Delphine Christin, Pablo Sánchez López, Andreas Reinhardt, Matthias Hollick, Michaela Kauer:
Share with Strangers: Privacy Bubbles as User-centered Privacy Control for Mobile Content Sharing Applications. In: Information Security Technical Report, 17(3), pp. 105-116, 2013.

Share with Strangers: Privacy Bubbles as User-centered Privacy Control for Mobile Content Sharing Applications☆

Delphine Christin[a], Pablo Sánchez López[a], Andreas Reinhardt[b], Matthias Hollick[a], Michaela Kauer[c]

[a]*Secure Mobile Networking Lab, Technische Universität Darmstadt, Mornewegstr. 32, 64293 Darmstadt, Germany*
Emails: firstname.lastname@seemoo.tu-darmstadt.de, Phone: +49 6151 16-70922, Fax: +49 6151 16-70921
[b]*Multimedia Communications Lab, Technische Universität Darmstadt, Rundeturmstr. 10, 64283 Darmstadt, Germany*
Email: andreas.reinhardt@kom.tu-darmstadt.de
[c]*Institute of Ergonomics, Technische Universität Darmstadt, Petersenstr. 30, 64287 Darmstadt, Germany*
Email: kauer@iad.tu-darmstadt.de

Abstract

A continually increasing number of pictures and videos is shared in online social networks. Current sharing platforms, however, only offer limited options to define who has access to the content. Users may either share it with individuals or groups from their social graph, or make it available to the general public. Sharing content with users to which no social ties exist, even if they were physically close to the places where content was created and witnessed the same event, is however not supported by most existing platforms. We thus propose a novel approach to share content with such users based on so-called *privacy bubbles*. Privacy bubbles metaphorically represent the private sphere of the users and automatically confine the access to the content generated by the bubble creator to people within the bubble. Bubbles extend in both time and space, centered around the collection time and place, and their size can be adapted to the user's preferences. We confirm the user acceptance of our concept through a questionnaire-based study with 175 participants, and a prototype implementation shows the technical feasibility of our scheme.

Keywords: Online content sharing, privacy, user control, user acceptance

1. Introduction

In recent years, the public interest for online social media has continuously increased and led to an unprecedented amount of content generated and shared by users. Picture and video sharing has become particularly popular, as shown by the estimated 135,800 pictures uploaded every minute to Facebook [2] and the approximated 72 hours of video shared on YouTube every minute [3]. In existing sharing platforms, users protect their privacy by confining the access to the uploaded content based on social distance. For example, users can share pictures with individuals, friends, friends of friends, or everyone on Facebook. The assumption behind this relationship-based access control is that the stronger the social tie between users, the lower the expected threat to their privacy. As a result, sharing content in a controlled fashion with individuals or a group of persons to which no social ties exist, is virtually impossible in existing platforms.

Let us however assume that two persons (Alice and Bob), who do not have any kind of social relationship to each other, attend the same event, e.g., a soccer match, a party, or a sightseeing tour. Using state-of-the-art solutions, Alice can only share the pictures she took with members of her social network or make them public. However, she cannot share them with Bob since they have no social ties. Sharing pictures with Bob may not pose a threat to Alice's privacy, though: both are likely to have observed the same scenes, because they have been to the same place concurrently. In this case, the perceived threat to Alice's privacy depends on the physical distance between Alice and Bob at the time the content was created as well as the time difference between Alice's and Bob's observations. If we further assume that Alice and Bob were situated close to each other, Alice might not feel that her privacy is endangered by sharing her pictures with Bob, while Bob can benefit from Alice's pictures.

☆This paper is an extended version of our paper "Privacy Bubbles: User-centered Privacy Control for Mobile Content Sharing Applications" [1], which appeared in the IFIP WISTP 2012 workshop and has been invited for ISTR journal publication.

We propose the use of *privacy bubbles* as a novel approach, which directly targets the aforementioned scenario, i.e., sharing content with strangers in a controlled manner. Note that our approach does not attempt to replace existing relationship-based access control mechanisms, but complements them by adopting a perspective which has received very little attention in the past. In order to share pictures with people in their physical vicinity, users create a privacy bubble by determining its radius and duration. The created privacy bubble is centered around the user and metaphorically represents his/her private sphere. The bubble sets spatiotemporal boundaries within which others users are granted access to the content created in the bubble. In particular, the radius of the bubble represents the maximal physical distance between the content creator and other users authorized to access the content. The duration of the bubble represents the maximal temporal difference between the time of capture and the presence of other users within the radius of the bubble. Users can customize both parameters depending on their privacy preferences. The smaller the radius and duration, the better the privacy protection. Note that users can still control which content is shared in the bubble. The access to content in the privacy bubbles of other users is transparently managed by the application. The applicability of the proposed concept is not confined to sharing pictures, and can easily be extended to additional user-generated contents such as video or audio recordings.

Our contributions can be summarized as follows:

1. We propose the concept of privacy bubbles, which enables sharing pictures between users having no social ties in a controlled manner.
2. We evaluate the viability of our concept by means of a questionnaire-based study involving 175 participants. Our evaluation focuses on: (1) the comprehensiveness of the concept, (2) the provided degree of user control, (3) the estimated management overhead, and (4) the user acceptance. We validate design drivers and design alternatives for the realization of privacy bubbles against the results of our study.
3. We present our proof-of-concept implementation of the privacy bubble concept, which takes the findings of our study into account.
4. As extension to [1], we extensively discuss potential attacks on privacy bubbles and derive countermeasures.

The paper is organized as follows. We explain the operation of the privacy bubbles using an example in Section 2 and describe the underlying concept in Section 3. In Section 4, we present the modalities and findings of our questionnaire-based study. We provide details about our prototype implementation in Section 5 and discuss attacks, countermeasures as well as possible extensions to our concept in Section 6. After summarizing existing work in Section 7, we make concluding remarks in Section 8.

2. Application Scenario

Let us examine the application of privacy bubbles in the realistic application scenario illustrated in Figure 1. Three tourists (Alice, Bob, and Carlos) are visiting London, where Alice and Bob join the same sightseeing tour, while Carlos prefers to visit the city's sights by foot. Although the tourists do not personally know each other, they are registered in the same photo sharing application which supports the concept of privacy bubbles.

