Poster: User Comfort with Android Background Resource Accesses in Different Contexts

Daniel Votipka, Seth M. Rabin, Kristopher Micinski*, Thomas Gilray, Michelle L. Mazurek, and Jeffrey S. Foster
University of Maryland, *Haverford College
dvotipka,srabin,tgilray,mmazurek,jfoster@cs.umd.edu; *kmicinski@haverford.edu

***Abstract*—Android apps ask users to allow or deny access to sensitive resources the first time the app needs them. Prior work has shown that users decide whether to grant these requests based on the context. In this work, we investigate user comfort level with resource accesses that happen in a *background* context, meaning they occur when there is no visual indication of a resource use. For example, accessing the device location after a related button click would be considered an interactive access, and accessing location whenever it changes would be considered a background access. We conducted a 2,198-participant fractional-factorial vignette study, showing each participant a resource-access scenario in one of two mock apps, varying what event triggers the access (*when*) and how the collected data is used (*why*). Our results show that both *when* and *why* a resource is accessed are important to users' comfort. In particular, we identify multiple meaningfully different classes of accesses for each of these factors, showing that not all background accesses are regarded equally.**

I. Introduction

Android apps potentially have access to a range of sensitive resources, such as location, contacts, and SMS messages. As a result, Android and similar systems face a critical privacy and usability trade-off: when should the system ask the user to authorize an app to access sensitive resources? Requesting permissions too often can overburden the user; requesting permission too infrequently can lead to security violations.

There has been significant research into this question, much of which shows that users' access-control decisions depend on the context, including when and why the access attempt is made [1]–[6]. However, this prior work has typically focused on individual aspects of context in isolation, such as app behavior at the point of resource-access [3], [5], [7], or the reason the app requires access to the sensitive resource [6]. Further, much of this work relies on a binary distinction between foreground and background accesses—sometimes defined as whether the app is visible on the screen [5], [7], and sometimes defined as whether the resource access is explicitly triggered by a specific user interaction [3], [8].

Using an online vignette study, we investigate more deeply how users understand resource uses that occur *in the background*, which we broadly define as not explicitly and obviously caused by a user interaction. We examine whether different kinds of background uses are viewed similarly, or whether more fine-grained distinctions are required for user comprehension. Participants were shown a mock app being used in a particular scenario, then asked whether they would be comfortable using an app that behaved similarly and whether they would recommend such an app to friends.

Fig. 1: Sample vignette. The orange boxes and arrows, and the gray circle, are not shown initially. They are added in the resource access description step, along with a textual description.

We found that both the *why* and *when* aspects of context played a significant role in users' expressed comfort with background accesses and differentiate important sub-types for each factor.

II. Method

We performed a 2,198-participant, between-subjects, fractional-factorial vignette study. Participants were recruited from the Amazon Mechanical Turk crowd-sourcing service. Participants were asked for their opinions regarding a given app's functionality and behavior, but we did not explicitly mention privacy or the possible sensitivity of specific resources. This study was approved by our university's Institutional Review Board.

We begin each survey by describing a mock app and how Jane, a fictional character, might use the app. We do so by showing a sequence of screenshots depicting Jane's use of the app throughout her day. Figure 1 shows an example. Note that in this first step, the orange boxes and gray circle in the figure are omitted from the vignette.

Participants are then informed of the app's "behind-the-scenes" access context where they are shown the same series of app screens with additional indicators showing *when* and *why* the resource access occurred.

After describing the app's access context, we ask the participants a series of five-point Likert-scale questions regarding

Variable	Value	Odds Ratio	CI	p-value
Why	*Personalize*	–	–	–
	Server	0.88	[0.72, 1.09]	0.240
	Analytics	**0.49**	**[0.37, 0.64]**	$<$ **0.001***
	Ads	**0.34**	**[0.28, 0.42]**	$<$ **0.001***
	N/A	**0.58**	**[0.43, 0.80]**	$<$ **0.001***
When	*Interactive*	–	–	–
	Prefetch	**0.64**	**[0.48, 0.87]**	**0.004***
	UI-Bg	**0.72**	**[0.55, 0.94]**	**0.014***
	Change	**0.34**	**[0.26, 0.44]**	$<$ **0.001***
Resource	*Location*	–	–	–
	Contacts	**0.33**	**[0.28, 0.39]**	$<$ **0.001***
	SMS	**0.12**	**[0.09, 0.16]**	$<$ **0.001***
Internet Skill	0	–	–	–
	+1	**0.95**	**[0.94, 0.97]**	$<$ **0.001***

*Significant effect – Base case (OR=1, by definition)

TABLE I: Summary of regression over participant comfort with different access contexts.

their comfort with the given access context and how likely it is that popular apps behave similarly.

Each participant was assigned round-robin to one of 52 conditions defined by four variables: the *app*, the *resource* being accessed, *why* the app accessed the resource, and *when* the resource was accessed. App and resource options included dating app and a rideshare app, as well as the user's location, list of contacts, and SMS messages.

