Understanding Sensor Notifications on Mobile Devices

Zongheng Ma, Saeed Mirzamohammadi, Ardalan Amiri Sani

University of California, Irvine
zonghenm@uci.edu, saeed@uci.edu, ardalan@uci.edu

Extended Abstract

Mobile devices use notifications to inform users of events, e.g., phone calls. Some security- and privacy-related notifications are *time-sensitive*: the user must be notified immediately. Examples are sensor notifications that inform the user of access to sensitive sensors (e.g., camera, microphone, and location) and AMBER alerts. In this paper, we set out to understand the properties of such notifications, i.e., *time-sensitive notifications*, through a user study. We perform the study in the context of sensor notifications but most of the results are applicable to other forms of time-sensitive notifications as well. Computers communicating with humans is one of the most fundamental aspects of computing. In this paper, we hope to understand how effectively a mobile device can grab its owner's attention using the means available to it. In an analogy to human-to-human communication, this is similar to one person seeking another person's attention in case of potential danger.

We focus on sensor notifications due to their increasing importance. Mobile devices incorporate a set of privacy-sensitive sensors, most importantly, camera, microphone, and location (i.e., GPS or cellular network-based sensing). These sensors can capture sensitive information of the user including photos, videos, conversations, and locations. A prior study has shown that most users consider this information to be private and sensitive [5]. Unfortunately, there have been several incidents where attackers have attempted to access these sensors without user's knowledge, e.g., to capture unauthorized video or audio of the user or to track her whereabouts [2–4]. And tools are available for attackers to remotely control infected device's camera and microphone [1]. To address this problem, *sensor notifications* are used, which use some notification channel, such as LED and display, to notify the user whenever one of these sensors is being used. For example, mobile operating systems, such as Android and iOS, show some form of notification on the screen when a location sensor is accessed (Figure 1 (a) and (b)). Moreover, some applications inform the user when they record audio in the background (Figure 1 (c)).

The rationale for sensor notifications is that users can distinguish between legitimate and illegitimate accesses to these sensors. For example, if the user is using a photography application, s/he expects the camera to be on. But when the user is watching a video, s/he expects the camera to be off. Therefore, a notification about the use of camera in the latter scenario is likely to indicate malicious activity.

We identify three important requirements for a sensor notification. First, it must be able to attract the user's attention successfully at all times and as fast as possible. This is important as it will minimize the window of vulnerability, i.e., the period of time that malware can access the sensor without user's knowledge. Second, it must be unambiguous. The user must be able to quickly understand the meaning of the notification and differentiate it from other sensor notifications and other notifications in the system, e.g., phone calls or application updates. Third, it should not cause annoyance to the user. This last property is indeed important because notifications are triggered whenever the sensor is accessed, even if the access is legitimate, e.g., user taking a photo with the camera.

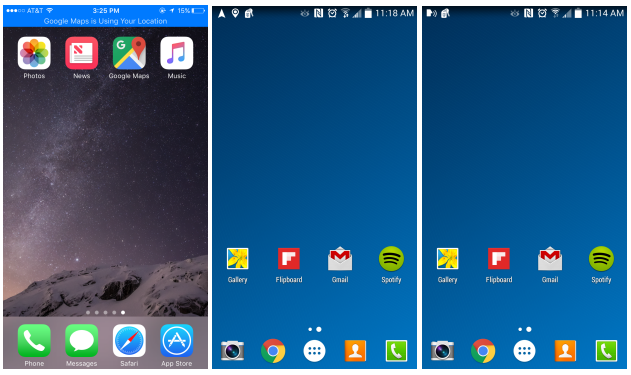

(a) (b) (c)

Figure 1: Sensor notifications in iOS and Android: (a) Location notification in iOS (the notification strip below the status bar on top of the screen, (b) location notification in Android (the icon in the top left corner), (c) microphone notification by the Samsung Voice Recorder application in Android (the icon in the top left corner).

We ask: *which one of existing notification channels, i.e., LED, vibration, sound, and display, best satisfy these requirements?* We answer this question with a user study performed on 40 participants[1]. To perform the user study, we first designed and implemented various types of sensor notifications for camera, microphone, and location, in the Android operating system. We then deployed our modified operating system on Google Nexus 5X smartphones and distributed them to the participants to use as their primary device for one week. In addition, we installed an application on the smartphone that emulates malware by using these sensors in the background a couple of times a day, which

[1]User study approved by UC Irvine's Office of Research, Human Research Protection, under IRB HS# 2016-3100.

would result in a notification to the user by the operating system. We then asked the participants to log the sensor that was illegitimately used as soon as they spotted the corresponding notification. This allows us to assess the first two requirements mentioned earlier. To assess the third requirement, we asked the participants to fill out a questionnaire right after the study. We have made the data collected in the study publicly available[2].

To better understand the notifications, we decided to evaluate the impact of the device's physical context, i.e., ambient light intensity and ambient noise, on the effectiveness of the notifications. To do this, we collected these physical context measurements whenever a notification was triggered.

We present several important findings about sensor notifications. First, on average, none of the existing channels achieve a success rate (in capturing user's attention) higher than 24%. Second, vibration achieves the best success rates at 24%, and LED achieves the worst at 4%. Third, we find that sound is almost always a bad choice as it incurs significant annoyance to the user, while being outperformed by vibration and even display notification. Fourth, we find that, quite counterintuitively, existing android notifications that only rely on textual content to convey their meaning are the hardest for users to recognize in a timely manner. Finally, we realize that physical context has important impact on the effectiveness of notifications. For example, we find that our new display notification (which turns on the display and shows a strip on top of the screen with textual content in it) can be very effective in light environments. Based on our findings, we suggest to use this display notification for camera since it is effective in light environments, which is where malicious access to the camera can be effective. For microphone and location sensor, we suggest to use different vibration patterns since vibration achieves high success rate and its different patterns are easy for users to recognize.

1. REFERENCES

[1] Dendroid: Android Trojan Being Commercialized. http://blog.trustlook.com/2014/03/20/dendroid-android-trojan-commercialized/.

[2] How the NSA can 'turn on' your phone remotely. http://money.cnn.com/2014/06/06/technology/security/nsa-turn-on-phone/.

[3] Man spies on Miss Teen USA. http://www.reuters.com/article/2013/10/31/us-usa-missteen-extortion-idUSBRE99U1G520131031.

[4] Men spy on women through their webcams. http://arstechnica.com/tech-policy/2013/03/rat-breeders-meet-the-men-who-spy-on-women-through-their-webcams/.

[5] I. Muslukhov, Y. Boshmaf, C. Kuo, J. Lester, and K. Beznosov. Understanding Users' Requirements for Data Protection in Smartphones. In *Proc. IEEE Int. Conf. on Data Engineering Workshops (ICDEW)*, 2012.

[2]http://www.ics.uci.edu/~ardalan/notifdata.html