Usable Privacy for Mobile Sensing Applications

Delphine Christin, Franziska Engelmann, Matthias Hollick

Usable Privacy for Mobile Sensing Applications

Delphine Christin[1,2], Franziska Engelmann[3], and Matthias Hollick[3]

[1] University of Bonn, Friedrich-Ebert-Allee 144, 53113 Bonn, Germany
[2] Fraunhofer FKIE, Fraunhoferstr. 20, 53343 Wachtberg, Germany
christin@cs.uni-bonn.de
[3] Technische Universität Darmstadt
Mornewegstr. 32, 64293 Darmstadt, Germany
firstname.lastname@seemoo.tu-darmstadt.de

Abstract. Current mobile applications gather an increasing amount of data about the users and their environment. To protect their privacy, users can currently either opt out of using the applications or switch off their mobile phones. Such binary choices, however, void potential benefit for both users and applications. As an alternative, finer control over their privacy could be given to users by deploying privacy-preserving mechanisms. However, it is unclear if users are able to perform the necessary configuration of such schemes. In this paper, we therefore investigate to which degree users can understand the underlying mechanisms as well as the resulting trade-offs in terms of, e.g., privacy protection and battery consumption. To this end, we have conducted a user study involving 20 participants based on user interfaces especially designed for this purpose. The results show that our participants would prefer deciding on the consequences and leave the system parameterizing the underlying mechanism.

1 Introduction

With over 6 billion subscriptions worldwide [14], mobile phones are ubiquitous and their technological advances have led to the emergence of millions of novel applications. However, most mobile applications require the collection of a wealth of information about the users [10]. This not only includes their current locations, but also data gathered by the sensors embedded in their mobile phones. For example, accelerometers can serve to monitor users' activity, while microphones can be leveraged to infer users' context. The information collected by the mobile phones can be further combined with, e.g., past users' search queries, agenda, or mails, in order to improve the application services and anticipate their next queries as proposed in Google Now [1]. Through the utilization of these applications, users' privacy is hence seriously put at risk.

Efforts to make the collection of location information transparent to the users have been recently undertaken, e.g., in the iOS 7 Beta 5 version [21] where users can consult their most frequently visited locations and the corresponding stay duration. While such transparency may increase user awareness about potential privacy issues, this still does not contribute to protect their privacy. On the

contrary, mobile phones fallen into wrong hands may reveal when users are usually not at home and thus help potential burglars. The most frequent solution offered to the users is to either disable such applications or even switch off their mobile phones in order to protect their privacy. Consequently, no fine-granular solutions exist. Such solutions could not only benefit to the users, but also to the applications. In other domains, it has been shown that providing control to users over their data and privacy protection increase their trust in the system [13]. Instead of completely opting out, privacy-conscious users may still benefit from limited application features, thus still providing information to the application but in a way that respects their privacy.

In this paper, we therefore investigate the feasibility of giving users control over their privacy protection and allow them to customize it according to their personal preferences. To this end, we select a noise monitoring application, in which users collect sound levels with their mobile phones. The collected sensor readings are then consolidated to build noise pollution maps. We further integrate the path jumbling scheme proposed in [7] into the noise monitoring application. In particular, our contributions are as follows:

1. We design privacy interfaces to provide users control over the underlying privacy-preserving mechanism and thus over their privacy protection. We base our design on a thorough analysis of the considered mechanism and its functional requirements. Simultaneously, our objective is to cater for *comprehension*, *transparency*, and *simplicity* in order to provide user interfaces with a high degree of usability.
2. We evaluate our proof-of-concept implementation by means of a user study involving 20 participants. In our study, the participants tested and evaluated the different privacy interfaces by completing both a guided and a free task, in which they had to configure the mechanism according to given settings and their personal preferences, respectively. The study highlights that most participants appreciated the additional control offered, but some of them were still overstrained by the overall complexity.

The paper is organized as follows. We first introduce and analyze the underlying privacy-preserving scheme in Sec. 2, before presenting our design drivers and design decisions in Sec. 3 and 4, respectively. We detail the results of our user study in Sec. 5 and summarize existing work in Sec. 6, before making concluding remarks in Sec. 7.

2 The Path Jumbling Concept

We assume that users are registered to a noise monitoring application. Their mobile phones automatically collect sensor readings, i.e., noise levels. The sensor readings are stamped with the collection time and location information. In order to protect their privacy, users leverage the collaborative path-hiding mechanism proposed in [7] instead of directly reporting the sensor readings to the application server. This means that their mobile phones swap their sensor readings when

they are in physical proximity in order to break the association between the spatiotemporal context of the sensor readings and the user's identity.

