

A Case for Battery Charging-Aware Power Management and Deferrable Task Scheduling in Smartphones

Salma Elmalaki, Mark Gottscho, Puneet Gupta and Mani Srivastava
University of California, Los Angeles

selmalaki@ucla.edu, mgottscho@ucla.edu, puneet@ee.ucla.edu, mbs@ucla.edu

Abstract

Prior battery-aware systems research has focused on discharge power management in order to maximize the usable battery lifetime of a device. In order to achieve the vision of perpetual mobile device operation, we propose that software also needs to carefully consider the process of battery charging. This is because the power consumed by the system when plugged in can influence the rate of battery charging, and hence, the availability of the system to the user. We characterize the charging process of a Nexus 4 smartphone and analyze the charging behaviors of anonymous Nexus 4 users using the Device Analyzer dataset. We find that there is potential for software schedulers to increase device availability by distributing tasks across the charging period. We estimate that approximately 53% of the users we examined could benefit from up to 18.9% improvement in net energy gained by the battery while charging. Accordingly, we propose new threads of research in charging-aware power management and deferrable task scheduling that could improve overall availability for a significant portion of smartphone users.

1 Introduction

In recent years, while the electronic devices in mobile phones have been experiencing dramatic improvements in performance and capability, battery energy and power density improvements have not kept pace. This widening “battery gap” [3] has resulted in tight constraints on device power consumption to achieve satisfactory battery life. To deal with this issue, there has been a flurry of battery-aware systems research [7, 2, 8, 4] over the past decade targeting improved system lifetime through discharge power management. However, we believe that maximizing system lifetime alone does not completely satisfy users’ needs.

It has been shown that user behavior has a major influence on overall battery life [6, 1]. In the ideal case, users should never need to explicitly manage battery life;




Figure 1: Illustration of typical battery charging current and voltage characteristics.

rather, devices should work in a perpetual manner with minimal user intervention. We propose that a metric to best capture this ideal is the availability of the system to meet user needs with the highest possible quality of service (QoS). We define *availability* in this context as the proportion of time the system can deliver the subjective user-desired functionality. At a high level, device availability is a function of the holistic battery charge/discharge processes over time. We consider the net energy stored in the battery as a proxy for availability.

In general, the energy gained from a charging event depends on (1) user’s behavior (e.g., how long they stay plugged in), (2) the battery-related hardware (e.g., the power *supply*, charge controller, and battery characteristics), and (3) the non-battery system hardware and software comprising the power *load*, whose energy consumption is also directly influenced by the user, e.g., by running applications.

We propose new threads of research in *battery charging-aware* (i) *power management* and (ii) *deferrable task scheduling*. The latter is the ability to defer, split, or otherwise reschedule a non-critical and/or non-real-time task at a macro time scale to prioritize power delivery to the battery while charging. We attempt to answer the following questions about hypothetical charging-aware software solutions:




Figure 2: Nexus 4 charging characteristics from 0% to 100% SOC. The smartphone was powered on and idle for each test.

- What opportunities exist to improve overall mobile system availability?
- Under what conditions would a user benefit?
- How many users would benefit?

2 Smartphone Charging Profile

As discussed in the previous section, the energy gained from charging depends on factors such as the power supply capability, the battery characteristics, etc. A representative illustration for the typical lithium-ion (Li-ion) battery charging process is shown in Fig. 1 [5, 9]. The charging process is typically divided into two primary phases from fully depleted 0% state-of-charge (SOC) to 100% SOC. In the first phase, the device’s charge controller circuit outputs a high constant current (CC) which delivers high power to the battery. Once the battery voltage has reached a certain threshold, typically 4.2 V, the charge management controller circuit transitions to the constant voltage (CV) phase, which maintains the voltage threshold and allows current drawn by the battery to fall off gradually. Once the drawn battery current has reached a minimum level, the charge controller terminates charging to prevent battery damage from overcharging. Note that the power delivered to the battery is typically highest in the CC phase, and drops off during the CV phase.

