones' battery was low. However, we observed neither dominant positive responses with Sandra (from Q4, Sandra: mean score of 3.1; nearly neutral) nor significant difference between Sandra and Sandra-lite (from Q4, Sandra-lite: mean score: 2.8, p-value = 0.236).

Still, 40% of the participants responded positively about Sandra's concern-relief effects. The most common case was a moment just before they fall asleep. (P4): "*My charger is far from my bedside, and I like playing with my phone on bed. About to fall asleep, it's really a big hassle to get out of the bed to plug my phone. I'm so tempted not to bother and just go sleep. (…) It is always a tricky bet when I see the battery halfway left. (…) [Sandra] answers exactly whether I have to do it or not.*" We observed an interesting case from P20 in her late pregnancy, *"Keeping my phone alive is extremely important, to make sure I can call [the husband] in emergency. I regularly take a walk for prenatal exercise, and I get so nervous about my battery. (…) With [Sandra], I felt a lot comfortable to know how long it will last, and even I could safely decide where to return.*" (Her walk path is a round trip with an arbitrary point of return.)

In contrast, about 30% of the Sandra users disagreed that the battery drain information alleviated their concern. Similar to Q3, once they get nervous with very low battery, they cannot extend its life. P11 mentioned, "*No matter I have [Sandra or Sandra-lite], it won't change anything.*"

Quantitative Behavioral Observation of Battery Usage

To observe the participants' behavioral change, we quantitatively analyzed the collected battery-related data such as charging frequency and smartphone use time. Statistically significant changes were not consistently observed between Sandra and Sandra-lite. From the interview, we speculated as to the reasons. First, they do not always face a low-battery situation in everyday life. (P16): "*I rarely ran low on battery during this study. Normally, I charge my phone as soon as I arrive at work or get back home.*" An average participant had battery issues a few times during the study. Despite infrequency, even a single low battery is a major inconvenience and Sandra was useful as shown in the previous subsection. Second, upon facing a battery problem, they had different strategies, making it hard to observe common behaviors; e.g., charging/replacing the battery, killing apps, or turning off Wi-Fi.

We also analyzed smartphone usage and Sandra (or Sandra-lite) usage logs for each participant. We gave them some pointers out of the logs that helped them reflect on specific use cases of Sandra. We report two notable cases.

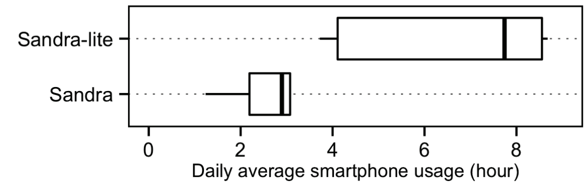

Figure 6. P2's daily average smartphone usage

Regarding changes in battery concern, P2 reported, "*When I was using [Sandra-lite], I was not too concerned about the remaining battery life. But [Sandra] showed notably lower expected standby time for conditions like walking. This got me worried about its life. I became more cautious about using my phone.*" In fact, analysis of the smartphone usage log supports her experience. As shown in Figure 6, her daily average duration of smartphone usage was 3.06 hours (std: 1.82) when using Sandra, which is significantly lower than 6.50 hours (std: 2.38) when using Sandra-lite. (t-test; p-value = 0.005) Another case is that, one day, P1 launched Sandra 8 times, significantly more frequent than the other days (mean: 2.8, std: 3.27). P1 recalled, "*There was a long distance business trip waiting for me. Since I had to drive a long way and it might get difficult to charge my phone once I get there, I frequently checked [Sandra] and regulated phone usage to make sure my phone alive during the trip.*"

Suggestions

The participants suggested features to enhance Sandra. We introduce a few interesting ones. First, some asked us to extend Sandra to consider their future mobility patterns, not a constant mobility condition. P16 stated that his mobility patterns are similar on weekdays and he wanted to see the expected battery life that considers the rest of the day. We discuss its feasibility in the following section. Second, two reported that it was difficult to see the per-mobility battery information at a glance and asked to display battery life for their current mobility condition only on a notification bar.

