3 W's of smartphone power consumption: Who, Where and How much is draining my battery?

Agrim Gupta*, Adel Heidari*, Avyakta Kalipattapu*, Ish Kumar Jain, Dinesh Bharadia
University of California, San Diego
{agg003,adheidari,rkalipattapu,ikjain,dbharadia}@ucsd.edu

Abstract

With 6.5 billion smartphones in use worldwide, each relying on a battery for key subsystems like display, compute, and cellular connectivity, previous studies on power consumption often used invalidated indirect estimates that failed to isolate specific hardware usage. We address this by utilizing Google's On Device Power Rails Monitor (ODPM) tool for precise power measurements of individual components. Our findings indicate that connectivity (Wi-Fi, 4G/5G) and screen display are the primary power consumers, as shown with the Google Pixel 7A. We also confirmed similar power consumption trends using an energy estimation method on the Samsung S23+. Given the prevalence of smartphones, we discuss the challenges and opportunities for optimizing power usage.

CCS Concepts

• **Hardware** → **Power and energy**; **Hardware test**; **Wireless devices**; • **Networks** → **Network components**.

Keywords

5G NR, Green communications, 6G, Cellular communications, Small cells, Mobile measurement

ACM Reference Format:

Agrim Gupta*, Adel Heidari*, Avyakta Kalipattapu*, Ish Kumar Jain, Dinesh Bharadia. 2024. 3 W's of smartphone power consumption: Who, Where and How much is draining my battery?. In *The 30th Annual International Conference on Mobile Computing and Networking (ACM MobiCom '24), November 18–22, 2024, Washington D.C., DC, USA.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3636534.3695905

1 Introduction

Smartphones have been one of the most important technical inventions of the past decade, which has almost led to universal adoption. As of 2024, the population of the entire world is about 8.1 billion, and the total number of smartphones estimated stands around 6.5 billion, about 80% of the entire world's population [1]. As shown in Fig. 1, this explosive growth has come at a large cost of embodied carbon footprint, as well as e-waste, to which smartphones are the highest contributors [2]. Although components like screen, memory, frame, can be repaired or recycled, batteries need to be 'replaced' as the smartphone life gets older as their batteries need to be charged on a daily basis, which has also recently shown to have an adverse effect on the environment [3]. Clearly, the battery is a crucial area of study if we aim to achieve sustainable growth in smartphones. Before suggesting any effective solutions to reduce battery power drain, it is essential to thoroughly understand **Who** is consuming the power at the component level, **Where** this power is being used at the application level, and **How much** amount each component is using.

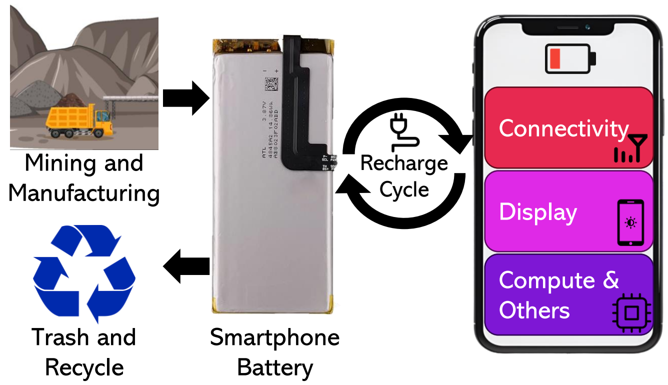

Figure 1: Each smartphone is powered with a rechargeable battery which powers various components like connectivity, display and computation.

Hence, it is consequential to understand what draws out battery on smartphone today. Smartphone consists of various hardware subsystems like, 'display' which powers up the smartphone screen, 'compute' which runs the operating system and various applications, 'connectivity' which consist of cellular/WiFi radios that connect the smartphone to internet. Almost every smartphone is able to estimate the total battery discharge, and some smartphones also enlist the battery consumed by individual 'applications'. However a clear breakdown of power due to the individual 'hardware components' has been difficult to measure.

Related work: There have been prior measurement studies [4–6] which measure power of these individual components via indirect estimation methods, which perform power subtraction between the total powers computed in a typical mode where the hardware component is operated normally, versus a baseline mode, where it is turned off. Due to the involved complexity, these indirect methods lack validation, and their accuracy is unclear. For instance, display power consumption has been traditionally attributed to be one of the highest contributors [7], however, moving from 4G to 5G, cellular modem power consumption has increased, and is more comparable to display power consumption [4].

