User Control Mechanisms for Privacy Protection Should Go Hand in Hand with Privacy-Consequence Information: The Case of Smartphone Apps

Position Paper

Gökhan Bal, Kai Rannenberg
Goethe University Frankfurt
Deutsche Telekom Chair of Mobile Business & Multilateral Security
Grüneburgplatz 1, 60323 Frankfurt, Germany
{goekhan.bal, kai.rannenberg}@m-chair.de

Abstract

User control for information privacy and informed decision-making are two societal values worth supporting. Nevertheless, there is a growing uncertainty in users' decision-making regarding their ability to have control over their privacy. Even when being provided with control mechanisms, very often users cannot make effective use of them. We argue that the users' lack of understanding with regard to the potential consequences of providers complying or not complying with user-control mechanisms is a strong inhibitor of using those mechanisms. We therefore call for a stronger integration of privacy-consequence information into users' decision-making processes. This aspect has so far been neglected in both research and practice. In this position paper, we discuss the challenges and some practical recommendations in the context of smartphone app ecosystems.

Introduction

Information privacy, "the interest an individual has in controlling, or at least significantly influencing, the handling of data about themselves" [1], is one of the most intensively discussed societal values today. There seems to be a common understanding between individuals, organizations, policy makers, and politics that information privacy – or user control for personal data – is something worth to support. However, regarding the question of "how", there is no satisfying answer yet. Looking at today's reality, we can observe a growing uncertainty of individuals regarding their ability to practice user control in the digital world. Either digital services do not provide sufficient means for user control or existing means exceed the users' capabilities to use them. Furthermore, the benefits that digital services such as social networking sites or mobile apps offer to users are what mainly drive them towards the use of these services. Privacy thoughts rather stay a "bitter aftertaste" that many users are willing to accept. This character of privacy being only a "supporting actor" in individuals' decision-making is further being supported by the lack of effective privacy transparency mechanisms in the online world. As a consequence, users are not able to grasp the actual extent of privacy risks caused by some services. As a second-order consequence, many users do not feel the need for additional privacy protection mechanisms. We believe that user control mechanisms for privacy protection should go hand in hand with useful risk information to support individuals' decision-making. In other words, the user should be able to understand the privacy *consequences* of (not) complying with user control mechanisms. In this position paper, we want to call for a stronger integration of privacy-consequence information into user-control mechanisms to foster informed decision-making. In the rest of this position paper, we will further support our call by examining in more detail the case of smartphone apps as one type of privacy-affecting service. Challenges and recommendations will be discussed within the context of smartphone app ecosystems.

The Nature of Privacy in Smartphone App Ecosystems

Smartphone apps became increasingly popular since Apple introduced its App Store for iOS devices in 2008. Meanwhile, there are more than 2.5 million apps available in the two major app markets, Apple's App Store and Google's Play Store. Apps are available for a wide range of use, whether it's gathering current information, using context-based services, tools, or entertainment apps. A specific enabler of apps' usefulness is platform APIs that provides apps with access to sensitive resources. There is a multiplicity of different resources that can be accessed by apps, e.g. positioning features, sensors, contacts databases, or the phone call history. Beyond the platform-specific APIs, the W3C is developing a series of standards to "create client-side APIs that enable the development of Web Applications and Web Widgets that interact with devices services such as Calendar, Contacts, Camera, etc."[1]. Those APIs that make apps particularly useful, at the same time naturally pose a threat to the user's information privacy. Often, sensitive information flow in the background without the users' awareness. In many cases even the legitimacy of access to some resources is not clear. Two real-world incidents that made a lot of negative publicity are caused by the two apps "Path", a social networking service, and "Brightest Flashlight". Both apps have been demonstrated to transmit highly sensitive user data to the application servers without proper user-consent mechanisms and – at first sight – also without legitimate grounds. A consequence of such behavior is that the users have no chance to capture the "true" privacy-invasiveness of their apps. A solution could be to provide users with better user-control mechanisms. However, our line of argumentation in this position paper is that those user-control mechanisms will not be very effective as long as the users do not "feel" the privacy threats. Current mechanisms for privacy risk communication used in smartphone app ecosystems are ineffective. Our research further showed that one of potentially many factors that lead to the ineffectiveness is the current conceptualizations of "privacy risks". Most often, the conceptualizations are limited to information about (atomic) information flows, e.g. "app X will have access to Y". While this undoubtedly *is* the core of the threats to information privacy, we argue that it uses a too limited conceptualization. The natural question that one might ask regarding the risks is "what does that mean for my privacy?". In other words, what are the consequences of sharing personal information in certain situations? We believe that if we can provide users with at least partial answers to this type of question whenever they can make a decision regarding information, we can more effectively support individuals' informed decision-making. It will give them a true belief of user control and power over their privacy. The obvious question now is: what *are* "privacy consequences"? The question is obvious, but at the same time not easy to answer. Here lies one of the biggest challenges for privacy research. Even though privacy research is well-developed regarding technologies, the aspect of effective privacy risk communication has almost been neglected regarding the contents of risk communication. We tried to foster research in this direction and have already developed some concepts and conducted a series user studies to examine the effectiveness of this approach.