Different strategies to exchange the sensor readings between users have been introduced in [7]. Users can swap all their sensor readings using the *realistic strategy*, while they can exchange a random number of them with the *random-unfair* and *random-fair* strategies. In the *random-fair* strategy, the users exchange the same number of sensor readings. As a result, the selection of an exchange strategy requires to balance the trade-offs between the expected jumbling degree (i.e., the percentage of exchanged versus own collected sensor readings), the reporting overhead (e.g., when users get more sensor readings as they initially collected and exchanged), and the degree of trust in other users (i.e., exchanging fewer sensor readings with less trusted users).

Depending on the user-meeting pattern, users may not be able to always exchange their sensor reading with others. In this case, the sensor readings can be either reported to the application or stored until the next encounter(s). In the former case, the original paths followed by users will be revealed to the application as the sensor readings could not be jumbled, while it will introduce additional latency for the application in the latter case. Depending on the application scenario, low latency may be preferred to allow a timely delivery of the collected sensor readings. Users can hence select and parameterize one of the following reporting strategies: *time-based*, *exchange-based*, and *metric-based*. Each strategy determines a particular condition needed to be fulfilled in order to trigger the reporting of the sensor readings to the application server.

In summary, users should be able to choose among the proposed exchange and reporting strategies based on the trade-offs between trust, overhead, reporting latency, and jumbling degree according to their preferences.

3 Design Drivers

In this paper, we aim at providing user interfaces that allow users to configure the path jumbling scheme presented in Sec. 2. Our first design driver is to increase the users' *consciousness* about potential privacy threats in mobile sensing applications as recommended in [20] in order to motivate the necessity of configuring and applying such a privacy-preserving scheme. Additionally, we intend to provide *control* to the users. They should be able to: (a) select one exchange strategy among the realistic, random-unfair, and random-fair strategies, (b) select a user reputation threshold above which users will be considered trustworthy enough to initiate an exchange of sensor readings, (c) select one reporting strategy among the time-based, exchanged-based, and metric-based strategies and customize the respective parameter. Once the path jumbling mechanism has been configured, users should be able to review the selected parameterization and consult the potential consequences. This caters for both *transparency* and *comprehension*. Through the whole configuration process, users should be assisted by different dialogues to support their comprehension of the overall mechanism. Further-

more, the required interactions should be kept to a minimum in order to enable fast configuration and reconfiguration and limit the burden for the users.

4 Designed Privacy Interfaces

Based on the drivers detailed in Sec. 3, we have designed and implemented the following privacy interfaces. Our proof-of-concept implementation is based on the iOS operating system (version 5.1). Our privacy interfaces are integrated into a noise monitoring application we called "Noisecapture". Similar to those proposed in [4] and [19], the application captures sound samples and extracts the corresponding noise levels. As illustrated in Fig. 1(a), users can access the application results in form of statistics or maps. When users select the "*Privacy*" option, an informative text about the nature of the collected data and the associated risks for their privacy is first displayed in order to increase user awareness (see Fig. 1(b)). A second view shown in Fig. 1(c) then explains the purpose and basic principles of the path jumbling concept. For both views, we have attempted to reduce the length of the texts to a minimum using as simple as possible wording and illustrate it with different icons to catch users' attention. Both descriptive views are only displayed when users access the privacy interfaces for the first time, except if the users explicitly require help using the corresponding button. The same principle is applied for the remaining interfaces: novice users are assisted by dialogues that explain the different process steps. Each dialogue follows the same structure and includes the goal of the current step, details about the mechanism to configure, and an explanation about the importance of configuring it.

