Towards PII-based Multiparty Access Control for Photo Sharing in Online Social Networks

Nishant Vishwamitra[†], Yifang Li[†], Kevin Wang[†], Hongxin Hu[†], Kelly Caine[†] and Gail-Joon Ahn[‡]

[†]Clemson University [‡]Arizona State University
{nvishwa, yifang2, kwang2, hongxih, caine}@clemson.edu, ahn@asu.edu

ABSTRACT

The privacy control models of current Online Social Networks (OSNs) are biased towards the content owners' policy settings. Additionally, those privacy policy settings are too coarse-grained to allow users to control access to individual portions of information that is related to them. Especially, in a shared photo in OSNs, there can exist multiple Personally Identifiable Information (PII) items belonging to a user appearing in the photo, which can compromise the privacy of the user if viewed by others. However, current OSNs do not provide users any means to control access to their individual PII items. As a result, there exists a gap between the level of control that current OSNs can provide to their users and the privacy expectations of the users. In this paper, we propose an approach to facilitate collaborative control of individual PII items for photo sharing over OSNs, where we shift our focus from *entire* photo level control to the control of individual PII items within shared photos. We formulate a PII-based multiparty access control model to fulfill the need for collaborative access control of PII items, along with a policy specification scheme and a policy enforcement mechanism. We also discuss a proof-of-concept prototype of our approach as part of an application in Facebook and provide system evaluation and usability study of our methodology.

KEYWORDS

Access control, privacy, PII, multiparty, online social networks

ACM Reference format:
Nishant Vishwamitra[†], Yifang Li[†], Kevin Wang[†], Hongxin Hu[†], Kelly Caine[†] and Gail-Joon Ahn[‡] [†]Clemson University [‡]Arizona State University {nvishwa, yifang2, kwang2, hongxih, caine}@clemson.edu, ahn@asu.edu . 2017. Towards PII-based Multiparty Access Control for Photo Sharing in Online Social Networks. In *Proceedings of SACMAT'17, June 21–23, 2017, Indianapolis, IN, USA, ,* 12 pages.
https://doi.org/http://dx.doi.org/10.1145/3078861.3078875

SACMAT'17, , June 21–23, 2017, Indianapolis, IN, USA

ACM ISBN ACM ISBN 978-1-4503-4702-0/17/06…$15.00.
https://doi.org/http://dx.doi.org/10.1145/3078861.3078875

1 INTRODUCTION

Online social networks (OSNs) have faced a tremendous growth in recent years and become a major aspect of the Internet for socializing and sharing information among hundreds of millions of users. Facebook, for example, claims that it has 1.86 billion current active users [28]. OSNs like Facebook allow sharing information such as photos, videos and text messages, which can possibly contain sensitive and private information. Especially in sharing of visual data, such as photos, users are likely to share private information with an unknown audience, due to limited control over sharing such visual data [18]. To protect such sensitive information, access control has received considerable attention as a central feature of OSNs [2].

A vast majority of current Internet users are also OSN users [1], which implies that more users are shifting to OSNs for information exchange. As a result, users themselves have emerged as the largest contributers of content towards OSNs. A critical implication of this is that users are now faced with the additional responsibility of managing the online content that is associated with them. A large part of the shared information on Facebook consists of photos [27]. Facebook allows users to share photos with other users, but the responsibility of managing the audience of the photo lies with the uploader of the photo [10]. Furthermore, in a group photo setting, also known as *multiparty* photo, users appearing in the photo have no control over who can view their personal information in the photo. Existing OSNs do not provide effective mechanisms to sufficiently address how users appearing in a multiparty photo can control the visibility of their individual private information.

Although it may appear that the main focus of a multiparty photo is user *faces* [12], there are numerous other private information of a user that can also appear in a multiparty photo. These private information points of a user are called as Personally Identifiable Information (PII) of the user. In the context of OSNs, PII can be defined as *"information which can be used to distinguish or trace an individual's identity either alone or when combined with other public information that is linkable to a specific individual"* [14]. There are a large number of PII items that can be leaked in a multiparty photo. For example, a user who has a very unique tattoo on her/his *body* can be used to identify the user in a photo. Similarly, a user who has a unique *belonging*, such as a uniquely colored vehicle, may be identified using the *belonging* in the multiparty photo. As research in the field of PII have pointed out [13, 14, 21, 26], there are numerous such PII items that can link a user with her/his identity in a multiparty photo.

