No technical understanding required: Helping users make informed choices about access to their personal data

Ilaria Liccardi[*†], Joseph Pato[*], Daniel J. Weitzner[*], Hal Abelson[*]
{ilaria, jpato, djweitzner, hal}@csail.mit.edu

David De Roure[†]
david.deroure@oerc.ox.ac.uk

[*]Computer Science and
Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, MA

[†]Oxford e-Research Center
University of Oxford
Oxford, UK

ABSTRACT

Many smartphone apps collect personal information used for a variety of purposes. Users, however, are often unaware of this kind of access even though they must grant the required permissions upon app installation. We have identified three reasons for this unawareness. First, relevant permissions can be missed in long lists of permissions. Second, apps that access personal information for functionality may appear suspicious even if they don't have the ability to disclose that information. Finally, updates to apps can lead to new permissions, accessing personal data, being granted.

We modified the Google Play permissions interface to include a quantitative measure (*sensitivity score*) of an app's ability to disclose personal information and to highlight the relevant permissions that contributed to this score in order to focus the user's attention on permissions that have the ability to access personal data. These improvements are easily integratable within the current structures and policies of the Android permissions interface and have been designed to allow inexperienced users to understand the permission interface and make *informed* and *conscious* decisions about access to their personal data.

We validated the effectiveness of this approach with a study of 125 Android smartphone users. We compared the current and improved versions of the interface and found that our improved permission interface led participants - especially inexperienced ones - to choose apps with less possible access to their personal data.

Keywords

Privacy, Android Apps, Personal Information

Categories and Subject Descriptors

K.4 [**Computers and Society**]: Privacy

General Terms

Theory, Human Factors

MOBIQUITOUS 2014, December 02-05, London, Great Britain

DOI 10.4108/icst.mobiquitous.2014.258066

1. INTRODUCTION

Millions of people use mobile phone apps for a variety of purposes, and both free and paid-for apps have become a lucrative business. Free and paid apps tend to add permissions that collect personal data which in some cases is not strictly required for functionality. Even if there is a legitimate need for access to users' personal data, that access may be reused for other purposes.

Unlike other operating systems, smartphone OSes have access controls that require authorization when personal data is accessed. This must be typically approved prior to installation of an app, or to an update (if permissions have changed). In our previous research [23] we found that 46% of apps had the ability to collect and transmit one or more types of personal information. When personal data is collected at this large scale, users should be made aware of access to this information in a simple and comprehensible manner.

The Android operating system provides more detailed transparency mechanisms than other smartphones, but the granularity and jargon of the displayed permissions is not easily understood by everyone [15]. In some instances the lists of permissions contain a variety of different types of access to the phone hardware itself, making the list long and confusing.

To help users better understand access to their personal data, we modified the Google Play permission interface. We use the *sensitivity score* [23] to display the number of sensitive permissions with the ability to access personal data when the app has the ability to disclose this information externally. The sensitivity score is zero when an app does not have any ability to disclose sensitive information and increases as the app gains more ability (i.e. permissions) to disclose information. This score is used to convey at glance, and in a clear and simple manner, how much information users might be giving away. This is used as an awareness mechanism rather than a quantitative measure. Users can investigate which permissions contribute to this score. We added a flag next to each permission that contributes to the score in order to allow users to make more informed decisions so to highlight permissions that threaten their privacy without having to read and understand every permission.

We conducted a user study to measure if these improvements can affect users' choices when selecting an app. Our results, described in section 6.2, show that improving the presentation of permission information in this way leads to the selection of apps with less possible access by users who do not have a clear understanding of an app's inner workings. Improving the presentation of permissions in this way helps users in several ways:

1. To identify safe apps[1] (apps that collect no personal data);

[1]We are only considering apps that can collect personal data, we do not imply that a

2. To help users focus their attention on the permissions with the ability to collect personal information;

3. To identify possible changes in the permission set when a new version is released: apps can change their permission requirements from one version to the next;

4. To make it easier for users to identify when an application transitions from being safe to having the potential to disclose data[2].

2. RELATED RESEARCH

Research from a wide variety of sources has shown how smartphone apps can both collect and infer a considerable amount of personal information about users. Patterns of mobile phone usage are valuable in detecting trends of behavior, especially for marketing [22] as well as customizing and personalizing services offered. It is possible to predict new app installations based only on information collected using the sensors found in smartphones [31], and it is also possible to infer the structure of friendship networks [7].

Past research [33] which analyzed applications gathered from the Google Play store showed that of 964 apps, 473 (49%) used nine different advertising libraries, and 110 requested multiple permissions that are only used for advertising [33]. The use of advertising within smartphones is widespread because sales of products which are advertised using mobile advertising increases compared to those which are not [27].

The collection of personal information allows marketers to increase the relevance of the adverts shown to users. However, while targeted advertising (using collected personal information) offers personalization and relevance to users, it also threatens their privacy: in order to provide such services, their preferences, behavior, and identity need to be tracked [40]. As a single example, the use of location information and user-preference data raises serious privacy concerns because these activities can be used to track users and infer their behavior [32], [39], [40].

Apps can intentionally or unintentionally [34] expose personal information to advertisers and expose personal data publicly, often without the user's knowledge [21]. Even when the app is in an 'idle' mode, it is not guaranteed that the app is not sending personal information [41]. Some apps do use personal information as a legitimate part of their operation [8]. Still, developers willingly or unwillingly often request more permissions than the app requires due to insufficient third-party API documentation [13]. Apps can also be malicious and leverage the privileges of another app through inter-process communication [6], [35].

Some developers provide free and paid versions of their apps, where the free version obtains revenue from advertising, while the paid version does not collect personal information. Users however tend not to buy apps even if they are as cheap as $0.99, in fact free apps are more popular and are downloaded far more than paid apps [23]. For developers it is often more lucrative to have a free app that uses advertising for revenue [25] than a paid app that requests no permissions, to target the right audience. However, users are often unaware of this access and when alerted to it, report surprise and the desire to remove the app [1], [24].

Obtaining personal information via mobile phone apps has become popular (due to its lucrative aspects) and because of this, privacy in mobile phones has become an important topic for research [36] to the extent that the European Commission [17], [29], the US Federal Trade Commission [12], [10], [11] and the US National Telecommunications and Information Administration [30] are analyzing and providing guidelines for app store markets and app developers to improve mobile privacy.

Despite concerns from users and policy makers about the privacy practices of mobile apps, existing approaches to evaluate and act on privacy practices have considerable shortcomings. Users generally have some awareness about mobile privacy issues; nonetheless, many still do not take steps to protect their privacy [16]. Researchers have tried to understand how people protect their personal data stored on mobile phones [4] and, where they do not, the reasons why [15], [5] and their perceptions of risks related to privacy leaks [14].

A 2012 Pew Internet & American Life Project report showed that more than half (57%) of users interviewed (2,254 adults age 18+) did not install apps when they *realized* that personal information could be collected, or removed apps from their phone if they found that personal information was collected [4]. However, not all users are able to recognize these situations, and apps which collect personal information are still extremely popular. We know that users have a difficult time understanding conventional privacy statements [26], often do not understand the technical jargon explained in the permissions [15], or are completely unaware of the personal information that they are sharing and need to be educated on the dangers posed to their privacy [37]. Others think that they have nothing to hide or that there is no danger to them [38].

