Analyzing the Power Consumption of the Mobile Page Load

Yi Cao Javad Nejati Pavan Maguluri Aruna Balasubramanian Anshul Gandhi
Stony Brook University
{yicao1, jnejati, smaguluri, arunab, anshul}@cs.stonybrook.edu

1. INTRODUCTION

A large portion of today’s Web traffic comes from mobile devices. User experience of mobile Web pages has become an important factor that content providers and service providers want to improve. One metric for evaluating user experience is the mobile Web performance, specifically the page load time (PLT). While improving mobile Web performance has drawn increasing attention [2], most optimizations tend to overlook an important metric, *energy*. PLT and energy are distinct metrics; we observe in our experiments that these metrics may not always be correlated. Given the importance of battery life for mobile users, we argue that Web page optimizations should also be evaluated for their impact on energy consumption.

However, examining the energy effects of a Web optimization is challenging, even if one has access to power monitors, for several reasons. First, Web page loads experience high variance due to network conditions and device scheduling. Our own experiments show that the energy consumption of a Web page can vary by about 30% under the same experimental setup. Second, Web pages can be highly complex since they are composed of distinct *components*, which are individual page load activities such as loading objects, parsing the page, or evaluating Javascript. There are also many levels of dependencies among those components during the page load process. Without detailed application-level semantics, it is hard to tease out the energy contribution of each component. Third, the page load process is short-lived, ranging from several milliseconds to a few seconds. Although different power models [3, 5] for mobile devices have been proposed, they do not work well for such short-lived applications. This is because such models require fine-grained *resource monitoring*, which incurs substantial overhead and impacts model accuracy.

2. RECON

We present RECON(RESource- and COMpoNent-based modeling), a model that addresses all of the challenges out-

Figure 1: Using WProf-M to decompose the components when loading the *instagram.com* Web page.

lined above to accurately predict the energy consumption of any mobile Web page load. The design of RECON is motivated by the fact that resource monitoring incurs significant overheads at fine-grained time scales, and does not provide insights into the page load process. Instead, RECON collects resource-level data at *coarse* granularity, and then exploits low-level *page load semantics*. RECON extracts page load semantics using WProf-M [4], a tool that decomposes the page load process into various components. In effect, RECON augments coarse-grained resource monitoring with component-level information, resulting in high prediction accuracy and providing benefits over individual resource-only or component-only models.

Our modeling approach breaks down the page load process into “segments”, where a segment is defined as an interval of page load activity during which the components of the Web page do not change. By definition, a segment is composed of at least one component. Further, the entire page load process can be partitioned into discrete (non-overlapping) segments. Figure 1 shows how we partition a page load into segments using WProf-M’s component-level decomposition. We then model the power consumption of a segment, s , as:

$$\hat{P}_s = w_0 + \sum_{i \in \text{Resources}} w_i R_i + \sum_{j \in s} w_j, \quad (1)$$

where i represents the various resources we monitor, including CPU utilization and frequency, bytes and packets transmitted through the WiFi interface; R_i represents the average value for that resource during s ; j represents a component, and thus $j \in s$ is the set of all components that make up segment s ; finally, w_i and w_j are coefficients (independent of s) representing the power contribution of the resources

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGMETRICS '16 June 14-18, 2016, Antibes Juan-Les-Pins, France

© 2016 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-4266-7/16/06.

DOI: <http://dx.doi.org/10.1145/2896377.2901491>

Figure 2: Actual vs. predicted energy consumption for 6 popular Web pages.

and components. w_0 is also a coefficient, and represents the baseline power consumption of the phone which could be attributed to background activities, screen brightness, etc. Our next task is to determine the values of these coefficients.

We use multiple linear regression to derive the weights vector, $\vec{w}$, that is indicative of the power contribution of each component and resource, and the baseline power. Once we derive the weights, we can compute the energy consumption of the segment by multiplying the predicted power with the segment length (obtained via WProf-M); summing over all segments gives us the Web page energy consumption.