Current OSNs, such as Facebook, do not provide any mechanisms for collaborative control of PII items in multiparty photos. In fact, Facebook does not provide any means of collaborative control of shared visual information. Facebook privacy policy allows the uploader of the photo to completely control photo sharing. Facebook has traditionally supported three levels of photo sharing: Public, Friends and Only Me. Recently, in an attempt to increase the granularity of photo sharing, Facebook has introduced *smart lists* [22]. Using these smart lists, an uploader can specify a subset of users from her/his friends list for sharing a photo, such as close friends and colleagues. However, numerous studies have shown that users struggle to adopt this feature for managing their friends and customizing their privacy settings [4, 7, 25], because of a non-trivial process [17, 30]. As a result, significant privacy violations and mismatched user expectations in OSNs have been identified [19, 20, 32].

The need of collaborative management for data sharing, especially photo sharing, in OSNs has been addressed by some recent research [5, 9, 11, 15, 29, 31]. However, all those solutions can only enable a collaborative control based on the *entire* photo level and lacks the support for the control of individual PII items in the shared photos. Face/Off [12] adopts a simple access control model to enable users to collaboratively control their *faces* in a shared photo in OSNs. When a viewer who does not have access to view the face of a user views a photo of the user, Face/Off uses the technique of *blurring* to hide the portion of the user's face in the photo, so that it is not visible to the unauthorized viewers. A main issue with this solution is that it does not enable specifying fine-grained access control policies for other crucial PII items, such as *body* and *belonging*, of a user. In addition, the nature of PII items differ substantially. For example, several PII items can be shared amongst multiple users, such as *location* information of a multiparty photo. A certain user may not wish to share the location of the photo, whereas another user might want to share the same location information with all her/his friends. Hence, this gives rise to additional *conflicts* in collaborative control of individual PII items in a multiparty photo. Therefore, it is essential to develop a more effective and flexible access control mechanism for multiparty photo sharing in OSNs, accommodating the special authorization requirements coming from multiple associated users for managing their individual PII items collaboratively.

In this paper, we propose an approach to enable collaborative management of shared visual data such as photos in OSNs, by enabling *fine-grained*, *PII-level* control. A PII-based Multiparty Access Control (PMAC) model is formulated to accommodate the core requirements of PII-level multiparty authorization in photo sharing in OSNs. We also provide a PII-based multiparty policy specification scheme and a policy evaluation mechanism. Since policy conflicts are inevitable in multiparty authorization enforcement, a conflict resolution method unique to PII-level privacy control is further introduced to deal with policy conflicts via balancing the need for privacy protection and information sharing. In addition, we provide a prototype implementation of our approach in the context of Facebook. Our experimental results based on comprehensive system evaluation and usability study demonstrate the feasibility and practicality of our solution.

The rest of the paper is organized as follows. In Section 2, we overview Facebook privacy management mechanism and evaluate the importance of PII level control in photo sharing in OSNs. We articulate our proposed PMAC model, including multiparty authorization specification and multiparty policy evaluation in Section 3. The details about prototype implementation and experimental results are described in Section 4. We overview the related work in Section 5. Section 6 concludes this paper and discusses our future directions.

2 PRELIMINARIES

2.1 Facebook's Privacy Model

Facebook allows its users to manage the privacy settings of their uploaded content, such as photos, videos, posts and comments. Currently, Facebook allows 4 levels of granularity for photo sharing: *Public, Friends, Only Me* and *Custom* [3]. The *Public* level allows all users of Facebook to access the shared content. The *Friends* level allows all users present in the user's friends list to access the shared content. In this level, a user can also specify if shared content should also be made available to *friends of friends* of the users tagged in the shared content. The *Custom* level allows users to specifically allow or deny a certain group of users to access the shared content. The default sharing setting is the *Public* level, i.e, if a user does not change her/his sharing settings, the shared content is accessible to everyone.

Users are provided with an option to create and maintain special lists, such as *Colleagues*, *Close friends* and *Family*, which offer more granular control of content sharing. Users can add/remove their friends to/from special lists. Users can manage trust levels by maintaining different privacy settings for different lists. The visibility of such lists is private, unless explicitly changed by users.

However, Facebook does not provide content stakeholders with any control over visibility of the shared content, allowing the content owner to be the sole controller of the shared content. In addition, Facebook does not provide any control over the visibility of context factors or mutual friends.

2.2 Importance of PII Privacy Control in Online Photo Sharing

Personally Identifiable Information (PII) can be defined as *any information about an individual maintained by an agency, including (1) any information that can be used to distinguish or trace an individual's identity, such as name, social security number, date and place of birth, or biometric records; and (2) any other information that is linked or likable to an individual, such as medical, educational, financial, and employment information* [21]. In addition to conventional PII items, there are several PII items that can be potentially used to identify a person, such as a person's location, belongings, certain distinguishing characteristics of their bodies and affiliations [14, 21]. Thus, the privacy leakage due to PII is a crucial privacy issue in OSNs [14]. Especially, due to large-scale photo sharing supported by OSNs, there is an immense compromise of user privacy in terms of users' PII items, since a problem of *linking* arises as a result, where a user in a photo can be associated with their identities by their PII items. However, current OSNs such as Facebook are not equipped with privacy models having PII level granularity.