To increase transparency and individual control, researchers have tried different approaches. Meurer & Wismuller [28] allow the users to filter apps by permission type [28] while Barrera et al. [2] propose a method to improve app permission expressiveness without increasing its overall complexity. Kelly et al. [18] proposed a "nutrition label" for privacy which has proven effective in allowing users to find relevant information quicker and more accurately.

Researchers have also enhanced Android itself in order to monitor the flow of information leaving the phone. Enck et al. [9] developed TaintDroid, a modified version of Android able of providing real-time analysis and tracking information that leaves the phone. The TaintDroid approach requires a modified version of the Android virtual machine to be installed on the phone by jailbreaking it. While it tracks information, it does not allow the user to stop the information from being distributed. Mockdroid [3] tries to tackle this problem by allowing users to revoke access to particular permissions at run-time, sacrificing functionality to stop collection of (and hence disclosure of) personal information. While these tools provide useful information and approaches to understand the inner working of apps, they are hard to set up and require specialized knowledge and technical skills.

3. PERMISSION ANALYSIS FAILS USERS

Apps request different permissions according to what access they need from the smartphone – both in terms of access to hardware and to personal information. Several studies have previously noted users not being able to understand privacy risks associated with mobile apps. The reasons were shown to be rooted in current transparency mechanisms, which ineffectively convey risk and level of access of different types of permission [19] and/or lack of attention to the permissions themselves. [15]. Other factors include the fact that recommendations from friends or family take precedence in choosing apps [5].

However, when users are guided in understanding permission access, research has shown a common reaction of "surprise and shock" in reaction to unexpected permission requests [24], and un-

safe app is free from malicious function

[2]Users can decide not to update the app and hence keep the *current* less intrusive version (with less access to their personal data) on their smartphone, which will still function for a time

foreseen patterns of collection of personal data [1] to the extent that some users consider removing apps that display this behavior. While conveying this information to users has proven successful within these studies, setting up how this information is displayed requires specialized knowledge and tools to monitor the behavior of apps. This requires deeper knowledge of the Android permission requests and operating system as well as rooting the device itself (when real time monitoring is used).

To analyze how permissions are used, we collected metadata about apps on the Google Play store[3] at two separate occasions[4]. The first dataset was collected from October 2012 to January 2013, gathering information about 563,528 apps. We then collected the second dataset from March 2013 to May 2013, collecting 635,264 apps. In analyzing the metadata we identified three reasons why users might have problems in understanding this information:

- **Too many permission requests**: Permissions that have the ability to access personal information can be embedded in otherwise innocuous permissions, making them difficult to notice.

- **Misunderstanding of legitimate apps**: Apps might request permissions that access personal information and appear suspicious, but do not have the ability to transmit data outside the phone (since communication-based permissions - either Internet connectivity or write permissions - are not requested).

- **Apps requesting new permissions – with access to personal data – in updates**: Apps can change permission requests between versions, so an app might not initially request access to personal information, but add that permission in a later version. Apps can also initially request access to personal data but withdraw that request later.

3.1 Too many permission requests

A large percentage of apps (52.8%) had the ability to access and collect varying levels of personal information stored on users' smartphones (sensitivity score ≥ 1). Free apps had the ability to access and collect on average more personal information that paid ones (Figure 1).

In apps where sensitivity score is ≥ 1, the sensitive permissions are often combined with indifferent permissions, making them harder to identify. Figure 2 shows all possible patterns of combination between sensitive permissions and indifferent permissions. App permission sets can vary from having a small number of different types of permission to varying numbers of permissions. For example, an app can have a small number of sensitive permissions together with a large number of indifferent permissions, or have a large number of sensitive permissions with a small number of indifferent permissions. Apps in the wild (Figure 2) were found to request up to 121 permissions (122 with network access).

[3] At the time of this study, the permission information was part of each app page. However, Google recently changed the way that Google Play works in the browser making the fetching of the data needed for this analysis more difficult and time-consuming. Permissions needed for each app are only reported when the *"install"* button is pressed and when the browser's user is logged in to an account associated with a smartphone compatible with the current app.

[4] We gathered different apps by performing searches for dictionary words on the Google Play website, and retrieving the page for each app that was found. The search results are split onto multiple pages, so we retrieved each page of search results; Google Play enforces a maximum limit of 20 pages of results for any given search. We used different dictionaries to collect the apps. A large English dictionary and dictionaries for French, Italian, and Spanish were combined to create different queries. The Google Play website enforces rate limiting if a large number of requests are made; we therefore included logic that would detect error messages, pause and retry.

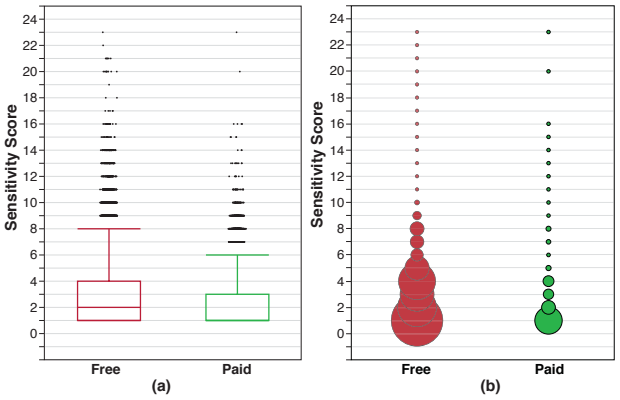

Figure 1: 335,869 apps (289,576 free and 46,293 paid) with $Sensitivity\ Score > 0$ showing (a) boxplot of sensitivity score values between free and paid and (b) bubble plot showing the concentration of apps in each sensitivity score value between free and paid apps.

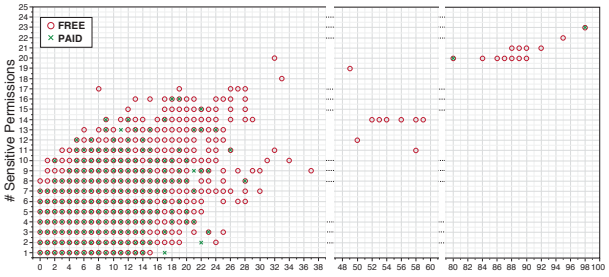

Figure 2: Number of sensitive permissions vs. number of indifferent permissions in each app with sensitivity score ≥ 1.

3.2 Misunderstanding of legitimate apps

Even though an app has access to sensitive permissions, it does not necessarily have communication-based permissions needed to disclose personal information.

Table 1: Number of sensitive permissions in apps with $Sensitivity\ Score = 0$

# SENSITIVE PERMISSIONS	# TOTAL	# FREE	# PAID
1	19,729	5,812	13,917
2	2,124	1,020	1,104
3	1,085	591	494
4	583	328	255
5	276	133	143
6	127	76	51
7	50	30	20
8	16	9	7
9	7	3	4
10	4	2	2
13	1	1	0

24,002 apps (from the most recent dataset) requested sensitive permissions, but did not request access to any communication permissions (either network or write access) that would allow data to be sent outside the device (*sensitivity score = 0*). A breakdown of the number of sensitive permissions with a corresponding number of apps is shown in Table 1, showing the counts for total, free and paid apps.

3.3 Apps requesting new sensitive permissions in updates

We compared the two datasets that we collected from Google Play. 411,101 apps were common to both sets, of which 302,823 were free and 108,278 paid. We identified 56,229 apps that changed their *sensitivity score* either by increasing or decreasing it:

Increase in Sensitivity Score: 52,430 apps (Figure 3 (a) & (c)) increased their sensitivity score. 20,051 apps increased from an initial value of 0 while 32,379 increased from a value already > 0.