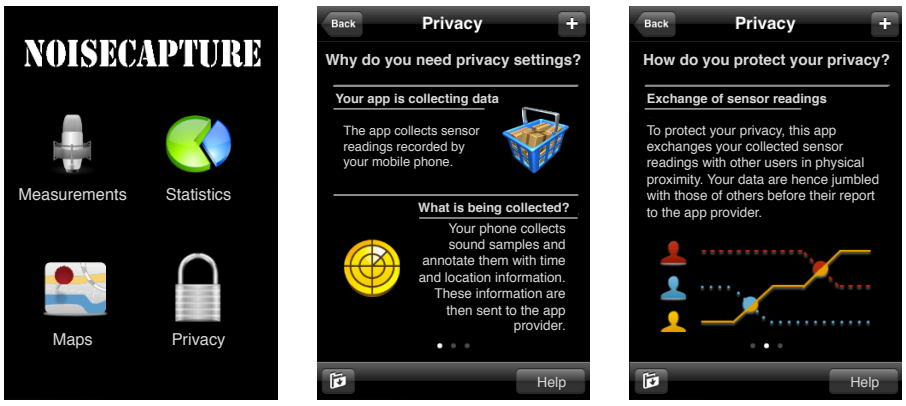

(a) Entry point of the noise monitoring application

(b) Informative view on data collection

(c) Informative view on the path jumbling concept

Fig. 1. Screenshots of the introductory interfaces

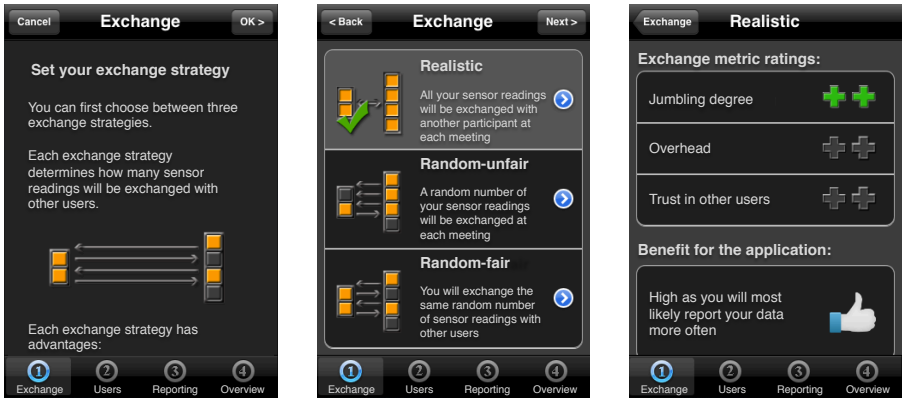

(a) Introductory dialogue on the exchange strategies

(b) Main screen to select the exchange strategy

(c) Details about the realistic exchange strategy

Fig. 2. Screenshots of the interfaces dedicated to the selection of the exchange strategies

In what follows, we detail the designed interfaces implementing the requirements defined in Sec. 3 and including the dialogues especially designed for novice users. We cluster these interfaces according to the following steps: (a) exchange strategy selection, (b) reputation-based user selection, (c) reporting strategy selection, and (d) setting review. Users can navigate through these steps either sequentially using the upper navigation bar or individually select the numbered views in the lower navigation bar.

4.1 Exchange Strategy Selection

If the dialogues are enabled, users first access an introduction on the exchange strategies illustrated in Fig. 2(a). Next, they can choose one of the proposed exchange strategies using the second screen displayed in Fig. 2(b). If the dialogues are disabled, users directly access this second screen. Each strategy is accompanied by an icon illustrating its main principle and a short description. By selecting the blue arrow, users obtain additional details about the corresponding exchange strategy (see Fig. 2(c)). These details include a rating of the strategy according to the resulting jumbling degree, overhead, and trust in other users and whether the strategy is more beneficial to the users (thumb-up icon), to the application (thumb-down icon), or both of them (thumb-middle icon). The more green crosses, the better the rating. While only the details of the realistic exchange strategies are displayed in Fig. 2, similar detail views are available for both random-fair and random-unfair exchange strategies. Consequently, users can see the consequences of the different exchange strategies and which parties benefit most from its application at a glance.

4.2 Reputation-based User Selection

After having selected the exchange strategy to apply, users can first inform themselves on the selection of users to exchange sensor readings with based on their reputation (see Fig. 3(a)). The reputation level is computed based on peer-based ratings about past exchanges as detailed in [8] and reflects the users' readiness to cooperate in this scheme. For example, dropping sensor readings or exchanging incorrect ones will result in low reputation scores. Users can choose the minimum reputation other users should have to initiate an exchange with them using the interface depicted in Fig. 3(b). The reputation level is computed based on peer-based ratings about past exchanges as detailed in [8]. The reputation levels are coded using a 5-star scale, each star differing in both size and color. The biggest green star is associated to the highest reputation level, while the smallest red star is for the smallest reputation level. By selecting a low reputation level, users take the risks that their sensor readings may not be reported to the application server by the concerned exchange partners. On the other side, the number of potential exchange partners may be limited when selecting a high reputation level.

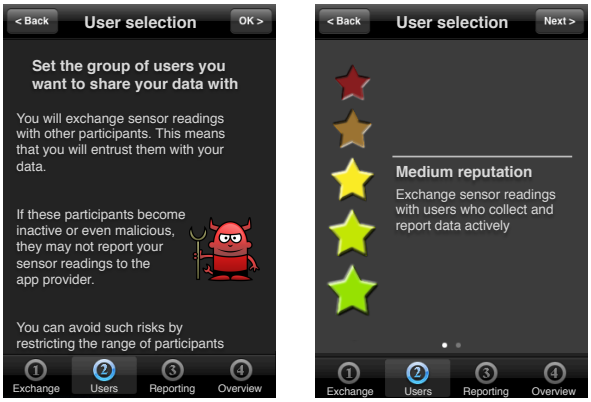

(a) Introductory dialogue on the reputation-based user selection

(b) Main screen to select the minimum users' reputation to exchange with

Fig. 3. Screenshots of the interfaces dedicated to the selection of users

4.3 Reporting Strategy Selection

Similarly to the exchange strategy selection, users first obtain basic information on the reporting principles as shown in Fig. 4(a) when using the dialogue-based configuration. Otherwise, they can directly select the desired reporting strategy in the screen represented in Fig. 4(b). Additionally, they can parameterize