DISCUSSION

Extensibility to Support Other Contexts

Currently, Sandra targets a limited group of CSAs whose main context is mobility condition; it was our design priority to deal with widely and commercially available CSA types. Still, the basic concept of Sandra can be applied to other CSAs using different types of contexts. To show such extensibility, we developed two research apps inspired by a conversation monitor [24] and a sound event detector [28], and measured their power consumption on Nexus 5 using a Monsoon power monitor. The former detects nearby people by periodic Bluetooth scans (110 mW) and conducts a conversation monitoring logic (233 mW) if detected. The latter detects non-silence frames (115 mW) and runs a classification logic for sound-related events (192 mW) only when detected. These CSAs' power behavior highly depends on contextual factors of encounter and ambient sound. Sandra can be easily extended to support these CSAs by providing battery information from such other contexts.

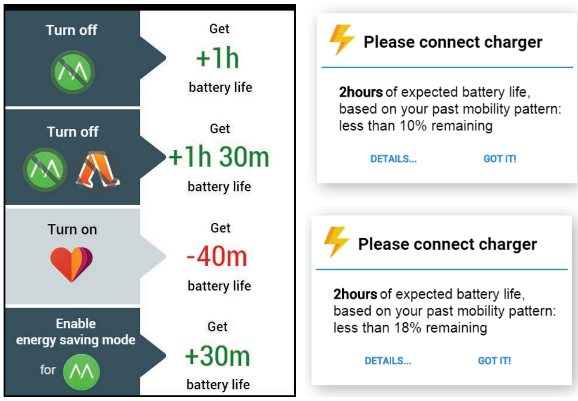

Figure 7. Mock-up UIs of Sandra Extension; (a) user guidelines (left) and (b) recharging alerts (right)

Study Limitation of Target Participants

Our participants were mostly limited to university students in their 20s-30s. Still, an interesting implication is that, although they are often regarded as "techy" users, they were unaware of the per-mobility battery impact of CSAs. In this light, those who are less familiar with smart devices may have even less knowledge and more troubles in managing battery. We expect Sandra would still be helpful for them.

Other Factors that Affect Nonlinear Battery Drain

We do not argue that CSAs are the only factor causing the nonlinear battery drain of smartphones. There are other known factors affecting phones' battery lives, e.g., app usage, screen brightness, and cellular strength. Throughout the years using smartphones, many users have developed an empirical understanding of the battery impacts from those factors to some extent. For example, users naturally expect that playing Angry Birds drains their phones' batteries faster. However, it is yet little known to users that their mobility condition affects their phones' battery lives.

Utility of Sandra

One may think that Sandra is only useful for a transient time until users get used to mobility-dependent battery drain. We believe users can benefit even for long-term use. First, we expect significant user convenience from a feature quantifying the battery impact from the mobility conditions automatically, rather than letting the user make a rough guess solely based on experiences. Second, more notably, CSAs are dynamic; users will install new CSAs or uninstall others. Different combinations of CSAs change contextual factors and their extent affecting the battery drains. For example, if a pedometer user installs an encounter-based conversation monitor, the battery impact of the mobility condition will change depending on her encounters. Sandra will be valuable to handle such dynamic situations.

Extending Sandra towards Richer User Utility

User guidelines for energy saving: A conventional way of providing user guidelines for energy saving is to notify about the expected increase of a phone's battery life if the user kills a specific, running application [9, 34, 42]. Martins et al. proposed diverse application modes of different utility and power costs, e.g., HD and SD streaming for video watching, and allowing users to select one considering the battery [31]. Similarly, Sandra can be extended for CSAs considering their mobility-dependent power consumption. For example, based on a user's current mobility condition, Sandra can notify about the expected increase of the battery life if she suspends a running CSA. Figure 7(a) illustrates a mock-up interface of such an extension.

Future mobility pattern-based battery advisor: Sandra assumes a constant user mobility condition when providing the expected standby time. If upcoming mobility patterns are predictable, Sandra can provide more personalized battery impact of CSAs. Suppose a user who usually stops by the gym after work, does incline walking for an hour, and drives home for an hour. When she is still at the office, Sandra can reasonably predict the battery level at the time she arrives home; if necessary, Sandra can notify her that she needs to recharge her phone before leaving the office. It is well known that human mobility has some predictable patterns [11]. In this light, we believe that it would be feasible to reasonably estimate the battery impact of CSAs for the rest of a user's day, especially for regular weekdays.

