

Battery-Aware Mobile Data Service

Liang He, *Member, IEEE*, Guozhu Meng, *Student Member, IEEE*, Yu Gu, *Member, IEEE*,
Cong Liu, *Member, IEEE*, Jun Sun, *Member, IEEE*, Ting Zhu, *Member, IEEE*,
Yang Liu, *Member, IEEE*, Kang G. Shin, *Life Fellow, IEEE*

Abstract—Significant research has been devoted to reduce the energy consumption of mobile devices, but how to increase their energy supply has received far less attention. Moreover, reducing the energy consumption alone does not always extend the device operation time due to a unique battery property — the capacity it delivers hinges critically upon how it is discharged. In this paper, we propose B-MODS, a novel design of battery-aware mobile data service on mobile devices. B-MODS constructs battery-friendly discharge patterns utilizing the *recovery effect* so as to increase the capacity delivered from batteries while meeting data service requirements. We implement B-MODS as an application layer library on the Android platform. Our experiments with diverse mobile devices under various application scenarios have shown that B-MODS increases the capacity delivery from the battery by up to 49.5%, with which an increase in the user-perceived data service utilities of up to 28.6% is observed.

Keywords—Mobile devices, battery management, data service, recovery effect

1 INTRODUCTION

Mobile devices, such as smartphones and tablets, are increasingly used for our daily lives and businesses. However, their limited operation time has been a major impediment to their usability [1]. As hardware components (e.g., displays or processors) that demand more power and more components (e.g., radios and sensors) are integrated in mobile devices, the rapidly increasing power demand has become an acute problem. Unfortunately, the battery technology — a common power source for mobile devices — has been advancing much slower than the increase of power demand [2]. To close this gap, efforts have been made to reduce the energy consumption of mobile devices at different layers of abstraction, including hardware [3], [4], operating system [5], [6], applications [7], [8], and human interactions [9], [10].

However, less energy consumption of mobile devices may not effectively extend their operation time without carefully considering the electrochemical battery characteristics [11], [12]. This is because the capacity the batteries deliver depends on how they are discharged [13]–[16], and the difference could be as much as $3\times$ [17].

The recent, wide adoption of smartphones has shifted their major functionalities from making/receiving real-time calls to diverse data services [18], such as cloud synchronization and video streaming, all with relatively soft timing requirements [7], [8]. This allows us to explore different battery discharge patterns to enhance their capacity delivery and thus to improve the data

service, especially in view of the fact that the communication modules are the dominating energy consumers besides the screens [8], [19]–[21], e.g., contributing to 18.5% and 22.3% of the overall power consumption of mobile devices, respectively [1].

In this paper, we propose a novel *Battery-aware MOBILE Data Service* (B-MODS) on mobile devices that improves the capacity delivery from batteries, thus extending the device operation time and more importantly, improving user-perceived data services such as downloadable file and streaming duration. B-MODS is inspired by two observations on batteries. First, the power-off voltage of the device — i.e., a dynamic set of voltages at which the mobile device powers itself off — is normally much higher than the cut-off voltage¹ of its battery. This voltage mismatch causes the battery to be deficiently discharged and the device to power off without utilizing the full battery capacity. Second, the *battery recovery effect* [17], [22]–[24] can be exploited to mitigate the voltage mismatch by drawing more capacity from the battery within the device operating voltage range. The recovery effect describes the phenomenon that the battery voltage can recover to a certain degree if the discharge current of the battery is reduced or interrupted, and is known to enhance the total capacity delivery from batteries. Beyond this conventional wisdom, we have made a new observation that the recovery effect can be exploited to shift the capacity delivery from (especially aged) batteries towards high voltage ranges within which the device operates, and it also reduces the energy loss on the resistance of batteries.

There are two key challenges in exploiting the recovery effect. First, the recovery effect depends heavily on dynamically changing battery conditions (e.g., SoC, resistance, and temperature) [25], making its modeling

• L. He and K. Shin are with The University of Michigan, Ann Arbor, MI, USA; G. Meng and Y. Liu are with Nanyang Technological University, Singapore; Y. Gu is with IBM Watson Health, Austin, TX, USA; C. Liu is with University of Texas at Dallas, Dallas, TX, USA; J. Sun is with Singapore University of Technology and Design, Singapore; T. Zhu is with University of Maryland at Baltimore County, Baltimore, MD, USA.

1. The voltage commonly defines the state of empty batteries.




Fig. 1: Energy vs power density [35].

either inaccurate or complex [26], [27]. Thus, it is non-trivial to identify the proper discharge patterns so as to exploit the recovery effect. Second, the utilization of recovery effect has to be transparent to mobile users so as not to degrade user experiences, because data services could be affected and/or interrupted under different battery discharge patterns. This becomes more challenging as we observe that the battery voltage requires tens of seconds or even minutes to fully recover, although a large part of it can be recovered in milliseconds as reported in the literature. B-MODS addresses these challenges by identifying the desired discharge pattern to exploit the recovery effect online, and scheduling the data-service tasks with the joint consideration of battery and workload requirements.

This paper makes the following contributions.


- We identify the fundamental discrepancy between energy supply and consumption, and a potential mitigation thereof (Sec. 3).
- We report new findings on the recovery effect via extensive measurements (Sec. 4).
- We propose B-MODS and implement it as an application library on the Android platform (Sec. 5).
- We evaluate B-MODS on diverse mobile devices under various application scenarios (Sec. 6).

The paper is organized as follows. Sec. 2 reviews the related literature. The motivation of this paper is discussed in Sec. 3. The basic idea of B-MODS is validated in Sec. 4 and its design is introduced in Sec. 5. B-MODS is evaluated in Sec. 6, and Sec. 7 sheds more insights. The paper concludes in Sec. 8.

2 RELATED WORK

Observing the limited operation time of mobile devices, extensive research has been done to identify their critical energy-consumption sources. The energy consumption of different hardware modules has been investigated [19], [28], [29], and those for different applications and services (and even bugs) of the devices have also been explored [20], [30]–[32]. Based on these identified critical energy-consumption sources, approaches have been proposed to reduce the device energy consumption at different layers, including hardware [3], [4], operating system [5], [6], [22], application layer [7], [8], and human interactions [9], [10], [33], [34].

Lithium-ion batteries are now the fastest growing and most promising battery chemistry thanks to their high



(a) Voltage traces

(b) Statistics

Fig. 2: Continuous vs. intermittent downloading: less consumption does not always lead to longer operation.

energy density. Fig. 1 compares various types of batteries in their energy and power densities, showing significant advantages of Lithium-ion batteries [35]. The unique properties of Lithium-ion batteries, however, have been studied far less in the energy management of mobile devices. For example, the battery voltage can increase to a certain degree if the discharge current of the battery is interrupted or reduced, commonly known as the recovery effect [17], [22]–[24]. These unique battery properties make their capacity delivery highly reliant on *how* they are discharged [13], [14]. An analysis of the discrepancy between the battery energy provision and the smartphone energy consumption is given in [11].

In this paper, we have explored the possibility of utilizing the battery recovery effect to extend the operation time of mobile devices. Although the recovery effect is known to help deliver more capacity during the discharge process, to the best of our knowledge, this is the first attempt to utilize its effect of shifting the distribution of the delivered capacities, significantly extending the device operation time. A few existing designs on mobile device energy management tackled the battery recovery effect. For example, the pulsed battery discharge for communication devices was proposed in [17], [23] based on a Markov model. However, besides the insufficient empirical model validation, there is no known joint consideration for the battery and application requirements.

B-MODS is orthogonal to existing battery-agnostic data service designs, with which can be integrated for further improvements. For example, B-MODS benefits from the communication framework *Bartendr* [8] by using the signal strength, besides the data rate, as another metric in the data service triggered state transitions. For 3G/LTE-based data services that show clear *tail effect*, B-MODS can be integrated with existing solutions on shortening the tail (e.g., *TailEndr* [29]), to jointly optimize energy efficiency and battery performance. The battery triggered transitions of B-MODS can be incorporated into the network interface assignment of *ATOM* to map user traffic across WiFi (i.e., low-discharge state) and LTE (i.e., high-discharge state) [36], and can also be incorporated into the RRC states in cellular networks [37].

