

Experiences in Logging Everyday App Use

Marek Bell¹, Matthew Chalmers¹, Lucas Fontaine¹, Matthew Higgs², Alistair Morrison¹,
John Rooksby¹, Mattias Rost¹, Scott Sherwood¹

¹School of Computing Science, University of Glasgow ²Department of Statistical Science, UCL.
¹{firstname.lastname}@glasgow.ac.uk ²M.Higgs@ucl.ac.uk

ABSTRACT

This paper discusses our experiences in logging app use on computers, mobile phones and tablets. We have created logging software to record app launches on iOS, Android and Mac OS X devices, and have used this in a study with 13 students over a period of one month. This paper discusses the practicalities of logging across multiple devices, the forms of enquiry suitable for log analysis, and the ethics of logging. We also discuss future work in which we will scale the study up to thousands of users.

Categories and Subject Descriptors

H.5.m [Miscellaneous]: Information Interfaces and Presentation (e.g. HCI) Miscellaneous.

General Terms

Measurement, Experimentation, Human Factors.

Keywords

Mobile devices, Logging, Apps.

1. INTRODUCTION

Mobile devices (including laptops, phones and tablets) are central to day-to-day life for many people. Many HCI studies have studied usage of individual apps (e.g. [4]), but such individual logging does not say a great deal about the holistic use of a device. Services such as Flurry (www.flurry.com) have been able to go some way to building pictures of overall device use by virtue of their logging software being built into many thousands of individual apps, but only have data on those apps that incorporate the Flurry service. More recently, several quantitative studies have been published [1][2][3] that create a fuller picture by examining all apps launched on mobile devices. For example, Böhmer et al. [1] have logged and analysed app launches on over four thousand Android devices, and Do et al. [2] have analysed logs provided by Nokia of 111 of their users.

In this paper we discuss our own experiences in holistically logging app use on mobile devices. Our work extends previous quantitative studies by logging app use across several operating systems. Ours is the first study, as far as we are aware, to log the use of all apps launched on iOS devices. We have also logged launches on Mac OS X and Android devices. This enables wider coverage, but also enables us to examine how people use devices

in tandem (for example how someone switches between an iPad, an iPhone and a MacBook). Our study can also be differentiated from prior work by our use of qualitative interviews in conjunction with quantitative analysis, and our concern with the user experience and the ethics of logging.

This paper reports on the development and use of a suite of logging applications known collectively as AppTracker. The paper discusses a study of thirteen students over a one-month period. This study is a precursor to a large-scale trial we will run with many thousands of users, and therefore this paper makes no attempt to generalise about populations. Instead it addresses issues associated with the practicalities of logging, directions for analysis, and ethics.

2. APPTRACKER

We have developed logging applications for three platforms: iOS (i.e. iPod, iPad, iPhone), Android, and OS X. AppTracker logs the time when apps are opened and closed, and when devices are locked and unlocked. We also log the connectivity of the device and when it is being charged. We do not log anything that is done within individual apps, or network activity – only when apps are opened and closed. We have taken a conscious decision not to log websites visited on the device or location information. Previous work by Böhmer et al. [1] and Do et al. [2] has shown such information to be useful, but we were uneasy about the ethical implications of this (and would ourselves not be comfortable revealing such information). The logs are stored on a database. To implement AppTracker, we had to resolve a number of technical and conceptual challenges. These included:

- Backgrounding is not readily supported in iOS. We therefore had to devise an approach that would enable a logger to run unobtrusively over the long term. To our knowledge ours is the first log based analysis of iOS devices.
- Logging inevitably has an effect on the performance of the devices. We therefore had to minimise battery consumption and data transmission.
- Working across platforms required us to address how a consistent log can be generated from diverse devices.
- Clearly, there are also ethical considerations. We minimised ethical concerns by choosing not to log location or any content. We treated ethics as an ongoing topic during the research (for example by discussing privacy issues with participants during our interviews).

3. PILOT STUDY

For the study we recruited 13 students at the University of Glasgow. We asked them to run AppTracker on one or more device for a minimum of 30 days, to complete a questionnaire and to sit for a semi-structured interview with a researcher. This was an exploratory study designed as a first foray for AppTracker in

Permission to make digital or hard copies of all or part 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. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Digital Economy '13, Nov 4–6, 2013, Salford, UK.

Copyright 2013 ACM 1-58113-000-0/00/0010 ...\$15.00.

the wild (beyond our own devices) and to refine our focus and approach for larger trials. Our key findings are as follows:

3.1 Patterns of Use

Our logs appear to confirm the point made by Böhmer et al. [1] and Do et al. [2] that there are patterns of app use (individually as well as in terms of ‘operational chains’). For example, it appears that social apps are used frequently and often, whereas other apps (particularly games) are used less frequently but for longer periods of time. Figure 1 gives a summary of how two participants used their devices during the study. It is important to acknowledge that our logs can never give the full picture (the foregrounded app on a laptop may not actually be the one in use, and mobile apps that work in the background, such as Spotify, are potentially misrepresented as the time spent in the app may be much shorter than the app is actually in use for). Our analysis so far has been simple. Figure 1 does not display chains of app use, or temporal issues over weeks and months. Our interviews point to many issues that our quantitative analysis has yet to address. We are currently developing statistical methods to look at app use in more detail.