(a) Introducing dialogue on the reporting strategies

(b) Main screen to select the reporting strategy

(c) Configuration of the jumbling threshold

Fig. 4. Screenshots of the interfaces dedicated to the reporting strategies

the selected reporting strategy according to their preferences. For example, they can determine the reporting frequency for the time-based reporting strategy, the number of exchanges for the exchange-based reporting strategy, the distance between the original paths, or the minimum jumbling percentage for the metric-based reporting strategies. Fig. 4(c) illustrates the parameterization of the minimum jumbling percentage. By moving the slider, the shares of personal and jumbled data are adjusted according to users' preferences.

4.4 Setting Review

After the configuration of the path jumbling mechanism, users can consult an overview of their selected settings and learn about the potential consequences as illustrated in Fig. 5. In Fig. 5(a), users can review which exchange strategy, reputation level for other users, and reporting strategy they have chosen in the upper part of the screen. In the lower part, they can see an estimation of the jumbling degree and the reputation level that could be reached when applying these settings. Moreover, the implications of their selection are displayed in a second view depicted in Fig. 5(b). In this view, users can see at a glance the influence of their settings with respect to privacy, trust in other users, reporting latency, and data completeness based on the different colors and associated icons. Data completeness refers to the reporting of consecutive sensor readings to the server. The better the completeness, the better the data processing at the server side, as results in the same area are available. By clicking on each cell, users can obtain additional information about potential risks and change the associated settings if those do not match their personal conception. Alternatively, they can navigate to the corresponding interface using the lower navigation bar.

(a) Setting overview (b) Setting consequences

Fig. 5. Screenshots of the interfaces dedicated to the review of the selected settings

4.5 Summary

By using the designed interfaces, users can *control* the path jumbling mechanism and take *informed* decisions based on the different proposed dialogues. Users can hence control both the exchange and the report of the sensor readings. They can also review their settings and their potential implications, thus catering for *transparency*. If the settings do not correspond to their personal conception, users can directly access them and update them.

5 Evaluation of the Designed Privacy Interfaces

In order to evaluate the usability of the privacy interfaces presented in Sec. 4, we have performed an empirical user study. Our study was advertised on different student forums at our university. In total, 20 participants volunteered to test and evaluated the designed interfaces. The participants were rewarded for their contribution with refreshments, no monetary remunerations were offered. In this section, we present the participants' demographics and provide details about the study settings, before commenting the obtained results.

5.1 Demographics

Our participant sample is composed of 20 undergraduate students aged between 20 and 25 years (μ=22.7, σ=1.87). They were predominantly male (n=12) and their fields of study were as follows: electrical engineering (30%), natural sciences (30%), computer science (25%), and humanities (15%). 62% of the participants owned a smartphone, among which 23% owned at least one iOS-based device. Their average experience level with such devices was rated with a score 4.25 with σ=1.58 on a scale from one (beginner) to seven (expert). While our sample

may not be representative for the whole population, we especially targeted this group of participants as they are more susceptible to contribute to mobile sensing applications than other socio-demographic categories as shown in [6].

5.2 Study Settings

The study was performed under supervised laboratory conditions. We distributed to each participant an iPhone 4 configured with the privacy interfaces detailed in Sec. 4. Additionally, each participant had a leaflet written in English including: (a) a brief introduction to mobile sensing applications and related privacy issues, (b) instructions for a guided task, (c) the same for a free task, and (d) a questionnaire. In the guided task, we asked the participants to conduct the following main steps:

1. Identify the strategy that requires the lower trust in other users and select the exchange strategy that guarantees the best jumbling degree,
2. Choose the reputation level that will allow them to exchange sensor readings with every encountered user,
3. Set the time-based and distance-based reporting strategies to a threshold of two days and 6 km, respectively. Select the metric-based reporting strategy and set the jumbling threshold to 75%,
4. Review the chosen settings and change those categorized as critical.

Next, the participants could freely customize their own privacy settings in the free task. In order to investigate their understanding of the existing trade-offs and the helpfulness of the review step, we first asked them to indicate whether their settings would benefit the application or their privacy protection. In average, the completion of the study took approximately one hour per participant.

5.3 Results

In this section, we present the results of our user study, including both our observations as well as the participants' answers to the distributed questionnaire. We first focus on the comprehensibility of the proposed dialogues, before addressing the different interfaces related to the selection of the exchange strategy, the minimum user reputation, the reporting strategy, and the setting review, respectively. We finally examine user acceptance.