3 BACKGROUND

Here we empirically show the necessity of battery-aware mobile data service, the theme of this paper.




Fig. 3: Power architecture in a mobile device.



Fig. 4: Regulator on Nexus S.



Fig. 5: NEWARE Battery Tester.

3.1 Motivation

The device operation time is a key consideration for mobile users. We now show that *less energy consumption of devices may not always lead to a longer operation time* — a counter-intuitive and surprising finding — due to the unique electrochemical characteristics of the batteries powering the devices.

To illustrate this finding, we use a fully charged Acro S phone to repeatedly download a file via our university WiFi until it powers off. Next, we stop the downloading process for 5s after every 10s of downloading, i.e., downloading intermittently.² The battery voltage traces during these two downloading processes are plotted in Fig. 2(a) and their statistics are summarized in Fig. 2(b). As expected, the intermittent downloading lasts longer than continuous downloading because of inserted rest durations. What is surprising, however, is that more files are downloaded with the intermittent downloading, albeit the frequent disconnections/reconnections that slow down its overall downloading rate. By integrating the discharge currents over time, we calculate the consumed battery capacities in milliamp hours (mAh) during the downloading processes.³ This way, about 50.3mAh more capacity is found to be consumed by the intermittent downloading. Similar measurements are taken on a Desire C phone, and similar observations are made — the intermittent downloading consumes 99.1mAh more capacity and downloads $\approx 5\%$ more files than the continuous case. These measurements show that the continuous downloading draws (and consumes) less capacity from the battery, leading to a shorter device operation and less user-perceived utility, e.g., downloaded files.

3.2 Why More Capacity Can Be Drawn?

Next, we explore the reasons for the counter-intuitive observations in Sec. 3.1, showing the need of a novel battery-aware design for mobile devices. Specifically, this design is motivated by the following observations.

- **Mismatched Voltage Ranges:** Fig. 3 illustrates the basic principle of how the battery — described by an ideal voltage source V_b and an internal resistance R_b connected in series [21] — powers a mobile device. The voltage regulator on the device motherboard accepts power

input from the battery and supply power to various device modules, such as screen, GPS, and WiFi. Its main purpose is to regulate the voltage to the required levels of these modules and isolate circuitry from the transient voltage changes of the battery. Fig. 4 highlighted the regulator on the motherboard of a Nexus S phone. However, the voltage regulator requires a minimum level of the input voltage V'_b from the battery. Note that by Ohm's law, $V'_b = V_b - I_b \cdot R_b$ where I_b is the discharge current of the battery. Taking the simple but widely-used linear voltage regulator as an example, its input voltage has to be at least V_{dropout} higher than the required output voltage, where V_{dropout} is called the *dropout voltage* (e.g., 350mV for the TPS76733 regulator by TI). For simplicity, assuming the same required voltage V_{device} for individual components, i.e., $V_1 = V_2 = \dots = V_{\text{device}}$ in Fig. 3, the device operates when

$$V'_b \in [V_{\text{device}} + V_{\text{dropout}}, V_{\text{full}}], \quad (1)$$

where V_{full} is the voltage upon fully charge.

On the other hand, most off-the-shelf batteries have a built-in safety switch to prevent their over-discharge (as illustrated in Fig. 3), which not only degrades their cycle life, but may also lead to overheating and unsafety. The switch is opened, thus preventing further discharge when the battery voltage decreases to the cut-off level V_{cutoff} . In other words, the battery can be discharged when

$$V'_b \in [V_{\text{cutoff}}, V_{\text{full}}]. \quad (2)$$

This cut-off voltage is set by the battery manufacturer and built in the protection circuit, potentially creating a mismatch between the device operating voltage range (Eq. (1)) and the voltage range within which the batteries can be discharged (Eq. (2)).

We collected the voltage traces of 6 mobile devices to show the mismatched voltage ranges, whose battery information is summarized in Table 1. These devices are first fully charged and then kept with their screens on (and off) until powering themselves off. The thus-collected device power-off voltages are also listed in Table 1, showing a power-off voltage of 3.1–3.5V. Note that the power-off voltage depends on the device operating conditions — higher if the devices are in active usage (e.g., with screen on) than in idle state (e.g., with screen off). Next, we investigate the voltage range within which the batteries can be discharged. We use the NEWARE Battery Tester (as shown in Fig. 5) to control and monitor the discharge of initially fully charged batteries. Fig. 6 shows the discharge of the Nexus S battery (the same Nexus S as in Table. 1) with a current of 500mA. The

2. These durations are only to show the finding.

3. The current information of Acro S can be found at a system file `sys/class/power_supply/ab8500_fg/current_now`. Instead of the unit of energy Joule (J), mAh is a unit of charge commonly used for batteries.

TABLE 1: Batteries of adopted devices and their power off voltages (V).

	Note 8.0	Note 2	Nexus S	Galaxy S5	Galaxy S6 Edge	Xperia Z
Battery Capacity (mAh)	4,600	3,100	1,500	2,800	2,600	2,330
Battery Type	Li-ion	Li-ion	Li-ion	Li-ion	Li-ion	Li-ion
Battery Removable	No	Yes	Yes	Yes	No	No
Power off Voltage (Screen Off)	3.364	3.305	3.343	3.358	3.131	3.156
Power off Voltage (Screen On)	3.498	3.411	3.418	3.423	3.323	3.203

battery can be discharged within voltage range 2.8–4.2V, indicating a cut-off voltage of 2.8V (agreeing with [12]). Discharge measurements on a Note 2 battery reveals a cut-off voltage of 2.5V (Fig. 8).

Comparison between Table. 1 and Fig. 6 reveals that due to the voltage requirements of various on-board chips of mobile devices, their power-off voltage is higher than the cut-off voltage of batteries. More importantly, since the voltage range between the battery cut-off level and device power-off level (e.g., 2.8–3.4V for Nexus S and 2.5–3.3V for Note 2) cannot be effectively utilized by mobile devices, the corresponding capacity cannot be used to power the devices either. For example, Fig. 6 indicates about 148.7mAh capacity is not effectively drawn from the battery to power the device (assuming a power-off voltage of 3.343V), accounting for 10.2% of the total battery capacity, and can support a 500mA discharge process for 17.8 minutes. In summary, the mobile devices operate within a much narrower voltage range than the battery can actually supply, and hence only part of the battery capacity can be used.

- **Battery Recovery Effect:** Our second observation is that the *battery recovery effect* can be exploited to mitigate the above-mentioned voltage mismatch by drawing more capacity from the battery within the device operating voltage range.

When batteries are rested for a short period of time after discharging, their voltage recovers to a certain degree, which is called the *recovery effect* [14], [17]. To show the recovery effect, we use the battery tester to discharge a Nexus S battery for 6 minutes and then let it rest for about 20 minutes, as shown in Fig. 7. The battery voltage drops from 4.036V to 3.995V during discharge, and then recovers to 4.030V during the subsequent rest, i.e., $\frac{4.030-3.995}{4.036-3.995} = 85.4\%$ of the battery voltage is restored. Fig. 8 shows the voltage traces when intermittently discharging (30s rest after every 10s discharge) a fully charged Note 2 battery until its cut-off voltage is reached. Note the recovery effect does not require the absolute rest of batteries (i.e., a discharge current of 0mA), and can be observed as long as the discharge currents change from a relatively higher to lower levels as commonly seen in mobile device operations. We will elaborate this more in Sec. 4.

In general, the recovery effect can be reasoned about with two facts. The first is the voltage drop due to the battery internal resistance (R_b in Fig. 3), commonly referred to as the *polarization potential* [38]. The second is the re-balancing process of active materials during the rest period. The voltage recovered due to these two

facts is highlighted in Fig. 7. The details of the recovery effect are not within the scope of this paper and are referred to [14]. The recovery effect keeps the battery voltage higher for a longer time and thus increases the capacity drawn within the operating voltage range of mobile devices, as we will validate in Sec. 4.

Conjecture: Combining the above two observations, the counter-intuitive phenomenon shown in Sec. 3.1 may be reasoned about as follows. In the case of intermittent downloading, a bursty discharge pattern of the battery is created, allowing the battery to recover. As a result, the battery can maintain higher voltages for a longer period. This, in turn, mitigates the voltage mismatch by enhancing battery capacity delivery within the device operating range, thus eventually extending the device operation and improving the data service.

