User Interactions and Permission Use on Android

Kristopher Micinski, Daniel Votipka, Rock Stevens, Nikolaos Kofinas, Michelle L. Mazurek, and Jeffrey S. Foster

University of Maryland, College Park
{micinski, dvotipka, rstevens, nkofinas, mmazurek, jfoster}@cs.umd.edu

Android has a *permission system* that asks users for authorization before an app uses sensitive resources such as contacts or GPS location. A key challenge in such authorization systems is balancing user interruptions with making sensitive resource use transparent. We hypothesize that Android's existing authorization systems (install-time permission lists or run-time dialog boxes, depending on the version) could achieve a better balance by integrating with the app's user interface (UI), because the UI deeply informs the user's mental model of the app's behavior, including security-relevant behavior.

In particular, in our work we ask whether *user interactions*—button clicks, page changes, dialog boxes, etc.—can be taken as evidence of authorization to use certain sensitive resources. If so, this could reduce the need for separate authorization requests. Conversely, we ask whether sensitive resource use without an associated interaction suggests a need for additional authorization requests. Note that while our studies are heavily influenced by Android, we believe our results will generalize to related mobile OS's and similar settings like web apps.

To answer these questions, we conducted two related studies. First, we reviewed 150 popular Android apps to determine whether sensitive resource uses are related to user interactions in existing apps. If so, an authorization mechanism integrated with the UI could work well with existing app designs. To carry out this study, we developed AppTracer, a dynamic analysis tool that instruments Android apps to log UI actions and resource uses, and then visualizes the logs as graphs that show temporal ordering of logged events. We used AppTracer to determine whether each observed resource use in each app was *interactive*, meaning either it was immediately preceded by a related UI event (e.g., accessing contacts after clicking a button marked "contacts"), or it was the main focus of the current screen (e.g., using location on a map screen).

We found that, across our subject apps, several resources (microphone, camera, external storage, and calendar) are used almost exclusively interactively; several others (including bluetooth and phone state) are used mostly non-interactively (which we call in the *background* even if the app itself is on screen); and several resources (most notably contacts and location) exhibit a mix of interactive and background uses. These results suggest interactive and background uses may call for different authorization mechanisms, and that these mechanisms cannot necessarily be divided strictly by resource.

These results informed the design of our second study, a 961-participant online survey investigating participants' expectations about interactive and background permission uses. Each participant viewed a slideshow of two usage scenarios for a mock mobile app, where each scenario shows a short interaction (e.g., launching the app, clicking a button, etc.) and then asks if the participant expects microphone, location, and/or contacts to be used after the interaction. We chose these resources to reflect a range of interactivity as measured in our app study. We aimed to gain insight into how different factors affect user expectations, and therefore which authorization mechanisms might be appropriate for different usage patterns. As we anticipated, we found that users are much more likely to expect resources to be accessed after a related interaction than in the background. However, we also found that seeing one interactive use does not prime the user to expect a future background use, indicating a potential weakness in the Android M request-on-first-use authorization model. In contrast, our findings show that an authorization request at launch does increase expectations for both interactive and background accesses, perhaps because it better conveys the idea that the resource could be accessed at any time.

Drawing on the results of our studies, we make three design recommendations. First, resource uses should be made after associated interactions as much as possible. Given the current makeup of apps, this should be achievable for many commonly used resources without extensive effort. Second, separate authorization dialogs might be unnecessary for resources that are accessed mostly interactively (see, e.g., [1]). Finally, authorization for background resource uses should be distinct from authorization for interactive uses, and these background authorizations may be most effective when the app is launched.

REFERENCES

1. Franziska Roesner, Tadayoshi Kohno, Alexander Moshchuk, Bryan Parno, Helen J. Wang, and Crispin Cowan. 2012. User-Driven Access Control: Rethinking Permission Granting in Modern Operating Systems. In *Proceedings of the 2012 IEEE Symposium on Security and Privacy (SP '12)*. IEEE Computer Society, Washington, DC, USA, 224–238. DOI: http://dx.doi.org/10.1109/SP.2012.24