2.1 Effect of Charger Type

To quantify the effect of different charger capabilities, we set up a testbed consisting of a Nexus 4 smartphone along with a programmable source measure unit (SMU). The power path from the supply (AC adapter or standard USB) to the battery consists of two parts. The first part is a 5 VDC path from the supply output to the power management integrated circuit (PMIC), which typically includes both the charge management controller circuit and the voltage regulators for the system’s VCC rails. The second part is a path from the charge controller to the battery. In our testbed, we utilize the SMU to measure the power in the first power path, while we query the internal charge controller circuit (Qualcomm PM8921) in the Nexus 4 to measure the power in the second power path. The chip reports the battery voltage and current to the operating system (OS) using the standard Hardware Abstraction Layer (HAL) in Android.

In our experiments, we charge the device from 0% state-of-charge (SOC) to 100% SOC via a standard 5 VDC USB plug (whose standard imposes a 500 mA current limit) and the AC to USB adapter included with the Nexus 4, which allows up to 1.2 A output at 5 VDC. The measured results are shown in Fig. 2.

2.1.1 Power Headroom

As can be seen from Fig. 1 and Fig. 2, the power drawn by the battery while charging depends on the phase of charge. Note that the maximum power of the 5 VDC supply is not drawn throughout the entire charging process. We define the instantaneous *power headroom* as the maximum power that the supply *can deliver* minus the maximum power that the battery *can absorb*. This headroom can be used to do useful work for the system load without impacting the energy gained by the battery during charging.

2.1.2 Charging Time

In order to understand how the charging process and duration depends on charger capability, we analyze the amount of time spent in the CC and the CV phases as well as the SOC at which the phase transition occurs. As shown in Fig. 2a, by charging the phone via USB, the current drawn is approximately 400 mA during the CC phase, and is fairly constant, being limited by the USB 500 mA restriction. When the CV phase starts after about 4.2 hours, the SOC is approximately 85%. The time spent in the CV phase is approximately 1.3 hours. In the AC charging experiment, there is no CC phase, due to the fact that the drawn power is limited by the battery’s ability to absorb current, not a limitation of the charger. In this case, the battery starts by drawing 800 mA and then the current decays to maintain a smooth rise in battery voltage. The battery is fully charged in 3.4 hours via AC compared to 5.5 hours in case of USB.

3 User Charging Behavior

The power headroom observed in Sec. 2 will play a significant role in the task deferral opportunities we explore in Sec. 4. Before we explore these opportunities, we now quantify how many users could benefit from charging-aware software techniques, which depends on user charging behavior. For example, if a specific user tends to




Figure 3: Mean SOC at plug-in events for each user.

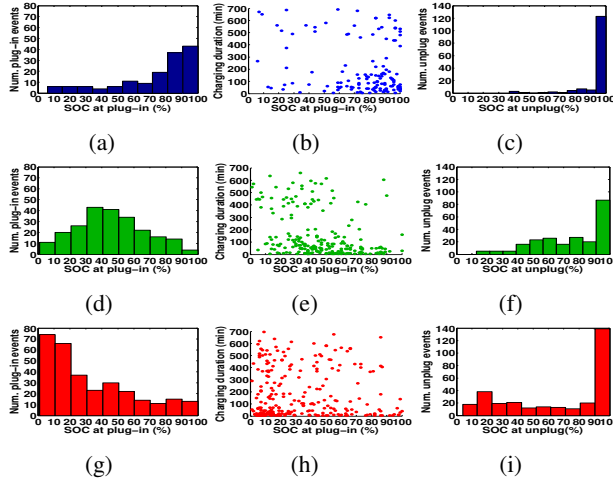


Figure 4: Charging behavior for three distinct users. Each row represents an exemplar user whose behavior follows different charging behavior trends from Classes 1, 2, and 3, which are colored accordingly as blue, green, and red.

charge their device such that it progresses through both the CC and CV phases during a single charging event, then this user might benefit from a task scheduler that accommodates time-varying power headroom. However, if a user tends to unplug the phone before entering the CV phase, then this user might not benefit from the proposed types of charging-aware techniques proposed in Sec. 4.