- **Apps that increased their sensitivity score from 0**: These apps did not have the ability to collect any personal data about users, However, they show an average increase of 1 (Figure 3).

- **Apps that increased their sensitivity score from n (where** $n > 0$**)**: These apps already had the ability to collect users' personal information. However, the requested permissions changed from a $median = 3$ with a maximum of 9, to $median = 5$ with a maximum of 13. This shows a significant increase in the information that these apps could potentially collect about users (Figure 3).

Decrease in Sensitivity Score: 3,799 apps (Figure 3) decreased their sensitivity score, i.e. reduced the amount of information that they could potentially collect about users.

- **Apps that decreased their sensitivity score**: These apps reduced the amount of data that they could collect about users from a $median = 3$ with a maximum of 7 to a $median = 1$ to a maximum of 5 with 1,422 reducing it to 0 (Figure 3).

4. IMPROVED INTERFACE

We enhanced the current Google Play permission request interface to allow users to better understand an app's ability to collect and transmit personal data (Figure 4) by providing an overall view of the different pieces of personal information accessed as well as focusing attention on the permissions that have the ability to read personal data. We first embed the sensitivity score [23] within the search page (Figure 4 (a)) to allow users to easily distinguish between apps with different types of access to their personal data. This score could also be displayed prior to installation (Figure 4 (b)) to make sure users are aware of this access. We draw attention to each permission contributing to the score (permissions that have access to the user's personal data) within the conventional permission list, (Figure 4(d)) both when choosing and updating an app. When a new update is available, this score alerts the user when an app is transitioning from being safe to potentially starting to collect their personal information, or when the app might be increasing their collection behavior (Figure 4(c)).

These improvements are minimalistic changes to the current permission interface and can be easily integrated within the current structures and policies of the Android permissions interface. Previous research has created more detailed permissions improvements [20] and privacy policy styles [18] which have been shown to be easily understood by users. These approaches, which describe usage, sharing and concrete access to users' personal data, would be very beneficial to users; however, they would require extensive restructuring and extensions to Android's underlying API, with checks that permissions really are used as specified.

Android groups permissions by access/functionality, e.g."personal permissions" includes access to contacts, bookmarks etc. Hence, creating a "Read personal permission" and "Third party connection permission" (e.g. Internet permission) or flagging permissions with

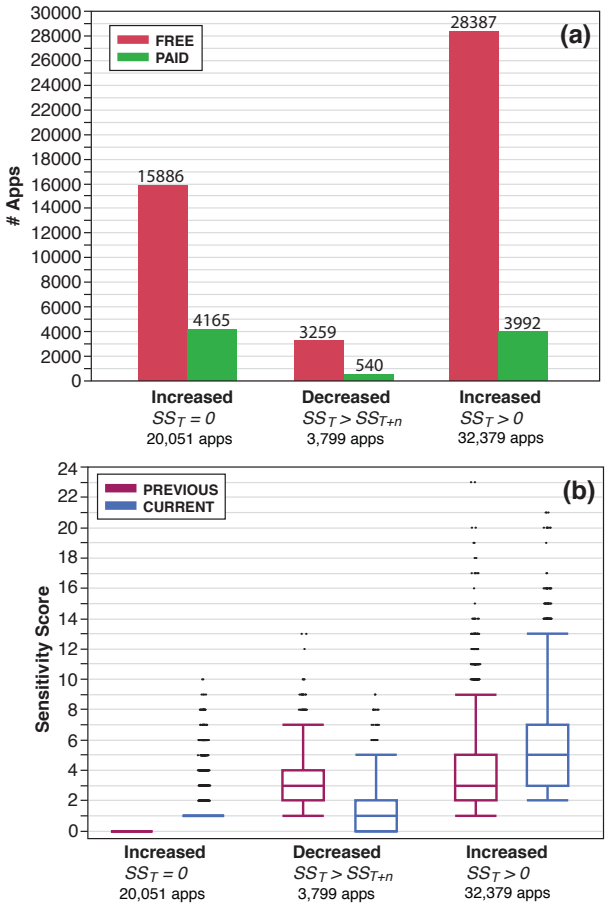

Figure 3: Showing (a) Distribution and (b) Box plots of 56,229 (free and paid) apps that have changed their sensitivity score; showing apps that increased their sensitivity score from 0 $(SS_T = 0)$**; apps that have decreased their sensitivity score** $(SS_T > SS_{T+n})$ **and apps that increased their sensitivity score greater than zero** $(SS_T > 0)$**.**

the ability to access personal data (when this data can also be transmitted) could be an attainable and feasible change. Given these kinds of awareness, users have the ability to choose apps with less access to their personal data, which not all currently have the informedexpertise to do [15]. Highlighting all possible access might make developers more inclined to disclose the reasons for that access.

5. USER STUDY

We conducted an online survey to gather attitudes regarding users' understanding and concerns about privacy when choosing an Android smartphone app. Survey participants were recruited through solicitations to mailing lists, CraigsList, and social media.

The survey consisted of 108 questions: 103 unique questions and 5 repeated questions. We informed participants that 30% would be randomly rewarded $20 if their answers were genuine and consistent with previous answers. We told them that we would check all the answers[5] prior to inserting them into the reward pool and that

[5]We recorded time for each question to be answered when "next" was pressed, total

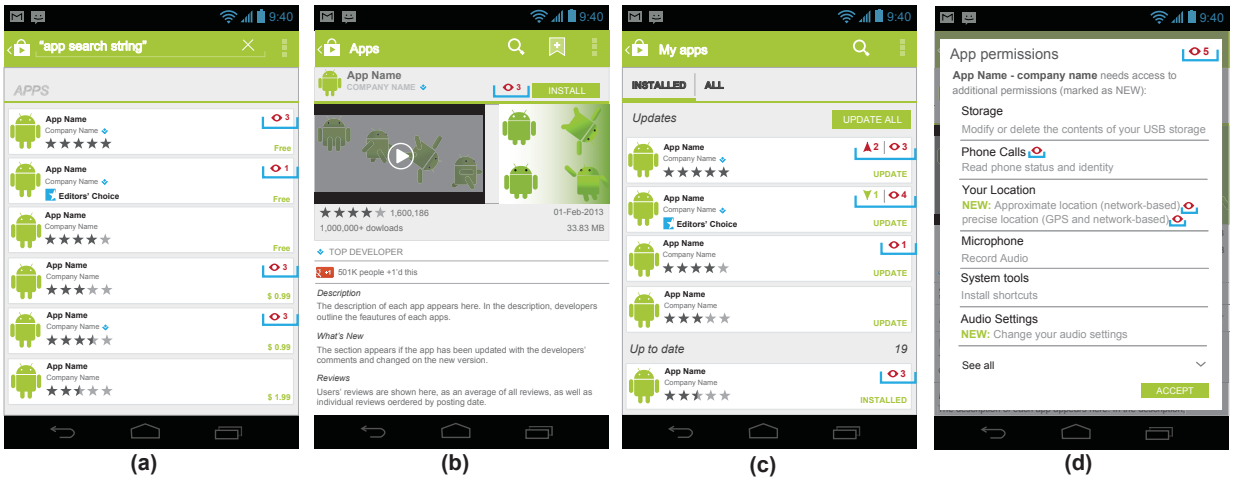

Figure 4: Presentation of improved interface when *choosing* an app, showing how the sensitivity score could be included within the Google Play store: (a) when users search for an app; (b) within each app's page prior to installation; when *updating*: (c) showing the increase (or decrease) in value of the sensitivity score next to the "update" button; (d) showing the total sensitivity score within the permission list and marking each permission (this improvement can be used both when installing and updating).

if one answer was not deemed genuine or did not match or was conflicting with previous answers that they would be automatically removed and not entered in the pool.