4 VALIDATION OF RECOVERY EFFECT

We have empirically investigated the impact of recovery effect on the capacity delivery of batteries, validating the above conjecture and shedding light on how to use it to improve the data service.

Before presenting our measurement methodology and the corresponding results, we first introduce a simplified notation to capture the battery discharge profiles. For any given discharge profile of batteries, we can approximate it with a piece-wise constant discharge profile with controllable accuracy [22], as illustrated in Fig. 9. This way, the discharge profile can be described by a sequence of tuples $\langle c_i, d_i \rangle$ ($i = 1, 2, \dots$), where c_i is the discharge current during the i -th phase of the discharge profile, and d_i is its corresponding duration.

In our measurements, batteries are discharged with a periodic profile of $\{ \langle c_h, d_h \rangle, \langle c_l, d_l \rangle \}$,⁴ where h and l denote the high discharge current (and duration) and low discharge current (and duration), respectively. Fig. 10 illustrates an example of this periodic discharge profile. The discharge currents and their durations are respectively measured in milliamp (mA) and second (s) in the rest of the paper, unless otherwise specified.

4.1 Efficacy of Utilizing Recovery Effect

We have taken measurements with different batteries — a Nexus S battery in use since 2010 and two Note 2 batteries in use since 2012 and 2015, respectively — to verify whether the recovery effect indeed enhances their capacity delivery. We discharge these initially fully

4. The periodic discharge profile is only to facilitate the measurements and is not required in B-MODS.




Fig. 6: Discharge a Nexus S battery.




Fig. 7: The voltage recovery process.




Fig. 8: Discharge a Note 2 battery.




Fig. 9: Approximating a given discharge profile with a piecewise constant current discharge profile.

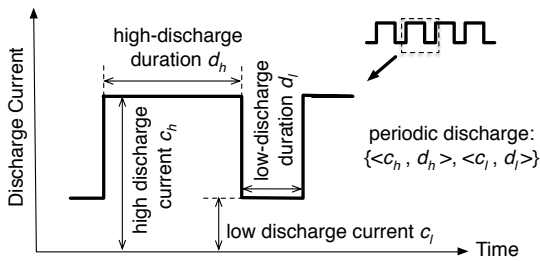


Fig. 10: The periodic discharge profile in our measurements.

charged (by the battery tester with the same charge profile) batteries according to the following two discharge profiles.⁵ First, the batteries are discharged according to a *continuous* discharge profile — $\{< 500, \infty \}$ for the Nexus S battery and $\{< 1,000, \infty \}$ for the two Note 2 batteries. Second, the batteries are discharged with a *bursty* discharge profile — $\{< 500, 10 \}$, $\{< 0, 1 \}$ for Nexus S and $\{< 1,000, 10 \}$, $\{< 0, 1 \}$ for Note 2. We make three important observations on recovery effect from these measurements.

• **Shifted Distribution of Delivered Capacity:** Discretizing the voltage range into a set of intervals of 0.1V each, Fig. 11 shows the delivered capacity in each interval (i.e., the distribution of delivered capacity with regard to voltages) obtained with these batteries. Note that the voltage ranges of 3.8–4.2V in Figs. 11(a) and 11(b) and 4.0–4.2V in Fig. 11(c) are missing from these results because of the voltage drop on the internal resistance of batteries. The bursty discharge increases the total capacity drawn from the battery, e.g., for Note 2 battery (2012), about 4% more capacity are delivered with bursty discharge when compared to the continuous discharge, agreeing with the conventional wisdom that recovery effect enhances the capacity delivery of batteries [17].

5. Again, these settings are only to show the finding.

More importantly, comparison between the capacity distributions obtained with continuous and bursty discharge profiles reveals that the bursty discharge *shifts* the distribution of delivered capacity towards higher voltages — the bars for the bursty discharge profile are shifted to the right as observed in Fig. 11. Again, taking the Note 2 battery (2012) as an example, about 1,232mAh capacity is delivered with voltage $\geq 3.4V$ (assuming a power-off voltage of 3.4V) with the bursty discharge profile, while only 823.9mAh capacity is delivered within the same voltage range when discharge continuously, indicating an increasing ratio of 49.5%. That is, by utilizing the recovery effect, more capacity can be drawn from batteries within the device operating voltage range, thus mitigating the voltage mismatch and extending the device operation time.

• **Pronounced with Aged Batteries:** Comparing Figs. 11(a), 11(b), and 11(c), we find the effect of bursty discharge on shifting the delivered capacity distribution is more pronounced with aged batteries — the Nexus S battery has been in use for ≈ 5 years and its shifting-effect is more pronounced than the two Note 2 batteries. (However, note that the bursty discharge draws 14.9% more capacity than the continuous discharge within the voltage range $\geq 3.4V$ even for the new Note 2 battery (Fig. 11(c)).) This is due to two reasons. First, the battery internal resistance increases as battery ages, leading to a larger polarization potential when the battery is rested. Second, aged batteries have more insufficiently activated material, leading to a more significant re-balancing process of active species. The stronger shifting-effect of aged batteries indicates that the mobile devices they power benefit more from exploiting the recovery effect for enhanced capacity delivery. This observation reinforces the necessity of a battery-aware design for users because (i) devices powered by aged batteries suffer more from the reduced operation time, and (ii) the replacement cycles of mobile devices are getting longer [39] and more devices are equipped with irreplaceable batteries, indicating more devices with aged batteries are in use.

• **Reduced Energy Loss on Internal Resistance:** Besides shifting the delivered capacity distribution, the recovery effect also helps reduce the energy loss on the internal resistance of batteries. The bursty discharge profile, during which the discharge current changes frequently, allows us to measure the internal resistance of batteries via $R = \Delta V / \Delta I$, where ΔV is the voltage changes when the discharge current changes by ΔI [38]. Fig. 12 shows the thus-measured internal resistance of a




Fig. 11: Bursty discharge shifts the distribution of delivered capacity towards high voltage intervals.




Fig. 12: The resistance of battery increases with lower voltage.

Nexus S battery and a Note 2 battery, where an increase of resistance can be observed with lower voltages, agreeing with the observations reported in [40]. Therefore, by keeping the voltage higher for a longer time, the recovery effect also keeps the battery resistance low and thus reduces its heating and internal energy loss.

More experiment results on the efficacy of utilizing recovery effect are provided in Appendix A due to space limit.

4.2 Factors Affecting Recovery Effect

We have shown that exploiting the recovery effect via the bursty discharge of batteries indeed increases their capacity delivery in the operating voltage range of mobile devices.⁶ Next, we investigate how the bursty discharge profiles affect the recovery effect. Specifically, we empirically study the impact of the high- and low-discharge durations on the recovery effect.

- **High-Discharge Duration:** We discharge the Note 2 battery (2015) with five profiles of various d_h s, as shown in Fig. 13. The bursty discharge increases its capacity delivery by exploiting the recovery effect, especially within the device operating voltage ranges. However, comparing the four bursty discharge processes, we find that an extremely long high-discharge duration is not desirable — the continuous discharge is actually the case of $d_h \rightarrow \infty$.

- **Low-Discharge Duration:** Next we discharge the battery with five profiles of different d_l s (Fig. 14). Again, the bursty discharge profiles draw more capacity from the battery within the device operating voltage range when compared to the continuous discharge. However, comparison among these bursty discharge processes reveals that (i) extending the low-discharge duration to the order of tens of seconds further pronounces the recovery effect on top of the milliseconds rest time as reported

in the literature [17], [23] (e.g., 9.4% more capacity is delivered within voltage range $\geq 3.4V$ with $d_l = 60s$ than that with $d_l = 0.1s$), which can also be inferred from Fig. 7 as the recovery process takes minutes to converge;⁷ (ii) an extremely large d_l has a diminishing effect in further strengthening the recovery effect, e.g., when comparing the cases with d_l s of 60s and 120s. These are of practical importance as we will see in Sec. 5.2.

Similar measurements (and observations) as in Figs. 13 and 14 with different batteries and discharge profiles are provided in Appendix A due to space limit.