A user’s charging behavior can be quantified as the answer to the following statistical questions:

1. What is the SOC when the device is plugged into the supply, irrespective of when it is unplugged?
2. What is the SOC when the device is unplugged, irrespective of when it was plugged?
3. What is the charging duration for each unique plug-to-unplug charging event?

To answer these questions, we study the user charging behavior of 40 randomly chosen and anonymous Nexus 4 users over a period of roughly six months using the Device Analyzer [10] dataset.

3.1 SOC at Plug-In Event

We start by calculating the arithmetic mean value for the SOC at plug-in for each individual user’s aggregated charging events, depicted in Fig. 3. We then calculate the global arithmetic mean (the mean of the individual user means) of SOC at plug-in, shown as a horizontal line in Fig. 3. For the 40 users under consideration, the global




Figure 5: Histogram of charging duration for all charging events across all users.




Figure 6: Charging duration vs. SOC when plugged in for all charging events across all users.

Blue * represents users of Class 1. Green • represents users of Class 2. Red + represents users of Class 3.

arithmetic mean for SOC when plug-in events occur is 47%. We use this global mean to classify the users into three groups: users who tend to plug-in (1) at high SOC (60-100%), (2) around the mean SOC (40-60%), and (3) at low SOC (0-40%). Fig. 4 shows a representative user from each of the respective categories from top to bottom. In particular, Fig. 4a, Fig. 4d, and Fig. 4g show the histograms of the SOC when plugged-in for each individual user. It is clear that these three users represent the three classes of users we defined: below, around and above the global mean SOC when plug-in events occur.

3.2 Charging Duration

The data shows that the global arithmetic mean of the charging durations across all users is 120 minutes (regardless of the SOC at plug-in time), but the median charging duration is less than that. The histogram of the charging duration for all the charging events aggregated across all users is shown in Fig. 5.

The correlation between the charging duration and the SOC level at plug-in could affect our choice of how to categorize user behaviors. Thus, we aggregate all the user data for charging events into one set, and calculate the correlation coefficient between the SOC at plug-in with the charging duration. We perform the same computation for each category of the users separately. We conclude that in general, the charging duration is weakly correlated to SOC (coefficient is below 0.06) at the time the user plugs-in the phone. Thus, we do not take charging duration into consideration for categorizing users based on their charging behaviors.

The previous conclusion can also be drawn by considering the representative users from the three classes. Fig. 4b, Fig. 4e, and Fig. 4h show the charging durations for these specific users versus the SOC level when the

phone was plugged-in. We observe that across all users, charging duration tends to be low. The same trend can be observed when considering all users across the three categories as shown in Fig. 6.

3.3 SOC at Un-plug Event

We extend our analysis by considering the SOC when the phone is un-plugged, irrespective of the SOC when it was plugged-in. The data for the representative users are shown in Fig. 4c, Fig. 4f, and Fig. 4i. We observe that typically either the users let their phone charge until complete, or it coincidentally completes because the charging duration happens to be long enough (as we observed in Sec. 3.2, the charging duration is not correlated with SOC when plugged-in, which implies that charge completion is not necessarily the primary goal for users). This observation can be generalized using the same correlation calculations done in the previous experiment from Sec. 3.2. We find that in general, all three user types have similar unplugging behavior. Hence, we conclude that using the SOC when un-plugged as a parameter does not affect the charging behavior classification of users.

3.4 User Distribution

From the previous discussions, we are able to classify the users based on their charging behavior. The main criterion is to consider their average SOC when they plug in their smartphone. It is important to understand how the users are distributed across the three categories.

From the 40 users in our data set, we observed that 44% of the users tend to charge their phone when the SOC is above 60% (above the global mean, i.e., Class 1), and 47% of the users charge their phone when the SOC level is below 40% (below the global mean, i.e., Class 3). Finally, 9% of the users charge their phone in the mid range between 40% and 60% (around the global mean, i.e., Class 2).

