Visualizing Privacy Risks of Mobile Applications through a Privacy Meter

Jina Kang[1], Hyoungshick Kim[1], Yun Gyung Cheong[1], and Jun Ho Huh[2]

[1] Department of Computer Science and Engineering, Sungkyunkwan University, Korea
[2] Honeywell ACS Labs, USA

Abstract. When it comes to installing mobile applications on Android devices, users tend to ignore privacy warning messages about permissions being requested. Warning messages are often shown too late and are hard to interpret for normal users. To improve users' awareness of potential privacy implications of installing an application, we designed a "privacy meter" that visualizes the risks (in a slider bar format) imposed by the types of permissions being requested. Interpreting and understanding privacy risks become quick and easy.

Our lab study shows that the privacy meter is the most effective solution compared to Google's existing permission screens and privacy fact feature. With the privacy meter in place, only about 26% of participants recommended applications that have high privacy risks to their friends and family members. That is a significant improvement from the 61% of participants who recommended high risk applications when Google's permission screens were used. The time taken to make recommendation decisions, on average, also dropped from 72 seconds to 26 seconds when the privacy meter was used.

Index Terms: Permission, Android, Mobile, Decision-making.

1 Introduction

When a user tries to install a mobile application (app) from Google Play (the official marketplace for Android apps), the user is asked to grant a set of *permissions* for the app to access information or use features from the user's smartphone. However, recent studies [7, 10, 11] have shown that the majority of users tend to ignore any warning messages shown on those permission requesting displays. For example, Felt et al. [7] showed that about 83% of their study participants did not pay attention to the permission information screens, and only about 3% of users correctly understood the meaning of given permissions. This is mainly because warning messages are shown when users have already decided to install an app; at that stage, most users just want to continue with installation without being interrupted [3]. Moreover, many users found it difficult to understand the terms and words used to describe permissions [11].

Although designing an effective warning message system for permissions seems challenging, Kelley et al. [10] proposed a warning display about permissions called *privacy facts* and showed that the use of privacy-focused warning messages do affect users' app selection behavior. This indicates that there is still room for improvement in

© Springer International Publishing Switzerland 2015
J. Lopez and Y. Wu (Eds.): ISPEC 2015, LNCS 9065, pp. 548–558, 2015.
DOI: 10.1007/978-3-319-17533-1_37

designing a warning display system for apps to highlight their privacy risks. A type of meter may be a reasonable candidate for providing provide warning to the users. Meter is a popular form of user interface for warning systems. In particular, password strength meter was intensively studied before (e.g., [15]). Since users are already familiar with such meter designs for warning systems, an app's privacy risks related to permissions might also be effectively visualized.

This paper proposes a new warning display mechanism called the '*privacy meter*' to warn users about potentially dangerous permissions (requested by an app). The privacy meter evaluates the risks associated with the permissions requested by an app, and visualizes the computed risk scores. In designing a prototype of the meter, we considered how many dangerous permissions (e.g., `INTERNET` and `READ_PHONE_STATE`) are included in the collection of permissions requested by an app. To evaluate the effectiveness of our prototype implementation, we conducted a user study involving 36 participants, asking participants to recommend apps to their friends or family members in a simulated role playing scenario. Each participant was assigned to one of the four experiment conditions: (1) our new privacy meter display, (2) the privacy facts display [10], (3) the current Google Play store permissions display, or (4) the previous Google Play store permissions display (before version 4.8.20). We compared the numbers of times a high risk app (an app that requests a high number of potentially dangerous permissions) was recommended under each condition, demonstrating that our privacy meter significantly outperforms all other mechanisms in affecting participants' recommendation decisions.

In summary, our contributions can be briefly described as follows: First, the proposed interface for app's permissions (*privacy meter*) is more effective than the existing methods in warning users before installing apps. Second, we found that the new Google Play store interface based on permission groups is significantly better than the previous Google one.

2 Related Work

We focus on Android permissions display due to its historically more detailed permissions system and its large user base. Prior work suggests that users tend to ignore the permission display on Android, mainly because the messages appear after users have already decided to download an app [11]. Furthermore, users who pay attention to permissions lists have trouble using them because the screens are jargon-filled, provide confusing explanations, and lack explanations for why the data is collected [7]. There were several different studies on user expectations about Android access control system (e.g., [11,12]). In particular, Felt and her colleagues have published a series of papers on the Android permission model, and how users understand it. They found that most users (83%) do not pay attention to the permission screens during installation. Also, only 3% of their surveyed users had a good understanding of what the permissions were actually asking in terms of accessing data on the phone [7].