When boarding the sightseeing bus, Alice creates a new privacy bubble, which has a validity duration of ±5 minutes and encompasses a radius of 50 meters. As a result, only persons located within 50 meters of Alice's location (the center of the bubble) are allowed to access her captured photos, and only do so if they have been at the location at most 5 minutes before or after the photo has been taken. As the bubble follows Alice's moves, the persons authorized to access her pictures are dynamically determined for each individual photo.

In contrast to Alice, Bob is more concerned about his privacy, and defines his own bubble to only include people within 10 meters around him when he takes a picture. In front of Westminster Abbey, Alice and Bob take a set of pictures, while Carlos is walking by in a distance of 20 meters from the bus after having taking photos of the sight. Back at home, Carlos is not fully satisfied with the quality of his pictures and is looking for better pictures on the picture sharing application that reflect the moment of his visit. As Carlos was located within Alice's bubble while she took pictures, he is able to access her pictures of the monument. However, he is not granted access to Bob's pictures since he was outside Bob's bubble.

3. Privacy Bubbles: The Concept

In this section, we highlight the design drivers of the concept of privacy bubbles and its principles. We detail their technical realization in Section 5.

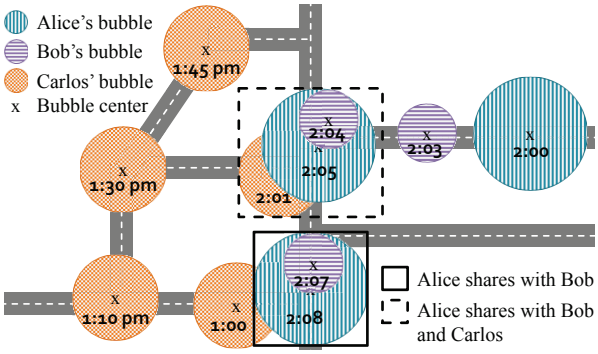

Figure 1: Representation of Alice's, Bob's, and Carlos' bubbles for each taken picture

3.1. Design Drivers

We aspire to develop an access control mechanism for sharing user-generated mobile content with people who were located in physical proximity to the content creator at the time of its creation. The designed access control mechanism should reflect the following design drivers:

1. **Comprehensiveness**: The mechanism should be intuitive and easy to comprehend, particularly for unexperienced users.
2. **User control**: Using this mechanism, the users should be able to control and customize the access to their generated content according to their personal preferences.
3. **Management overhead**: The required user interactions should however be kept to a minimum in order to limit the associated management overhead and foster its usage by potential users.
4. **User acceptance**: We believe that the users need to enjoy and feel comfortable with the proposed approach to adopt and accept it.
5. **Privacy protection**: Ultimately, the privacy of the users should be respected. This includes the control of the users over the pictures released to others and the selection of the bubble parameters according to their personal preferences. Furthermore, the collection of sensitive information by the sharing application should be kept to the minimum.

3.2. Concept and Principles

The concept of privacy bubbles serves as a metaphorical representation of the privacy spheres of the users. The user occupies the center of its bubble and can share information (we have chosen to design our prototype for the sharing of pictures) with other users located in his bubble in a protected manner. In contrast, users located outside his bubble are not allowed to access the shared pictures. Privacy bubbles can be dynamically created by the user that shares the content, who selects its radius and duration. The radius of the bubble determines the maximal distance at which other users should be from the bubble creator at the time of capture of the picture to be able to later access the picture. The duration of the bubble determines the maximal time range during which others users should be included in the bubble (i.e., at a distance inferior to the bubble radius) to access the picture. Let us assume that Alice has a bubble with a radius of 5 meters and a duration of 2 minutes and takes a picture at time t. Every user located at a distance of up to 5 meters from Alice in the time interval [t–2 min, t+2 min] will be able to access the picture taken by Alice if she decides to share it. Alice controls which pictures she shares in her bubble. She can therefore deselect pictures, which potentially compromise her privacy. These users are granted access to the picture until Alice decides not to share the picture anymore. The access authorization does not depend on the current location of the users when they search for shared pictures, but only on their location around the time of the capture of the picture. Moreover, the access authorization is not symmetric. This means that Alice can access the pictures of others if she was included in their bubbles, while they cannot access hers. In our solution, users do not share pictures according to a tit-for-tat mechanism, but the individual privacy preferences of each user (expressed by means of the bubble parameters) are respected. Note that the concept of privacy bubbles does not replace existing access control mechanisms but it complements these by a new sharing paradigm.

4. Evaluation of the Privacy Bubble Concept

We have conducted a questionnaire-based study in order to investigate how potential users perceive the concept of privacy bubbles. Since this study focuses on online picture sharing applications, we have specifically approached participants who could be potential users of such applications. We have recruited them by posting announcements on multiple forums and mailing lists at our university and partner universities. The study was conducted using an online questionnaire in order to collect responses from a broad spectrum of participants. The questions were written in English and their completion took approximately 15 minutes. In total, 175 participants anonymously answered our online questionnaire. In this section, we first present demographic informa-

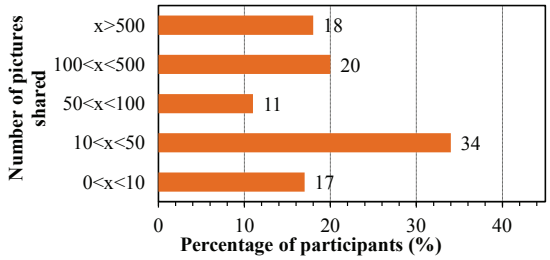

Figure 2: Overall number of pictures shared

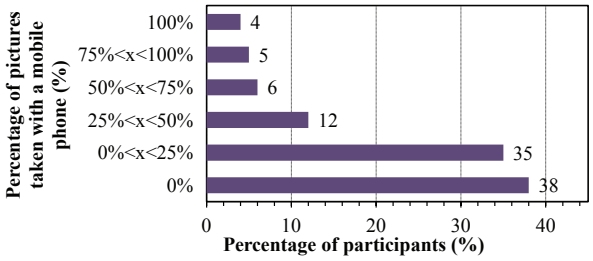

Figure 3: Estimated percentage of shared pictures taken with a mobile phone

tion about our participants, before highlighting the findings of the study.