The survey posed questions in three segments. We first collected demographic profile information about each participant. This included age, gender, origin, current country of residence, educational level and native language(s). In the second segment we collected information about participants' current approaches to app usage. Finally, the third segment required participants to choose between pairs of applications based on the interface for presenting permission information. Within each of the second and third segments, questions were randomized except as outlined below.

5.1 Approach of users to app selection

In the second segment of questions, we looked at the participant's usage of smartphones and collected the number and names of apps they have installed and frequency of usage. Participants' attitudes towards the types of app they search for on the Google Play store was also recorded. In particular we were interested in motivations for searching for free, paid or either type of app.

We then looked at what information participants most value and focus on when trying to decide which app to choose. Using a 5-point Likert scale we asked participants to rate how often they view particular information available for each app (Table 3). Subsequently we asked them to rank this information in order of priority.

We then asked them to rate their comprehension of permission information (Table 4), and followed with questions related to permission information. In particular we were interested if participants look at or value having this information prior to installation and their motivations behind these actions (positive or negative).

We also asked participants to rank information they would be most likely to use to compare apps when undecided about which to choose (Table 3). In this question we also made a distinction between fewer permissions and fewer permissions that read/access personal data.

We also enquired about their update practices. In particular we asked participants about update settings, whether automatic update was on or off (participants could also specify uncertainty about this option). We then enquired about the reasoning behind this choice. We asked participants who do not enable automatic updates when they decide to update an app, and what kind of information they are most likely to check when an update is available.

We also asked all participants if they look at permission information when they are prompted, and if not, the reason for this choice. We asked them to select an option that would describe how they update an app when permissions have changed (Table 6). We asked all participants to rank in order of priority what kinds of information they look at when a new update is available (if permissions have changed). We enquired about removing an app and motivations behind this choice.

After all other questions in this segment were posed, we enquired about attitudes towards privacy policy. In particular we probed to determine if participants check if an app has a privacy policy and if they read it when it is present as a link or part of the description of an app and motivations for their answers.

5.2 Effect of improved permission interface

In the final segment, we posed 52 questions designed to compare our improved permission interface with a conventional interface and gather motivations for participants' choices. We created 13 questions which compared permission lists of two apps. The compared app lists were chosen from apps of the same category and similar purpose (games, wallpaper etc.), hence the disparity in access needed to be highlighted when searching for apps of similar functionality.

The list of permissions presented different patterns of aggregation of sensitive permissions (permissions that have the ability to read personal data), indifferent permissions (permissions that do not have the ability to read personal data), communication permissions (either network or write access that allow connection or communication with other apps and transmission of information from the phone) and indifferent permissions with write access to personal information. The length of the list of permissions varied between apps in eight questions but was kept constant – either short (2 questions) or long (3 questions) – in the remaining five. The

time to take the survey, inserted five repeated questions to ensure participants were consistent in their answers, and checked that participants' motivations reflected the chosen app features.

presence of communication permissions (we used network access within our study) alternated between apps in seven questions, and was constant in the six remaining questions. The description of the permission itself as presented in the Google Play store was also included. When app lists of equal length were displayed in the study the description of the permission was considered to make sure that visually the two permission lists looked of similar length. A list of all questions with corresponding types and number of permissions for each app is shown in Table 7.

We asked participants to select which app they would choose if their decision was based solely on the permissions each app requested. Participants could decide to choose either one of the apps (app 1 or app 2) or select the "either one" choice. We presented the permission list in two ways: 1) using the conventional permission interface and 2) using an improved interface. The improved interface is based on the sensitivity score for each app. We explained in each question the reasoning behind the sensitivity score and we inserted it at the top of the permission list as a numeric value, also highlighting the permissions that contributed to the score (Figure 4(d)). Participants first viewed all questions using the conventional permission interface and then viewed the same questions using our improved interface (question order was randomized within each presentation interface). After selecting one app – in both the conventional and improved interface – each participant also had to provide rationale for their choice by either selecting one of the provided answers or by choosing to specify their own in free-from text. For each question, the "correct" response was the app (or both) with ability to access and transmit the least personal data (Table 7). In particular we are interested in:

1. Can our improved interface help users who are not familiar to *simply understand* the complexity of permission requests?
2. Did our improved interface *affect* users' choices in *selecting* apps?
3. Does awareness of access to personal data change users' choices in selecting an app? If so, do users tend to select apps with less access to personal data?

6. RESULTS

We recruited 170 participants using CraigsList and mailing lists. We discarded 45 participants: 27 participants did not complete the entire survey, 8 participants took less than the minimum measured time required to complete the survey[6], 10 participants were discarded because their answers were not consistent ie. motivation for answers did not match the answers chosen[7].

Participants (125) were 48 females (average age of 29) and 72 males (average age of 36), with five participants who opted not to disclose gender (average age of 36). 63.2% of participants had more than 10 apps currently installed on their phone, 33.6% had between 5-9 apps, and only 3.2% had 1-2 apps (these apps are in addition to apps already installed by default on participants' smartphones). The level of education of participants was broadly distributed from having completed elementary school to advanced graduate degree (Table 2). The survey was completed online and it took between 33 minutes and 2 hours to be completed with an average of 53 minutes across all participants.

6.1 Approach of users to app selection

[6]We piloted the study with 10 participants; The average time to complete the pilot was between 30 and 55 minutes with an average of 41 minutes.

[7]For example participants in these groups selected apps with the longest permission lists but selected as motivation the app with fewer permissions.

Table 2: Participants' distribution of Educational Level

	PARTICIPANTS (125)	
EDUCATION LEVEL	**#**	**%**
Elementary school only	10	8
Some high school, but did not finish	7	5.6
Completed high school	20	16
Some college, but did not finish	11	8.8
Two-year college degree / A.A / A.S.	6	4.8
Four-year college degree / B.A. / B.S.	15	12
Undergraduate student	12	9.6
Some graduate work (e.g. master student)	3	2.4
Graduate student	20	16
Completed Masters or other professional degree	12	9.6
Advanced graduate work or Completed Ph.D.	9	7.2

In this section we report participants' responses to current approaches for choosing and updating apps as well as highlights into which of the app's information is most valued and used by participants during their decision making process.

6.1.1 What information do users value when choosing apps?

Participants presented different preferences which influence their decisions when choosing apps. Reviews and ratings from other users take precedence either in choosing, or when undecided between two apps (Table 3). In particular participants reported looking within the reviews to see how other users rate the app, focusing on if the app presents *"excessive use of advertisements"* which could compromise the actual app experience. Price and description are also considered high priorities (Table 3).

Table 3: Priority of available types of information about each app showing when choosing (with 8 being the highest priority) and when undecided between apps (with 9 being the highest priority) prior to installation.