Dialogue Comprehensibility After having read the first introductory dialogues displayed in Fig. 1, the users answered a set of multiple choice questions about potential risks to their privacy caused by contributions to mobile sensing applications, the basic principle of the path jumbling mechanism, and the objective of the proposed interfaces. Based on these dialogues, 90% of the participants correctly answered all questions, meaning that they fully understood the

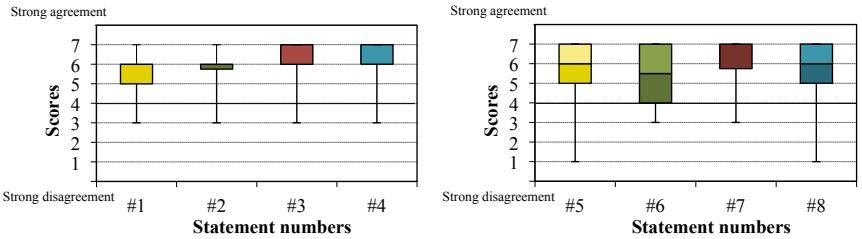

(a) Comprehensibility (b) Exchange strategy selection

Fig. 6. Minimum, quartiles, and maximum score attributed to the statements focused on the comprehensibility and the exchange strategy selection

motivation for these interfaces and the key principle of the underlying privacy-preserving mechanism. The remaining participants had a majority of correct answers, but did not select all possible correct answers.

Additionally, we submitted the following different statements to the participants: "*The first view [in Fig. 1(b)] helped me to understand the risks of mobile sensing applications*" (#1), "*The second view [in Fig. 1(c)] helped me to understand the goals of these privacy interfaces*" (#2), "*The second view clearly described what I had to do next*" (#3), and "*The second view clearly described the goal of the next step*" (#4). The participants rated them using a seven point Likert scale. A score of 1 indicates a strong disagreement with the statement, 4 is neutral, and 7 indicates a strong agreement. Figure 6(a) shows the minimum, quartiles, and maximum scores attributed to these statements. With the exception of one participant, all participants agreed with the proposed statements. This confirms that the first views contribute to the comprehensibility of the privacy threats, the motivations behind the interfaces, as well as the different steps of the configuration process. Globally, the second view about the path jumbling concept obtained better scores than the first view describing the potential privacy threats. Participants may be more willing to have a detailed information about possible risks when contributing to such applications, as participant P_3 commented that "you should also indicate what providers can do with your personal data: location tracking, habit analysis,...".

Exchange Strategy Selection In a second step, we asked the participants to rate the interfaces designed for the selection of the exchange strategies introduced in Fig. 2. With the exception of three participants, all participants agreed that "*the icons appropriately illustrate the exchange strategies*" (#5), and "*selecting an exchange strategy is easy*" (#6), as shown in Fig. 6(b). Moreover, they found that "*the table describing the pros and cons of the exchange strategies is clearly structured*" (#7), and it "*helped [them] to find the exchange strategy that best fits [their] preferences*" (#8).

Concerning the disagreeing participants, the participant P_2 did not find the proposed icons appropriate (#5), but did not comment on how to improve them.

For #8, the participant P_4 strongly disagreed as he preferred using the textual descriptions rather than the summary table "[...] because they provide more information". In comparison, the participant P_{20} thought that the table is not useful as "reporting strategies can change the pros and cons of the exchange strategies again". Her reasoning is due to a confusion between the achieved jumbling degree and the jumbling-based reporting strategy. Despite these three strong disagreements, the scores selected by the participants however remain positive. By comparing the results of #5 to #8, the scores given to #6 are globally lower. Based on our observations, this difference may not be exclusively due to the design of the main interface (see Fig. 2(b)), but also to the navigation complexity between the interface itself and both the introductory dialogues (see Fig. 1) and the table displaying the setting consequences (see Fig. 5(b)).

When observing and discussing with the participants, we noticed an important variation in their degree of comprehension of the exchange strategies. Some participants perfectly understood the principles and consequences of the different strategies, whereas others had only a vague idea. Hence, this indicates that additional efforts should be provided to further increase the comprehensibility of the configuration process. Moreover, we noted that several participants interpreted the consequences of each exchange strategy based on their descriptions, instead of using the table showing the setting consequences as shown in Fig. 5(b). This may suggest that the design of the table is still not optimal and can still be improved to better help all users. In both cases, understanding and selecting an exchange strategy was time-consuming and required concentration. While we attempted to keep the amount of text to the minimum, our observations showed that other alternatives should be found to reduce the burden for the users. For example, videos or cartoons, could be investigated in the future.