4.1. Demographic Information

The participants of our study were predominantly male (n=118) and aged between 21 and 55 (m=28, SD=5). Table 1 illustrates the distribution of the most represented nationalities, current jobs, and fields of occupation among the participants. Our sample of participants includes diverse profiles of potential users with various fields of occupation such as theology, law, or business. Among the participants, 81% indicated to have already shared pictures online (n=142). The estimated number of pictures shared by the participants is visualized in Figure 2, which shows that only 17% of the participants do not share pictures online. Furthermore, Figure 3 shows that more than 60% of the participants have shared photos that were taken with their mobile phones.

4.2. Results

In this section, we present the findings of our study classified by design drivers (cf. Section 3.1). We especially analyze whether the participants estimate that the design driver is reflected in the proposed concept of privacy bubbles. Moreover, we assess the suitability of different design alternatives for the implementation of our proof-of-concept presented in Section 5.

4.2.1. Comprehensiveness

The first design driver aims at providing for a solution which is easy to comprehend and intuitive for potential users. After a textual description of the privacy bubble concept, we first submitted the following statement to the participants: "The concept of privacy bubble is easy to comprehend". The participants indicated their degree of agreement with this statement on a seven point Likert scale. A score of 1 indicates a strong disagreement, 4 is neutral, and 7 indicates a strong agreement. Figure 4 illustrates the distribution of the resulting scores

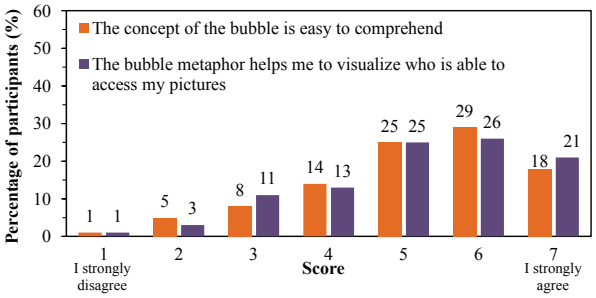

Figure 4: Distribution of the answers about the comprehensiveness and intuitiveness of the privacy bubbles

and shows that 72% of the participants agreed with the submitted statement, i.e., 72% of the participants found the privacy bubble concept easy to comprehend. Moreover, 72% of the participants found that "the bubble metaphor helps [them] to visualize who is able to access [their] pictures as shown in Fig. 4, thus catering for intuitiveness.

4.2.2. User Control

The second design driver targets at allowing the users to tailor the access control to their individual preferences. In our solution, the users customize the radius and duration of their bubble to control the persons able to access their pictures. Figure 5 shows that 71% of the participants confirmed that "being able to determine the radius of the bubble is important for [them]", while 70% of the participants indicated that "being able to determine the duration of the bubble is important for [them]".

Furthermore, we have investigated different control options for the design of our prototype implementation in order to tailor its features to the feedback of the participants. Firstly, 88% of the participants wish to review their pictures before their release to other users (cf. Figure 6) — a feature easily integrable in our proof-of-concept implementation. Secondly, we examined if the participants wish reciprocal relationships with people authorized to access their pictures. Since 39% of the

Table 1: Demographics of the participants (n_{total}=175)

Nationality	*n*
German	108
French	22
Spanish	9
Romanian	3
Indian	3
Ukrainian	3
Other	27

Current job	*n*
PhD student	72
Undergraduate student	59
Postdoctoral researcher	18
Professor	6
Administrative staff	5
Technical staff	4
Other	11

Field of occupation	*n*
Computer science	99
Electrical engineering	35
Psychology	5
Biology	5
Physics	4
Mechanical engineering	4
Other	23

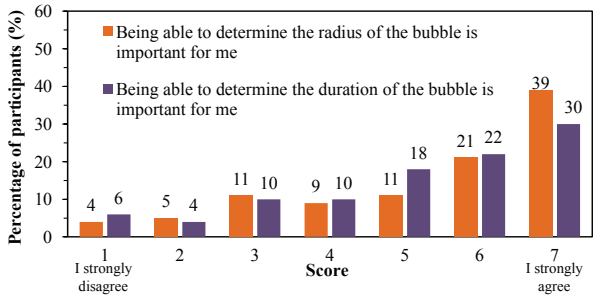

Figure 5: Distribution of the answers about the importance of the control over the radius and the duration of the privacy bubbles

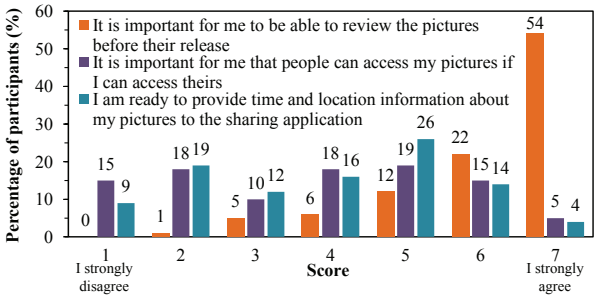

Figure 6: Distribution of the answers about the importance of reviewing pictures before their release, the importance of reciprocal relationships, and the participants' readiness to provide spatiotemporal information

participants agreed that "it is important for [them] that people can access [their] pictures if [they] can access theirs", 18% remained neutral, and 43% disagreed, no trend can be clearly identified from the participants' answers (see Figure 6). We therefore have introduced this feature as an option in our prototype, which can be optionally activated by the users depending on their preferences. We finally asked the participants if "[they are] ready to provide time and location information about [their] pictures to the sharing application". As a result, 44% of the participants indicated to be ready to do it, 16% remained neutral, and 40% indicated not to be ready (see Figure 6). Again, no trend can be clearly identified from the given answers. Consequently, we have integrated two different mechanisms in our prototype, one is transmitting spatiotemporal information to the sharing applications, while the other one does not transmit any such data.

4.2.3. Management Overhead

The third design driver aims at limiting the management burden for the users to the minimum. In our solution, the users only have to select the duration and radius of their bubble and the access control is automatically and transparently managed by the application. The participants confirm the viability of this approach since 56% of the participants indicated that "[they] find it appreciable that the persons authorized to access [their] pictures are dynamically defined", while 57% stated that "[they] find it appreciable that they do not have to select each person individually" as depicted in Figure 7.