Factors influencing users' choices	**CHOOSING**	**COMPARING**
Other users' reviews	6.88	7.37
Price	6.99	6.01
# of downloads	6.33	6.51
App description	6.20	6.34
App screen shots	5.08	4.79
Content ratings	3.48	2.78
Permission list	3.37	
Fewer permissions	-	3.38
Fewer personal permissions	-	2.92
Privacy policy	2.02	2.21

While the use of advertisements plays a role in choosing the app itself, this does not include the use of their personal information. In fact, information regarding which permissions the app requires is not rated a high priority on average for participants. When trying to decide between two apps, participants on average reported choosing apps with fewer permissions in total than apps with fewer permissions that access personal data (Table 3). The presence of privacy policies was rated lowest priority on average.

The reason why permissions have not been rated high on a list of priorities is probably due to the participants' level of understanding of the permissions themselves. We asked participants to rate their comprehension and understanding of permissions, by providing four pre-filled selections and an additional option (free-form text) where they could specify their understanding in their own words (Table 4). Only 20% reported having a clear understanding of each of the different permission requests, while 24.8% of

participants reported not understanding permissions at all. The remainder of participants reported selected a moderate understanding of permissions (Table 4).

Table 4: Sensitivity Score (mean) and number of apps (excluding the ones already pre-installed in the phone) showing average and median for participants grouped according to their self description of permission understanding. The percentage of participants in each self description category is also shown.

	SUMMARY OF LEVEL OF PERMISSION UNDERSTANDING[1]	S.S.[2]	NUMBER OF APPS AVG.	MED.
1	I am very familiar with permission requests (20%)	1.5	8.64	8
2	I am somewhat familiar with permission requests (24%)	2.6	10.28	9
3	I understand some of the permission requests (32.8%)	4.7	12.72	11
4	I do not understand permission requests (24.8%)	6.5	14	13

[1]For a detailed description of the provided descriptions of permission understanding shown to participants of refer to Table **??**

[2]This value represent the mean of the Sensitivity Score (section 3.1). The sensitivity score measures the amount of personal information that could be collected by apps.

The difference in understanding of permissions is correlated with the number (0.94) and type (0.99) of apps currently (at the time of this study) installed by participants - less or more possible access to personal data measured as the sensitivity score (section 3.1). Participants who are more familiar and understand permissions in detail (1) tend to have fewer apps in general with a lower sensitivity score, which means less possible access and ability to collect personal data (Table 4). Participants with a lower (self-measured) level of understanding of permissions present a higher number of apps with more ability to collect personal data.

The number of permissions and ability to collect personal data increases as the level of understanding of permission decreases (Table 4). Participants who have declared not understanding permission requests had apps with the highest sensitivity score (6.5) and had downloaded (currently present on their smart phone at the time of the study) on average more apps (14) than the other participants with better understanding of permissions.

6.1.2 Do users look at Privacy Policies?

While permissions might provide users who are able to understand them with an overview of the kinds of access the app requires to function, the only way to understand the inner working of the apps, with respect to how user data is used, is with a privacy policy. For example a weather app might request "location information" and "Internet access" to provide users with the correct weather for that particular location. However, in the background it might also be sending that information to the advertisements system in order to produce targeted advertisements tailored for the specific location to be displayed to users. This behavior cannot be predicted just by looking at the permission list on its own. This kind of information and behavior of the app should be disclosed within the app's privacy policy.

Including a privacy policy within an app's page is not common or required in the Google Play store (or other markets). Only 8.7% of the apps in the collected dataset had a privacy policy listed on the app's page (55,598 apps (21,078 companies) of the 635,264 apps recently collected compared to 38,568 apps (14,788 companies) in the 563,528 apps previously collected).

While *privacy policies* can provide invaluable information about the inner workings of the apps and users' personal data collection and usage, its presence or usefulness has been rated on average to be the lowest priority in the decision making process (Table 3). When we asked participants whether they check if apps have a privacy policy, 81.6% reported not looking, while 18.4% reported looking to see if the app has a privacy policy, even though only 11.2% reported that they read it prior to installation (Table 5).

Table 5: Attitudes of participants towards privacy policies for apps. Participants could select only one choice.

Motivations for participants (81.6%) not looking for privacy policies	
Privacy policies are too long to read	43.13%
Privacy policies are full of incomprehensible legal jargon	24.31%
I would not know how to use a privacy policy	7.05%
Others	5.4%
Motivations for participants (18.4 %) looking for privacy policies	
To understand how the developers plan to use my data before installing the app.	11.2%
To ensure that a privacy policy exists.	2.4%
I believe apps with privacy policies are more careful with my data	2.4%
I feel more secure downloading apps with privacy policies	2.4%
Other	0

When we enquired about motivations for participants' decisions, they reported that privacy policies are usually very long, vague, and not easily understood (Table 5). 7.05% of participants reported additional reasons for not reading or even looking for them within the app's page, with two participants expressing defeat about safeguarding their own privacy and one relying on popularity to influence his decision:

> P32: *Well, I sometimes do, depending on the kind of app. But usually it's too much bother, since I presume that at least some of my apps on Android are leaking information to advertisers anyway.*

> P86: *I just assume that if it has enough downloads its fine with most people and thus fine for me. I have given up trying to keep things on my phone private.*

Another participant underlined the imprecision and vagueness of privacy policies, especially when disclosing collected personal information to undefined third parties.

> P5: *I have tried to read privacy policies but it is impossible to understand, sometimes they say we will disclose your information "only" to associated third parties without actually telling who they are.*

Privacy law and practice depends on the ability to make informed decisions about how personal information is used. Few smartphone applications currently provide a privacy policy. Of those that do, many fail to provide users with useful information. Policies attached to applications should be concise, easily understood and reasonably viewed within the constraints of a smartphone screen.

6.1.3 Choosing between Free or Paid Apps?

77.6% of participants highlighted that they always choose free apps, and 3.2% always choose paid apps, while 19.2% tend to

choose either one. Participants always choosing free apps do so because they do not want to pay for apps. They reported always being able to find a free version of an app they require, and either did not understand that free apps make money with advertisements, or they did not seem to care. We have seen in our collected data set that there are three times as many free apps as paid apps within the Google Play store. Only 3.2% of participants tended to always pay for apps due to not wanting to have advertisements or to let developers collect their personal data. The remainder of participants reported choosing either one depending on the app. However they reported first installing the free version and then changing to the paid version: in some cases to *reward the developer* or to *access more features or functionality* that the free version did not possess.

6.1.4 Updating Patterns

62.4% of participants had "automatic update" turned on while 27.2% did not. The remainder (10.4%) showed uncertainty about this setting on their smartphone.

Table 6: Update practices of participants when an update requests additional permissions. Participants could select only one choice.

	AUTOMATIC UPDATE SETTING		
	ON	OFF	Uncertain
Even though permissions have changed	40	–	10
Only if the app offers new features	22	6	3
Only if I understand the new permissions.	10	7	–
Only if the app does not add permissions with access to personal data.	6	12	–
Only if the new features and the new permissions are reasonable*	–	6	–
I never update apps	–	3	–

*This option was reported by 6 participants as a free form text option.

However when new permissions are added, even users who have automatic update turned on would be notified of the changes. 59.2% of participants reported checking the new permission list once they are prompted about the new update (all participants who have automatic updates off belong to this group). 28% reported not looking at the permission list, while the remainder reported that it depends on the app. They stated that they tend to look at and review the permission list for apps that they use less frequently, while not checking and updating immediately for apps that they use often. More than 50% of participants who had automatic updates turned on reported not looking at the new permission list once they are prompted (Table 6) with only 20% basing their choices on permission information. Of participants who had automatic updates off, 73% based their decision on the new list of permissions, with 2.4% of participants never updating apps which change their permissions.