In this paper, we first showcase measurements from the newly released Google's On Device Power Rails Monitor (ODPM) tool, which is able to measure 'direct' fine-grained power measurements. Hence, using ODPM tool, we are able to obtain power measurements of all different components like compute, display, camera, sensor systems, modem and memory access as well (**Who in 3 W's**). Using ODPM, we observe that across most of the popular applications and smartphone use-cases (**Where in 3 W's**), that the connectivity and display components consistently consume the highest power (**How much in 3 W's**). Further, for the first time, we are also able to validate, as well as profile the accuracy of previously studied indirect estimation methods by baselining them against the ODPM oracle measurements on Google Pixel phone by considering cellular power usage as the most significant contributor. Then, we extend these measurements to another smartphone (Samsung Galaxy S23), to show how these

*These authors contributed equally to this work
Paper website: https://wcsng.ucsd.edu/ue-power/

ACM MobiCom '24, November 18–22, 2024, Washington D.C., DC, USA

ACM ISBN 979-8-4007-0489-5/24/11
https://doi.org/10.1145/3636534.3695905

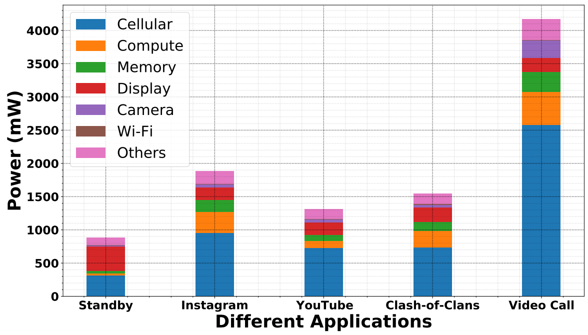

Figure 2: Average power consumption computed over 1-minute uses for different activities, like watching a 720p video on YouTube, doing a video call via Google Meet, scrolling on Instagram feed, playing Clash-of-Clans and keeping the phone idle in standby.

observations around connectivity power consumption and display generalize to different chipsets. We conclude with a brief discussion on further trends/opportunities in power consumption breakdown of smartphones.

2 Power Consumption Study

In Section 2.1, we explain how the On-Device Power Measurement (ODPM) system on Google Pixel phone is used to obtain fine-grained power data from individual hardware components. This allows us to analyze power consumption trends for popular applications (e.g., standby, Instagram, YouTube) and break down consumption by hardware. We identify two primary sources of power consumption: the cellular modem and the display. To generalize these findings to other devices, we profile an indirect measurement approach using total battery discharge with ODPM, demonstrating its accuracy in Section 2.2. In Section 2.3, we present measurements from the Samsung S23+, which features different chipsets compared to the Google Pixel.

2.1 Pixel 7A ODPM Power Measurements

To capture information from these ODPM rails[8] on Google Pixel phones, we can run a system trace (using the open source Perfetto tool) to collect the power measurement for a certain amount of time. While the Perfetto trace runs in background, one can use the smartphone as usual for different applications. With the trace running in the background for a set time of 1 minute, we use different applications, like YouTube, Video Call using Google Meet, Instagram, Clash-of-Clans, and also let the phone be idle in standby. As observed from Fig. 2, ODPM is able to capture different subsystems power, since the camera power consumption increases during video call, compute power consumption increases while playing Clash-of-Clans game (compared to YouTube), and overall the power reduces in the idle standby mode. For all these usecases, we observe that cellular and display stand out as the two largest power consuming components.

2.2 Measuring Cellular Power w/o ODPM

Unfortunately, ODPM tool is only available on Google Pixel phones 6 and above, and they all have the same cellular modem (Samsung Exynos 5300), and hence, we need to make similar measurements on a different modem to confirm the observations. This requires a measurement technique that can estimate the modem power consumption without relying on ODPM.

We do this by measuring net battery discharge, which is reported by almost every android phone. To isolate modem and display powers from the net battery discharge, we take baseline measurements in airplane mode (where modem is turned off). Then, for individual measurements, like uplink/downlink data movement using iPerf, we perform subtraction of the baseline power to obtain the estimates for modem power. However, a key difference between this work and prior works which used similar baseline suppression method [4] is that we also show accuracy of this method,