Human context-dependent recharging alert: Many phones today pop up a low-battery alert upon the battery level dropping below a predefined value, e.g., 15%. With CSAs involved, this 15% remaining battery would mean quite different battery lives depending on a user's mobility conditions. Sandra may help improve this alert, so that a battery alert is triggered when the likely remaining battery lifetime is below a predefined time as in Figure 7(b); the underlying percentage value may vary depending on the current or predicted user mobility conditions.

CONCLUSION

Emerging CSAs introduce new major factors governing phones' overall battery consumption behaviors: (1) added nontrivial persistent battery drain and (2) different battery drain rates depending on the different mobility contexts. To address these factors outdating users' existing battery model, we explored an initial approach to help users understand the new causality. We proposed *Sandra*, a novel battery advisor highlighting the impacts of users' different mobility conditions. We conducted a real deployment study for 30 days with 24 users. Our findings report what they essentially learned, in which situations they found Sandra particularly helpful, and the lessons learned to help in the design of future mobility-aware battery advisors.

ACKNOWLEDGMENTS

We thank the anonymous reviewers and shepherd for their valuable comments to improve the quality of this paper. This work was supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIP) (No. 2011-0018120).

REFERENCES

1. Athukorala, K., Lagerspetz, E., Kugelen, M.v., Jylha, A., Oliner, A.J., Tarkoma, S., and Jacucci G. How carat affects user behavior: implications for mobile battery awareness application. In *Proc. CHI* 2014, ACM Press(2014), 1029-1038.
2. Azizyan, M., Constandache, I. and Choudhury, R.R. SurroundSense: Mobile Phone Localization via Ambience Fingerprinting. In *Proc. MobiCom 2009*, ACM Press (2009), 261-272.
3. Banerjee, N., Rahmati, A., Corner, M. D., Rollins, S., and Zhong, L. Users and batteries: Interactions and adaptive energy management in mobile systems. In *Proc. UbiComp 2007*, Springer-Verlag, 4717 (2007), 217-234.
4. Chu, D., Lane, N., D., Lai, T., T., Pang, C., Meng, X., Guo, Q., Li, F., and Zhao, F. Balancing energy, latency and accuracy for mobile sensor data classification. In *Proc*. *SenSys* 2011, ACM Press (2011), 54-67.
5. Consolvo, S., McDonald, D. W., Toscos, T., Chen, M. Y., Froehlich, J., Harrison, B., Klasnja, P., LaMarca, A., LeGrand, L., Libby, R., Smith, I. and Landay, J. A. Activity sensing in the wild: A field trial of UbiFit Garden. In *Proc. CHI 2008*, ACM Press (2008), 1797-1806.
6. Falaki, H., Mahajan, R., Kandula, S., Lymberopoulos, D., Govindan, R. and Estrin, D. Diversity in smartphone usage. In *Proc. MobiSys 2010*, ACM Press (2010), 179-194.
7. Feldman, D., Sugaya, A., Sung, C. and Rus, D. iDiary: From gps signals to a text-searchable diary. In *Proc. SenSys 2013*, ACM Press (2013), 6.
8. Ferreira, D., Dey, A.K. and Kostakos, V. Understanding Human-Smartphone Concerns: A Study of Battery Life. *Pervasive 2011, 6696*, Springer-Verlag (2011), 19-33.
9. Ferreira, D., Ferreira, E., Goncalves, J., Kostakos, V., and Dey, A. K. Revisiting human-battery interaction with an interactive battery interface. In *Proc. UbiComp 2013*, ACM Press (2013), 563-572.