- **Current Gap:** We discharge the battery with different current gaps to verify their impact on the recovery effect. Fig. 31 plots the capacity delivery during four discharge processes with different current gaps. First, the recovery effect occurs, and thus more capacity can be delivered, so long as the current transits from a higher to lower levels, not requiring the absolute rest of battery. Furthermore, a larger current gap, and thus a burstier discharge pattern, further improves the capacity delivery, e.g., Fig. 31 shows 15% more capacity is delivered when discharging the battery with the profile of $\{<1000, 30>, <0, 30>\}$ than that with $\{<1000, 30>, <700, 30>\}$, assuming again a power-off voltage of 3.4V as indicated in Table 1.

4.3 Validation Summary

Although the recovery effect has been known for years, we gained new insights based on the above measurements, including (i) exploiting the recovery effect shifts the capacity delivery of batteries towards higher voltage; (ii) the recovery effect helps reduce the energy loss on the internal resistance of the battery; (iii) the recovery effect is strengthened for aged batteries; and (iv) the battery voltage needs tens of seconds of low-discharge duration to fully recover.

Two conclusions can be drawn from the above observations. First, exploiting the recovery effect to extend (especially aged battery powered) device operation time is *feasible*. Second, we need to identify proper bursty discharge profiles to fully utilize the recovery effect. This is difficult because the recovery effect depends on dynamically changing battery parameters (such as SoC and resistance [22], [24], [25]) and real-time discharge currents, while at the same time we need to ensure the

6. The bursty discharge also reduces the overall discharge intensity of the battery, slowing down its capacity fading [12].

7. However, it is true that a large part of the voltage can be recovered in a short time, as shown in Fig. 7 and also in [17], [24].




Fig. 13: Impact of high-discharge duration on recovery effect.




Fig. 14: Impact of low-discharge duration on recovery effect.




Fig. 15: Impact of current gap on recovery effect.




Fig. 16: Discharge current of mobile device is inherently bursty.

user experience does not degrade. In the next section, we introduce B-MODS that exploits recovery effect to improve the data services on mobile devices.

5 BATTERY-AWARE MOBILE DATA SERVICE

As battery voltage needs tens of seconds or even minutes to fully recover, the inherent variance in the power draw of mobile devices (as illustrated in Fig. 16), normally in milliseconds [12], [17], is not enough for voltage recovery. However, the data service on mobile devices allows to take a full advantage of recovery effect with its relatively soft real-time requirements.

5.1 Data Service Property

Data-oriented services have become the dominant applications on mobile devices [18], [36]. Unlike traditional real-time services on mobile phones such as making/receiving calls, data service usually has softer real-time requirements [29], i.e., there is a certain slack in completing these data-service tasks.

On the other hand, a bursty discharge pattern needs to be formed to utilize the recovery effect. For mobile devices, this means some of their modules (or services) must be operated intermittently. However, the interruption of any device module or service means its temporary unavailability, which in turn delays the completion of the corresponding task. To guarantee the timely task completion and thus prevent any degradation in user experience, there must be a certain feasible time cushion for the task to be completed. The data service has such a time cushion because of its soft real-time requirements.

The soft real-time requirements of data service are also the reason for the lack of, or insufficient coverage of the topic of B-MODS in the literature, although the recovery effect itself is known for years — even though mobile devices (especially phones) have been in use for over 20 years, their traditional functions (i.e., make/receive calls) clearly do not possess the soft real-time property to leverage the recovery effect.

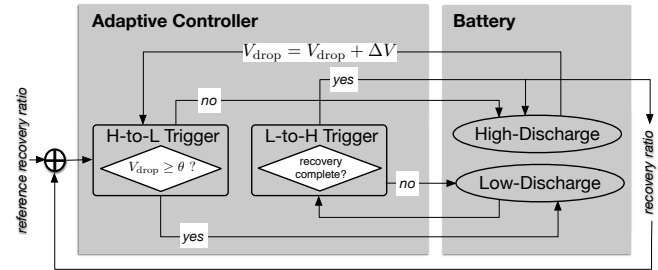


Fig. 17: Controller for battery triggered transitions.

5.2 Apply Recovery Effect to Data Service

Here we discuss how to apply the recovery effect to improve the mobile data service without jeopardizing user experience — how high and low discharge currents are formed and how long they should last.

5.2.1 Identify High and Low Discharge Currents

As mentioned earlier, we can alternatively turn on and off certain modules/services of the device to create high and low discharge currents. For example, in a WiFi-based downloading scenario, we can create the two currents by continuously switching between performing and not performing the downloading task. Specifically, the high and low discharge currents can be captured by

$$c_h = c_{bg} + c_{on} \quad \text{and} \quad c_l = c_{bg} + c_{off},$$

where c_{on} and c_{off} are the required current draw of the WiFi module when performing and not performing the downloading task, respectively, and c_{bg} is the sum of the required current by other modules. Note c_{bg} may not be constant over time.

After creating high and low discharge currents, we need to determine how the corresponding service transits between on and off states, i.e., determining the high and low discharge durations. This must be tackled based on the requirements of both the battery (to fully exploit the recovery effect) and the data service (to ensure no degradation in user experience).

5.2.2 Battery-Triggered Transitions

We first investigate the transitions between high and low discharge states from the perspective of batteries. Intuitively, a battery discharged with high current should be allowed to rest (i.e., switch to the low-discharge state) if it can restore a significant voltage level, and the battery being discharged with low current should transit to the high-discharge state if/when its voltage recovery has




Fig. 18: Recovered voltage ratios during discharge.

been completed. In B-MODS, we determine the battery-triggered transitions in real time. This is because the recovery effect not only depends on real-time battery parameters, but also varies with batteries (as will be revealed in Sec. 6.2) or even with the same battery at different cycles (e.g., the battery resistance increases as ages [38]), preventing a unified offline solution. Fig. 17 presents the controller architecture of the battery triggered transitions, details of which are given below.

- **High-to-Low Transitions:** We use the voltage drop θ during the high-discharge period as the transition trigger — the high-to-low transition is triggered once the voltage has dropped θ in the high-discharge period. To observe the impact of θ on the recovered voltage, Fig 18 plots the recovered voltage ratios (similar to Fig. 7) when discharging a battery with profiles of different d_h , where the x -axis is the voltage at which the recovery period starts. Two observations are made: (i) the recovered voltage ratio changes during discharge even with a given d_h , which again indicates the proper high-discharge duration should be identified in real time; (ii) a shorter high-discharge period, and thus a smaller θ , leads to a larger recovered ratio. However, too small a θ causes frequent state transitions, increasing the switching overhead (e.g., our measurements show that the WiFi-based HTTP connection takes at least 0.1–0.3s to be stable and could be longer) and reducing the energy-efficiency due to the *tail effect* [29], [41]. Combining these observations, we find that making θ as large as possible while guaranteeing a high enough voltage recovery ratio is desired for the high-to-low transitions.

We borrow the slow-start algorithm in TCP congestion control to identify the proper θ online. Specifically, we initially use a small θ and examine how much of the dropped voltage can be recovered. If the recovered voltage ratio is large enough, e.g., $\geq 85\%$, we increase θ to 2θ until the recovered voltage ratio is less than 85% . Then, we use a binary search to find the largest θ guaranteeing an 85% recovery ratio.

- **Low-to-High Transitions:** Fig. 14 shows an over-long rest time of the battery has diminishing effect on further increasing the delivered capacity. To shed more light on the voltage recovery process, the recovered voltage of the Nexus S battery after being discharged with 500mA current are plotted in Fig. 19. The voltage recovery process can be captured by a piece-wise linear model consisting of a set of line segments. Specifically, we define *turning points* as the points in the voltage recovery curve with a considerable slope change, which




Fig. 19: Voltage recovery process after discharge.

Fig. 20: Goodness when fitting β as exponential.




Fig. 21: The voltage recovery process of the Nexus S battery after its receiving WiFi module entering the idle state.

are used to determine the start and end points of each line segment. By denoting $\{v_1, v_2, v_3, \dots\}$ as the sequence of recovered voltages at the turning points of the segments and $\{t_1, t_2, t_3, \dots\}$ as the corresponding time instances, we can capture the voltage recovery process as

$$v(x) = \beta_i \cdot x + v_i \quad (t_i \leq x < t_{i+1}), \quad (3)$$

where β_i is the slope coefficient defining each segment and quantifies the voltage recovery speed. Fig. 19 also shows the sequence of β corresponding to the same voltage recovery trace, which visually conforms to an exponential decay process — the sequence of β follows $\beta_i = a \cdot e^{b \cdot i}$ for certain a and b .