In addition to the control over the radius and duration of the bubble, the participants wish additional features as shown in the above section, which complete the original concept described in Section 3.2. The integration of these features in the prototype implementation may introduce additional management overhead for the users. This overhead remains however limited and the additional features contribute to the acceptance of our approach by potential users.

4.2.4. User Acceptance

In addition to the analysis of three design drivers, we finally investigated whether the participants would accept this novel approach for controlling the access to their pictures. The results presented in Figure 8 show that 61% of the participants would "[...] feel comfortable that people can access pictures [they] took when they were in [their] bubble". Note that only 4% strongly

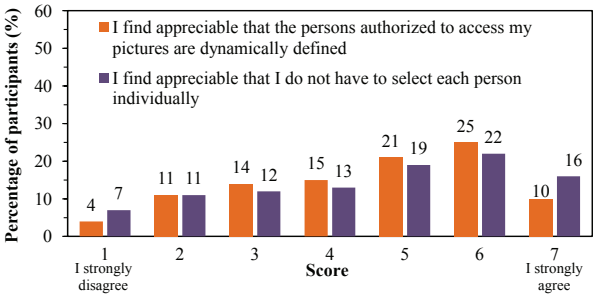

Figure 7: Distribution of the answers about the appreciation of the dynamical and automatic nature of the privacy bubbles

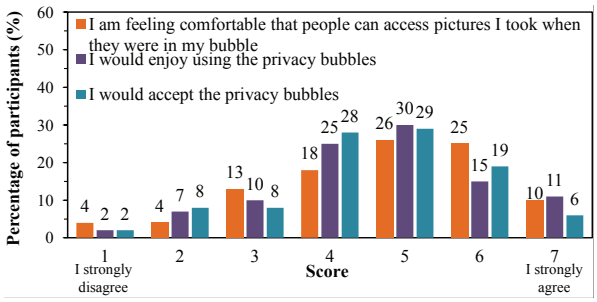

Figure 8: Distribution of the answers related to the acceptance of the privacy bubbles by the participants

disagreed with this statement. Furthermore, Figure 8 shows that 56% of the participants "would enjoy using the privacy bubbles" (vs. 15% who would not) and 54% of the participants "would accept the privacy bubbles" (vs. 18% who would not). These results have been confirmed by the following comments left by the participants: "Privacy bubbles seem to be an easy process for sharing photos", "Interesting concept. I guess this would make things much easier", "It sounds like a great idea", "It sounds like an interesting new concept to share pictures with others based on their whereabouts when the picture was taken", or "Where could I access and test it?".

In summary, the participants have confirmed that the four design drivers are reflected in the proposed concept of privacy bubbles. Additionally, they have provided valuable insights about different design alternatives for the implementation of our prototype detailed in Section 5.

5. Proof-of-Concept Implementation

Based on the findings of the aforementioned questionnaire-based study, we have prototypically implemented the concept of privacy bubbles. We first present an overview of the architecture of our implementation, before describing the underlying mechanisms in detail.

5.1. Overview

The architecture comprises mobile phones, and an application server, modeling a online sharing platform. Mobile phones are particularly adapted to the implementation of the concept of privacy bubbles, since 61% of the participants of our study have already adopted them to take the pictures they share (cf. Figure 3). All mobile phones run an application that allows their users to (1) create and configure their bubbles, (2) take pictures, (3) review pictures prior to their upload, (4) upload and share pictures, and (5) search for pictures taken by others. Moreover, the application collects contextual information about the users in the background. The client application is implemented for Nexus S mobile phones running the Android operating system. The server stores the uploaded pictures and manages the privacy policies and access rights to each picture according to the user's privacy preferences. The communication between the mobile phones and the server are secured using TLS/SSL, and the server is secured against fraudulent access using well-established mechanisms.

5.2. Underlying Mechanisms

We herein present the steps conducted by the users and the associated mechanisms from the creation of privacy bubbles to the release of pictures.

5.2.1. Bubble Creation

Users start the creation of a new bubble via the main interface illustrated in Figure 9(a). They first choose the radius of their bubbles ΔT by choosing between the options of (1) city, (2) area, (3) building, or (4) nearby. If the *area* option is selected, the users have to determine the radius of the area centered around their own location using the interface depicted in Figure 9(b). Note that the proposed values for the radius of an area can be customized by the users in order to reflect their personal preferences as good as possible. Finally, users achieve the creation of the bubble by setting its duration of existence ΔT as one of the options of (1) unlimited, (2) same day, (3) a custom time interval, or (4) the exact time of the capture of the picture.

5.2.2. Taking Pictures

Once a privacy bubble has been created, the users access the picture management interface shown in Figure 9(c) and can take pictures as usual. A background

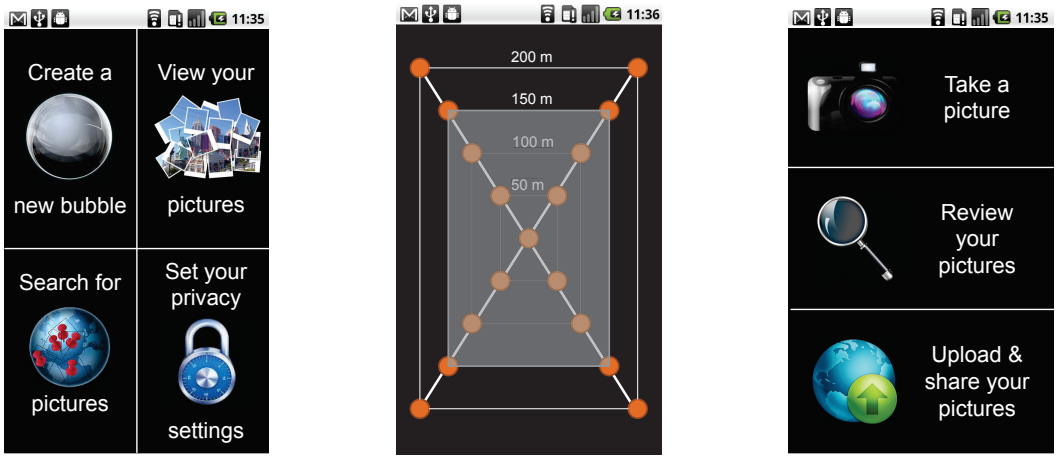

(a) Main functions (b) Bubble personalization (c) Picture management

Figure 9: Screenshots of selected user interfaces

mechanism captures information about the user's current context while pictures are being taken in order to later determine which other users were included in the current privacy bubble.