Table 1: Example of PII Items and Leaked Private Information.

PII	Leaked Private Information
Face/Body	Gender Sexual orientation Relationship status
Affiliation	Groups affiliated to Job/Occupation School information
Belonging	Official documents Relationship status Religion Interests Favorite music Favorite books
Location	Physical address Hometown

Table 1 gives examples about the kind of private information that is leaked when PII items are viewed by unauthorized users [16, 21]. These examples of PII items are especially relevant to photo sharing in OSNs. For example, let's assume that an OSN user uses the Face/Off solution [12] to blur her/his face from unknown users in all her/his shared photos. Let's assume this user has a unique tattoo on her/his hand that many users who are not her/his friends know about. Even though this user uses face blurring to hide her/his identity from unknown users, the unknown users may guess her/him in the photo through her/his tattoo. As a second example, assume that an OSN user works at a firm, which does not allow its employees to disclose their association with the firm to anyone outside the firm. Let's assume that this user is photographed at a party, wearing her/his firm's uniform, with the firm logo clearly seen on the front of her/his shirt. In this case, there are two possibilities. First, even if the user chooses to blur her/his face, people who met her/him at the same party may recognize her/him by her/his firm logo. In addition, her/his firm might come to know of her/his violation of the firm's policy of non disclosure.

Besides, in a multiparty photo, an individual's privacy can be compromised by the presence of mutual friends of the individual and the viewer [12]. For example, assume Bob and Jane are photographed together. Suppose Bob does not want to be seen with Jane by users who are not his friends, hence he uploads a photo where his face is not clearly visible. But, since Jane's friends know that Bob and Jane are friends, they conclude that the other person in the photo with Jane is Bob. Since not all friends of Jane and Bob are friends with each other, users who are not Bob's friends end up viewing the photo and learning about Bob in the photo. Hence, a privacy policy related to mutual friends in a photo must be enforced as well.

3 PII-BASED MULTIPARTY ACCESS CONTROL FOR OSNS

To enable collaborative control of individual PII items for photo sharing in OSNs, we formalize the PMAC model (Section 3.1), along with a policy scheme (Section 3.2) and a policy evaluation mechanism (Section 3.3) for the specification and enforcement of PMAC policies in OSNs.

3.1 PMAC Model

An OSN system, such as Facebook, typically contains a set of users, a set of user profiles, a set of user visual data, and a set of user relationships (called friends lists in Facebook). *User profile* indicates who a user is in the OSN, including identity and personal information, such as name, birthday and interests. *User visual data* represents visual information, such photos and videos, that the user has in the OSN, created through various activities in the OSN. *User relationship* shows who a user knows in the OSN, representing user connections with friends, mutual friends, family, coworkers, colleagues, and so on.

Existing OSNs including Facebook do not provide effective mechanism to support collaborative privacy control of PII items over shared visual data. Several access control schemes [6, 9, 10, 33, 36] have been introduced that propose collaborative access control in OSNs. Unfortunately, these schemes only allow coarse-grained control of the whole visual data and do not offer any solutions for fine-grained control of PII items. An effective access control mechanism should allow fine-grained, collaborative control of individual PII items associated with a user.

One exception to the above schemes is the Face/Off model [12]. This model enables collaborative control of a user's $face$ in multiparty photos. However, several previous work [14, 16, 21] have discussed the importance of PII items in compromising privacy of an individual. In a multiparty photo, there are numerous PII items apart from *face*, that can compromise the privacy of an OSN user as well. In addition, PII items may be co-owned by several users, hence we also need a robust mechanism to address potential conflicts caused by the collaborative control of PII items. A flexible access control mechanism in a multi-user environment like OSNs should allow multiple controllers, who are associated with the shared visual data, to specify access control policies that can control individual PII items. As we have discussed in Section 2.2, in addition to the *owner* of content, *stakeholder* (the tagged user associated with the content) need to govern the access of the shared data as well due to possibly different privacy concerns. Additionally, every controller must be allowed to govern access control to their PII items to minimize the compromise of privacy.