To verify this hypothesis, we empirically collect 107 voltage recovery traces with various discharge currents and durations, and then we apply exponential fit to their corresponding β sequences with 95% confidence level. The goodness of the exponential fit (in both *root-mean-square error* (RMSE) and *adjusted-R²*) is summarized in Fig. 20. As a close-to-0 RMSE and a close-to-1 adjusted- R^2 indicate good fitting accuracy, the goodness-of-fit clustered in the right-bottom corner of the figure, as observed in Fig. 20, validates the hypothesis on the exponential decay process.

In B-MODS, we estimate β_i based on the sampled voltages with linear regression and predict β_{i+1} with exponential fit. As an example, Fig. 21(a) plots the voltage recovery process of a Nexus S battery after a receiving WiFi module enters the idle state. In general, the voltage recovery process converges as the idle process continues, agreeing with Fig. 19 where the discharge current is explicitly controlled. However, two differences can be observed. First, the recovered voltage is smaller and the time for the recovery process to converge is shorter when compared with Fig. 19, due to the existence of background current. Second, the voltage recovery process in Fig. 21(a) is not as smooth as that in Fig. 19 because the background current is not constant in practice, leading to a less smooth trace of β . To reliably identify the convergence of the recovery process, we use the moving

average filter to smooth the sequence of β before applying the exponential fit, and high level of goodness-of-fit can still be obtained (e.g., the adjusted- R^2 is 0.8512 for the recovery process shown in Fig. 21(b)). The recovery process is concluded if the accuracy of the exponential fit is high (e.g., adjusted- $R^2 > 0.8$) and the predicted β_{i+1} is small (e.g., $< 10\text{mV}$ per second).

5.2.3 Data Service Triggered Transitions

Next, let us consider how the transitions between high and low discharge states should be determined according to the data service requirements. The time to complete the data service may vary with time due to the dynamic link quality, leading to diverse slack times. These dynamics in data service can be captured by the *variable rate execution* model [42], and we determine the data-service triggered transitions based on whether the available slack time is sufficient or not.

For the high-to-low transition, we determine whether the slack is sufficient based on the service rates (e.g., downloading rates for a downloading task) during the previous high-discharge period.⁸ Specifically, let T and t be the deadline of the task and the current time (e.g., the user will arrive home in one hour and she wants the system updates to be completed by that time), and T can be either specified by the user or estimated based on historic usage information [43]. We compute the distribution of the service rate during the previous high-discharge period, and adopt its p -percentile $r_p(t)$ to determine whether the slack is sufficient: the slack is deemed to be sufficient *iff*

$$r_p(t) \cdot (T - t) \geq B - f(t), \quad (4)$$

where B is the total data volume associated with the task, and $f(t) \in [0, B]$ is the completed data volume until time t . This condition can be interpreted as even if the service rate would be slow throughout the remaining time (i.e., as slow as $r_p(t)$), the data service task can still be completed within its time cushion.

Similarly, we say the slack is insufficient *iff* the q -percentile, $r_q(t)$, of the service rate during the previous high-discharge period satisfies

$$r_q(t) \cdot (T - t) < B - f(t), \quad (5)$$

indicating that even the service rate would keep high throughout the remaining time (i.e., as high as $r_q(t)$), the data service task cannot be completed within the time cushion if the service is not resumed immediately.

5.2.4 Joint Transition Design

Fig. 22 summarizes the transitions between high and low discharge states. Note that the high-to-low transition

⁸ With the battery-awareness as the theme of the work, we use the service rate — the most direct factor determines whether the data service tasks can be completed in time — to simplify the design. More in-depth approaches involving detailed communication factors exist in the literature, e.g., *Bartendr* [8] and *ATOM* [36], with which can be potentially integrated for further improvement.




Fig. 22: Transition diagram of B-MODS.




Fig. 23: Battery-aware data service library architecture.

occurs if *both* A.1 and B.1 are met, while *either* of A.2 or B.2 triggers the low-to-high transition. This way, the condition B.1 imposes a constraint on the shortest duration for each high-discharge period, and the condition B.2 limits the longest duration the lower discharge period can last. As a result, B-MODS exploits the recovery effect to extend device operation whenever possible, otherwise it seamlessly reverts back to continuously provide the data service to ensure user-experience.

5.3 Implementation

We implement B-MODS as an application layer library on the Android platform, which accepts data-service (i.e., downloading or streaming) requests/tasks from other applications and schedules them accordingly. In the library, the battery voltage is sampled and made accessible via a system file located at different places for heterogeneous mobile devices. Specifically, for Nexus S, it is located at `/sys/class/power_supply/battery/voltage_now`, while for Acro S, Xperia Z and Desire C, it is at `/sys/devices/i2c-3/3-0055/power_supply/bq27520/voltage_now`. Also, the interruptible data service is realized by disconnecting (re-connecting) an HTTP connection during low (high) discharge periods. Use of an HTTP connection implies a TCP-based implementation and provides reliable delivery of octets stream. This is adopted by most data-service applications, such as video streaming, cloud storage, and file transmission [36]. The library can be found at [44], whose overall architecture is shown in Fig. 23, consisting of the manager, the battery-aware scheduler, and the executor.

- *Manager*, running in the background, monitors the AIDL interface to accept and handles the data service tasks from other applications. When a task is received, the manager first stores the task information (e.g., file size and the time cushion to complete the task) to avoid task duplication, and then it passes the task to the battery-aware scheduler and launches the executor to

perform the task. The manager automatically starts boot once by registering a broadcast receiver to listen to the boot complete event.

- *Battery-aware scheduler*, which can be further decomposed into *battery monitor* and *schedule unit*, is responsible for deciding the status of executor (i.e., working or resting) based on the continuously sampled battery voltages and the task information. The battery monitor captures real-time battery voltages via a system file containing the voltage information, and passes them to the schedule unit. Based on these sampled voltages, the schedule unit determines whether the executor should be in the working or resting states.

- *Executor*, consisting of a downloader and an uploader, actually performs the data service. It supports the interruptible data service to allow the battery to rest and recover — disconnects the HTTP connection during the low-discharge period and reconnects for the next high-discharge period. The downloader writes the response to a file specified by the downloading task, and the uploader reads the local file and writes it into a network stream specified by the uploading task.

6 EXPERIMENTAL EVALUATION

6.1 Evaluation Settings and Scenarios

We evaluate B-MODS under different application scenarios with various mobile devices. We use the following default settings unless otherwise specified. We sample the battery voltage at 10Hz. A recovered voltage ratio of at least 85% is adopted when identifying θ for the battery-triggered high-to-low transitions. Linear regression is used to estimate β based on the previous 5 samples. The recovery process is concluded to be completed if the adjusted- R^2 of the exponential fit is larger than 0.8 and the predicted recovery speed is slower than 10mV per second. $p = 20$ and $q = 80$ percentiles of the downloading-rate distribution are used to estimate if the slack is sufficient. Besides B-MODS, we also implement a battery-agnostic data service for comparison, under which the devices perform the tasks continuously with best effort.

Observing the wide availability of WiFi to which the devices are connected for $\approx 70\%$ of the time [1], we evaluate B-MODS under 3 WiFi-based application scenarios: (i) background downloading, (ii) downloading with (emulated) human interactions, and (iii) online streaming with commercial APPs.

6.2 Background Downloading

We first evaluate B-MODS under the scenario of WiFi-based background downloading, in which the screen is kept off throughout the downloading process. This background downloading scenario not only allows us to better observe the advantages of B-MODS because of smaller and relatively stable background currents, but is also common in the daily usage of mobile devices, as

TABLE 2: Voltage drop (mV) and duration (s).

	Initial V	End V	Duration
B-MODS (Nexus S)	4,087.5	3,932.5	2,669
Continuous (Nexus S)	4,097.5	3,937.5	2,043
B-MODS (Acro S)	4,164	4,102	635.4
Continuous (Acro S)	4,098	4,033	524.8




Fig. 24: Voltage traces of the two downloading processes.

most users keep their phones in the *connected standby* mode for a large fraction of time [7], accounting for $>40\%$ of the total energy drain [1]. A Nexus S and an Acro S phone are used in this set of experiments.