Indirect vs. Direct Localization Mechanism. We have designed and implemented two mechanisms that both capture the user's context, but differ in the modality of the collected location information. In the first mechanism, referred to as the *indirect* mechanism, the mobile phone collects spatiotemporal information (denoted by L_P) for each taken picture P. This includes GPS coordinates (referred to as GPS(L_P)), scanned Wi-Fi access points (APs) (referred to as Wi-Fi(L_P)), and scanned Bluetooth devices (referred to as Bluetooth(L_P)). Depending on the degree of location granularity selected by the users, different types of data are collected as summarized in Table 2. Note that in indoor environments, the application solely relies on Wi-Fi APs and Bluetooth devices to determine the location of the users, therewith inherently confining the radius of the bubble. The collected information L_P as well as the parameters of the bubble (i.e., the duration ΔT and the radius ΔL) are then appended as metadata to the picture.

In comparison to the indirect mechanism, the *direct* mechanism does not collect location information from GPS or other wireless communication technologies, but instead annotates pictures with the IDs broadcasted by nearby users. For the duration of the bubble, the mobile phone therefore periodically broadcasts messages advertising the ID of its user and at the same time listens for such messages sent from other phones. In our prototype, we have used the discovery service provided by

Table 2: Mapping of degrees of location granularity with collected data

Granularity	Collected data
City	GPS
Area	GPS
Building	Wi-Fi APs
Nearby	Bluetooth devices

the AllJoyn technology [4] to broadcast the advertisement messages. This technology supports ad-hoc communication between mobile devices in physical proximity. It is agnostic to the operating system running on the phones and the underlying wireless interfaces (Wi-Fi or Bluetooth). In contrast to Bluetooth, it does not requires any user interactions and runs in the background. In the direct mechanism, the radius of the bubbles is thus determined by the range of the wireless technology. Users who select to use this mechanism in the "Privacy settings" (illustrated in Figure 9(a)) hence only configure the duration of their bubble and do not access the bubble personalization interface (shown in Figure 9(b)), which is exclusively used in the indirect mechanism.

In summary, the indirect mechanism allows the users to freely define the radius of their bubbles. This freedom however comes at the cost of reduced location privacy, since users provide spatiotemporal information to the server. On the other side, the direct mechanism does not reveal spatiotemporal information about the users to the application server, but restricts the bubble radius to the range of the underlying wireless technology. Our

prototype implementation includes both the indirect and the direct mechanism in order to foster the acceptance of our approach by potential users. The decision to implement both mechanisms is based on the fact that roughly half of the participants were ready to provide spatiotemporal information to the server, whereas the remaining users were reluctant to provide information about the location in which their pictures were taken.

Picture-based vs. Periodic Location Detection Mechanism. In the above mechanisms, the spatiotemporal information and the collected user IDs are transmitted along with each uploaded picture. This implies that the pictures serve as grant for accessing further pictures, and that users thus need to take and share pictures in order to access pictures of other users. While this tit-for-tat aspect has been identified as important by 39% of the participants of our study, 43% judged it as an unimportant feature (cf. Figure 6). We therefore propose extended versions of both the direct and indirect mechanisms with relaxed sharing conditions. In the extended indirect mechanism, the mobile phone periodically provides spatiotemporal information to the sharing application. Similarly, the mobile phones periodically broadcasts the identity of its user in the extended direct mechanism, even if no picture is taken. While these extended versions may increase the number of pictures accessible by each user, they come at the cost of providing additional information to the sharing platform, such as the locations visited by the user or the identities of the users encountered. The choice between the regular and the extended mechanisms is up to the user, as both mechanisms depend on their personal privacy conception and their willingness to access more content. An evaluation of the impact of both original and extended versions of the direct and indirect mechanisms on the sharing behavior of users as well as their respective acceptance by potential users is considered as future work.

5.2.3. Reviewing Pictures

The users can review the taken pictures and decide which pictures they are willing to share with the persons who were included in their bubbles. After having reviewed the pictures, the users upload the pictures to share to the application server, which stores and clusters them by user ID or spatiotemporal information to facilitate the later verification of the inclusion of potential retrievers in the bubbles.

5.2.4. Accessing Pictures from Other Users

In addition to sharing pictures, a user can also query the server for pictures taken by other users. These pictures are however only accessible to a requesting user if he was included in the privacy bubble defined by the photographer at the time of capture of each picture.Verifying the inclusion in an existing privacy bubble is performed at the server side in the two steps shown in Algorithm 1. The server first verifies if the picture includes the user ID of the requesting user (noted ID_R) in its metadata. Next, the server compares the timestamp T_R (step 4 of Algorithm 1 and Algorithm 2) as well as the location L_R (step 5 of Algorithm 1 and Algorithm 3) included in the picture of the requester to those of the stored pictures. Note that the spatiotemporal information of the requester can be either included in his pictures or periodically delivered if he used the extended version of the indirect mechanism. Positive search results (i.e., pictures to release) are then displayed on a map, which can be browsed by the users on their mobile phone.

Algorithm 1 Search for pictures

Require: ID_R, T_R, L_R
1: **for all** pictures **do**
2: **if** $ID_R \in$ list of authorized IDs **then**
3: add picture to pictures to release
4: **else if** $T_R \in \Delta T$ **then** ▷ see Algorithm 2
5: **if** $L_R \in \Delta L$ **then** ▷ see Algorithm 3
6: add picture to pictures to release
7: **end if**
8: **end if**
9: **end for**
10: **return** pictures to release

6. Security Analysis, Discussion and Future Work

Based on the encouraging findings of our study and the proposed prototype implementation, we performed a security analysis focusing on (1) the protection of location information and (2) the protection of user IDs. We discuss potential attacks on location information and derive countermeasures against these attacks. Subsequently, we investigate the threat due to the falsification of user IDs. As an improvement to the original concept, we propose various extensions of both the design and the realization of the privacy bubbles.