In the context of OSNs that allow photo sharing, we have identified three kinds of PII items that can be associated with a user:

- **Unique PII.** A user's unique PII items uniquely identify the user in an OSN. For example, a user's face, body and belongings are unique PII items;
- **Shareable PII.** A user's shareable PII items can be linked to the user's identity and are shared with other users in an OSN. For example, a user's location is a shareable PII, as other users present in a photo share the location information in a photo; and
- **Relational PII.** The PII items that can be indirectly used to identify a user, based on the user's relationships with other users in an OSN are relational PII items. For example, *mutual friendship* in Facebook is a relational PII, because

in a multiparty photo, a viewer can guess the identity of a friend's friend in the photo, by the knowledge of the friend's identity.

Next, we formally define *Unique PII* and *Sharable PII* for PMAC model (see Figure 1) as follows:

Definition 1. ***(Unique PII)**. Let d be a visual data in the social network. Let u be a user identified in d. Unique PII (UP) of u constitutes the portion of the visual data that uniquely identifies u and it is owned by u. In PMAC model, a user's face, body and belonging are Unique PII items.*

Definition 2. ***(Shareable PII)**. Let d be a visual data in the social network. Let u be a user identified in d. Shareable PII (SP) of u constitutes the portion of the visual data that is collaboratively owned by u and a set of m users $\{u_1, \ldots, u_m\}$, $m \geq 1$. In PMAC model, a user's affiliation and location are Shareable PII items.*

Three types of controllers are identified in PMAC model. We define these types of controllers as follows:

Definition 3. ***(Visual Data Owner)**. Let d be a visual data item in the space of user $u \in U$ in the social network. The user u is called the Visual Data Owner of d.*

Definition 4. ***(PII Item Owner)**. Let d be a visual data item in the social network. Let $p \in I$ be a PII item identified in d, where I is a set of PII items in a set of Unique PII items, UP. Let u be a user who is linked to p. The user u is called the PII Item Owner of p.*

Definition 5. ***(PII Item Stakeholder)**. Let d be a visual data item in the social network. Let $p \in I$ be a PII item identified in d, where I is a set of PII items in a set of Shareable PII items, SP. Let T be the set of users who can be linked to p. A user u is called the PII Item Stakeholder of p, if $u \in T$.*

Different PII items tend to have different levels of importance for different user privacy concerns. Users tend to distinguish visibility of their PII items from various friends groups and also assign them different priorities called *sensitivity levels*. For example, a user may want to share her/his face with only her/his close friends. Therefore, her/his face has a *high* sensitivity level. In addition, the same user would not want her/his colleagues to learn about her/his personal belongings. Facebook introduced the concept of "smart lists" to enable sharing different content with different online social relationships. Using these smart lists, users can classify their friends into separate groups with different sharing settings, such as close friends, colleagues, and school. However, Facebook does not provide fine-grained sharing of PII items with separate smart lists, forcing users to share *entire* visual data with different friends lists. Of course, users in OSNs would assign different degrees of sensitivity to different PII items, and a PII item's sensitivity level can be leveraged to determine who are authorized to access the PII item. Several existing approaches [9, 10] have discussed how *sensitivity levels* can be utilized in OSNs. The concept of sensitivity level is also applicable to our PII-based collaborative sharing scenario. Therefore, in our model, we make the assumption that users can explicitly specify how sensitive their PII items are to their respective privacy concerns by assigning each PII item a sensitivity level when they specify their policy.

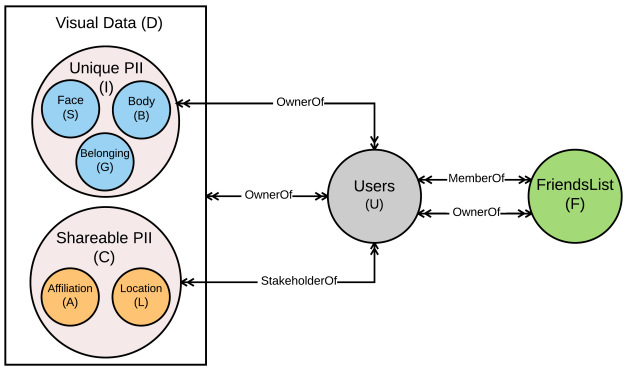

Figure 1: PMAC Model: Components and Relations.