3.2 Pictures of Life

Mobile devices are used in routine or patterned ways to the extent that simple plots of when a device is in use give an indication of people’s lives. For example, a comment made by nearly all participants about plots of their phone data was that it is possible to see when they were sleeping (Figure 1 gives an example of this). Comments about tablet computers included that it is possible to see when the participant was in bed or getting home from work (situations in which tablets are seemingly used). Participants’ comments about plots of computer use included that you can see when they were working (and not working). Combined with other information, such as when a device is being charged, when Wi-Fi networks are joined and left, and when particular apps (such as alarms or travel apps such as maps, bus or rail timetables, and taxi-booking services) are used, then it is possible to build compelling pictures of life. For this reason, we concur with Morrison et al.’s [4][5] point that the ethics of logging mobile devices requires serious attention. The issue is not that these are accurate pictures of life, but that they are interpretable (and misinterpretable) as speaking of people’s lives. We showed our participants visualisations of their data and asked whether they would be happy to share these with friends, family and staff at university. All participants were less troubled by this than we expected, but perhaps had not had time to think through the implications (and had of course volunteered in the first place to be logged).

3.3 Limitations of Logging

It is important to temper arguments about ‘patterns of use’ and ‘pictures of life’ with the observation that our logs are not logs of human action but logs of the device. An app being open does not mean it is in use, and it does not mean that people are performing activities representative of the category by which apps are often discussed. Firstly, things like docking a device or installing flavours of Android such as Facebook Home can mean apps can remain open for long periods while the device is put down. Computers can also remain on for long periods while the user is away from it. Even where an app is used for seconds or minutes it is not necessarily the case that it is being used. It is also apparent that categories such as “communication”, “social networking” and “gaming” serve to order app stores but do not describe use. In particular, it is clear people happily use several apps of one

category (e.g. social networking apps including WhatsApp, Facebook, Facebook Messenger, and Snapchat) and that taking these as indicative of a category of use fails to address their individual patterns. We suggest therefore that statistical methods are used with caution, and that qualitative and conceptual work is also important in this area.

3.4 Impact of Logging

Logging can affect the battery consumption of a device, and to a lesser extent has implications for speed and data transmission. It became clear during our study that performance is very much an issue of perception, and that different people hold different expectations about this. The participants predominantly discussed performance in terms of time between charges, and would complain if apps appear slow to load. Different participants seemed to have different needs and expectations regarding performance. Some participants, notably computer science students, would charge their devices frequently. Others wanted two to three days between charges. Some participants would often be on Wi-Fi (and thus data could be transmitted at no cost) whereas others would not join Wi-Fi even where this was freely available to them (e.g. on campus). It was also apparent that not everyone made a direct association between individual pieces of software and performance (not just AppTracker, but, for example, the use of fitness apps). They would see performance as something attributable to the device itself rather than apps. We became aware in our study that the ways in which we framed questions about performance, or revealed possible performance issues related to AppTracker, would affect what people said on this topic. This raises interesting questions, and makes it difficult to address the issue of whether logging has an effect on use.

4. FURTHER WORK

This paper reports a relatively small study with 13 participants over one month. We now intend to run large-scale deployments of the logging software, releasing it via app stores. In these trials, we will begin by supporting self-reflection by enabling users to explore visualisations of their own data. We are currently working with statisticians to develop methods and, ultimately, online services for the analysis of AppTracker data. As we have noted, it is appropriate to integrate quantitative and qualitative methods and therefore we propose to follow Morrison et al.’s [6] hybrid approach to mass participation trials.

Further down the line we propose that AppTracker can be incorporated into several kinds of service, including 1) tools for self-reflection and behaviour change; 2) social tools for awareness among friends or co-workers; 3) diagnostic frameworks for troubleshooting and reasoning about performance. We also feel it is important to address the ethical issues in logging, these being issues that other studies seem to have overlooked or played down.

5. REFERENCES

- [1] Bohmer, M., Hecht, B., Schoning, J., Kruger, A., and Bauer, G. Falling asleep with Angry Birds, Facebook and Kindle: a large scale study on mobile application usage. *Proc. MobileHCI '11*, 47–56 (2011).
- [2] Do, T.-M.-T. and Gatica-Perez, D. By their apps you shall understand them: mining large-scale patterns of mobile phone usage. *Proceedings of the 9th International Conference on Mobile and Ubiquitous Multimedia*, 27:1–27:10 (2010).

[3] Hang, A., De Luca, A., Hartmann, J., and Hussmann, H. Oh app, where art thou?: on app launching habits of smartphone users. *Proceedings of the 15th international conference on Human-computer interaction with mobile devices and services*, 392–395 (2013).

[4] Morrison, A. and Chalmers, M. SGVis: Analysis of Data from Mass Participation UbiComp Trials. *International Journal of Mobile Human Computer Interaction (IJMHCI)*, 3(4), 36–54 (2011).