6.1. Tampering with Location Information and Countermeasures

In the current version of our prototype implementation, malicious users can tamper with the location information included in the uploaded pictures when using

Algorithm 2 Verification of temporal inclusion

Require: T_R, T_P, ΔT
if ΔT = unlimited **then**
  return true
else if ΔT = day **then**
  if day(T_R)= day(T_P) **then**
    return true
  end if
else if ΔT = interval **then**
  if $T_R \in [T_P - \frac{\Delta T}{2}, T_P + \frac{\Delta T}{2}]$ **then**
    return true
  end if
else if ΔT = exact **then**
  if $T_R = T_P$ **then**
    return true
  end if
else
  return false
end if

Algorithm 3 Verification of spatial inclusion

Require: L_R, L_P, ΔL,
if ΔL = city **then**
  if city(GPS(L_R))=city(GPS(L_P)) **then**
    return true
  end if
else if ΔL = area **then**
  if $GPS(L_R) \in$ circle of radius $\frac{\Delta L}{2}$ around $GPS(L_P)$ **then**
    return true
  end if
else if ΔL = building **then**
  if $\exists$ Wi-Fi(L_R) $\in$ {Wi-Fi(L_P)} **then**
    return true
  end if
else if ΔL = nearby **then**
  if $\exists$ Bluetooth(L_R) $\in$ {Bluetooth(L_P)} **then**
    return true
  end if
else
  return false
end if

the indirect mechanism. In this section, we thus list potential attacks by malicious users before discussing possible countermeasures.

6.1.1. Attacks on the Location Information

Malicious users may attempt to tamper with the GPS location information by conducting the attacks described in [5], which we list as follows:

Tampering with GPS Modules. Malicious users may tamper with the GPS module embedded in the mobile phone. This can be realized by physically modifying the GPS hardware, or simulating a GPS device [5]. In the latter case, malicious users may introduce a simulator modeling an external GPS receiver interfaced via Bluetooth, such as [6] or [7]. In a subsequent step, malicious users configure the mobile phone to deliver the coordinates generated by the simulator to the application, instead of those provided by the on-board GPS module. Using the simulator, malicious users can hence directly control the generation of fake coordinates.

Manipulating GPS APIs. Malicious users may also modify the source code of the GPS API in open source operating systems, such as Android, in order to generate fake location coordinates. For debugging purposes, the Android operating system even provides means to supply arbitrary location information to the system. Thus, instead of relying on the coordinates provided by the on-board GPS receiver, fake coordinates from a server or a local file can be injected in the application.

Leveraging Device Emulators. Instead of only simulating a GPS receiver, malicious users may leverage mobile phone emulators in order to create virtual devices. They can install the application on the created devices and control the coordinates provided by the emulated GPS module in a similar fashion as presented above.

Manipulating Transmitted Data. In addition to the attacks presented in [5], malicious users may attempt to modify the location information during its transmission to the application server. Our prototype implementation, however, establishes secured communication between the mobile phones and the server based on the TLS/SSL protocols. In absence of key leakage, such manipulations on the communication path are hence strongly complicated.

6.1.2. Countermeasures

In order to counter these attacks, our prototype implementation could be extended by the following measures.

Additional Location Information. In our current implementation, the GPS coordinates are completed by scanned Bluetooth devices and Wi-Fi APs. While this

additional location information can also be manipulated, the more information provided by the users about their positions, the more difficult for malicious users to find the right combination of information to fake in order to fraudulently access shared data. Furthermore, we plan to improve the precision and reliability of the location information provided by the mobile phones by completing the positioning information by additional sensing modalities (such as microphone and light sensor). Enhanced precision will refine the granularity of the bubbles and allow users to define even smaller bubbles, e.g., at room level. The reliability of the access control will also be improved since it currently only depends on precision of the GPS coordinates and the scanned Bluetooth and Wi-Fi APs. Note that introducing additional sources of location information inherently reduces the radius of the corresponding bubbles, but simultaneously reduces the risk of manipulation.

Verified Location Information. By relying on location information uniquely provided by the mobile phones of the users, there always exists a risk that malicious users succeed in manipulating them. This risk could be reduced, if a third party could be able to verify and confirm that users were actually located in the area of the city they pretend to be. One imaginable third party could be the network providers since they already have access to the location of users based on the identity of the GSM cell they are currently located in. In order to preserve the location privacy of the participants, the bubble sensing application could provide the spatiotemporal information provided by the users to the network providers for verification, but would not be able to request location information about particular users. In addition to verifying the location of the users, introducing this out-of-band verification would allow the application to detect and eliminate emulated devices potentially created by malicious users. The precision of the verification would, however, be limited to the dimension of the GSM network cell. Hence, a verification with a finer degree of granularity would not be achievable.

Introduction of Detection Mechanisms. In addition to verify location information, additional mechanisms could be implemented at the server side in order to detect malicious users attempting to fraudulently access shared data. For example, such mechanisms could monitor the frequency at which users upload data, the content of their upload, and the coherence of the route they follow. Indeed, malicious users need to guess the combination of both the location and temporal information verifying the bubble parameters of a picture. Since the probability to guess a right combination increases with the number of fake uploaded pictures, malicious users may be tempted to upload a large number of pictures, but this simultaneously increases the risk to be detected by the application server. Moreover, the application server could verify if an identical version of an uploaded picture has been previously uploaded with different metadata. Finally, the application server could analyze consecutive uploads of the same users in order to verify the coherence and the realism of the associated location information. In case of fake coordinates arbitrarily generated by malicious users, the application server could identify incoherent paths or implausible motion speeds. By generating realistic paths, malicious users increase the chance not to be identified by the server, but also decrease the chance to get access to shared pictures, as it limits the range of possible trials.

In order to reduce the success probability of malicious users to the minimum and ideally to zero, we plan to examine applicability and performances of the aforementioned countermeasures in order to develop an adequate and efficient solution for our proof-of-concept implementation.