We apply this user classification to determine the proportion of users that tend to plug in their devices at medium or high SOC. The behavior of these users tends to progress through both the CC and CV phases. Those users could benefit from deferring some tasks to CV phase where greater power headroom typically exists. According to the previous user distribution, 53% from the examined users fall into this category.

4 Opportunities for Task Deferral

From our experiments described thus far, we demonstrated the existence of power headroom during certain phases of the charging process, and concluded that 53% of users likely have their devices experience significant time-varying power headroom while plugged in, but


are currently not able to exploit it. In this section, we propose simple task deferral policies that exploits this power headroom in order to enhance the device availability. This is done by attempting to increase the net energy stored in the battery at the end of the charging event and task completion, whichever occurs later. It is given that the task must begin running after plug-in occurs. The charger used in this experiments is USB cable for Nexus 4 phone running Android 4.2.

We evaluate the simple proposed policies by manually launching an application during different phases of charge, emulating the ability of the OS to defer the task automatically. We assume that the OS would know which tasks are deferrable, and which are not, without affecting the user experience. Tasks that are interactive or otherwise time-sensitive should generally not be deferred, as this would severely affect the highly subjective device availability to the user. Modification of the OS scheduler to bin tasks based on their tolerance for deferral or other rescheduling techniques is part of our future work.

4.1 Schedule Tasks After Unplugging

One simple scheduling possibility is to run the task just after the phone is unplugged. This leaves the phone to charge at the maximum rate while plugged in, without being affected by the power consumption from the task.

To evaluate this policy, we fix the charging duration to 120 minutes (which we found as the global arithmetic mean for charging duration in Sec. 3.2) and the initial SOC to 25%. The phone runs an app that uses power-hungry GPS and Wi-Fi for ten minutes. In the first sub-experiment, the app is run while the phone is charging, while in the second sub-experiment, we manually defer the same app until the phone is unplugged. For both cases, we calculate the energy gained by the battery by observing the current and the voltage of the battery over time. A representative run of the experiments are shown in Fig. 7a and Fig. 7b. Fig. 7c shows the net energy gained by the battery in both cases. We observe that deferring the task to after the phone is un-plugged performs worse compared to running the task while it is charging.



(a) App runs while phone is charging. (b) App runs after unplugging phone. (c) Energy gained by the battery.

Figure 7: Effect of running an app during vs. after the charging period.

4.2 Schedule Tasks Within the Constant Current Phase

In this experiment, we schedule the same task at different stages of the CC phase in order to determine whether it will affect the net energy gained by the battery by the end of the charging period.




Figure 8: Effect of running an app early vs. late in the CC phase.

As in the previous experiment, we run sub-experiments to compare between the two scenarios. We fix the initial SOC to 70% SOC and the charging duration to 70 minutes. In the first sub-experiment, we run the app beginning at ten minutes after the plug-in event as shown in Fig. 8a. In the second sub-experiment, we start the same app ten minutes before the beginning of the CV phase as shown in Fig. 8b. We find a small improvement in net energy gain that occurs due to deferring the task to later in the CC phase, as shown in Fig. 8c.

4.3 Schedule Tasks in the Power Headroom

A third possibility is to defer the task to the CV phase, where the greatest power headroom is present. In this experiment, we fix the initial SOC to 78 (CC phase) and the charging duration to 120 minutes (which ensures that the phone hits the CV phase). We run the app once in the CC phase (Fig. 9a) and once on the CV phase (Fig. 9b). The CV phase starts once the peak battery voltage is reached. From Fig. 9c, we observe an 18.9% increase in the energy gained by the battery in the latter case.




Figure 9: Effect of running an app in the CC vs. CV phases.

This can be explained if we look carefully in Fig. 9a and Fig. 9b. In the former case the current drops approximately from an average of 400 mA to 150 mA. However, in the latter case the current drops from an average of 300 mA to 100 mA. The difference in the drop between

the two cases can be explained by the power headroom present in the CV phase at this time. Assuming that the maximum deliverable power from the supply and the app power demand are constant, this means that a greater part of the load current was drawn from the supply and not from the battery during the CV phase.