Reasons for resource access. We selected five variations for *why* based on reverse-engineering usage patterns in popular apps. These included to provide personalized features either within the device (*Personalize*) or by sending data to the app's own server (*Server*); to support debugging and analytics via a third-party service (*Analytics*); or for targeted advertising, also via a third-party service (*Ads*). We also include a case where the participant is not given a reason for data collection (*N/A*).

Triggers for resource access. We considered four variations in *when* the app requests resources, designed to target different levels of interactivity. These included access directly after a related button click (*Interactive*); access prior to a related UI event in order to prefetch data that will soon be needed (*Prefetch*); access after an unrelated UI event (*UI-Bg*); and directly after a resource has been modified (e.g., the user changes locations or adds a contact), regardless of whether or not the app is on screen (*Change*).

III. Results

We found that both *why* and *when* resource accesses occurred had a significant effect on user comfort. Additionally, we found that there are several meaningful classes of accesses for each part of the access context. Table I shows the results of our logistic regression over user comfort.

With respect to *why* the access occurred, we observed that users were more comfortable when data was shared with the app developer (*Personalize* and *Server*) than a third-party (*Analytics* and *Ads*). Further, within third-party sharing, users are more comfortable when data is shared for app analytics (to improve the functionality of the application) as opposed to sharing data for advertising. Additionally, if no reason for access was provided, we found that users were less comfortable than they would be if told the data never left their device (*Personalize*), but slightly more comfortable than having their data shared with advertisers (*Ads*).

For *when*, as expected, users are the most comfortable when accesses occur interactively, directly after a UI event (*Interactive*). Non-interactive (background) accesses can further be divided into two classes: participants were more comfortable if the access occurred when the app was on-screen (*Prefetch* and *UI-Bg*) compared to off-screen (*Change*).

Acknowledgment

This research was supported in part by NSF CNS-1064997, a UMIACS contract under the partnership between the University of Maryland and DoD, and a Google Research Award.

References

[1] B. Bonné, S. T. Peddinti, I. Bilogrevic, and N. Taft, "Exploring decision making with android's runtime permission dialogs using in-context surveys," in *Proceedings of the 13th Symposium on Usable Privacy and Security*, ser. SOUPS '17. Santa Clara, CA: USENIX Association, 2017, pp. 195–210. [Online]. Available: https://www.usenix.org/conference/soups2017/technical-sessions/presentation/bonne

[2] C. Thompson, M. Johnson, S. Egelman, D. Wagner, and J. King, "When it's better to ask forgiveness than get permission: Attribution mechanisms for smartphone resources," in *Proceedings of the 9th Symposium on Usable Privacy and Security*, ser. SOUPS '13. New York, NY, USA: ACM, 2013, pp. 1:1–1:14. [Online]. Available: http://doi.acm.org/10.1145/2501604.2501605

[3] K. Micinski, D. Votipka, R. Stevens, N. Kofinas, M. L. Mazurek, and J. S. Foster, "User interactions and permission use on android," in *Proceedings of the 35th ACM on Human Factors in Computing Systems*, ser. CHI '17. New York, NY, USA: ACM, 2017. [Online]. Available: http://cs.umd.edu/~micinski/apptracer-2016.pdf

[4] I. Shklovski, S. D. Mainwaring, H. H. Skúladóttir, and H. Borgthorsson, "Leakiness and creepiness in app space: Perceptions of privacy and mobile app use," in *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems*, ser. CHI '14. New York, NY, USA: ACM, 2014, pp. 2347–2356. [Online]. Available: http://doi.acm.org/10.1145/2556288.2557421

[5] P. Wijesekera, A. Baokar, A. Hosseini, S. Egelman, D. Wagner, and K. Beznosov, "Android permissions remystified: A field study on contextual integrity," in *Proceedings of the 24th USENIX Security Symposium*, ser. USENIX Security '15. Washington, D.C.: USENIX Association, Aug. 2015, pp. 499–514. [Online]. Available: https://www.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/wijesekera

[6] J. Lin, B. Liu, N. Sadeh, and J. I. Hong, "Modeling users' mobile app privacy preferences: Restoring usability in a sea of permission settings," in *Proceedings of the 10th Symposium On Usable Privacy and Security*, ser. SOUPS '14. Menlo Park, CA: USENIX Association, 2014, pp. 199–212. [Online]. Available: https://www.usenix.org/conference/soups2014/proceedings/presentation/lin

[7] P. Wijesekera, J. Reardon, I. Reyes, L. Tsai, J.-W. Chen, N. Good, D. Wagner, K. Beznoso, and S. Egelman, "Contextualizing privacy decisions for better prediction (and protection)," in *Proceedings of the 36th ACM on Human Factors in Computing Systems*, ser. CHI '18. New York, NY, USA: ACM, 2018. [Online]. Available: https://blues.cs.berkeley.edu/wp-content/uploads/2018/01/chi18-android.pdf

[8] T. Ringer, D. Grossman, and F. Roesner, "Audacious: User-driven access control with unmodified operating systems," in *Proceedings of the 23rd ACM Conference on Computer and Communications Security*. Vienna, Austria: ACM, oct 2016. [Online]. Available: http://tlringer.github.io/pdf/audacious.pdf