6.2. Falsification of User IDs

Malicious users using the direct mechanism may attempt to either (1) register the IDs of other users to the application server or to (2) broadcast the IDs of other users. While the former attack is prevented by well-established authentication mechanisms at the application server, additional mechanisms could be integrated into our prototype implementation in order to protect the users against the latter attack. These mechanisms may include (1) the signature of the broadcasted IDs by the clients and their verification by either the application server or the clients, or (2) the introduction of ID-based encryption in our architecture [8]. By applying the first approach, the application server would verify that the broadcasted ID matches with the identity of the client who broadcasted it before releasing the potential shared data, while the clients could verify it upon reception of the broadcasted IDs. By applying the second approach (i.e., introducing ID-based encryption), an independent and trusted key generator should be introduced in order to generate a master key pair as well as a private key for each client. Upon registration, each client would be assigned a key pair composed of the public master key and the client's own private key. The clients would use their private key to sign the ID they broadcast. Other users receiving the broadcasted ID would compute the public key of the associated client based on the public master key and the identity of the broadcasting client.

As a result, the clients could verify that the broadcasted ID actually corresponds to the ID of the sender. The key generator, however, represents a single point of failure and all keys need to be revoked if it is compromised. Moreover, the signature of the broadcasted IDs does not prevent malicious users to replay the signed IDs. Additional solutions should therefore be found to prevent this attack, such as introducing certified timestamps in the broadcasted IDs or integrating distance-bounding protocols. While the discussed solutions would allow us to prevent malicious users from falsifying the broadcasted IDs, it would seriously increase the complexity of the architecture and introduce major overhead for limited gains. In fact, malicious users falsifying broadcasted IDs need to collude with the user having this ID or create multiple accounts to get access to pictures, which they could directly access by broadcasting their user ID.

6.3. Location Privacy

Privacy bubbles require the disclosure of information about the users to the application server in order to match the persons included in the privacy bubbles and their creator. The more content users are willing to access, the more location information should be provided and hence, the more threats to location privacy arise. In our prototype implementation, we have proposed different mechanisms allowing users to choose both the type of information released to the application server and the corresponding frequency. The indirect mechanism leverages spatiotemporal annotations, while the direct mechanism monitors nearby user IDs (cf. Section 5.2.2). Both mechanisms can collect the information of interest either at the time of the capture of the picture or periodically. If users want to protect their location privacy, they can choose to use the direct mechanism, which only reveals the IDs of nearby mobile phones. Location privacy may however only be endangered if other users collude with the application server and reveal their location and thus the location of their victims. The likeliness of this attack is limited since it requires the physical proximity of the attackers to their victims. We envision to protect the location privacy of users using the indirect mechanisms by adding a trusted middleware to our current implementation and applying obfuscation mechanisms. For example, mechanisms based on the k-anonymity principle [9], such as tessellation [10] or microaggregation [11], can be applied. In the tessellation mechanism, the geographic area is divided into multiple tiles, each of them containing at least k users. The exact coordinates of the users are then replaced by either the geographical boundaries or the center of the current tile, which are then reported to the application server. Since k users are included in the same tile, they become indistinguishable. In comparison, the microaggregation scheme replaces the exact coordinates of the users by the average location of the k nearest users and similarly protects the location privacy of the k users. While both mechanisms increase the location privacy of the users, they simultaneously prevent the definition of fine-grained bubbles and lower the precision at which the inclusion of other users in bubbles can be verified. Consequently, further mechanisms should be examined to provide enhanced location privacy while still supporting the realization of the privacy bubbles.

6.4. Modular and Malleable Bubbles

The proposed bubbles are currently spheric and centered on the users. In the future, malleable bubbles could be used, which can dynamically adapt themselves to the form of a room where the users could freely move without modifying their bubbles in order to provide enhanced privacy protection.

6.5. Multimedia Contents

In this paper, the feasibility of privacy bubbles has been studied for picture sharing applications. Its applicability is, however, not confined to sharing pictures, and should be further investigated for additional user-generated content such as videos or audio recordings. By extending the scope of the shared content to other media, we plan to examine, e.g., if the interest of potential users in using the privacy bubbles would increase and if it would influence their acceptance. Moreover, we will analyze the impact of the nature of the shared content on the decisions of the users in terms of bubble dimensions and selected mechanisms (see Section 5.2.2 and Section 5.2.2).

6.6. Long-term Evaluation

Once the above enhancements will be integrated, we will deploy our approach for a long-term user study. A set of users will evaluate the privacy bubbles under real-world conditions. This will allow us to gain precious insights about the utilization and acceptance of the privacy bubbles by real users. For example, we are planning to investigate in which contexts users make use of the privacy bubbles (e.g., concerts, sport events, etc), which mechanisms they prefer to apply among the direct and indirect mechanisms and their extended versions. We expect to brain feedback about the offered and potentially missing functions as well as their usability with a fine degree of granularity. Finally, a long-term study will also reveal how the behavior of the users evolves with their experience in using the bubbles.

7. Related Work

A wide range of existing work, such as [12, 13], focuses on defining policies, rules, or semantics for access control mechanisms. They mainly contribute technical solutions, which remain invisible to the users and obscure for non-experts. Within the scope of this work, we however concentrate on existing mechanisms directly controlled by the users. Among the existing solutions, most of the mechanisms rely on individual authorizations managed by the users, who manually select individuals (or build groups of individuals) authorized to access their pictures. The way how groups are defined varies from an application to another, but the underlying principle remains the same. For example, Facebook utilizes scrolling lists, while Google+ proposes "circles" to visualize the groups of individuals formed. In contrast to these solutions, our concept differs in two dimensions: (1) the authorization to access pictures is delivered based on spatiotemporal conditions and (2) this authorization is dynamically and automatically managed by the system based on the radius and duration of the bubbles defined by the users. The "geofences" introduced in Flickr [14] allow users to define geographical zones on a map and select the persons able to access the pictures taken in theses zones. Even if the geofences includes a spatial component, the proposed solution remains static and the users need to set up each fence and select the authorized users individually. Moreover, our concept not only considers the location of the photographers at the time of capture of the pictures, but also the location of the persons able to access these pictures at the same time. The Color application [15, 16] shares a number of similarities with our approach since people located in proximity of the photographers can directly access their pictures. Color does however not only limit their access to the nearby persons, but considers each picture as public, which endangers the users' privacy.