First Experiences with Communicating Privacy Consequenecs

In our research, we first started to get a first sense for the concept of "privacy consequences". We have consulted privacy literature and came to the conclusion that beyond information such as "*app X will have access to Y*", there is a bigger threat to information privacy when adding the long-term perspective to the analysis. The following example shall further illustrate this. Location information is dynamic, thus it is a function of time. Consequently, when location information is shared *once*, the full capabilities for knowledge extraction are not exploited. However, if location information is shared frequently and regularly, more private information is potentially revealed since a location profile can be built with the collected data. More complex examples could be thought of. In more general terms, a long-term privacy risk is the *implicit* revelation of private information due to long-term profiling or data-mining capabilities. We call this type of privacy risk the "second-order privacy risk" in order to differentiate it from the first type, which we call the "first-order privacy risk" perspective (cf. Figure 1). We took the second-order privacy risks as a conceptualization of privacy consequences, developed some proof of concepts for the Android smartphone platform and the Google Play Store and conducted some experiments with users in

[1] http://www.w3.org/2009/dap/

order to measure the effectiveness of this approach regarding the communication of privacy risks and influencing privacy behavior [2], [3]. The results were positive. Users were better able to grasp the actual privacy risk severity of apps and made better app choices when consequences were made explicit. However, the nature of scientific studies requires researchers to create controlled environments by introducing a very limited number of "changes to reality" in order to keep the study manageable and the results interpretable. Thus, we were only able to develop a small set of consequence information for a small set of apps with only limited context information. There is still much work to do in this area.

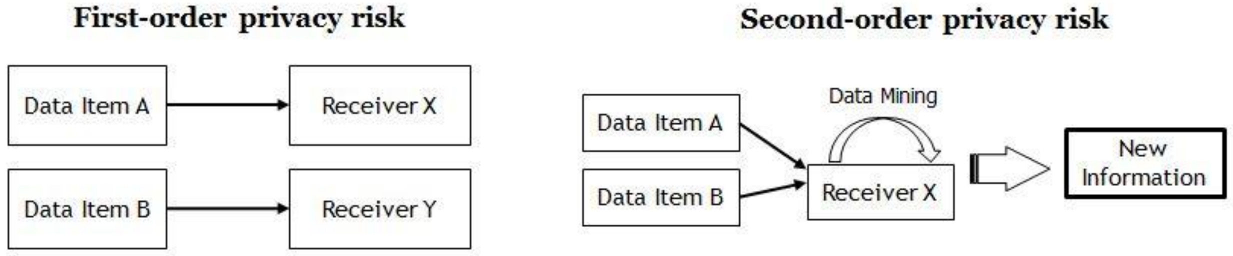

Figure 1: First-order Privacy Risks vs. Second-order Privacy Risks

Communicating Privacy Consequences for Better User Control

With this position paper, we want to call for a stronger integration of consequence information into privacy risk communication and user control mechanisms. We further propose the second-order privacy risk perspective as one potential way to communicate consequences of privacy behavior. We need a more in-depth discussion about the design of consequence information and how they can be interweaved with user control mechanisms. The challenges and some recommendations will be discussed in the following.

Challenges

The following list of challenges is not complete and also not 100% mature. Our aim is to identify interesting points for discussion for the future of user control for privacy and informed decision-making.

Challenge 1: Conceptualization of Privacy Consequences

One of the biggest challenges is to conceptualize privacy consequences. It is about the question how the user's privacy is (potentially) affected by sharing personal data in a specific case and how to communicate this to the users. So far, research and practice has used a very simple conceptualization of privacy consequences, namely the fact that some personal data is shared with some party. This information is not sufficient for individuals to perform risk analyses to make more informed judgments. We have introduced the second-order privacy risk perspective as one potential way to conceptualize privacy consequences. However, it is maybe not realistic to believe that an exhaustive set of potential consequences can be defined. Also, the consequences might vary across different scenarios or in different contexts. We believe that the challenge is to at least identify a small and manageable - and maybe generic - set of second-order privacy risk information that represents a broader set of potential consequences. Even though they might not be 100% accurate, they at least would provide useful information for decision-making. Another important aspect to be considered is the nature of consequences to be communicated. One can think of **negative consequences** (e.g. "*the app provider will know where you live if you share your location*") or **positive consequences** (e.g. "*the app provider will not learn where you live if you block access to your location*"). While the first approach describes consequences of sharing personal information (i.e. not making use of user-control mechanisms), the second approach describes consequences of not sharing personal information (i.e. making use of user-control mechanisms). Both approaches should be considered and tested. The real solution might be a hybrid one as known from many other technologies.