Those studies showed the problems of Google permission lists. To solve those problems, some researchers developed automated tools to detect overprivileged and malicious apps. Stowaway [5] was created to detect overprivileged Android apps by checking whether an app requests for excessive permissions. Taintdroid [4] developed a

mechanism for inspecting whether a running app requires the user information it has access to. Papamartzivanos et al. [13] proposed a cloud-based system that runs on a crowdsourcing logic, providing a privacy-flow tracking service in real-time. Felt et al. [6] found that about one third of 940 Android apps they experimented were considered overprivileged.

The problem is that users often ignore the permission lists and warnings messages. To improve users' awareness of privacy implications of permissions being requested by apps, Felt et al. [10] suggested a new warning display interface called *privacy facts* that shows a checklist of permissions requested by an app. Privacy facts display eight different types of information an app can collect: personal information, contacts, location, calendars, credit card or financial, diet or nutrition, health or medical, and photos. These items are displayed with a checkbox next to it, indicating whether each information type is being requested by that app. Privacy facts help users avoid apps that suspiciously require unnecessary permissions.

To help users make more privacy-conscious decisions, Harbach et al. [9] proposed to leverage many personal data available on smartphones by providing customized examples. Providing private information examples can help users pay more attention to the relevant, important information and make better decisions. Our goals are inline with their research goals. We want users to be able to decide for themselves the risks associated with an app they are about to install. The key difference, however, is that they use images to visualize the meaning of permissions. On the other hand, our privacy meter visualizes the privacy risk information through an one-dimension slider bar.

Some researchers have tried grading apps. PrivacyGrade [12, 14] aims to improve users' awareness of app behaviors that may compromise users' privacy. PrivacyGrade provides detailed information about an app's privacy-sensitive behaviors. Such behaviors are summarized as grades, ranging from "A+" (being the most privacy sensitive) to "D" (being the least privacy sensitive). Gates et al. [8] proposed a scoring system that assigns a risk score to each app and displays summary of the risk scores to users. Again, the key difference between their scoring system and our meter is the way the risks are visualized. Recently, Biswas et al. [2] proposed metrics to quantify privacy risks associated with an app accessing user data. There are some similarities between privacy panel and our privacy meter, but their primary focus is to track privacy-flows. Moreover, user study is missing from their work.

3 Privacy Meter Design

This section presents the design of a *privacy meter* that measures the risks associated with permissions requested by an app by counting the number of privacy-sensitive, potentially dangerous permissions. Our privacy meter was designed and prototyped with a simple focus on visualizing the safety (privacy risks) of apps. With the initial work presented in this paper, our goal was not to develop very accurate risk scoring functions, but to first demonstrate the effectiveness of visualizing privacy risks through a meter type display. Once we can show the effectiveness of a simple meter design, we will then focus on building more precise risk scoring functions as part of future work. We used a circular slider thumb to indicate where the privacy risks for an app lie on

Table 1. Top 10 popularly requested permissions for Android malware [1]

Rank	Permission name	Description
1	INTERNET	Full network access
2	READ_PHONE_STATE	Allows read only access to phone state
3	SEND_SMS	Allows an application to send SMS messages
4	WRITE_EXTERNAL_STORAGE	Allows an application to read from external storage
5	RECEIVE_SMS	Allows an application to receive SMS messages
6	READ_SMS	Allows an application to read SMS messages
7	ACCESS_COARSE_LOCATION	Allows an app to access approximate location derived from network location sources
8	READ_CONTACTS	Allows an application to read the user's call log
9	ACCESS_FINE_LOCATION	Allows an app to access precise location from location sources such as GPS, cell towers, and Wi-Fi
10	WRITE_SMS	Allows an application to write SMS messages

the slider bar (see Fig. 1), which is the default slider design available in Android SDK. The next three sections explain the meter features we carefully considered in the meter design.