Figure 1 represents the core components and relationships of our PMAC model. Note that in our model, a *user* can be the *owner* of her/his face, body and belonging, which is constituted by *PII item owner*, but a *stakeholder* of other types of *PII items*, such as *affiliation* and *location*, constituted by *PII item stakeholder*. Users are *owners* of their friends lists and can be *members* of other users' friends lists. A user can be a *viewer* of a visual data item, but a user who uploads a visual data item is the *owner* of the visual data item. We now formally define our model as follows:

- $U = \{u_1, \ldots, u_n\}$ is a set of users of the OSN. Each user has a unique identifier;
- $F = \{f_1, \ldots, f_m\}$ is a set of friends lists created by users in the OSN. Each friends list is identified by a unique identifier;
- $D = \{d_1, \ldots, d_p\}$ is a set of visual data items in the OSN. Each visual data item is identified by a unique identifier;
- $S = \{s_1, \ldots, s_l\}$ is a set of user faces in the OSN. Each user face is a <*u: sl: face-id*> tuple, $s_i =< u_i : sl_i : sid_i >$, where u_i is a face owner identifier, sl_i is a sensitivity level identifier and sid_i is a face identifier;
- $B = \{b_1, \ldots, b_t\}$ is a set of user body in the OSN. Each user body is a <*u: sl: body-id*> tuple, $b_i =< u_i : sl_i : bid_i >$, where u_i is a body owner identifier, sl_i is a sensitivity level identifier and bid_i is a body identifier;
- $A = \{a_1, \ldots, a_o\}$ is a set of user affiliations in the OSN. Each user affiliation is a <*u: sl: affiliation-id*> tuple, $a_i =< u_i : sl_i : aid_i >$, where u_i is an affiliation stakeholder identifier, sl_i is a sensitivity level identifier and aid_i is an affiliation identifier;
- $G = \{g_1, \ldots, g_s\}$ is a set of user belongings in the OSN. Each user belonging is a <*u: sl: belonging-id*> tuple, $g_i =< u_i : sl_i : gid_i >$, where u_i is a belonging owner identifier, sl_i is a sensitivity level identifier and gid_i is a belonging identifier;
- $L = \{l_1, \ldots, l_w\}$ is a set of user locations in the OSN. Each user location is a <*u: sl: location-id*> tuple, $l_i =< u_i : sl_i : lid_i >$, where u_i is a location stakeholder identifier, sl_i is a sensitivity level identifier and lid_i is a location identifier;
- $UF = \{uf_1, \ldots, uf_r\}$ is a collection of user friends lists, where $uf_i = \{uf_{i1}, \ldots, uf_{is}\}$ is a set of friends lists created by a user $i \in U$, where $uf_{ij} \in F$;

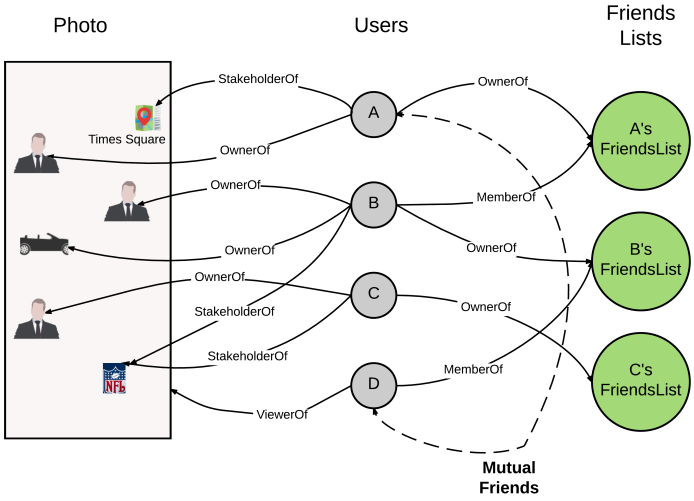

Figure 2: An Example of PII-based Multiparty Social Network.