Challenge 2: Consider Functionality & Context of Data Access

Different applications have different demands regarding sensitive information. It is absoulute vital for an LBS application to have regular access to location information, while for a weather forecast application, it is sufficient to access location information only when a new favorite location is setup by the user. On the other hand, some applications require location information only for side functionality and some don't need access location information at all. The challenge is to consider this aspect of "demand level" when assessing the privacy consequences and communicating them to the users. This will help to make fair judgements about the privacy-intrusiveness of apps. Also important is to consider the context of access to sensitive resources. Apps can access data while the app is running in the background, while the app is active or when the smartphone is in stand-by mode or similar. Users' awareness of data access might be depending on this factor as well.

Challenge 3: Monitor Data-Access Behavior of Apps

In order to get an accurate picture of an app's privacy intrusiveness, not only the *potentials* for data access are important to consider. Rather, the actual privacy-impacting *behavior* of that app is important to know about. The concept of privacy-impacting behavioral patterns had been proposed by one of the authors of this paper in the 2012 IEEE Mobile Security Technologies (MoST) workshop [4]. It is about the dynamic data-access behavior of an app that influences how privacy intrusive that app is (what resources does it access? how frequent? what combinations? interactions with other apps?). To determine the privacy-impacting behavioral patterns of apps, there is the need for an information-flow monitoring system as proposed in the workshop paper. Systems such as TaintDroid [5] demonstrate the feasibility of such solutions. We believe that a privacy monitoring of some kind should be an integral part of a smartphone platform. It should keep track of sensitive information flows *for* the user.

Challenge 4: Consider Privacy Transparency of App Providers

There exist multiple ways for app providers to be transparent about the processing of personal data. One prominent example is privacy policies in which they explain what data they collect and how they process it. But the ineffectiveness of privacy policies is well-known, "no one reads them". In any case, they potentially provide valuable information about how privacy impacting a service is. Therefore, we suggest that these information should be considered when determining privacy consequences. Since users are not able to consume privacy policies, we need to find alternative approaches to translate or communicate those information and consider them in the privacy-consequence analysis for the user.

Challenge 5: Automatization

In terms of efficiency, effectiveness, scalability and costs, it is vital to have processes to automize all the necessary steps to make judgments about privacy properties of apps. Therefore, we need an automated monitoring, automated mechanisms to determine apps' privacy-impacting behavior, an automated way to process transparency and privacy policy information, and automated mechanisms to determine privacy consequences.

Recommendations

In the following, we discuss some ideas about potential market players that can contribute to the goal described in this position paper.

What Can Smartphone Platform Providers Do?

In order to facilitate true user control over privacy, smartphone platforms are the obvious "entities" to **keep track of sensitive information flows** on the devices. Furthermore, smartphone platforms should have mechanisms to reason about the privacy intrusiveness of apps based on the apps' data-access profiles. Third, any information regarding the privacy-impacting behavior of an app of a user is valuable for other users as well. Thus, findings about an app should flow into a system where information about specific apps from different sources can be aggregated, e.g. app marketplaces.

What Can App Market Places Do?

App markets should include more efficient and simpler privacy information about apps, specifically about privacy consequences. This will help users in making more informed decisions regarding app selection before they even install an app. Also, app marketplaces should provide app providers with a standardized and easy way to explain their data-access requests. **App markets should add some type of privacy profile information about apps**. Another option could be to include 5-star based privacy-rating mechanisms just as for the overall rating of apps. We have tested such a design in an experimental user study with the Google Play Store [2] and it showed positive results.

What Can App Developers Do?

Most app developers don't have bad intentions when they request permissions to access sensitive resources. Therefore, in order to avoid inaccurate judgments about apps and app providers, developers could use any option to **explain their apps' privacy impact behavior**. This will definitely increase trust towards their applications. Right now, the only option in smartphone app ecosystems is to use some space in the app descriptions or privacy policies. However, we believe that there is a **need for a simpler and standardized way to be transparent** about this.

What Can the WC3 Do?

In the context of the W3C Device API specifications, the W3C can **support app developers in being more transparent about their privacy-impacting practices**. The privacy requirements could further define standardized mechanisms for app developers to e.g. justify their data-access requests or to map their requests to specific features of their apps. In that way, it will be easier for the user to understand the (positive and negative) consequences of sharing or not sharing the required data.

References

[1] R. Clarke, "Internet privacy concerns confirm the case for intervention," *Commun. ACM*, vol. 42, no. 2, pp. 60–67, Feb. 1999.

[2] G. Bal, "Designing Privacy Indicators for Smartphone App Markets: A New Perspective on the Nature of Privacy Risks of Apps," in *Proceedings of the 20th Americas Conference on Information Systems (AMCIS 2014)*, 2014.

[3] G. Bal, K. Rannenberg, and J. Hong, "Styx: Design and Evaluation of a New Privacy Risk Communication Method for Smartphones," in *29th IFIP TC 11 International Conference, SEC 2014, Marrakech, Morocco, June 2-4, 2014. Proceedings*, 2014, pp. 113–126.

[4] G. Bal, "Revealing Privacy-Impacting Behavior Patterns of Smartphone Applications," in *MoST 2012 - Proceedings of the Mobile Security Technologies Workshop 2012*, 2012.

[5] W. Enck, P. Gilbert, B. 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 USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 2010.