[5] McMillan, D., Morrison, A., and Chalmers, M. Categorised ethical guidelines for large scale mobile HCI. *Proceedings of the 2013 ACM annual conference on Human factors in computing systems*, 1853–1862 (2013).

[6] Morrison, A., McMillan, D., Reeves, S., Sherwood, S., and Chalmers, M. A Hybrid Mass Participation Approach to Mobile Software Trials. *Proc. ACM CHI 2012*, 1311–1320 (2012).

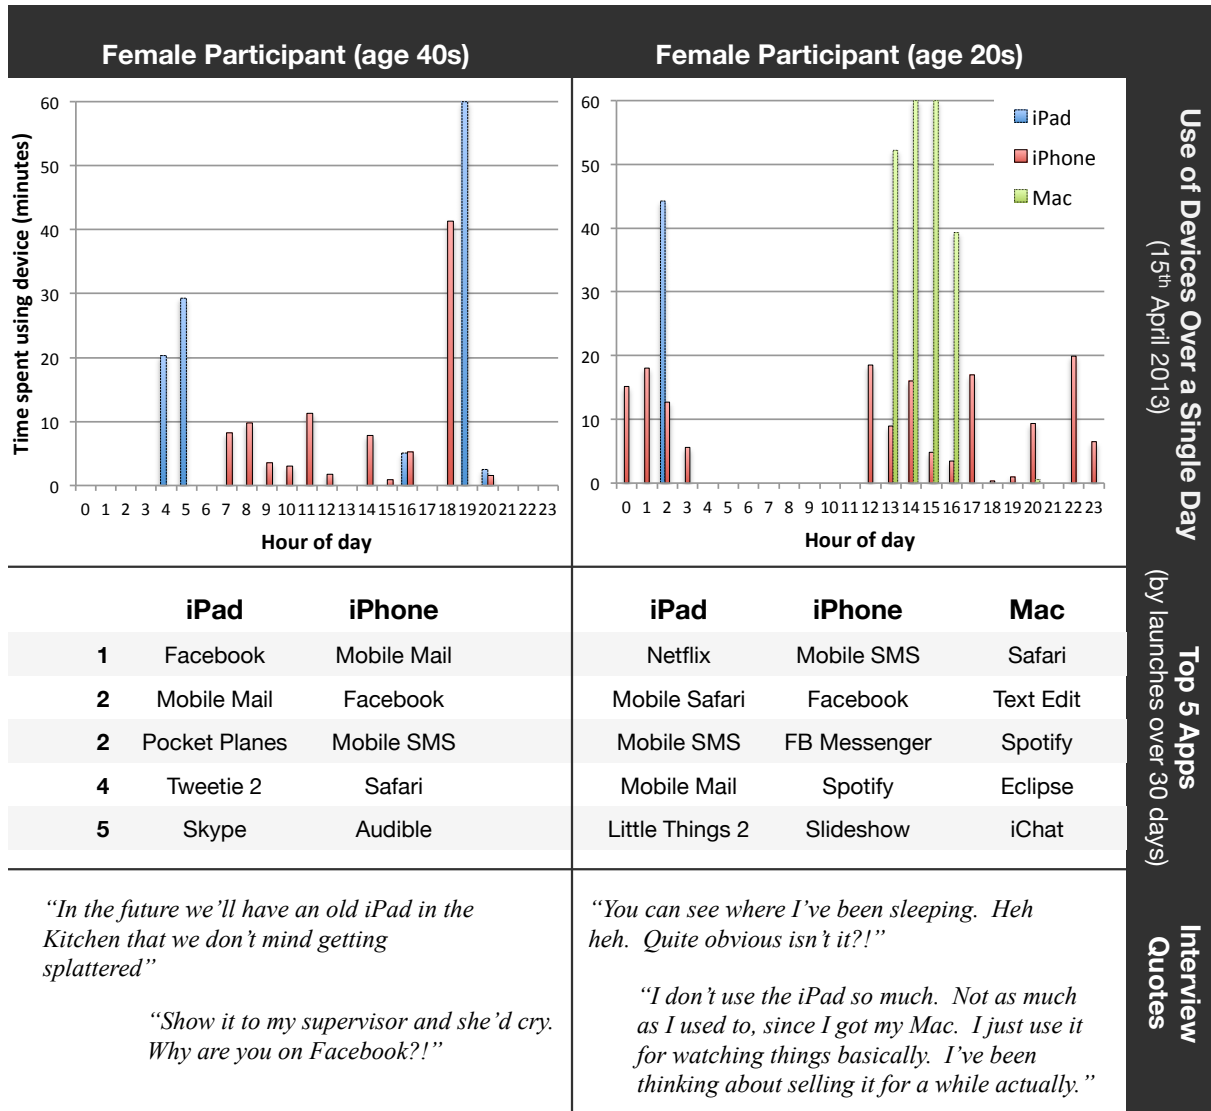


Figure 1: Data Concerning Two Study Participants