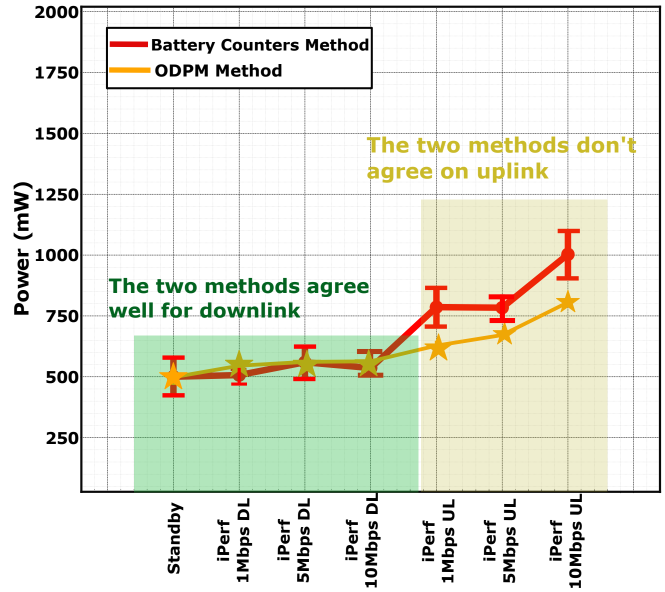

Figure 3: Cellular power consumption of Google Pixel 7A using both battery counters and ODPM methods across various experimental categories. These categories encompass standby mode, uplink and downlink data transmission over a 5G base-station setup.

by comparing the results with ground truth values obtained from ODPM measurements. The estimates of modem/display power and the ground truth from ODPM are visually illustrated in Fig. 3. From Fig. 3, we can see that the estimation strategy works fairly well for downlink and standby modem power measurements. However, as observed, for uplink the battery discharge estimates for modem power are not accurate, and a reason for this could be larger instantaneous current draws for power amplifiers during uplink that might not get captured via total battery discharge. Similar to airplane mode for modem power, to compute display power, we capture battery counter measurements when the screen is turned off. Hence, using battery discharge, we can estimate the downlink modem and display powers by base lining the airplane and screen off power measurements respectively.

2.3 Samsung S23+ Power Measurements

In this section, we utilize the battery discharge method to estimate modem and display components power for Samsung S23+, and compare with measurements on Google Pixel 7A. We perform this experiment for standby mode, and for streaming a 720p youtube video at the same nominal display brightness. To ensure fairness in the comparison, both phones were equipped with the same commercial SIM card and operated on the same wireless band. Additionally, both phone experiments were conducted at the same location to maintain same quality of connection to the base-station. As depicted in Fig. 4, the results show that 5G does consume a higher (about two times) standby power consumption than 4G. Also, we can observe that Samsung S23+ with Qualcomm X70 modem has lower power consumption than Google Pixel 7A with Samsung Exynos 5300. Since S23+ has a larger AMOLED screen compared to Pixel 7A's smaller OLED screen, the display power consumption is higher for S23+.

Next, we show youtube power consumption results in Fig. 5. Note that YouTube data consumption is majorly downlink, for which the battery discharge approach works reasonably accurate. We can also see here that S23+ with improved modem chipsets for both Wi-Fi and Cellular connectivity has lower power consumption as compared to Pixel 7A. Overall, across both standby and YouTube measurements, we could see that the improved modems in newer phones do help reduce the connectivity power consumption. However it still forms a major bulk of battery draw and is comparable to that of display.

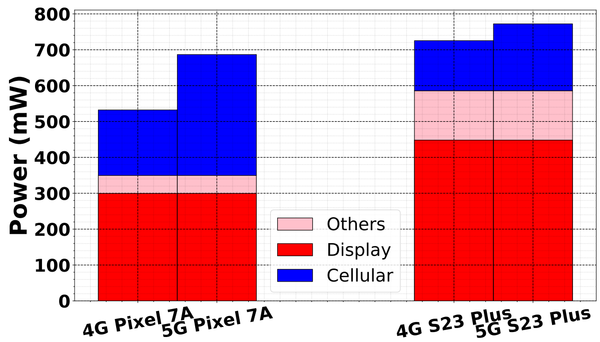

Figure 4: Assessing the efficacy of the battery discharge method for two distinct Google Pixel 7A and Samsung S23+ phones maintained in standby mode, connected to 4G-LTE and 5G networks, and delineating the power usage into display, cellular, and other components.