6.1.5 Deleting apps

97.6% of participants reported having deleted an app on their phone. However only 5.1% of these participants removed an app because of unwanted collection of personal data. Four major reasons were reported by participants for removing an app: not needing or using it anymore, excessive use of advertisements, malfunctioning, or misleading description.

6.2 Effect of improved permission interface

In this section we report on how participants selected apps based solely on permission information. Participants were asked to choose between two apps, with the additional option of "either one". They were then asked to provide a reason for their choice. Each pair of apps was presented using both the conventional and the improved interface. The permission lists shown to participants in each question are shown in Table 7.

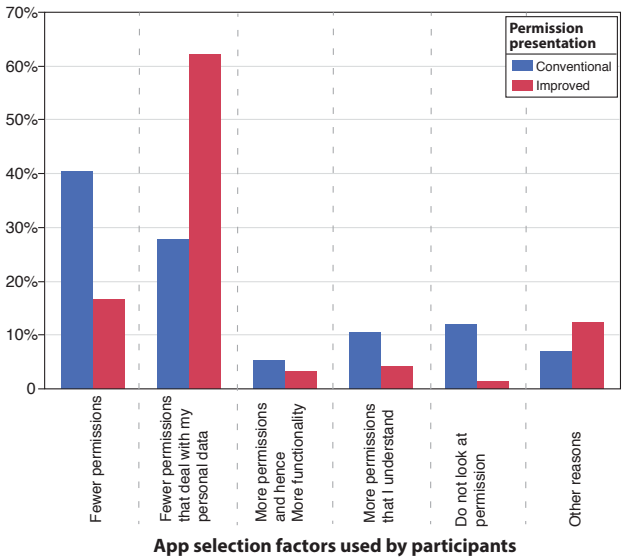

Figure 5: Factors used by participants when selecting apps according to interface used

Figure 5 compares the motivations behind the selections made by participants when presented with the two permission interfaces. We can see that when using the current interface, participants were prone to choosing apps with fewer permissions regardless of the type of permission itself: the length of the list of permissions is a factor influencing participants' choices. However, when using our improved interface (Figure 4 (d)), participants were more likely to choose apps with fewer personal permissions - even though these permissions were included with other (indifferent) permissions that do not access personal information - which made the lists longer (Figure 5). Also 15 (12%) participants initially stated that they do not look at permissions, however when presented with the improved interface this dropped to only 2 (1.6%) participants.

Table 7 & 8 shows the responses of participants to each question. Table 7 shows the overall responses by all participants for each question, displaying the type of permissions in each app within each question, desired answers as well as number of participants who changed their choice to apps with either less or more access to personal data. Table 8 shows the responses to each question grouped by participants' self-comprehension evaluation level of permission understanding, showing the total number of participants in each comprehension level as well as the number of participants choosing the desired answer with the conventional and improved interface and corresponding changes to the app with either less or more access to personal data. We can see from Table 7 & 8 an effect on participants' choices between the two interfaces.

The way the information is presented with our improved interface can allow people who are unclear about permissions (especially self described comprehension levels from 2-4) to visually gather more details about an app's access to their personal information and hence make informed choices focusing their attention only on permissions that grant access their personal data (Table 8).

Q1, *Q5*, *Q8* show smaller improvements than the rest of the questions, but this is due to the fact that the desired app had fewer total permissions, which is how participants tend to make choices when using the conventional interface.

However, when the desired answer is part of a longer list of permissions, regardless of the type of permission, participants overwhelmingly changed their choices to prefer apps with less access to personal data (*Q4*, *Q6*, *Q9*). For example *Q4* and *Q6* both presented lists of permissions of varying length. In both questions, the shorter permission list grants access to three types of personal information and one network access, which allows the app the ability to transmit information.

However while in *Q4* the longer list contains nine sensitive permissions, *Q6* contains 11 indifferent permissions and 1 network access. In responses to both questions, when using the conventional interface, participants overwhelmingly chose the app with fewer permissions, which had the ability to collect and transmit personal data (Table 7). Only participants who self-described themselves as having a deeper understanding of permission information (level 1 in Table 8) could choose the app with less access to their personal data. Participants with less understanding of permission information (level 2 to 4 in Table 8) chose the app with the shorter list.

However, when participants were asked the same questions again using our improved interface, 93 (91 initially chose the less desirable app and 2 selected 'either one' choice) and 86 (83 initially chose the less desirable app and 4 selected 'either one' choice) participants changed their choices by selecting the apps with more permissions but less ability to access and transmit their personal information (Table 7). This result shows that when users - especially users who are unfamiliar with or do not understand permission information (Table 8) - are made aware of possible access to their personal data, their choice changes preferring apps which are less likely to collect their data.

Participants were better able to distinguish between sensitive and indifferent permissions when presented with shorter permission lists (*Q3*). This ability to distinguish was shown to be reduced when longer lists of permissions were used that combined sensitive and indifferent permissions (*Q11*, *Q12*, *Q13*).

In *Q3*, participants selected the less risky app (the app with less access to personal data) when using the conventional interface, but were less successful when presented with a longer list of permissions (*Q11*, *Q12*, *Q13*). They particularly seemed to have trouble understanding the difference between the presence of sensitive permissions and the ability to transmit this information.

For example in *Q11*, the length of the permission list was kept constant while varying the type of permissions and whether network permission was present (app 1 presented 11 SPs & app 2 presented 3 SPs + 1NP + 8 IP). App 1, while suspicious, does not have the ability to transmit data, while app 2 requests network access and three sensitive permissions that grant access to personal data. When using the conventional interface, only 45 participants (the majority of whom are participants who understand permission information (1) in Table 8) selected the less risky app (app 1), 61 chose app 2, while 19 chose 'either'. Participants selecting app 2 were unaware that app 1 could not transmit any data and hence mistakenly selected app 2 under the impression that it presented less of a threat to personal information (97% of these participants reported this motivation for their choice). However, when presented with our improved interface, 62 participants (50 initially chose app 2 and 12 selected the either one option) were able to distinguish between the two set of apps and chose the less risky one (Table 7).

Participants (especially participants with less understanding of permissions (3 & 4) in Table 8) were also confused about the difference between read and write permissions. *Q13* presented two apps requesting either read access or write access: both apps requested network permissions. 60 participants chose either the riskier app (40) or selected 'either' (20) when using the conventional interface. 45 of these participants (33 of whom initially chose the riskier app and 12 of whom initially selected 'either one' choice) changed their answer to the app with less access using our improved interface.

Participants chose apps with fewer permissions because of a limited understanding of permission requests (Table 8). From Table 8 we can see that participants with better understanding of permission requests tended to choose apps with less access to personal data in both types of interface. Our improved interfaces allow participants with less understanding of permission information (2,3,4) to choose apps with less access to their personal data.

Our interface has shown improvement when apps present no access (hence possible disclosure) to personal information. This is the case whether the permission list is constant (*Q2*) or varies in length (*Q7*, *Q10*). Participants showed an improvement in understanding this difference and changed their choice to the 'either' option – giving reasons related to the inability to access personal data in both apps. However since neither app had any access to personal data, while we record the right choice to be "either", participants answering this question would always select the right choice.