8. Conclusions

In this paper, we have presented a complementary approach to the relationship-based access control mechanisms applied in most current online picture sharing platforms. We have defined design drivers for a novel concept called privacy bubbles, which allow users to share pictures with other users to which no social ties exist. Users control the bubbles, i.e., the sharing spatiotemporal boundaries, as well as the pictures shared within the bubbles. The privacy bubble paradigm is thus centered around the users and takes into account their individual privacy conception. We have hence thoroughly investigated the feasibility of our concept by submitting it to the 175 participants of our study for evaluation. The results show that a majority of the participants would feel comfortable using our approach and would be ready to accept it. We have further implemented a proof-of-concept of our approach to examine its technical feasibility.

Acknowledgments

The authors would like to thank the participants of the study as well as Stanislaus Stelle and Stefan Hartung for their contributions to this paper. This work was supported by CASED (`www.cased.de`).

[1] D. Christin, P. Sánchez López, A. Reinhardt, M. Hollick, M. Kauer, Privacy Bubbles: User-Centered Privacy Control for Mobile Content Sharing Applications, in: I. Askoxylakis, H. Pöhls, J. Posegga (Eds.), Information Security Theory and Practice. Security, Privacy and Trust in Computing Systems and Ambient Intelligent Ecosystems, Vol. 7322 of Lecture Notes in Computer Science, Springer Berlin / Heidelberg, 2012, pp. 71–86.

[2] A Snapshot of Facebook in 2010, Online: `http://www.facebook.com` (accessed in 01.2012).

[3] YouTube Statistics, Online: `http://www.youtube.com/t/press_statistics` (accessed in 06.2012).

[4] AllJoyn Peer-to-Peer, Online: `http://developer.qualcomm.com` (accessed in 01.2012).

[5] W. He, X. Liu, M. Ren, Location Cheating: A Security Challenge to Location-Based Social Network Services, in: Proceedings of the 31st IEEE International Conference on Distributed Computing Systems (ICDCS), 2011, pp. 740–749.

[6] Skylab GPS Simulator, Online: `http://www.skylab-mobilesystems.com` (accessed in 06.2012).

[7] Zyl Soft, Online: `www.zylsoft.com` (accessed in 06.2012).

[8] A. Shamir, Identity-based Cryptosystems and Signature Schemes, in: Proceedings of CRYPTO 84 on Advances in Cryptology, 1985, pp. 47–53.

[9] L. Sweeney, K-anonymity: A Model for Protecting Privacy, International Journal of Uncertainty, Fuzziness, and Knowledge-Based Systems 10 (5) (2002) 557–570.

[10] K. L. Huang, S. S. Kanhere, W. Hu, Preserving Privacy in Participatory Sensing Systems, Computer Communications 33 (11) (2010) 1266 – 1280.

[11] J. Domingo-Ferrer, J. Mateo-Sanz, Practical Data-oriented Microaggregation for Statistical Disclosure Control, IEEE Transactions on Knowledge and Data Engineering 14 (1) (2002) 189–201.

[12] J. Joshi, E. Bertino, U. Latif, A. Ghafoor, A Generalized Temporal Role-based Access Control Model, IEEE Transactions on Knowledge and Data Engineering (2005) 4–23.

[13] D. Kulkarni, A. Tripathi, Context-aware Role-based Access Control in Pervasive Computing Systems, in: Proceedings of the 13th ACM Symposium on Access Control Models and Technologies (SACMAT), 2008, pp. 113–122.

[14] D. Leung, Introducing Geofences on Flickr!, Online: `http://blog.flickr.net` (accessed in 01.2012).

[15] Color Application, Online: `http://www.color.com` (accessed in 09.2011).

[16] B. Upbin, Color, a Twitter for Photo and Video, Launches with $41 Million, Online: `http://www.forbes.com` (accessed in 01.2012).

Vitæ

Delphine Christin graduated in Electrical Engineering from Ecole Nationale Supérieure de l'Electronique et ses Applications, France, and Technische Universität Darmstadt, Germany, in 2009. She is a research assistant at the Center for Advanced Security Research Darmstadt (CASED) since April 2009 and a member of the Secure Mobile Networking Lab (SEEMOO) at Technische Universität Darmstadt since October 2009. Her research interests include privacy schemes for participatory sensing scenarios, privacy interfaces, and privacy metrics.

Pablo Sánchez López graduated in 2011 in Telecommunication Engineering as well as in Information and Communication Engineering from Universidad Politécnica de Madrid (ETSIT UPM), Spain, and Technische Universität Darmstadt, Germany, respectively. Since September 2011, he is working as a radio optimization engineer at Ingenia Telecom, Spain. His current research interests are 3G and 4G networks optimization, efficient usage of the radio resources, and mobile networks from design to deployment.

Andreas Reinhardt received his Dipl.-Ing. (M.Sc. equivalent) in 2007 and his doctoral degree in 2011, both in Electrical Engineering and Information Technology from Technische Universität Darmstadt, Germany. Since 2012, he is head of the Distributed Sensing Systems group at the Multimedia Communications Lab at TU Darmstadt. His research interests lie in the area of ubiquitous computing, with a special focus on wireless sensor and actuator networks and their energy-efficient operation.

Matthias Hollick is heading the Secure Mobile Networking Lab (SEEMOO) at the Computer Science Department of Technische Universität Darmstadt, Germany. He received his Ph.D. degree in 2004 from the TU Darmstadt. He has been researching and teaching at TU Darmstadt, Universidad Carlos III de Madrid (UC3M), and the University of Illinois at Urbana-Champaign (UIUC). In 2005, for his research, he has received the Adolf-Messer Foundation award. His research focus is on secure and quality-of-service-aware communication for mobile and wireless ad hoc, mesh, and sensor networks.

Michaela Kauer finished her studies of psychology at the Technische Universität Darmstadt, Germany, in September 2008. Since November 2008, she is research associate at the Institute of Ergonomics at the Technische Universität Darmstadt, Germany. She is head of the research group usability at the Institute of Ergonomics since November 2010. Her research interests are user acceptance of technical products with focus on usable security and privacy and semi-automated driving.