Reputation-based User Selection Based on their experience in the guided and free tasks, the participants next evaluated both the dialogue (cf. Fig. 3(a)) and the main interface (cf. Fig. 3(b)) used to set the minimum reputation level that other users should have to initiate an exchange with them. As shown in Fig. 7(a), the distribution of the scores attributed to the corresponding statements are slightly higher than for the previous results. Most participants agreed that "*the illustration clearly indicates the minimum reputation score of [their] exchange partners should have*" (#9). Moreover, "*the combination of color and size of the stars [helped them] to recognize the corresponding reputation score*" (#10) and "*the text [helped them] to understand the characteristics of the users having the respective reputation score*" (#11). This means that the participants are more positive about the control provided to select the minimum reputation level for their exchange partners than that for the exchange strategy selection.

Most participants were able to understand and explain the consequences of exchanging data with users having either low or high reputation scores. For example, P_3 explained the implications of choosing very high reputation scores as follows: "The network of exchange partners shrinks as you are excluding many [users] this way". Participants having initial doubts indicated that the text had

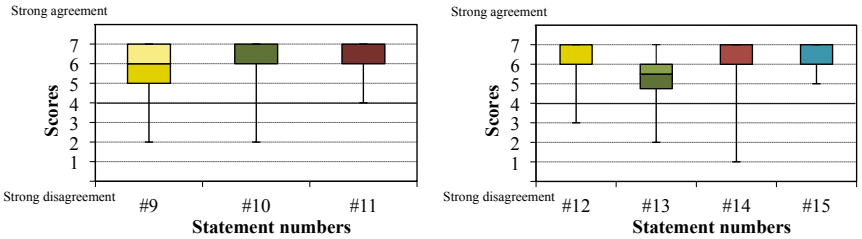

(a) Reputation-based user selection (b) Reporting strategy selection

Fig. 7. Minimum, quartiles, and maximum score attributed to the statements focused on the reputation-based user selection and the reporting strategy selection

been useful to select the appropriate reputation level. P_3, however, commented that the labeling could be improved as "[it] is not intuitively clear what means low and high". Again, P_{10} particularly disagreed on #9 and #10. While he understood the interface objective as shown by his comments, the reputation attribution remained unclear to him.

Reporting Strategy Selection The participants next rated the interfaces dedicated to the selection of the reporting strategy detailed in Fig. 4. By comparing the obtained results shown in Fig. 7(b) with those for the exchange strategy selection in Fig. 6(b), we can observe that the participants found that the respective icons better illustrate the reporting strategies (#12) than the exchange strategies (#5). Moreover, fewer participants strongly agreed that "*selecting a reporting strategy is easy*" (#13) compared to the exchange strategy selection (#6). This may be due to the additional interaction required to customize the reporting strategy parameter, e.g., the reporting frequency in the time-based reporting strategy. At the same time, more participants globally agreed with this statement based on a comparison of the first quartiles. Our observations show that the degree of comprehension not only varied between participants as for the exchange strategy selection, but also between strategies. According to our expectations, the time-based reporting strategy was relatively easy to understand while the distance-based reporting strategy was the most difficult. With the exception of P_2, all participants, however, rated "*the animations used to configure the metrics of the reporting strategies are comprehensible*" (#14) and "*the animations used to configure the metrics of the reporting strategies are illustrative*" (#15) with a score of either six or seven. This means that the proposed interactions were appreciated by the participants, but the navigation and the overall comprehension could be generally improved.

Setting Review Fig. 8(a) shows that all participants agreed that "*information on [their] configuration are clearly arranged in the overview*" (#16). A wide spread of scores is however observed for #17 about the intuitiveness of the scrolling between the overview and consequence table introduced in Fig. 5(a) and

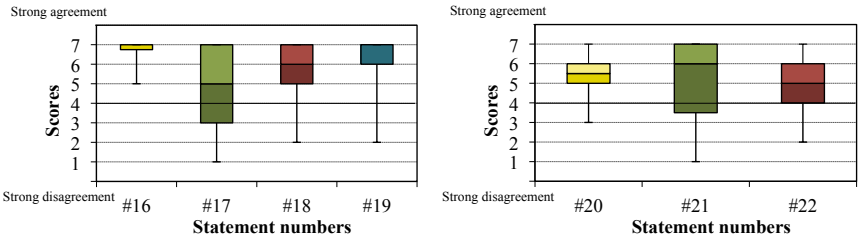

(a) Setting overview (b) User acceptance

Fig. 8. Minimum, quartiles, and maximum score attributed to the statements focused on the setting overview and the user acceptance

Fig. 5(b), respectively. Our observations confirm that participants had difficulty to find the consequence table because of the implemented sideways scrolling. As a result, the sideways scrolling should be replaced by a more transparent interaction in order to address this issue. Moreover, almost all participants claimed to have understood the objective of the consequence table as shown by the score distribution of #18. Some participants, however, needed to read the provided explanations in order to fully understand it. While most participants agreed that the color mapping "*helped them to quickly recognize critical aspects of [their] settings*" (#19), some of them indicated that the color mapping could be improved to provide additional levels, instead of the current binary classification between critical and uncritical. The participant having attributed the lowest score commented that "it is not always clear what the colors mean" (P_5). These encouraging results are confirmed by our observations, as most participants needed only one to two attempts to correctly modify their settings when those were identified as critical. Few participants were even able to immediately identify which strategy and parameter needed to be changed.