- **B-MODS vs. Continuous Downloading:** With a downloading task of file size 1GB and a time cushion T of 1 hour, we record the B-MODS-based and continuous downloading processes, as summarized in Table 2. B-MODS takes longer to complete the task than the continuous downloading, but still within the 1 hour cushion. More importantly, the accumulated voltage drop after completing the task with B-MODS is 5mV and 3mV less for Nexus S and Acro S respectively, as a result of exploiting the recovery effect.

- **Downloading Until Power-Off:** While the previous results show smaller voltage drops with B-MODS, we still need to quantify how much the device can actually benefit from B-MODS — what does the 5mV gain mean for users? To this end, we examine the scenario where the to-be-downloaded file is sufficiently large such that the task cannot be completed with a single charge of the phones and the time cushion is large enough to complete the downloading. This way, the downloading processes continue until the phones power off, allowing us to quantify the effect of B-MODS on improving the data services. Similar method was used in [29] and widely adopted in mobile device specifications, e.g., a talk time of 18+ hours for Galaxy S6 Edge [45] and 8 hour LTE time for Nexus 5X [46].

The voltage traces during the two downloading processes are plotted in Figs. 24. The first observation is that B-MODS significantly extends the downloading process, which is not surprising because the low-discharge state reduces the overall load intensity. Furthermore, the voltage drop of the continuous downloading processes is sharp when the voltage is low (e.g., below 3.6V), while by utilizing the recovery effect, this voltage drop




Fig. 25: Effective download time and downloaded file.




Fig. 26: Distributions of the receiving and idle periods.

becomes much slower with B-MODS. The intermittent downloading also helps address the storage deficiencies in mobile devices [47].

To check whether this extended downloading process indeed improves the data service, Fig. 25 shows the statistics of the two downloading processes, in terms of the accumulated receiving periods of the WiFi module and the downloaded file size. By utilizing the recovery effect, not only the downloading process is extended, but more importantly, the effective downloading time increases as well. For example, the effective downloading time of Nexus S is extended by 2,074s, or 16.2% compared to the continuous downloading. The extended effective downloading time and the corresponding ratio for Acro S are 6,361s and 28.6%, respectively. Similar observations can be made on the downloaded file — B-MODS downloads 15.3% (with Nexus S) and 12.3% (with Acro S) more data as compared to the continuous downloading. The improvement in downloaded file is smaller than that in effective downloading time due to the HTTP reconnections, slowing down its overall downloading rate. These improvements are significant especially in view of the fact that the battery density has only been doubled over the past 15 years [2].

- **Receiving and Idle Periods Distributions:** To shed more light on the B-MODS-based downloading, Fig. 26 plots the distributions of the receiving and idle time period durations obtained from Nexus S and Acro S. The two devices show significant differences in the periods distributions. For example, the average receiving and idle periods for Nexus S are about 204.3s and 730.7s, respectively, yielding a receiving-to-idle-ratio of ≈ 0.28 . On the other hand, the average receiving and idle periods for Acro S are 461.6s and 260.3s, and the ratio between them is ≈ 1.77 . These results indicate that different mobile devices (and batteries) favor different discharge patterns to exploit the recovery effect, e.g., the Nexus S battery needs more time to recover than the Acro S battery according to the results shown in Fig. 26. This, in turn, calls for the online determination of the transitions between high- and low-discharge states — as is done in B-MODS — instead of attempting to identify a *single* discharge pattern for all devices offline.

6.3 Downloading with Human Interactions

To verify if B-MODS benefits the downloading process in more general scenarios, we introduce a control parameter $\phi \in [0, 1]$ in the second set of our experiments, reflecting the probability for the screen to stay on during downloading. Specifically, dividing time into unit




Fig. 27: Download time with human interactions.




Fig. 28: Online streaming durations with commercial APPs.

intervals of 1 minute, the screen of the phone stays on during each time unit with probability ϕ . This way, we emulate the randomized human interactions during the downloading processes, and the thus-obtained results help us see how B-MODS performs in the presence of other battery-agnostic device operations. A Nexus S and an Xperia Z are used for this set of experiments.

The accumulated effective downloading time with various ϕ are summarized in Fig. 27 (averaged over 5 runs). B-MODS achieves longer effective downloading time in all the cases we studied, but the advantage of B-MODS diminishes as ϕ gets larger. Taking Nexus S as an example, when compared to the continuous downloading, B-MODS makes 10.3%, 5.2%, and 2.5% improvements with $\phi = 0.1, 0.5,$ and 0.9 , respectively. This can be reasoned about as follows: as the background current increases because of a larger probability for the screen to stay on, the current gap formed by adjusting the WiFi operation accounts for a decreasing portion of the overall discharge current of the phone, thus degrading the recovery effect. We will elaborate more on this in Sec. 6.5.

6.4 Online Streaming with Commercial APPs

Last but not the least, we evaluate the battery-aware design for online streaming with three off-the-shelf APPs: *Youtube* [48], *Repeat on Youtube* [49], and *LiveNow!tv* [50].⁹ These well-encapsulated APPs pose new experimental challenges as we (i) have no access to the size of buffered data and (ii) have no control on whether the downloading requests should be sent to the WiFi module, making the data service library not applicable. Therefore, instead of scheduling the downloading requests, we build *WiFi-Switch*, an application that directly turns on/off the WiFi module, to emulate a B-MODS-based streaming process. Specifically, *WiFi-Switch* consecutively turns on/off the WiFi module according to two user-customized *on* and *off* durations $\langle x, y \rangle$ — turning on the WiFi module for x s before tuning it off for y s, and then repeat. After

9. We use *Youtube* to stream a 10-hour video *Nyan Cat* [51]. For *Repeat on Youtube*, a video *One Man Les Miserables* [52] is played repeatedly until the phone powers itself off. *LiveNow!tv* is used to stream an on-demand audio *Comedy Bang Bang*.

TABLE 3: WiFi state distribution (%).

WiFi State	Disabling	Disabled	Enabling	Enabled
Ratio	0.3178	4.5264	1.0229	94.1329




Fig. 29: Two-peak discharge current distribution.




Fig. 30: The peak currents under three scenarios.

testing various settings of $\langle x, y \mid x \in \{100, 200\}, y \in \{5, 10, 20\} \rangle$, we observe the setting of $\langle 100, 5 \rangle$ is able to show the advantages of the battery-aware design while ensuring a smooth streaming process in our experimental environment. A video recording part of such a streaming process can be found at [53].

Fig. 28 summarizes the streaming durations. The battery-aware streaming significantly extends the streaming processes when compared to the continuous case, and even the accumulated duration with an on-WiFi module (e.g., $\approx 9,518.2 \times \frac{100}{100+5} = 9,065.0$ s for Nexus S) is longer than that with continuous streaming. Similar observations can be made from the results collected with the other two APPs and also those collected with Xperia Z. Table 3 summarizes the WiFi states distribution during a battery-aware streaming process with Nexus S. Enabling WiFi model is found to take longer than disabling it, as reported in [29].

6.5 Cross-Application Comparison

The above experiments over three application scenarios show that exploring battery-awareness in the data service not only extends the device operation, but also improves user-perceived utilities, such as file downloading and streaming duration. However, comparing the improvement ratios across the three scenarios, we find the background downloading benefits most from the battery-awareness (e.g., 15.3% more downloaded files with Nexus S), while the improvement ratio of downloading with emulated human interactions is limited, especially with a high screen-on probability. To investigate a potential reason for this, Fig. 29 plots the current distribution of the Xperia Z phone when operating under the downloading with human interaction with $\phi = 1$. Not surprisingly, two peaks are observed, corresponding to the cases when performing (i.e., the right peak) and not performing (i.e., the left peak) the downloading tasks. A similar two-peak pattern of the current distribution is observed under the other two application scenarios, as shown in Fig. 30. An interesting finding is the advantage of battery-awareness decreases as the ratio between the right and left peak currents get smaller, implying a *burstier* discharge pattern may further improve the capacity delivery from battery, which we will empirically verify in the next section.

