Privacy Bubbles: User-Centered Privacy Control for Mobile Content Sharing Applications

Delphine Christin¹, Pablo Sánchez López¹, Andreas Reinhardt²,
Matthias Hollick¹, and Michaela Kauer³

¹ Secure Mobile Networking Lab, Technische Universität Darmstadt
Mornewegstr. 32, 64293 Darmstadt, Germany
`firstname.lastname@seemoo.tu-darmstadt.de`

² Multimedia Communications Lab, Technische Universität Darmstadt
Rundeturmstr. 10, 64283 Darmstadt, Germany
`andreas.reinhardt@kom.tu-darmstadt.de`

³ Institute of Ergonomics, Technische Universität Darmstadt
Petersenstr. 30, 64287 Darmstadt, Germany
`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 user study with 175 participants, and a prototype implementation shows the technical feasibility of our scheme.

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 [1] and the approximated 48 hours of video shared on YouTube every minute [4]. 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 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.

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 user 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 user study.

3. We present our proof-of-concept implementation of the privacy bubble concept, which takes the findings of our user study into account.

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 user study. We provide details about our prototype implementation in Section 5 and list possible future 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.

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

Fig. 1. Representation of Alice's, Bob's, and Carlos' bubbles for each taken picture

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 $[t-2 \text{ min}, t+2 \text{ min}]$ 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 user 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 information 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,

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

Nationality	n	Current job	n	Field of occupation	n
German	108	PhD student	72	Computer science	99
French	22	Undergraduate student	59	Electrical engineering	35
Spanish	9	Postdoctoral researcher	18	Psychology	5
Romanian	3	Professor	6	Biology	5
Indian	3	Administrative staff	5	Physics	4
Ukrainian	3	Technical staff	4	Mechanical engineering	4
Other	27	Other	11	Other	23

Fig. 2. Overall number of pictures shared

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 User Study Results

In this section, we present the findings of our user 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.

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 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.

Fig. 3. Estimated percentage of shared pictures taken with a mobile phone

Fig. 4. Distribution of the answers about the comprehensiveness and intuitiveness of the privacy bubbles

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 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

Fig. 5. Distribution of the answers about the importance of the control over the radius and the duration of the privacy bubbles

Fig. 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

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.

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

Fig. 7. Distribution of the answers about the appreciation of the dynamical and automatic nature of the privacy bubbles

Fig. 8. Distribution of the answers related to the acceptance of the privacy bubbles by the participants

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.

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 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 user study, we have prototypically implemented the concept of privacy bubbles. Our proof-of-concept is based on Android Nexus S mobile phones and an application server, modeling an 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). Moreover, they enable an easy collection of contextual information about the users. In this section, we present the different steps conducted by the users and the corresponding mechanisms from the creation of a privacy bubble to the upload of the pictures and their access by other users. Note that the application server is secured against fraudulent access using well-established mechanisms and its different functions can be easily integrated into existing sharing applications.

5.1 Bubble Creation

Users start the creation of a new bubble via the main interface illustrated in Figure 9(a). They determine its radius using the second interface depicted Figure 9(b). Moreover, the proposed values for the radius can be customized by the users in order to reflect their personal preferences as good as possible. The creation of the bubble is completed by the selection of its duration.

5.2 Taking Pictures

The users then access the picture management interface shown in Figure 9(c) and can take pictures as usual with their mobile phone. While taking pictures, a mechanism transparent for the users captures information about the user’s context in order to later determine which other users were included in the current privacy bubble.

Fig. 9. Screenshots of selected user interfaces

Indirect vs. Direct Localization Mechanism. We have first designed and implemented two mechanisms, which 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 about each taken picture. This includes the GPS coordinates, scanned Wi-Fi access points, and scanned Bluetooth devices. The collected information as well as the parameters (i.e., the duration and radius of the bubble) are then appended as metadata to the picture.

In comparison to the indirect mechanism, the *direct* mechanism does not collect absolute location information, but instead collects the IDs broadcasted by nearby users added to the picture’s metadata. For the duration of the bubble, the mobile phone therefore periodically broadcasts messages advertising the ID of its user and listening for similar messages coming from other phones. In our implementation, we have used the AllJoyn technology [2], which supports Wi-Fi ad-hoc communication between Android phones. In the direct mechanism, the radius of the bubbles is determined by the range of the wireless technology. Users, who select to use this mechanism in the “Privacy settings” illustrated in Figure 9(a), may 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 need to provide spatiotemporal information to the sharing platform. 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 employed wireless technology. We have included both indirect and direct mechanisms in our prototype implementation in order to foster the acceptance of our approach by potential users, since roughly half of

the participants were ready to provide spatiotemporal information to the sharing platform, while the other 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 unimportant (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.