3.1 Counting the Number of Dangerous Permissions

Most importantly, the meter was designed to reflect on how many privacy-sensitive, potentially dangerous permissions are being requested by an app. We simply counted the number of privacy-sensitive permissions in the requested permission set based on an existing list of dangerous permissions available in [1], and computed the privacy risk score by dividing that counter value by the maximum number of dangerous permissions. Table 1 shows the top 10 of the most popularly requested permissions by Android malware. For more intuitive visualization, the number of dangerous permissions were scaled and normalized to values ranging from 0 to 1.

3.2 Computing the Privacy Risk Score

Since we used the list of 30 most popularly used dangerous permissions by malware samples [1], the maximum number of dangerous permissions in an app would always be 30 in our design. However, previous research has shown that for top 234 popular apps, the number of dangerous permissions requested by the apps is much smaller than 30 (95% confidence interval, 3.059 to 3.540) [reference removed for anonymous submission]. Therefore, given a set of dangerous permissions p, we calculated the privacy score S_p as follows: $S_p = \min(|p|, 4)/4$ where $|p|$ is the number of permissions in p. Based on the observed confidence interval, we assumed that an app can be dangerous if the number of requested dangerous permissions is greater than or equal to 4. For example, if an app requests two dangerous permissions such as INTERNET and

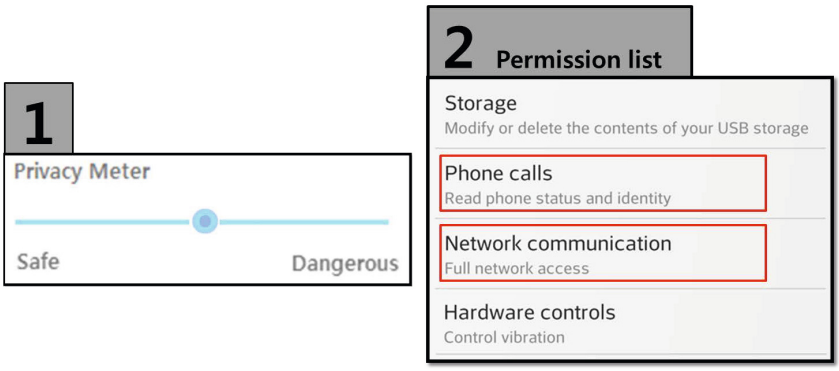

Fig. 1. Our prototype implementation of *privacy meter*. "1" is privacy meter; and "2" is the requested permission list (dangerous permissions are highlighted in red box) which can additionally be shown when the slider bar is clicked.

READ_PHONE_STATE, S_p would be 0.5, and the slider thumb would be located at the half way of the slider bar (see Fig. 1).

3.3 Display Style

As indicated above, we simply used the default interactive slider available on Android SDK to minimize any visualization bias that could be introduced by designing our own proprietary slider design. Our prototype, shown in Fig. 1, is based on an Android interactive slider with a height of 90 pixels and a width that stretches the full width of a given Android phone. The circular slider thumb visualizes the level of privacy risks associated with an app. By clicking on the slider bar, users can see the details of the requested permissions.

4 Methodology

This section describes our hypotheses and describes the user study, which was designed based on the hypotheses to help evaluate the effectiveness of the proposed privacy meter.

4.1 Hypotheses

The primary goal of this work is to design a highly effective visualization mechanisms for privacy risks that can effectively protect users from installing (potentially) privacy-compromising apps. Based on that goal, the following hypotheses were defined.

1. The privacy meter interface is more effective than existing warning displays such as the privacy facts interface [10] in visualizing the privacy risks associated with an app.
2. Current Google design is more effective than previous Google design in presenting the privacy risks associated with an app.

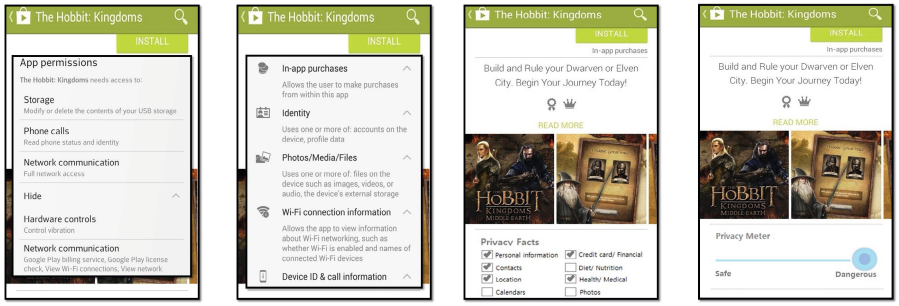

(a) Google (previous) **(b)** Google (current) **(c)** Privacy facts **(d)** Privacy meter

Fig. 2. The warning displays used for user study

4.2 Existing Warning Displays about Permissions

We compared the effectiveness of the privacy meter with the following two representative warning displays.