7 BATT.-AWARENESS BEYOND DATA SERVICE

We have shown B-MODS to improve data services, but also noticed its limited advantage when the ratio between the two peak currents is small. So, we next explore the reasons for this phenomenon and suggest its mitigation — extending battery-awareness to multiple device modules to form burstier discharge patterns.

7.1 Burstier Discharge Does Help

We discharge two batteries with different current gaps to verify whether the burstier discharge indeed pronounces the recovery effect. Fig. 31 plots the capacity delivery during these discharge processes. First, the recovery effect occurs, and thus more capacity can be delivered, so long as the current transits from a higher to lower levels, not requiring the absolute rest of battery. Furthermore, as expected, a larger current gap, and thus a burstier discharge pattern, does further improve the capacity delivery, e.g., Fig. 31(b) shows 11.6% more capacity is delivered with the profile of $\{ \langle 500, 30 \rangle, \langle 0, 30 \rangle \}$ than that with $\{ \langle 500, 30 \rangle, \langle 300, 30 \rangle \}$, assuming again a power-off voltage of 3.4V.

This observation implies a degraded recovery effect when the formed current gap is small, e.g., when intermittently download via 3G/LTE which shows a clear tail-effect. However, this also indicates the battery-aware energy management can be improved further by extending battery-awareness to other device modules and jointly scheduling them to create larger current gaps — e.g., via voltage scaling (up to 641mA current gap can be formed by adjusting the GPU frequency for Galaxy S4 [1]) and adaptive display (as we will show next). This, in turn, allows for a thorough evaluation of battery-aware energy management with real-life device usage beyond mobile data service.

7.2 Example: Battery-Aware Display

We explore the battery-aware display as a preliminary study of applying battery-awareness to other device modules. Specifically, we built an application that gradually adjusts screen brightness between 10% to 100% of its maximum level, with a step of 10% that changes every 2s. We ran the application on an Xperia Z until it powers off. We recorded the process and compared it with the case when the screen displays with the maximal brightness, as shown in Fig. 32. The device operation is significantly extended with this adjustment of brightness, i.e., not only by reducing the brightness level but also by receiving $2,723.5 - 2,506.8 = 216.7$ mAh more capacity from the battery, showing the feasibility of applying battery-awareness to adaptive screen display. The challenge, again, is to make the brightness adjustment transparent to users. Exemplary approaches include adjusting screen brightness during the scrolling operations [9] or based on user-screen distance [33] and dimming display areas of less interest to users [34].




Fig. 31: A larger current gap improves battery capacity delivery further.




Fig. 32: Battery-aware display.

8 CONCLUSIONS

Observing the mismatch between the consumed energy on mobile devices and the supplied energy from batteries, we have shown the necessity of battery-aware energy management for mobile devices. Focusing on the prevalent data-oriented services seen on mobile devices, we have proposed a novel design of B-MODS for enhancing the capacity delivery of the batteries by utilizing the battery recovery effect and improving user-perceived data service utilities. We have evaluated B-MODS with heterogeneous mobile devices under different application scenarios. The results have shown B-MODS to improve data service by up to 28.6%.

Acknowledgment: The work reported in this paper was supported in part by NSF under Grants CNS-1329702, CNS-1446117, and CNS-1503590.

REFERENCES

- [1] X. Chen, N. Ding, A. Jindal, Y. C. Hu, M. Gupta, and R. Vanithamby, "Smartphone background activities in the wild: Origin, energy drain, and optimization," in *MobiCom'15*, 2015.
- [2] "Microsoft Aims for Smartphones That Run for a Week," <http://www.technologyreview.com/news/528201/>.
- [3] C. Shen, S. Chakraborty, K. R. Raghavan, H. Choi, and M. B. Srivastava, "Exploiting processor heterogeneity for energy efficient context inference on mobile phones," in *HotPower'13*, 2013.
- [4] R. LiKamWa, B. Priyantha, M. Philipose, L. Zhong, and P. Bahl, "Energy characterization and optimization of image sensing toward continuous mobile vision," in *MobiSys'13*, 2013.
- [5] J. Flinn and M. Satyanarayanan, "Managing battery lifetimes with energy-aware adaptation," *ACM Transaction on Computer Systems*, vol. 22, no. 2, pp. 137 – 179, 2004.
- [6] A. Roy, S. M. Rumble, R. Stutsman, P. Levis, D. Mazières, and N. Zeldovich, "Energy management in mobile devices with the Cinder operating system," in *EuroSys '11*, 2011.
- [7] F. Xu, Y. Liu, T. Moscibroda, R. Chandra, L. Jin, Y. Zhang, and Q. Li, "Optimizing background Email sync on smartphones," in *MobiSys'13*, 2013.
- [8] 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 *MobiCom'10*, 2010.
- [9] H. Han, J. Yu, H. Zhu, Y. Chen, J. Yang, G. Xue, Y. Zhu, and M. Li, "E³: energy-efficient engine for frame rate adaptation on smartphones," in *SenSys'13*, 2013.
- [10] M. Martins and R. Fonseca, "Application modes: a narrow interface for end-user power management in mobile devices," in *HotMobile'13*, 2013.
- [11] M. Kim, Y. G. Kim, S. W. Chung, and C. H. Kim, "Measuring variance between smartphone energy consumption and battery life," *Computer*, vol. 47, no. 7, pp. 59–65, July 2014.
- [12] D. Panigrahi, C. Chiasserini, S. Dey, R. Rao, A. Raghunathan, and K. Lahiri, "Battery life estimation of mobile embedded systems," in *VLSI Design, 2001. Fourteenth International Conference on*, 2001, pp. 57–63.
- [13] S. Park, A. Savvides, and M. B. Srivastava, "Battery capacity measurement and analysis using Lithium coin cell battery," in *ISLPED'01*, 2001.

- [14] T. L. Martin, "Balancing batteries, power, and performance: System issues in CPU speed-setting for mobile computing," *Ph.D. thesis, Carnegie Mellon University*, 1999.
- [15] L. He, L. Gu, L. Kong, Y. Gu, C. Liu, and T. He, "Exploring adaptive reconfiguration to optimize energy efficiency in large-scale battery systems," in *RTSS'13*, 2013.
- [16] L. He, Y. Gu, C. Liu, T. Zhu, and K. G. Shin, "Share: Software-aware reconfiguration to enhance deliverable capacity of large-scale battery packs," in *ICCPs'15*, 2015.
- [17] C. F. Chiasserini and R. Rao, "Pulsed battery discharge in communication devices," in *MobiCom'99*, 1999.
- [18] "Cisco visual networking index: Global mobile data traffic forecast update, 2013–2018," http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white_paper_c11-520862.html.
- [19] A. Carroll and G. Heiser, "An analysis of power consumption in a smartphone," in *USENIXATC'10*, 2010.
- [20] N. Ding, D. Wagner, X. Chen, A. Pathak, Y. C. Hu, and A. Rice, "Characterizing and modeling the impact of wireless signal strength on smartphone battery drain," in *SIGMETRICS'13*, 2013.
- [21] L. Zhang, B. Tiwana, R. Dick, Z. Qian, Z. Mao, Z. Wang, and L. Yang, "Accurate online power estimation and automatic battery behavior based power model generation for smartphones," in *CODES+ISSS'10*, 2010.
- [22] D. Rakhmatov, S. Vrudhula, and D. A. Wallach, "A model for battery lifetime analysis for organizing applications on a pocket computer," *TVLSI*, vol. 11, no. 6, pp. 1019 – 1030, 2003.
- [23] C. F. Chiasserini and R. R. Rao, "Improving battery performance by using traffic shaping techniques," *IEEE Journal on Selected Areas in Communications*, vol. 19, no. 7, pp. 1385–1394, Jul 2001.
- [24] C. Chiasserini and R. Rao, "A model for battery pulsed discharge with recovery effect," in *Wireless Communications and Networking Conference, 1999. WCNC. 1999 IEEE*, 1999, pp. 636–639 vol.2.
- [25] S. Hong, "Battery-aware real-time task scheduling in wireless sensor networks," *Master thesis, KAIST*, 2005.
- [26] P. Rong and M. Pedram, "An analytical model for predicting the remaining battery capacity of lithium-ion batteries," *IEEE Transactions on Very Large Scale Integration Systems*, vol. 14, no. 5, pp. 441–451, 2006.
- [27] S. J. Moura, J. L. Stein, and H. K. Fathy, "Battery-health conscious power management in plug-in hybrid electric vehicles via electrochemical modeling and stochastic control," *IEEE Transactions on Control Systems Technology*, vol. 21, no. 3, pp. 679–694, 2013.
- [28] 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 *MobiSys'12*, 2012.
- [29] N. Balasubramanian, A. Balasubramanian, and A. Venkataramani, "Energy consumption in mobile phones: a measurement study and implications for network applications," in *IMC'09*, 2009.
- [30] A. Pathak, Y. C. Hu, and M. Zhang, "Fine grained energy accounting on smartphones with Eprof," in *EuroSys'12*, 2012.
- [31] R. Mittal, A. Kansal, and R. Chandra, "Empowering developers to estimate App energy consumption," in *MobiCom'12*, 2012.
- [32] M. Dong, T. Lan, and L. Zhong, "Rethink energy accounting with cooperative game theory," in *MobiCom'14*, 2014.
- [33] H. Z. Songtao He, Yunxin Liu, "Optimizing smartphone power consumption through dynamic resolution scaling," in *MobiCom'15*, 2015.
- [34] K. W. Tan, T. Okoshi, A. Misra, and R. K. Balan, "Focus: A usable and effective approach to oled display power management," in *UbiComp '13*, 2013.
- [35] "Storage, distribution, electrical, convertors," http://www.joules-project.eu/Joules/technologies/storage_distribution_electrical_convertors.