5.3 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 later facilitate the verification of the inclusion of potential retrievers in the bubbles.

5.4 Accessing Pictures from Other Users

In addition to sharing pictures, a user can also query the application 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. The verification of the potential inclusion in privacy bubbles happens at the server side in two steps. First, the server searches for pictures including the user ID of the requesting user in their metadata. Next, the server compares the spatiotemporal information provided by the requester and potential content providers by taking the corresponding bubble parameters into account. The spatiotemporal information of the requester can be included in his pictures or periodically delivered if he used the extended

version of the indirect mechanism. Positive search results are then displayed on a map, which can be browsed by the users on their mobile phone.

6 Discussions and Future Work

Based on encouraging findings of our study and the proposed prototype implementation, we envision to improve both the concept and the realization of the privacy bubbles in the following dimensions:

6.1 Tampering with Spatiotemporal Information

In the current version of our prototype implementation, malicious users can tamper with the spatiotemporal information included in the uploaded pictures when using the indirect mechanism. For example, malicious users may upload the same fake picture with different metadata to fraudulently gain access to additional pictures. To be successful, attackers need to guess the combination of both the location and temporal information verifying the bubble parameters of a picture. The probability to guess a right combination increases with the number of fake uploaded pictures, but this simultaneously increases the risk to be detected by the application server, thus limiting the expected impact of this type of attack. In order to reduce the success probability of the attacker to zero, we however envision to examine different alternatives, such as the utilization of Trusted Platform Modules or Location Proofs [10] with regard to their applicability and acceptance in order to develop an adequate solution for our proof-of-concept implementation.

6.2 Falsification of User IDs

Malicious users using the direct mechanism may attempt to register the IDs of other users to the application server. This attack is however prevented by the authentication mechanisms at the application server. Next, an attacker can broadcast an user ID different from their own, granting the access to pictures taken in their proximity to another user. This attack is however useless, as an attacker needs to collude with the attacker having the broadcast ID to get access to pictures, which he could directly access by broadcasting his own 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). 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 [11], such as tessellation [6] or microaggregation [5], 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 Reliability of Location Information

We further plan to improve the precision 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 access points.

6.5 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.

6.6 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 contents such as videos or audio recordings.

6.7 Long-Term Evaluation

Once the above enhancements will be achieved, we will deploy our approach for a long-term user study. A set of users will evaluate the privacy bubbles under real-world conditions and provide additional feedback for their improvement.

7 Related Work

A wide range of existing work, such as [7,8], 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 [9] allow users to define geographical zones on a map and select the persons able to access the pictures taken in these 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 [3,12] 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 privacy of the users.

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 user 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 user study and Stanislaus Stelle for their contributions to this paper. This work was supported by CASED (www.cased.de).

References

1. A Snapshot of Facebook in 2010, <http://www.facebook.com> (accessed in January 2012)
2. AllJoyn Peer-to-Peer, <http://developer.qualcomm.com> (accessed in January 2012)
3. Color Application, <http://www.color.com> (accessed in September 2011)
4. YouTube Statistics, http://www.youtube.com/t/press_statistics (accessed in January 2012)
5. Domingo-Ferrer, J., Mateo-Sanz, J.: Practical Data-oriented Microaggregation for Statistical Disclosure Control. *IEEE Transactions on Knowledge and Data Engineering* 14(1), 189–201 (2002)
6. Huang, K.L., Kanhere, S.S., Hu, W.: Preserving Privacy in Participatory Sensing Systems. *Computer Communications* 33(11), 1266–1280 (2010)
7. Joshi, J., Bertino, E., Latif, U., Ghafoor, A.: A Generalized Temporal Role-based Access Control Model. *IEEE Transactions on Knowledge and Data Engineering*, 4–23 (2005)
8. Kulkarni, D., Tripathi, A.: Context-aware Role-based Access Control in Pervasive Computing Systems. In: *Proceedings of the 13th ACM Symposium on Access Control Models and Technologies (SACMAT)*, pp. 113–122 (2008)
9. Leung, D.: Introducing Geofences on Flickr! <http://blog.flickr.net> (accessed in January 2012)
10. Luo, W., Hengartner, U.: Proving Your Location Without Giving up Your Privacy. In: *Proceedings of the 11th Workshop on Mobile Computing Systems and Applications (HotMobile)*, pp. 7–12 (2010)
11. Sweeney, L.: K-anonymity: A Model for Protecting Privacy. *International Journal of Uncertainty, Fuzziness, and Knowledge-Based Systems* 10(5), 557–570 (2002)
12. Upbin, B.: Color, a Twitter for Photo and Video, Launches with \$41 Million, <http://www.forbes.com> (accessed in January 2012)