- $CT = \{VO, PO, PS, VW, MB\}$ is a set of controller types, indicating *VisualDataOwnerOf*, *PIIItemOwnerOf*, *PIIItem-StakeholderOf*, *ViewerOf* and *MemberOf*, respectively;
- $CD = \{CD_{ct_1}, \ldots, CD_{ct_x}\}$ is a collection of binary user-to-PII item relations, where $CD_{ct_i} \subseteq U \times D$ specifies a set of $< user, visual\ data\ item >$ pairs with a controller type $ct_i \in CT$;
- $SL = \{sl_1, \ldots, sl_y\}$ is a set of supported sensitivity levels, which are assumed to be in the closed interval [0,1] in our model;
- $FU \subseteq F \times U$ is a set of 2-tuples $< FriendsList, user >$ representing user-to-friends list ownership relations;
- $controllers : D \xrightarrow{CT} 2^U$, a function mapping a visual data item $d \in D$, to a set of users who are the controllers of the visual data item with the controller type $ct \in CT$:
  $controllers(d : D, ct : CT) = \{u \in U \mid (u, d) \in CD_{ct}\}$;
- $visual\_data\_items : U \xrightarrow{CT} 2^D$, a function mapping each user $u \in U$ to a set of visual data items, where the user is a controller of the visual data items with the controller type $ct \in CT$:
  $visualDataItems(u : U, ct : CT) = \{d \in D \mid (u, d) \in CD_{ct}\}$;
- $user\_own\_friends\_lists : U \rightarrow 2^F$, a function mapping each user $u \in U$ to a set of friends lists created by this user:
  $user\_own\_friends\_lists(u : U) = \{f \in F \mid (\exists uf_u \in UF)[f \in uf_u]\}$;
- $friends\_list\_contain\_users : F \rightarrow 2^U$, a function mapping each friends list $f \in F$ to a set of users who are the members of this friends lists:
  $friends\_list\_contain\_users(f : F) = \{u \in U \mid (c, u) \in FU\}$;
- $user\_belong\_friends\_lists : U \rightarrow 2^F$, a function mapping each user $u \in U$ to a set of friends lists to which this user belongs:
  $user\_belong\_friends\_lists(u : U) = \{f \in F \mid (f, u) \in FU\}$;
- $UPS \subseteq U \times P \times S$ is a set of 3-tuples $< User, PIIItem, SensitivityLevel >$ representing user assigned sensitivity levels to PII items of the user;
- $sensitivity\_level : U, p \rightarrow SL$, a function returning the sensitivity level of a user-to-PII-item relation:
  $sensitivity\_level(u : U, p : (S \cup B \cup A \cup G \cup L)) = \{sl \in SL \mid (u, p, sl) \in UPS\}$;
- $all\_friends\_users : U \rightarrow 2^U$, a function mapping a user $u \in U$ to a set of users who are the members of the user's friends list:
  $all\_friends\_users(u : U) = \{u^{'} \in U \mid (\exists f \in user\_own\_friends\_lists(u))[u^{'} \in friends\_list\_contain\_users(F)]\}$; and
- $mutual\_friends\_list : F \rightarrow 2^U$, a function mapping a pair of users $< u, u' >, \{u, u'\} \in U$ to a set of users who belong to friends lists of both u and u'.
  $mutual\_friends\_list(u : U, u' : U) = \{all\_friends\_users(u) \cap all\_friends\_users(u')\}$.

Figure 2 depicts an example of PII-based multiparty OSN representation. It contains four individuals, Alice (A), Bob (B), Carol (C) and Dave (D), along with their relations with visual data items and friends lists. Note that a user may be related to more than one friends list, thus forming complex relationships. For example, in Figure 2, Bob is a *memberOf* Alice's friends list and also the *ownerOf* his own friends list. Dave is a *memberOf* Bob's friends list. Hence Alice and Dave are mutual friends, through Bob. This example depicts that a collaborative visual data item has multiple controllers. Since the photo depicts Alice, Bob and Carol, all three of them are controllers of the photo. The controller types are depicted in the example. In addition, a visual data item can have multiple *stakeholders*. For example, Bob and Carol are both stakeholders of the "NFL" logo. In our model, each user has complete ownership

of her/his face, body and belongings, as shown by the *ownerOf* relationship in Figure 2.

3.2 PMAC Policy Specification

To achieve authorization requirements with respect to the multiparty privacy concerns owing to multiple PII items, it is essential for access control policies to be in place to regulate access over individual PII items contained in a shared visual data associated with multiple controllers. Our policy specification scheme is constructed based on the proposed PMAC model. In our model, each controller of a shared visual data can specify one or more rules, as her/his policy governs who can view the PII items associated with them, contained in the shared visual data.

Viewer Specification: Viewers are a set of users who are granted/ denied access to the visual data item.We formally define the viewer specification as follows:

Definition 6. *(**Viewer Specification**). The viewer specification of a user $u \in U$ is defined as a set, $\{a_1, \ldots, a_n\}$, where each element is a user friends list $uf_u \in UF$, a set of users, $\{u_1, \ldots, u_m\}$, where $u_i \in U$, or everyone (*).*

For example, Alice can specify her colleagues, a particular set of users from her friends list, as viewer specification in her rule.