User Acceptance We finally investigated the participants' acceptance and show the results in Fig. 8(b). With the exception of one participant, all participants globally agreed that "*the concept of path jumbling is easy to understand*" (#20). This is not fully aligned with our observations, as some participants required additional information from the study supervisor. Overall, the more technical backgrounds the students had, the easier it was for them to provide fast and precise answers. However, there were some exceptions. For example, a student in physics performed better than one in mechatronic. Additional efforts are hence still needed to improve the overall scheme comprehensibility. Asked if "*[they] would like to configure the mechanism [themselves] if an application would offer it*" (#21), 50% of the participants strongly agreed, despite the observed time and concentration required. The remaining remained neutral or disagreed, thus showing that the proposed control and associated interfaces did not gain the full acceptance of our participants. Those participants however indicated that "*[they] would rather like to directly adjust the consequences according to [their]*

preferences than configuring the mechanism in detail" (#22) by selecting higher scores compared to #21. However, the participants having strongly agreed with #21 indicated to be less interested in controlling the mechanism as compared to selecting the consequences.

In summary, the majority of our participants understood the path jumbling mechanism and configured it wisely. This shows that potential users are able to excerpt fine-granular control over the protection of their privacy. Some of them considered their privacy and the associated control as important, but were overwhelmed by all scheme details. They would prefer only deciding on the consequences and leave the system parameterizing the underlying mechanism.

By putting the configuration of the privacy settings in the foreground and conducting the study in a laboratory setting, we were able to evaluate our design decisions based on the participants' comments and reactions. However, the chosen methodology cannot fully capture normal user behaviors. We hence plan to conduct an additional long-term study in order to investigate, i.e., whether and how privacy settings are updated over time under real-world conditions and how many interactions do the users actually need in absence of guidance.

6 Related Work

In recent years, designing privacy interfaces and analyzing privacy concerns and behaviors have attracted increasing attention in a wide range of application domains. Generic guidelines to design privacy user interfaces have been provided in [20, 22] and recommendations to avoid common pitfalls have been made in [16]. Moreover, enhanced privacy interfaces for online social networks have been proposed, e.g., in [18], while the impact of the related information exposure on privacy concerns and behaviors have been investigated in [2, 17]. Additionally, interfaces for peer-to-peer file sharing systems and website privacy policies have been designed and evaluated in [12] and [11], respectively. In the former, existing interfaces have been leveraged, whereas new concepts, such as the Privacy Bird, have been introduced in the latter. Users' privacy decisions have also been examined in picture sharing applications [3]. These solutions, however, focus on application domains orthogonal to participatory sensing applications.

Concerning mobile sensing applications, few user studies have been conducted. Users' privacy concerns contributing to a mobile sensing application have been explored in [15], while the authors of [5] have analyzed how users understand, choose, and feel comfortable with different location privacy-preserving schemes. No dedicated user interfaces have, however, been proposed. This work shares more similarities with our previous work [9], in which different privacy interfaces allow users to select the granularity degree at which their sensor readings are released. Similarly to this work, a user study based on a proof-of-concept implementation have been conducted. Their focuses however differ. In [9], we explore the users' preferences in terms of visualization of privacy settings, while we build upon this work and focus on investigating to which degree users can understand and configure complex technical schemes to protect their privacy.

7 Conclusions

We have designed and implemented privacy interfaces that provide control over a privacy-preserving scheme to users of mobile sensing applications. By using our interfaces, users can select and customize different strategies according to their personal preferences. We have evaluated our interfaces by means of a user study involving 20 participants and shown that most of our participants were able to comprehend the underlying mechanism and the associated trade-offs based on our interfaces despite their complexity. While some users would prefer an assisted version where the system would configure the settings based on their chosen consequences, others would be ready to invest time and manually configure each setting according to their preferences. In addition to providing insights about future design improvements, the outcomes of our study therefore demonstrate that users have more than a binary choice between either renouncing to their privacy or not using the application at all, thus laying the first stones on the path to usable and controllable privacy protection for mobile applications.

8 Acknowledgments

The authors would like to thank the participants of the user study. This work was supported by CASED (`www.cased.de`).