Having this improved interface will allow users to better understand permissions information, especially users who are unable to do so already. Users can then willingly decide whether to give their personal information in exchange for the app rather than be unaware and oblivious to this access due to their lack of understanding of the inner workings of permission information.

7. CONCLUSION

We have shown that people are often unaware of possible collection of personal information. When faced with decisions based only on permission information, users tend to choose apps with fewer permissions (Figure 5) regardless of type and/or access. This lack of awareness and understanding makes permission information a low importance factor when deciding to choose (Table 3) or update an app (Table 6). We have introduced an improvement to the Google Play store that allows users, especially those who do not have a deep understanding of permission information, to better understand possible access, and be guided to the permissions that have the ability to read personal data when an app has the ability to transmit it (Figure 4). Our results show that when users are made aware of possible access to their personal data, they favor choosing apps with less access (fewer sensitive permissions) rather than basing their selection on the total number of permissions requested. We understand that various other factors are considered (and in same case preferred) when choosing an app such as functionality, popularity etc. rather than just the amount or the type of access to personal information. However users should be made aware of this access in a simple manner rather than being confused and unknowingly granting access. This study has shown that current transparency mechanisms are not suitable for all users and should be improved to alert users of such access. These improvements suggested are minimalistic to the current permission interface and can be easily integrated within the current structures and policies of the Android permissions interface. Applying these improvements might also encourage developers to describe the *purpose* of access to personal permissions or add a simple and concise privacy policy. This is particularly relevant if users' personal information is used for functionality and/or fed to the advertisement system or other purposes, since having personal permission requests (as we have shown in the survey) might have an effect on users' choices.

Table 7: Survey's results showing permission lists of both app in each question (section 5.2), the desired app choice (app with less access), the number of participants selecting the desired choice using the conventional and improved permission interfaces. The effects of using our improved interface is reported using the number of participants who changed their choice of apps.

	TYPES OF PERMISSIONS[1]			# PARTICIPANTS SELECTING DESIRED ANSWER		EFFECT OF IMPROVED PRESENTATION	
QUESTIONS	APP 1	APP 2	Desired Answer	Conventional	Improved	Less Access	More Access
Q (1)[4]	9 SPs + 1 NP	3 SPs + 1 NP.	APP 2	111	117	9	3
Q (2)[2]	5 IPs	5 SPs	EITHER	5	72	67	20
Q (3)[2]	5 IPs	4 SPs + 1 NP	APP 1	96	106	21	11
Q (4)[4]	9 SPs	3 SPs + 1 NP	APP 1	10	102	93	1
Q (5)[4]	9 SPs + 1 NP	3 SPs	APP 2	118	119	5	4
Q (6)[4]	11 IPs + 1 NP	3 SPs + 1 NP	APP 1	24	108	85	1
Q (7)[4]	11 IPs + 1 NP	3 SPs	EITHER	4	72	71	3
Q (8)[4]	3 SPs +11 IPs + 1 NP	3 SPs + 4 IPs	APP 2	108	117	15	6
Q (9)[4]	3 SPs + 12 IPs	3 SPs + 3 IPs + 1NP	APP 1	14	103	91	2
Q (10)[4]	9 IPs + 1NP	3 IPs + 1NP	EITHER	7	34	28	1
Q (11)[3]	12 SPs	3SPs + 1 NP + 8 IPs	App 1	45	105	62	2
Q (12)[3]	11 SPs + 1 NP	3 SPs + 9 IPs	App 2	81	108	34	7
Q (13)[3]	7 SPs + 1 NP	7 WPs + 1NP	App 2	65	107	45	3

[1]**SPs**: Sensitive Permissions (permissions with the ability to read personal data); **NP**: Network Permission (ability to transmit data over the Internet); **IPs**: Indifferent permissions (permissions with no access to personal data); **WPs**: Write Permissions to write to personal information; [2]Permission list is constant and short; [3]Permission list is constant and long; [4]Permission list varies length between apps.

Table 8: Survey's results divided according to the self-described level of expertise (Table 4), showing the number of participants choosing apps with less or more access to personal data using the conventional and improved permission interfaces. The effects of using our improved interface is reported using the number of participants who changed their choice of apps.

	SELF RATING BY PARTICIPANTS ABOUT PERMISSION COMPREHENSION (Table 4) & NUMBER OF PARTICIPANTS IN EACH LEVEL															
	COMPREHENSION :1; #P.: 25				COMPREHENSION :2; #P.: 30				COMPREHENSION: 3; #P.: 39				COMPREHENSION :4; #P.: 31			
	INTERFACE		ACCESS		INTERFACE		ACCESS		INTERFACE		ACCESS		INTERFACE		ACCESS	
QUESTIONS	C.	I.	Less	More	C.	I.	Less	More	C.	I.	Less	More	C.	I.	Less	More
Q (1)[4]	25	25	0	0	27	29	3	1	31	34	5	2	28	29	1	0
Q (2)[2]	4	17	13	0	1	17	16	0	0	25	25	0	0	13	13	0
Q (3)[2]	25	25	0	0	18	27	11	2	24	27	9	6	29	27	1	3
Q (4)[4]	10	24	15	1	0	26	26	0	0	29	29	0	0	23	23	0
Q (5)[4]	25	25	0	0	28	30	2	0	34	34	3	3	31	30	0	1
Q (6)[4]	24	24	1	1	0	25	25	0	0	33	33	0	0	26	26	0
Q (7)[4]	4	17	16	3	0	14	14	0	0	25	25	0	0	16	16	0
Q (8)[4]	22	25	3	0	26	30	4	0	31	34	6	3	29	28	2	3
Q (9)[4]	9	25	16	0	2	25	24	1	2	30	29	1	1	23	22	0
Q (10)[4]	7	15	9	1	0	5	5	0	0	4	4	0	0	10	10	0
Q (11)[3]	25	25	0	0	10	26	16	0	6	30	25	1	4	24	21	1
Q (12)[3]	25	25	0	0	25	27	5	3	20	29	13	4	11	27	16	0
Q (13)[3]	22	25	3	0	23	30	7	0	16	32	18	2	4	20	17	1

C. = Conventional Interface; I. = Improved Interface;[2]Permission list is constant and short; [3]Permission list is constant and long; [4]Permission list varies length between apps.

8. ACKNOWLEDGMENTS

Ilaria Liccardi was supported by the European Commission Marie Curie International Outgoing Fellowship grant 2011-301567 *Social Privacy*. Daniel J. Weitzner acknowledges support from National Science Foundation grant CNS-0831442 *CT-M: Theory and Practice of Accountable Systems*, from the Ford Foundation for the MIT Information Policy Initiative and the Department of Homeland Security grant N66001-12-C-0082 *Accountable Info Systems*.

9. REFERENCES

[1] R. Balebako, J. Jung, W. Lu, L. F. Cranor, and C. Nguyen. "little brothers watching you": raising awareness of data leaks on smartphones. In *Proc. of ACM SOUPS '13*, pages 12:1–12:11, 2013.

[2] D. Barrera, H. G. Kayacik, P. C. van Oorschot, and A. Somayaji. A methodology for empirical analysis of permission-based security models and its application to android. In *Proc. of ACM CCS '10*, pages 73–84, 2010.