PII Item Specification: In OSNs, users can share their visual data, such as photos, with others. To facilitate effective policy conflict resolution for multiparty access control (Section 3.3.1), we introduce *sensitivity levels* for PII item specification, which are assigned by the controllers to the shared PII items. A user's judgment of the sensitivity level of the PII item is not binary (private/public), but multi-dimensional with varying degrees of sensitivity. Formally, the PII item specification is defined as follows:

Definition 7. *(**PII Item Specification**). Let $dt \in D$ be a data item. Let P be the set of PII items of a user $u \in U$ in dt. Let sl be a sensitivity level, which is a rational number in the range [0,1], assigned to PII items in P. The PII item specification is defined as a tuple $< p, sl >$, where p is a PII item .*

Access Control Policy: To summarize the above-mentioned policy elements, we give the definition of PMAC access control rule as follows:

Definition 8. *(**PMAC Rule**). A PMAC rule is a 5-tuple $R =< controller, viewer, shared\ visual\ data, PII\ items, effect >$, where*

- *$controller \in U$ is a user who can regulate the access of $data$;*
- *$viewer$ is a set of users to whom the authorization is granted/ denied, representing with an access specification defined in Definition 6.*
- *$shared\ visual\ data$ is a specific photo or all photos (*) where in the user is identified.*
- *$PII\ items$ is a set of PII items that the user wants to regulate access to, representing with an PII item specification defined in Definition 7.*
- *$effect$ is a tuple defined as $< p, share/blur >$ where p represents a PII item associated with the controller and $share/blur$ represents the authorization effect of the rule regarding p.*

Suppose a controller can leverage five sensitivity levels: 0.00 (*none*), 0.25 (*low*), 0.50 (*medium*), 0.75 (*high*), and 1.00 (*highest*) for the shared visual data, the following is an example rule:

Example 3.1. Alice denies users who are in her "Colleagues" friends list, from viewing her face, body and location in all photos that she is tagged in, where Alice considers her location with a high sensitivity level:

$r_1 = (Alice, \{< Colleagues >\}, *, \{< face : 1 >, < body : 1 >, < location : 0.75 >\}, \{< face : blur >, < body : blur >, < location : blur >\})$.

We apply this rule to the example social network shown in Figure 2. Let us assume Bob is a colleague of Alice. Bob satisfies this rule, since he is in "Colleagues" friends list of Alice. Let us assume Bob comes across a photo, "party.jpg", taken at a popular downtown bar near Times Square, "Times Square Bar" depicting Alice. As an effect of this rule, Alice's face, body and location ("Times Square Bar") would be blurred out, so that they are not visible to Bob.

Furthermore, a PII item stakeholder may define more than one rule in her/his policy for a shared visual data. In this case, users who satisfy *any* rule in the policy are considered as authorized users for the resource. The following is another example rule:

Example 3.2. In addition to the rule defined in Example 3.1, let's consider another authorization requirement from *Alice*, where she wants to disclose all her PII items in all photos to users in her "Close Friends" list:

$r_2 = (Alice, \{< CloseFriends >\}, *, \{*\}, \{* : share\})$.

Example 3.3. There are cases where relational PII items, such as mutual friends, can leak the privacy of a stakeholders, as shown in Example 3.1. Let's consider an authorization requirement from *Alice* where, in addition to Example 3.1, she wants to deny one of her colleagues, *Dave*, to view their mutual friends' faces:

$r_3 = (Alice, Dave, *, \{< face : 1 >, < body : 1 >, < location : 0.75 >, < mutual\_friends\_list(Alice, Dave).face : 0.75 >\}, \{< face : blur >, < body : blur >, < location : blur >, < mutual\_friends\_list(Alice, Dave).face : blur >\})$.

When we apply this rule to the example social network (Figure 2), Dave will not be able to see Bob's face, in addition to Alice's specified PII items.

3.3 PMAC Policy Evaluation

In our PMAC model, we adopt two steps to evaluate a viewer request over multiparty access control policies as shown in Figure 3. In the first step, we first perform some pre-evaluation procedures that involve the detection of controllers and *PII items* in the photo. Then, the privacy policies of the controllers are retrieved, following their detection. In PMAC model, a controller can leverage a positive rule to define a set of viewers to whom the controller's *PII items* are visible, or/and a negative policy to exclude some specific viewers from whom the *PII items* should be blurred. A *PII item owner* has complete ownership of the *unique PII items* that they own. As a result, only the *PII item owner's* policy is used for determining the effect on *unique PII items*. However, in case of *shareable* and *relational PII items*, several co-owners, known as *PII item stakeholders* can have different privacy policies associated with them, based on different privacy needs and concerns.

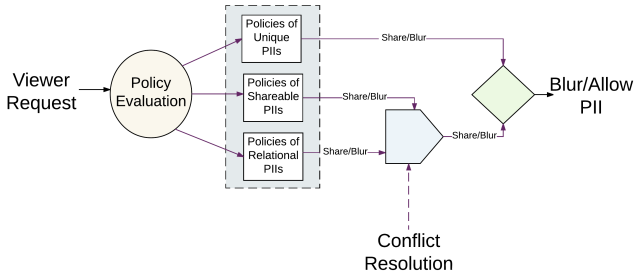

Figure 3: PMAC Policy Evaluation Process.