- *Google permission displays.* A warning display is currently used in the official app store in Google. When a user tries to install an app through the store, it shows a list of permissions required to use the app. We used two different interfaces: Fig. 2(a) and (b) show the warning displays before and after version 4.8.20, respectively. After version 4.8.20, permissions have been organized into groups.
- *Privacy facts.* A warning display was proposed by Kelley et al. [10] to help users become more aware about privacy risks upon installing apps; their warning display shows personal information that can be accessed by an app. Their study results show that privacy facts are more effective in influencing users' app selection behavior than Google's existing permission lists [10].

4.3 User Study Design

We used a mixed-method design, an experiment and a follow-up questionnaire, to study participants' behaviors in understanding the the privacy risks of requested permissions, and test our two hypotheses. To test the hypotheses, we designed a between-subject study in which each participant was given only one of the four warning display mechanisms (the previous Google warning display, the current Google warning display, the *privacy facts* checklist, and the *privacy meter* warning display). The participants were sequentially allocated to one of the four warning displays (See Fig. 2) in the order in which they were recruited.

4.4 Procedure

We advertised and conducted the user study as a *application recommendation* role playing experience. The study was carried out in an on-campus lab.

For each participant, we provided an Android smartphone with a pre-installed app that we developed for simulating a "custom Android market" and explained a brief overview of the experiment (without mention of the privacy or security issues associated with the app). Before performing the app selection tasks, each participant was asked to complete a short pre-survey about demographic information, which allowed us to find out whether that a participant currently uses an Android phone. After the pre-survey, through the pre-installed app, we asked the participant to (1) select one between two similar apps but each requiring a totally different permission set, and (2) recommend that app to his or her friends. We repeated this exercise six times with different six app pairs that belong to the six different categories (word game, documents scanning, Facebook, nutrition, review movie, travel apps) that we carefully selected. Once the participant finished the app recommendation tasks, he or she was interviewed to explain his or her decisions; the interviews were audio-recorded. On average, the participants spent about 15 minutes to complete the entire process and received a two-dollar cafe voucher for completing it.

4.5 Tested Apps

For the app selection task, we carefully chose 2 apps for each of six representative categories ('document', 'social networks', 'movie', 'health', 'travel', and 'word game') in Google Play where one app (*safe* app) requested less access to permissions compared with the other app (*dangerous* app). All of the apps we used were real apps which can be found in Google Play, but those apps were chosen in a manner that they are similar in terms of functionality and popularity (e.g., with 100 to 1,000 download range); we made our best efforts to eliminate bias occurring in the app selection task.

We implemented an Android app for the user study, which simulated a "custom Android market" to achieve a reasonable amount of ecological validity. That app sequentially asked a participant to select one app in each of the six app pairs that we defined. The order of the six categories presented to a participant and the two apps selected for each category were randomly defined.

5 Results

This section presents the data collected through the user study and the statistical analysis results.

5.1 Demographics

We recruited a total of 36 Android phone users. 10 participants (27.78%) were female and 26 participants (72.22%) were male. All participants were aged between 23 and 32 and had a computer science bachelors degree. We wanted to find out the security knowledge and awareness levels of participants when it comes to using Android phones. Therefore we asked the participants whether they are currently using a security lock (e.g., PINs, patterns, etc.) to protect their mobile phones. Most participants (77.78%) did not use any security lock because they found entering patterns or passwords annoying.