[36] R. Mahindra, H. Viswanathan, K. Sundaresan, M. Y. Arslan, and S. Rangarajan, "A practical traffic management system for integrated LTE-WiFi networks," in *MobiCom'14*, 2014.

[37] S. Rosen, H. Luo, Q. A. Chen, Z. M. Mao, J. Hui, A. Drake, and K. Lau, "Discovering fine-grained RRC state dynamics and performance impacts in cellular networks," in *MobiCom'14*, 2014.

[38] D. Andrea, *Battery Management Systems for Large Lithium-Ion Battery Packs*. Artech House, 2010.

[39] "The smartphone upgrade cycle gets longer," <http://www.businessinsider.com/the-smartphone-upgrade-cycle-gets-longer-2013-9>.

[40] H. Zhang, R. Lu, C. Zhu, and Y. Zhao, "On-line measurement of internal resistance of lithium ion battery for ev and its application research," *Computer*, vol. 7, no. 4, pp. 301-310, July 2014.

[41] W. Hu and G. Cao, "Quality-aware traffic offloading in wireless networks," in *MobiHoc'14*, 2014.

[42] S. Goddard and X. Liu, "Variable rate execution," *Technical Report, University of Nebraska-Lincoln*, p. 2004, 2004.

[43] H. Falaki, R. Mahajan, S. Kandula, D. LyMBERopoulos, R. Govindan, and D. Estrin, "Diversity in smartphone usage," in *MobiSys'10*, 2010.

[44] "B-MODS Library," <https://sites.google.com/site/bamds2015/>.

[45] "Galaxy S6 Edge Full Specifications," http://www.gsmarena.com/samsung_galaxy_s6_edge-7079.php.

[46] "Nexus 5X," https://store.google.com/product/nexus_5x.

[47] H. Kim, N. Agrawal, and C. Ungureanu, "Revisiting storage for smartphones," in *FAST'12*, 2012.

[48] "Youtube," <https://play.google.com/store/apps/details?id=com.google.android.youtube&hl=en>.

[49] "Listen on Repeat for Youtube," <https://play.google.com/store/apps/details?id=com.youtuberepeatfree>.

[50] "LiveNow!tv," <https://play.google.com/store/apps/details?id=com.filmon.android.aff38906Te&hl=en>.

[51] "Nyan Cat," <https://www.youtube.com/watch?v=wZZ7oFKsKzY>.

[52] "One Man Les Misérables," <http://www.youtube.com/watch?v=R9IWAxSpgKA>.


[53] "A video for B-MODS-based streaming with Youtube," <https://youtu.be/RwUwnSPeE5w>.



Cong Liu is an assistant professor in the Department of Computer Science at the University of Texas at Dallas. He received his Ph.D. degree in Computer Science from the University of North Carolina at Chapel Hill in July 2013. His research interests include real-time systems and GPGPU. He has published over 30 papers in premier conferences and journals. He received the Best Paper Award at the 30th IEEE RTSS and the 17th RTCSA.



Jun Sun received Bachelor and PhD degrees in computing science from National University of Singapore (NUS) in 2002 and 2006. In 2007, he received the prestigious LEE KUAN YEW postdoctoral fellowship in School of Computing of NUS. Since 2010, he joined Singapore University of Technology and Design (SUTD) as an Assistant Professor. Jun's research interests include software engineering, formal methods, software engineering, program analysis and cyber-security. He is the co-founder of PAT.



Ting Zhu is an assistant professor in the Department of Computer Science and Electrical Engineering at the University of Maryland, Baltimore County. He received his Ph.D. from the Department of Computer Science and Engineering at the University of Minnesota in 2010. He received his B.S. and M.E. degree from the Department of Information Science and Electronic Engineering at the Zhejiang University in 2001 and 2004, respectively. His research interests include wireless networks, renewable energy, embedded systems, and security.



Yang Liu graduated in 2005 with a Bachelor of Computing in the National University of Singapore (NUS). In 2010, he obtained his PhD and continued with his post-doctoral work in NUS. Since 2012, he joined Nanyang Technological University as a Nanyang Assistant Professor. His research focuses on software engineering, formal methods and security. Particularly, he specializes in software verification using model checking techniques. This work led to the development of a state-of-the art model checker, PAT.



Kang G. Shin (LF'12) is the Kevin & Nancy O'Connor Professor of Computer Science in the Department of Electrical Engineering and Computer Science, The University of Michigan, Ann Arbor. His current research focuses on QoS-sensitive computing and networking as well as on embedded real-time and cyber-physical systems.

He has supervised the completion of 77 PhDs, and authored/coauthored more than 830 technical articles, a textbook and more than 30 patents or invention disclosures, and received numerous best paper awards, including the Best Paper Awards from the 2011 ACM International Conference on Mobile Computing and Networking (MobiCom'11), the 2011 IEEE International Conference on Autonomic Computing, the 2010 and 2000 USENIX Annual Technical Conferences, as well as the 2003 IEEE Communications Society William R. Bennett Prize Paper Award and the 1987 Outstanding IEEE Transactions of Automatic Control Paper Award. He has also received several institutional awards, including the Research Excellence Award in 1989, Outstanding Achievement Award in 1999, Distinguished Faculty Achievement Award in 2001, and Stephen Attwood Award in 2004 from The University of Michigan (the highest honor bestowed to Michigan Engineering faculty); a Distinguished Alumni Award of the College of Engineering, Seoul National University in 2002; 2003 IEEE RTC Technical Achievement Award; and 2006 Ho-Am Prize in Engineering (the highest honor bestowed to Korean-origin engineers).

He was a co-founder of a couple of startups and also licensed some of his technologies to industry.



Liang He (S'09-M'12) is currently a research fellow at University of Michigan, Ann Arbor, MI. He worked as a postdoc fellow and then research scientist at Singapore University of Technology & Design during 2012-2014, and as a research assistant at University of Victoria during 2009-2011. His research interests include CPSes, battery management, and networking. He has been a recipient of the best paper awards of QShine'14, WCSP'11 and GLOBECOM'11.



Guozhu Meng received Bachelor and Master degrees in school of computer science and technology from Tianjin University, China in 2009 and 2012. He worked in Temasek lab in National University of Singapore for one year as an Associate Scientist. Since 2013, he started his Ph. D study in school of computer science and engineering of Nanyang Technological University, Singapore. His research interests include mobile security, software engineering and program analysis.



Yu Gu is currently a research scientist at IBM Research, Austin. He was an assistant professor at the Singapore University of Technology and Design between 2010-2014. He received the Ph.D. degree in the Department of Computer Science and Engineering at the University of Minnesota, 2010. He is the author and co-author of over 100 papers in premier journals and conferences. He has received several prestigious awards from the University of Minnesota.