In the training period, we monitor the ground-truth power consumption by using the Monsoon Power Monitor [1]. We train our regression model on nine instantiations of a given Web page and test on the tenth instantiation. Our error results are based on a 10-fold cross validation.

3. RESULTS

We validate RECON on the Samsung Galaxy S4 and S5 devices using 50 popular Web pages chosen from *alexa.com*. Our results show that RECON’s average prediction error for energy consumption across these Web pages is 4.48%, and the average segment-level energy consumption prediction error is less than 8.88%. Figure 2 shows the actual and predicted mean energy consumption for 6 popular Web pages, and the confidence intervals around the predicted mean.

4. APPLICATIONS: OPTIMIZATION

An important application of RECON is studying the energy effects of Web optimizations. To that end, we consider four popular optimizations: *Compression*, *Minification*, *Inlining* and *Ad-block*. Our mean prediction error for Web page energy consumption under these optimizations is a low 5.45%. In addition to accurately predicting the energy effects of an optimization, RECON also provides Web developers with several useful applications in optimization:

4.1 Breaking down total power consumption

RECON decomposes the total energy consumption of a Web page into component-level energy consumption, so that we know each component’s contribution to the total energy consumption. For example, Figure 3 (left) shows the predicted per-component energy usage for *alexa.com*. Further, RECON allows Web developers to visualize how an optimization affects component-level energy consumption. Figure 3 (right) shows the change in predicted per-component energy usage under compression for *fico.com*. We see that, after compression, the energy consumption of evalhtml and Javascript increases, indicating that compression is not useful for these components. On the other hand, the energy

Figure 3: Left: predicted per-component power consumption for *alexa.com*; Right: changes in per-component energy consumption for *fico.com* due to compression.

consumption for text decreases, suggesting that text compression is useful.

4.2 Analyzing why energy consumption changes

With the help of RECON, we can explore not only *how* but also *why* energy consumption changes due to an optimization. For example, the inlining optimization embeds all the external Javascripts and CSS in the root HTML file. Inlining avoids small downloads, so the energy consumption due to downloading decreases. However, inlining lengthens the HTML evaluation, which, according to RECON, is a power-hungry component. In our experiments with *netflix.com*, we found the total energy consumption increases on inlining although the PLT decreases.

5. CONCLUSION

RECON provides Web developers with quick and accurate power predictions of the mobile page load. Moreover, RECON helps users explore how and why energy consumption changes due to an optimization. Our end-goal in this research is to enable and promote simultaneous energy- and performance-analysis of Web pages; RECON is an important step in this direction.

Acknowledgements

We gratefully acknowledge the anonymous reviewers for their feedback. This work was supported by the National Science Foundation through the grant CNS-1551909 and a Google Research Grant R2-2015-839.

6. REFERENCES

- [1] Monsoon Power Monitor. <http://msoon.github.io/powermonitor/>.
- [2] M. Butkiewicz, D. Wang, Z. Wu, H. V. Madhyastha, and V. Sekar. Klotski: Reprioritizing Web Content to Improve User Experience on Mobile Devices. In *NSDI 2015*, pages 439–453.
- [3] X. Chen, N. Ding, A. Jindal, Y. C. Hu, M. Gupta, and R. Vannithamby. Smartphone energy drain in the wild: Analysis and implications. In *SIGMETRICS 2015*, pages 151–164.
- [4] X. S. Wang, A. Balasubramanian, A. Krishnamurthy, and D. Wetherall. Demystifying Page Load Performance with WProf. In *NSDI 2013*, pages 473–485.
- [5] L. Zhang, B. Tiwana, Z. Qian, Z. Wang, R. P. Dick, Z. M. Mao, and L. Yang. Accurate online power estimation and automatic battery behavior based power model generation for smartphones. In *CODES+ISSS 2010*, pages 105–114.