Table 2. Proportion of participants who decided to recommend potentially dangerous applications to their friends and family members, and average time take to make recommendation decisions

Design	# recommended	Time
Google (previous)	41/54 (75.92%)	70.67 sec
Google (current)	33/54 (61.11%)	71.50 sec
Privacy facts	24/54 (44.44%)	53.54 sec
Privacy meter	14/54 (25.92%)	26.11 sec

5.2 Recommending Potentially Dangerous Applications

First, we look at the proportion of participants who decided to recommend to their friends potentially dangerous applications that contain 3 or more privacy-compromising permissions. As shown in Table 2, with the privacy meter being present, only 14 participants (25.92%) decided to recommend potentially dangerous applications to their friends. In contrast, 33 participants (61.11%) decided to recommend potentially dangerous applications when Google's current permission lists were used, which is a huge jump.

The Fisher's Exact Tests (FET) were used to compare the proportion of participants in a statistically significant manner. Post-hoc comparisons were corrected for multiple-testing using the Benjamini-Hochberg method. The test results of all the warning displays appeared to be significantly different (corrected FET, all $p < 0.05$).

Most importantly, there is a statistically significant difference between the proportions of participants recommending potentially dangerous applications when the privacy meter was present and when the privacy facts were used (corrected FET, $p < 0.05$). It is also interesting to note that Google's current permission screens are more effective than the previous screens (corrected FET, $p < 0.05$), showing that the newer permission screens have indeed improved users' experience in making more informed decisions. From those observations, we confidently conclude that the privacy meter is more effective than privacy facts in visualizing the privacy risks associated with the permissions requested by an app. We also conclude that current Google privacy display design is better than the previous Google privacy display design.

5.3 Decision Time

We now look at the average time taken to make recommendation decisions. As Table 2 shows, the participants who were given the privacy meter took the shortest time on average to make decisions. We used unpaired t-test to show statistical significance. Google's current and pervious permission screens did not show any statistical significant difference in the average decision times (corrected unpairwise t-test, $p = 0.773$). The average decision time for privacy meter (26.11 sec) did show statistically significant improvement over privacy facts (53.54 sec) though (corrected unpairwise t-test, $p < 0.05$). From those results, we confidently conclude that the privacy meter can help users become more efficient in making privacy-conscious decisions. The participants' decision time for the privacy meter was much shorter than that of the privacy facts.

5.4 Interview Results

This section presents the interview results. We asked the following two questions.

Q1. When you are installing mobile applications, are you concerned about permission lists? Only 3 participants said that they are concerned about permission lists and make decisions based on them (8.33%). This result is similar to the results presented in Felt et al [7]. For those who are not concerned with permission lists, they provided several reasons. 14 of those participants (38.89%) believe that most applications request similar list of permissions anyway. 10 of those participants (27.78%) trust Google to inspect applications for privacy and security. 5 of those participants (13.89%) ignored the list due to long descriptions. 4 of those participants (11.11%) ignore permission lists because by that stage they would have already made decisions to install an application. 5 of those participants (13.89%) mentioned that these interfaces cannot be trusted since Android system can be compromised — Interestingly, all of them are using a security lock, which implies that they are more concerned with the security of their mobile phones.

Q2. Why did you not recommend a mobile application? 6 participants (16.67%) said that they did not recommend an application that contains version information in its name (e.g., FreeTravel ver.2) because it looks suspicious. 5 participants (13.89%) did not recommend an app that has a name of an unusual manufacturing company. Some participants wanted to know the relation between the requested permissions and their privacy impact, which are lacking in the privacy facts. In comparison, our privacy meter is capable of visualizing the privacy impact. The interview results also indicate that 11 participants pay attention to video files and images (31%), and popularity ratings (31%) when making decisions. 8 participants (22%) said that they read the application descriptions before downloading an application.

6 Limitations

This study was limited in several ways. First, all of our participants are from a single pool of users who major in Computer Science and Engineering. This may make it difficult to generalize the findings to all Android users with various backgrounds. Second, the main task (*application recommendation* task) in the user study was artificially designed through a simulation of the app installation process. Inherently, using such a simulation model might lower the ecological validity of the experiment compared with conducting real empirical studies.

7 Conclusions and Future Work

To improve users' awareness of privacy risks associated with installing mobile applications, we proposed a novel privacy measuring mechanism called the privacy meter. The

privacy meter looks at different types of permissions requested by a mobile application to gauge the level of privacy risks associated with it. The outcomes are visualized through a slider bar, making it quick and easy for users to interpret the associated risks and make decisions.