[3] A. R. Beresford, A. Rice, N. Skehin, and R. Sohan. Mockdroid: trading privacy for application functionality on smartphones. In *Proc. of ACM HotMobile '11*, pages 49–54, 2011.

[4] J. L. Boyles, A. Smith, and M. Madden. Privacy and data management on mobile devices. *Pew's Report: Mobile*

Identity, pages 1–19, September 5th 2012.

[5] E. Chin, A. P. Felt, V. Sekar, and D. Wagner. Measuring user confidence in smartphone security and privacy. In *Proc. of ACM SOUPS '12*, pages 1:1–1:16, 2012.

[6] M. Dietz, S. Shekhar, D. S. Wallach, Y. Pisetsky, and A. Shu. Quire: Lightweight provenance for smart phone operating systems. In *Proc. of Usenix Security Symposium*, pages 347–362, 2011.

[7] N. Eagle, A. Pentland, and D. Lazer. Inferring friendship network structure by using mobile phone data. *Proc. of the National Academy of Sciences (PNAS)*, 106(36):15274–15278, 2009.

[8] S. Egelman, A. Porter Felt, and D. Wagner. Choice architecture and smartphone privacy: There's a price for that. In *Workshop on the Economics of Information Security (WEIS)*, 2012.

[9] W. Enck, P. Gilbert, B.-G. Chun, L. P. Cox, J. Jung, P. McDaniel, and A. N. Sheth. Taintdroid: an information-flow tracking system for realtime privacy monitoring on smartphones. In *Proc. of OSDI'10*, pages 1–6, 2010.

[10] Federal Trade Commission. Mobile apps for kids: Current privacy disclosures are disappointing. pages 1–23, February 2012.

[11] Federal Trade Commission. Mobile apps for kids: Disclosures still not making the grade. pages 1–21, December 2012.

[12] Federal Trade Commission. Mobile privacy disclosures: Building trust through transparency. pages 1–29, February 2013.

[13] A. P. Felt, E. Chin, S. Hanna, D. Song, and D. Wagner. Android permissions demystified. In *Proc. of ACM CCS '11*, pages 627–638, 2011.

[14] A. P. Felt, S. Egelman, and D. Wagner. I've got 99 problems, but vibration ain't one: a survey of smartphone users' concerns. In *Proc. of ACM SPSM '12*, pages 33–44, 2012.

[15] A. P. Felt, E. Ha, S. Egelman, A. Haney, E. Chin, and D. Wagner. Android permissions: user attention, comprehension, and behavior. In *Proc. of ACM SOUPS '12*, pages 3:1–3:14, 2012.

[16] A. Gahran. Survey: Most cell phone users don't protect mobile privacy. *CNNTech*, September 2012.

[17] GSM Association. Gsma mobile and privacy, privacy design guidelines for mobile application development. pages 1–27, February 2012.

[18] P. G. Kelley, J. Bresee, L. F. Cranor, and R. W. Reeder. A "nutrition label" for privacy. In *Proc. of ACM SOUPS '09*, pages 4:1–4:12, 2009.

[19] P. G. Kelley, S. Consolvo, L. F. Cranor, J. Jung, N. Sadeh, and D. Wetherall. A conundrum of permissions: Installing applications on an android smartphone. In *Proc. of USEC '12*, pages 1–12, 2012.

[20] P. G. Kelley, L. F. Cranor, and N. Sadeh. Privacy as part of the app decision-making process. In *Proc. of ACM CHI'13*, pages 3393–3402, 2013.

[21] B. Krishnamurthy and C. E. Wills. On the leakage of personally identifiable information via online social networks. *SIGCOMM Comput. Commun. Rev.*, 40(1):112–117, January 2010.

[22] N. D. Lane, E. Miluzzo, H. Lu, D. Peebles, T. Choudhury, and A. T. Campbell. A survey of mobile phone sensing. *Communications Magazine, IEEE*, 48(9):140 –150, 2010.

[23] I. Liccardi, J. Pato, and D. J. Weitzner. Improving Mobile App selection through Transparency and Better Permission Analysis. *Journal of Privacy and Confidentiality: Vol. 5: Iss. 2, Article 1.*, pages 1–55, 2014.

[24] J. Lin, S. Amini, J. I. Hong, N. Sadeh, J. Lindqvist, and J. Zhang. Expectation and purpose: understanding users' mental models of mobile app privacy through crowdsourcing. In *Proc. of ACM UBICOMP '12*, pages 501–510, 2012.

[25] J. Manoogian. How free apps can make more money than paid apps. *TechCrunch*, August 2012.

[26] A. McDonald and L. F. Cranor. Beliefs and behaviors: Internet users' understanding of behavioral advertising. In *TPRC*, pages 1–31. 2010.

[27] M. Merisavo, J. Vesanen, A. Arponen, S. Kajalo, and M. Raulas. The effectiveness of targeted mobile advertising in selling mobile services: an empirical study. *International Journal of Mobile Communications*, 4(2):119–127, 2006.

[28] S. Meurer and R. Wismüller. Apefs: An infrastructure for permission-based filtering of android apps. In *Security and Privacy in Mobile Information and Communication Systems*, volume 107, pages 1–11. 2012.

[29] Mobile Europe. Gsma urges european mobile industry to adopt new privacy framework. *Press Wires*, January 2013.

[30] National Telecommunications Information Administration. Privacy multistakeholder process: Mobile application transparency. March 2013.

[31] W. Pan, N. Aharony, and A. S. Pentland. Composite social network for predicting mobile apps installation. In *Proc. of ACM AAAI '11*, pages 821–827, 2011.

[32] T. Parka, R. Shenoya, and G. Salvendya. Effective advertising on mobile phones: a literature review and presentation of results from 53 case studies. *Journal of Behaviour Information Technology*, 27(5):355–373, 2008.

[33] P. Pearce, A. P. Felt, G. Nunez, and D. Wagner. Addroid: privilege separation for applications and advertisers in android. In *Proc. of ACM ASIACCS '12*, pages 71–82, 2012.

[34] N. Perlroth and N. Bilton. Mobile apps take data without permission. *New York Times*, February 2012.

[35] A. Porter Felt, S. Hanna, E. Chin, H. J. Wang, and E. Moshchuk. Permission re-delegation: Attacks and defenses. In *Proc. of Usenix Security Symposium*, pages 331–346, 2011.

[36] K. Shilton. Four billion little brothers?: privacy, mobile phones, and ubiquitous data collection. *Communications of the ACM*, 52(11):48–53, 2009.

[37] D. J. Solove. *Understanding Privacy*. Harvard University Press, March 30 2010.

[38] D. J. Solove. *Nothing to Hide: The False Tradeoff between Privacy and Security*. Yale University Press, May 31 2011.

[39] I. Trestian, S. Ranjan, A. Kuzmanovic, and A. Nucci. Measuring serendipity: connecting people, locations and interests in a mobile 3g network. In *Proc. of ACM IMC '09*, pages 267–279, 2009.

[40] H. Xu, H.-H. Teo, B. C. Y. Tan, and R. Agarwal. The role of push-pull technology in privacy calculus: The case of location-based services. *Journal of Management Information Systems*, 26(3):135–174, 2009.

[41] F. Zhang, F. Shih, and D. Weitzner. No surprises: Measuring intrusiveness of smartphone applications by detecting objective context deviations. In *Workshop on Privacy in the Electronic Society (ACM CCS' 13)*, pages 1–6, 2013.