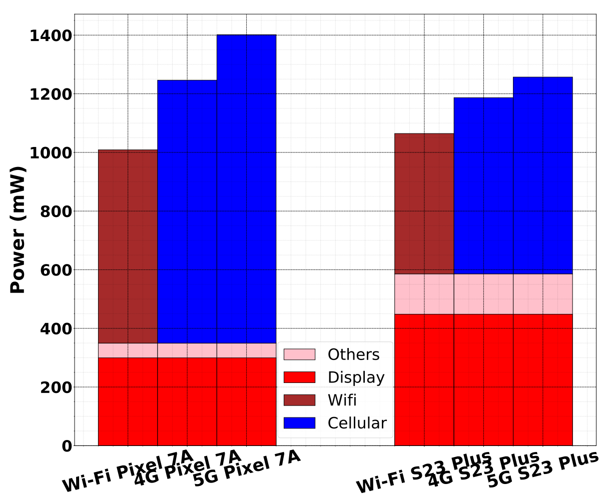

Figure 5: Power consumption during 720p YouTube video streaming over Wi-Fi, 4G-LTE and 5G cellular networks for both the Google Pixel 7A and Samsung S23 Plus phones. Additionally, replicating the measurements for the same activity via Wi-Fi.

3 Discussion and Future Work

(a) Implications on Edge Compute: The measurements showcased in this work have a huge implication on edge compute. With rise of AI services, like Galaxy AI on S24 and Gemini on Pixel 8, it might be possible that an online server inference requiring connectivity may be worse for energy than an on-device inference with edge compute. This also requires further work on profiling compute energy for more complicated tasks like AI inferences, and exploring tradeoffs between the communication and computing energy costs. Some initial studies assuming WiFi communication to cloud have already started to surface up [3], however, a much careful look at cellular power consumption is needed.

(b) Devising Energy Currency unit, enforcing ODPM rails across different phones: We emphasize again that smartphone as a technology has achieved an unprecedented scale (6.5 billion smartphones [1] vs 8 billion world population today). Hence, ODPM profiling needs to be mandated across smartphones, and concept of energy currency needs to be devised for applications.

(c) Future smartphone innovations targeted towards modem energy optimizations: Wi-Fi still consumes lesser power than cellular networks while they transfer same amount of data. With upcoming innovations like 5G redcap [9], smart network configurations [10], this differential can be reduced, and these research efforts can also benefit from more detailed smartphone power breakdown studies.

Acknowledgment

This work was supported in part by the U.S. National Science Foundation (NSF) under Award No. 2030245. Further, the authors acknowledge the anonynmous reviewers and WCSNG group members for their feedback.

References

[1] Total number of smartphones, big data. https://whatsthebigdata.com/smartphone-stats/.

[2] Jennifer Switzer, Gabriel Marcano, Ryan Kastner, and Pat Pannuto. Junkyard computing: Repurposing discarded smartphones to minimize carbon. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, pages 400–412, 2023.

[3] David Patterson, Jeffrey M Gilbert, Marco Gruteser, Efren Robles, Krishna Sekar, Yong Wei, and Tenghui Zhu. Energy and emissions of machine learning on smartphones vs. the cloud. *Communications of the ACM*, 67(2):86–97, 2024.

[4] Dongzhu Xu, Anfu Zhou, Xinyu Zhang, Guixian Wang, Xi Liu, Congkai An, Yiming Shi, Liang Liu, and Huadong Ma. Understanding operational 5g: A first measurement study on its coverage, performance and energy consumption. In *Proceedings of the Annual conference of the ACM Special Interest Group on Data Communication on the applications, technologies, architectures, and protocols for computer communication*, pages 479–494, 2020.

[5] Monsoon power monitor. https://www.msoon.com/high-voltage-power-monitor.

[6] Matteo Varvello, Kleomenis Katevas, Mihai Plesa, Hamed Haddadi, and Benjamin Livshits. Batterylab, a distributed power monitoring platform for mobile devices. In *Proceedings of the 18th ACM Workshop on Hot Topics in Networks*, pages 101–108, 2019.

[7] Noor Zaman and Fatimah Abdulaziz Almusalli. Smartphones power consumption & energy saving techniques. In *2017 international conference on innovations in electrical engineering and computational technologies (ICIEECT)*, pages 1–7. IEEE, 2017.

[8] Google odpm power rails and power profiler. https://developer.android.com/studio/profile/power-profiler.

[9] Sandeep Narayanan Kadan Veedu, Mohammad Mozaffari, Andreas Höglund, Emre A Yavuz, Tuomas Tirronen, Johan Bergman, and Y-P Eric Wang. Toward smaller and lower-cost 5g devices with longer battery life: An overview of 3gpp release 17 redcap. *IEEE Communications Standards Magazine*, 6(3):84–90, 2022.

[10] Anum Ali, Yuqiang Heng, Vutha Va, Priyabrata Parida, Boon Loong Ng, and Jianzhong Charlie Zhang. Ue power saving with traffic classification and ue assistance. *IEEE Access*, 2023.