Our user study and statistical analysis show that the privacy meter outperforms other existing warning systems like Google's permission screens and the privacy facts. When only Google's permission screens were used, about 61% of the participants recommended mobile applications that have high privacy risks to their friends and family members. In comparison, that proportion went down dramatically to 26% when our privacy meter was present.

As part of future work, we will conduct more user studies to optimize the privacy meter design in terms of its visualization effects and accuracy in measuring risks.

Acknowledgements. This work was partly supported by the ICT R&D program (2014-044-072-003 , 'Development of Cyber Quarantine System using SDN Techniques') of MSIP/IITP and was also supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (No. 2014R1A1A1003707).

References

1. AV-Comparatives. Mobile security review. Technical report (2012)
2. Biswas, D., Aad, I., Perrucci, G.P.: Privacy Panel: Usable and Quantifiable Mobile Privacy. In: Proceedings of the 8th IEEE International Conference on Availability, Reliability and Security (2013)
3. Egelman, S., Tsai, J., Cranor, L.F., Acquisti, A.: Timing is Everything?: the Effects of Timing and Placement of Online Privacy Indicators. In: Proceedings of the 27th SIGCHI Conference on Human Factors in Computing Systems (2009)
4. Enck, W., Gilbert, P., Han, S., Tendulkar, V., Chun, B.-G., Cox, L.P., Jung, J., McDaniel, P., Sheth, A.N.: TaintDroid: an Information-Flow Tracking System for Realtime Privacy Monitoring on Smartphones. In: Proceedings of the 9th USENIX Conference on Operating Systems Design and Implementation (2014)
5. Felt, A.P., Chin, E., Hanna, S., Song, D., Wagner, D.: Android Permissions Demystified. In: Proceedings of the 18th ACM Conference on Computer and Communications Security (2011)
6. Felt, A.P., Greenwood, K., Wagner, D.: The Effectiveness of Application Permissions. In: Proceedings of the 2nd USENIX Conference on Web Application Development (2011)
7. Felt, A.P., Ha, E., Egelman, S., Haney, A., Chin, E., Wagner, D.: Android Permissions: User Attention, Comprehension, and Behavior. In: Proceedings of the 8th Symposium on Usable Privacy and Security (2012)
8. Gates, C.S., Chen, J., Li, N., Proctor, R.W.: Effective Risk Communication for Android Apps. IEEE Transactions on Dependable and Secure Computing 11(3), 252–265 (2014)
9. Harbach, M., Hettig, M., Weber, S., Smith, M.: Using Personal Examples to Improve Risk Communication for Security & Privacy Decisions. In: Proceedings of the 32nd SIGCHI Conference on Human Factors in Computing Systems (2014)
10. Kelley, P.G., Cranor, L.F., Sadeh, N.: Privacy as Part of the App Decision-Making Process. In: Proceedings of the 31st SIGCHI Conference on Human Factors in Computing Systems (2013)

11. Kelley, P.G., Consolvo, S., Cranor, L.F., Jung, J., Sadeh, N., Wetherall, D.: A Conundrum of Permissions: Installing Applications on an Android Smartphone. In: Blyth, J., Dietrich, S., Camp, L.J. (eds.) FC 2012. LNCS, vol. 7398, pp. 68–79. Springer, Heidelberg (2012)
12. Lin, J., Amini, S., Hong, J.I., Sadeh, N., Lindqvist, J., Zhang, J.: Expectation and Purpose: Understanding Users' Mental Models of Mobile App Privacy through Crowdsourcing. In: Proceedings of the 2012 ACM Conference on Ubiquitous Computing (2012)
13. Papamartzivanos, D., Damopoulos, D., Kambourakis, G.: A Cloud-Based Architecture to Crowdsource Mobile App Privacy Leaks. In: Proceedings of the 18th Panhellenic Conference on Informatics (2014)
14. Lin, J., Sadeh, B.L.N., Hong, J.I.: 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 (2014)
15. Ur, B., Kelley, P.G., Komanduri, S., Lee, J., Maass, M., Mazurek, M.L., Passaro, T., Shay, R., Vidas, T., Bauer, L., et al.: How Does Your Password Measure Up? The Effect of Strength Meters on Password Creation. In: Proceedings of the 21st USENIX Conference on Security Symposium (2012)