5 Conclusion and Future Work

We present a case for battery charging-aware power management and deferrable task scheduling to improve overall device availability. In particular, we propose to utilize the power headroom during certain phases of battery charging to run these tasks, rather than starve the battery of energy during its most power-intensive charging time. Increasing the energy delivered to the battery during the charging period, or conversely, decreasing the required charging duration to reach full SOC would improve overall device availability to the user.

Our study on Nexus 4 smartphone user charging behavior shows that most users tend to charge their phone for less than 120 minutes, and that the charging duration is largely independent of the SOC when the smartphone is plugged in or unplugged. We estimate that around 53% of users could benefit from battery charging-aware software policies that maximize energy delivered to the battery while plugged in.

We observe mixed results for the different simple proposed charging-aware task deferral policies. Positively, we find that deferring tasks to the CV phase can improve the net energy gained by the battery by approximately 18.9%. In contrast, we observe that deferring tasks until the end of the charging period or deferring the tasks within the CC phase leads to no significant net battery energy increase. All of the proposed scheduling schemes could be implemented in a smartphone OS via online monitoring of the power headroom, and they represent only a small portion of the possible scheduling policies.

Our future work seeks to pursue the ideas around power headroom and user charging behavior further. Quantifying power headroom based on the battery characteristics and the stage of the charging process is essential to determine the number and type of tasks to be deferred based on their predicted energy requirements. A task could also be split between two phases based on the amount of headroom available and the energy requirement of the task. Furthermore, charging-aware software might have even greater potential to improve availability in systems with heterogeneous energy storage architectures that require careful energy management. Finally, it would be useful to predict whether a given user at during some charging event is likely to reach a period with greater power headroom, using factors beyond those explored in this paper. Accordingly, our future work may include a study of user-specific models in this direction.

References

- [1] FALAKI, H., MAHAJAN, R., KANDULA, S., LYMBEROPOULOS, D., GOVINDAN, R., AND ESTRIN, D. Diversity in Smartphone Usage. In *Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services (MobiSys)* (New York, New York, USA, June 2010), ACM Press, p. 179.
- [2] FLINN, J., AND SATYANARAYANAN, M. Managing Battery Lifetime With Energy-Aware Adaptation. *ACM Transactions on Computer Systems* 22, 2 (May 2004), 137–179.
- [3] LAHIRI, K., RAGHUNATHAN, A., DEY, S., AND PANIGRAHI, D. Battery-Driven System Design: A New Frontier in Low Power Design. In *Proceedings of the Asia and South Pacific Design Automation Conference (ASP-DAC)* (Jan. 2002), IEEE Computer Society, p. 261.
- [4] LI, J., QIU, M., NIU, J.-W., AND CHEN, T. Battery-Aware Task Scheduling in Distributed Mobile Systems With Lifetime Constraint. In *Proceedings of the Asia and South Pacific Design Automation Conference (ASP-DAC)* (2011), IEEE, pp. 743–748.
- [5] PANASONIC CORPORATION. Lithium Ion Batteries Technical Handbook, 2007.
- [6] RAHMATI, A., AND ZHONG, L. Human-Battery Interaction on Mobile Phones. *Pervasive and Mobile Computing* 5, 5 (Oct. 2009), 465–477.
- [7] RAO, R., VRUDHULA, S., AND RAKHMATOV, D. Battery Modeling for Energy-Aware System Design. *IEEE Computer* 36, 12 (2003), 77–87.
- [8] RAVI, N., SCOTT, J., LU, H., AND IFTODE, L. Context-Aware Battery Management for Mobile Phones. In *Proceedings of the International Conference on Pervasive Computing and Communications* (2008), IEEE, pp. 224–233.
- [9] SONY CORPORATION. Lithium Ion Rechargeable Batteries: Technical Handbook, 2014.
- [10] WAGNER, D. T., RICE, A., AND BERESFORD, A. R. Device Analyzer: Understanding smartphone usage. In *Proceedings of the International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services* (Tokyo, Japan, 2013), ACM.