References

1. Google Now. The right informaiton at just the right time. Online: `http://www.google.com/landing/now/` (accessed in 08.2013)
2. Acquisti, A., Gross, R.: Imagined Communities: Awareness, Information Sharing, and Privacy on the Facebook. In: Proceedings of the 6th Workshop on Privacy Enhancing Technologies (PET). pp. 36–58 (2006)
3. Ahern, S., Eckles, D., Good, N.S., King, S., Naaman, M., Nair, R.: Over-exposed? Privacy Patterns and Considerations in Online and Mobile Photo Sharing. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI). pp. 357–366 (2007)
4. Bilandzic, M., Banholzer, M., Peev, D., Georgiev, V., Balagtas-Fernandez, F., De Luca, A.: Laermometer: A Mobile Noise Mapping Application. In: Proceedings of the 5th ACM Nordic Conference on Human-Computer Interaction (NordiCHI). pp. 415–418 (2008)
5. Brush, A., Krumm, J., Scott, J.: Exploring End User Preferences for Location Obfuscation, Location-based Services, and the Value of Location. In: Proceedings of the 12th ACM International Conference on Ubiquitous Computing (Ubicomp). pp. 95–104 (2010)
6. Christin, D., Büchner, C., Leibecke, N.: What's the Value of Your Privacy? Exploring Factors That Influence Privacy-sensitive Contributions to Participatory Sensing Applications. In: Proceedings of the IEEE Workshop on Privacy and Anonymity for the Digital Economy (LCN Workshop) (2013)

7. Christin, D., Guillemet, J., Reinhardt, A., Hollick, M., Kanhere, S.S.: Privacy-preserving Collaborative Path Hiding for Participatory Sensing Applications. In: Proceedings of the 8th IEEE International Conference on Mobile Ad-hoc and Sensor Systems (MASS). pp. 341–350 (2011)
8. Christin, D., R. Pons-Sorolla, D., Hollick, M., Kanhere, S.S.: TrustMeter: A Trust Assessment Framework for Collaborative Path Hiding in Participatory Sensing Applications. In: Proceedings of the 9th IEEE International Conference on Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP) (2014)
9. Christin, D., Reinhardt, A., Hollick, M., Trumpold, K.: Exploring User Preferences for Privacy Interfaces in Mobile Sensing Applications. In: Proceedings of 11th ACM International Conference on Mobile and Ubiquitous Multimedia (MUM). pp. 14:1–14:10 (2012)
10. Christin, D., Reinhardt, A., Kanhere, S.S., Hollick, M.: A Survey on Privacy in Mobile Participatory Sensing Applications. Journal of Systems and Software 84(11), 1928–1946 (2011)
11. Cranor, L.F., Guduru, P., Arjula, M.: User Interfaces for Privacy Agents. ACM Transactions on Computer-Human Interacteration (TOCHI) 13, 135–178 (2006)
12. Good, N.S., Krekelberg, A.: Usability and Privacy: A Study of Kazaa P2P File-sharing. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI). pp. 137–144 (2003)
13. Hansen, M., Fischer-Hübner, S., Pettersson, J., Bergmann, M.: Transparency Tools for User-controlled Identity Management. In: Proceedings of the 17th eChallenges Conference (e-2007). pp. 1360–1367 (2007)
14. International Communication Union: The World in 2013: ICT Facts and Figures. Online: `http://www.itu.int` (accessed in 05.2013) (2013)
15. Klasnja, P., Consolvo, S., Choudhury, T., Beckwith, R., Hightower, J.: Exploring Privacy Concerns about Personal Sensing. Pervasive Computing pp. 176–183 (2009)
16. Lederer, S., Hong, I., Dey, K., Landay, A.: Personal Privacy Through Understanding and Action: Five Pitfalls for Designers. Personal Ubiquitous Computing 8(6), 440–454 (2004)
17. Lipford, H., Besmer, A.: Users' (Mis)Conceptions of Social Applications. In: Proceedings of the 36th Graphics Interface Conference (GI). pp. 63–70 (2010)
18. Lipford, H., Besmer, A., Watson, J.: Understanding Privacy Settings in Facebook with an Audience View. In: Proceedings of the 1st Conference on Usability, Psychology, and Security (UPSEC). pp. 1–8 (2008)
19. Maisonneuve, N., Stevens, M., Niessen, M.E., Steels, L.: NoiseTube: Measuring and Mapping Noise Pollution with Mobile Phones. In: Proceedings of the 4th International Symposium on Information Technologies in Environmental Engineering (ITEE). pp. 215–228 (2009)
20. Patrick, A.S., Kenny, S.: From Privacy Legislation to Interface Design: Implementing Information Privacy in Human-Computer Interactions. In: Proceedings of the 3rd Workshop on Privacy Enhancing Technologies (PET). pp. 107–124 (2003)
21. Warzel, C.: This Is What It Looks Like When Your Phone Tracks Your Every Move. Online: `http://www.buzzfeed.com` (accessed in 08.2013) (2013)
22. Yee, K.P.: User Interaction Design for Secure Systems. In: Proceedings of the 4th International Conference on Information and Communications Security (ICICS). pp. 278–290 (2002)