10. Gao, L., Liu, S., and Dougal, R., A. Dynamic Lithium-Ion Battery Model for System Simulation. IEEE Transaction on Energy Conversion, Vol 17, No. 2, pp. 285-294, 2002.
11. Gonzalez, M. C., Hidalgo, C. A., and Barabasi, A. L. Understanding individual human mobility patterns. *Nature, 453(7196)*, 779-782
12. Hao, T., Xing, G. and Zhou, G. iSleep: unobtrusive sleep quality monitoring using smartphones. In *Proc. SenSys 2013*, ACM Press (2013), 4.
13. Hemminki, S., Nurmi, P. and Tarkoma, S. Accelerometer-based transportation mode detection on smartphones. In *Proc. SenSys 2013*, ACM Press (2013), 13.
14. Hwang, I., Yoo, C., Hwang, C., Yim, D., Lee, Y., Min, C., Kim, J., and Song, J. TalkBetter: family-driven mobile intervention care for children with language delay. In *Proc*, *CSCW* 2014, ACM Press (2014), 1283-1296.
15. Ju, Y., Lee, Y., Yu, J., Min, C., Shin, I., and Song, J. SymPhoney: a coordinated sensing flow execution engine for concurrent mobile sensing applications. In *Proc*. *SenSys* 2012, ACM Press (2012), 211-224.
16. Jung, W., Chon, Y., Kim, D. and Cha, H. Powerlet: An Active Battery Interface for Smartphones. In *Proc. UbiComp 2014*, ACM Press (2014), 45-56.
17. Kang, S., Kwon, S., Seo, S., Yoo, C., Park, K., Song, J., and Lee, Y. Sinabro: Opportunistic and Unobtrusive Mobile ECG Monitoring System. In *Proc*. *HotMobile* 2014, ACM Press (2014), 11.
18. Kang, S., Lee J., Jang, H., Lee H., Lee Y., Park, S., Park, T., and Song, J. SeeMon: scalable and energy-efficient context monitoring framework for sensor-rich mobile environments. In *Proc*. *MobiSys* 2008, ACM Press (2008), 267-280.
19. Kang, S., Lee, Y., Min, C., Ju, Y., Park, T., Lee, J., Rhee, Y., and Song, J. Orchestrator: An active resource orchestration framework for mobile context monitoring in sensor-rich mobile environments. In *Proc*. *PerCom* 2010, IEEE Computer Society Press (2010), 135-144.
20. Kansal, A., Saponas, S., Brush, A., J., B., McKinley, K., S., Mytkowicz, T., and Ziola, R. The latency, accuracy, and battery (LAB) abstraction: programmer productivity and energy efficiency for continuous mobile context sensing. In *Proc*. *OOPSLA* 2013, ACM Press (2013), 661-676.
21. Kim, D. H., Kim, Y., Estrin, D., and Srivastava, M. B. Sensloc: sensing everyday places and paths using less energy. In *Proc. SenSys 2010*, ACM Press (2010), 43-56.
22. Lee Y., Ju Y., Min, C., Kang, S., Hwang, I., and Song, J. CoMon: cooperative ambience monitoring platform with continuity and benefit awareness. In *Proc*. *MobiSys* 2012, ACM Press (2012), 43-56.
23. Lee, S., Min, C., Yoo, C., and Song, J. Understanding customer malling behavior in an urban shopping mall using smartphones. In *Proc*. *MCSS* 2013 (UbiComp Adjunct), ACM Press (2013), 901-910.
24. Lee, Y., Min, C., Hwang, C., Lee, J., Hwang, I., Ju, Y., Yoo, C., Moon, M., Lee, U. and Song, J. SocioPhone: Everyday Face-to-Face Interaction Monitoring Platform Using Multi-Phone Sensor Fusion. In Proc. *MobiSys 2013*, ACM Press (2013), 375-388.
25. Lester, J., Hartung, C., Pina, L., Libby, R., Borriello, G. and Duncan, G. Validated caloric expenditure estimation using a single body-worn sensor. In *Proc. UbiComp 2009*, ACM Press (2009), 225-234.

26. Liu, Y., Xu, C., and Cheung, S. C. Where has my battery gone? Finding sensor related energy black holes in smartphone applications. In *Proc. PerCom 2013*, IEEE (2013), 2-10.
27. Lu, H, A. Brush, A.J., Priyantha, B., Karlson, A.K. and Liu, J. Speakersense: Energy efficient unobtrusive speaker identification on mobile phones. *Pervasive 2011, 6696*, Springer-Verlag (2011), 188–205.
28. Lu, H., Pan, W., Lane, N, Choudhury, T. and Campbell, A. SoundSense: scalable sound sensing for people-centric applications on mobile phones. In Proc. *MobiSys 2009*, ACM Press (2009), 165–178.
29. Luo, C. and Chan, M.C. Socialweaver: Collaborative inference of human conversation networks using smartphones. In *Proc. SenSys 2013*, ACM Press (2013), 20.
30. Ma, X., Huang ,P., Jin, X., Wang, P., Park, S., Shen, D., Zhou, Y., Saul, L.K., Voelker, G.M. eDoctor: Automatically Diagnosing Abnormal Battery Drain Issues on Smartphones, In *Proc. NSDI 2013*, USENIX (2013), 57-70.
31. Martins, M. and Fonseca R. Application modes: a narrow interface for end-user power management in mobile devices. In *Proc*. *HotMobile* 2013, ACM Press (2013), 5.
32. Metri, G., Shi, W., Brockmeyer, M., and Agrawal, A. BatteryExtender: an adaptive user-guided tool for power management of mobile devices. In *Proc. UbiComp 2014*, ACM Press (2014), 33-43.
33. Nath, S., ACE: exploiting correlation for energy-efficient and continuous context sensing. In *Proc. MobiSys* 2012, ACM Press (2012), 29-42.
34. Oliner, A. J., Iyer, A. P., Stoica, I., Lagerspetz, E., and Tarkoma, S. Carat: Collaborative energy diagnosis for mobile devices. In *Proc. SenSys 2013*, ACM Press (2013), 10.
35. Paek, J., Kim, J. and Govindan, R. Energy-Efficient Rate-Adaptive GPS-based Positioning for Smartphones. In *Proc. MobiSys 2010*, ACM Press (2010), 299-314.
36. Pathak, A., Hu, Y., C. and Zhang, M. Where is the energy spent inside my app?: fine grained energy accounting on smartphones with Eprof. In *Proc. EuroSys 2012*, ACM Press (2012), 29-42.
37. Rachuri, K.K., Musolesi, M., Mascolo, C., Rentfrow, P.J., Longworth, C. and Aucinas, A. EmotionSense: A Mobile Phones based Adaptive Platform for Experimental Social Psychology Research. In *Proc. UbiComp 2010*, ACM Press (2010), 281-290.
38. Rahmati, A., and Zhong, L. Human–battery interaction on mobile phones. *Pervasive and Mobile Computing, 5(5)*, (2009), 465-477.
39. Rahmati, A., Qian, A. and Zhong, L. Understanding human-battery interaction on mobile phones. In *Proc. MobieHCI 2007*, ACM Press (2007), 265-272.
40. Ravi, N., Scott, J., Han, L., and Iftode, L. Context-aware battery management for mobile phones. In *Proc. PerCom 2008*, IEEE Computer Society Press (2008), 224–233.
41. Strauss, A. and Corbin, J. *Basics of Qualitative Research Techniques and Procedures for Developing Grounded Theory*. Sage Publications, (2007).
42. Truong, K.N., Kientz, J.A., Sohn, T., Rosenzweig, A., Fonville, A. and Smith, T. The Design and Evaluation of a Task-Centered Battery Interface. In *Proc. UbiComp 2010*, ACM Press (2010), 341-350.
43. Yoon, C., Kim, D., Jung, W., Kang, C. and Cha, H. AppScope: application energy metering framework for android smartphones using kernel activity monitoring. In *Proc. USENIX ATC 2012*, USENIX (2012), 36-36.
44. Zhang, L., Tiwana, B., Qian, Z., Wang, Z., Dick, R.P., Mao, Z.M. and Yang, L. Accurate online power estimation and automatic battery behavior based power model generation for smartphones. In *Proc. CODES/ISSS 2010*, ACM Press (2010), 105-114.
45. Zhou, P., Zheng, Y., Li, Z., Li, M. and Shen, G. IODetector: A generic service for indoor outdoor detection. In *Proc. SenSys 2012*, ACM Press (2012), 113-126.