In the second step, decisions from all controllers corresponding to each *PII item* in the photo pertaining to the viewer request are first aggregated. Since those controllers may generate different decisions (`share` and `blur`) with respect to the *shareable PII items* and *relational PII items* for the viewer request, conflicts may occur. Conflict resolution is used in our PMAC policy evaluation process in this step. We will address our approach for resolving such conflicts in detail subsequently, where we use a strategy called *privacy adjustment*. As can be seen in Figure 3, policies concerning *unique PIIs* are directly used to determine the effect, whereas policies concerning *shareable PIIs* and *Relational PIIs* are first aggregated to perform conflict resolution, followed by the effect decision.

3.3.1 Conflict Resolution in PII-based Multiparty Access Control. Due to collaborative control of photo sharing in PMAC model, we have two main areas of conflicts as follows:

- **Relational PII Conflict.** Mutual friends are users who are common friends of two users in the OSN. There may exit a conflict between a mutual friend's own policy for her/his PII items and a user's policy applied to the mutual friend's PII items. For example, assume Bob and Jack are friends and John is a friend of Jack but not a friend of Bob. Bob wants his friend Jack's face to be blurred to all users who are not his friend. But, Jack as a mutual friend of Bob and John wants all his friends to view his face. Thus, there is a conflict about whether or not to blur Jack's face, when John is the viewer.
- **Shareable PII Conflict.** As shown in Figure 2, there can exist multiple stakeholders for a single PII item. Stakeholder conflict occurs when the privacy concerns of the collaborating stakeholders for the same PII items do not match. For example, assume Jack and Bob both work for the *same* company. While Jack wants to share his affiliation with the company with all his friends, Bob does not want anyone except his close friends to learn about his affiliation with his company. Jack and Bob are photographed together with the company logo in the background. The stakeholder conflict arises for viewers who are in both Jack's and Bob's target viewers list: whether to blur the company logo according to Bob's policy or to show the company logo according to Jack's policy.

The process of privacy conflict resolution makes a decision to allow or deny the viewer access to view *conflicted* PII items. In general, allowing a viewer to view a conflicted PII items may cause *privacy leakage*, but denying a viewer access to the conflicted PII items may result in *sharing loss*. Our privacy conflict resolution approach employs a *privacy adjustment* mechanism that ensures that there is minimum privacy leakage for users who want to blur PII items, but at the same time allow users to share PII items. We employ this mechanism differently to ensure optimum conflict resolution for *Relational PII Conflict* and *Shareable PII Conflict*.

Relational PII Conflict Resolution Through Privacy Adjustment. We consider a multiparty photo containing a user and a mutual friend of the user. Let v be a viewer of the photo who is a friend of the mutual friend but not a friend of the user. Table 2 depicts the conflicting PII items and resolution strategy for the user and the mutual friend. In the last scenario, conflict occurs because the user wishes to blur her/his own face and blur the mutual friend's face from v, but the mutual friend wishes to share her/his face and body with v. In this case, we use *privacy adjustment* to resolve this conflict, defined as follows.

- **Policy Restrictiveness**: The restrictiveness of a policy p defined by a user i is the level of restriction imposed by p on the visibility of PII items of i, denoted by R_i. For example, a user's policy that is set to blur *both* face *and* body has higher *policy restrictivenss* than a user's policy that is set to only blur her/his face. In PMAC model, the restrictiveness of a policy is computed by summing the pre-defined weights (w) of PII items in PMAC model. These weights are defined based on the degree of identification for a user that the PII item provides. For example, a user's face can reveal the identity of the user much more effectively, when compared to the location information of the user. The computation of policy restrictiveness in shown in Equation 1.

$$R_i = \sum_{j \in PIIs_{User}(i)} w_j \tag{1}$$

  Where the function $PIIs_{User}(i)$ returns a set of PII items associated with the user i.
- **Privacy Adjustment**: Automatically adjust the privacy policy of a user in case of a conflict by increasing the *policy restrictiveness* of the user, so that a higher level of privacy is achieved by restricting visibility of additional PII items of the user. For example, let's assume Bob is a user who has set his policy to blur his own face and blur his mutual friend's face, when John is the viewer. Let's say Jack is a mutual friend of Bob and John. Conflict occurs if Jack has set his policy to share his face with all his friends. In this case, *Privacy Adjustment* mechanism increases the *policy restrictiveness* of Bob's policy by blurring both his body and face, but allows Jack to share his face with his friends.

The ability to control individual PII items in PMAC model enables us to use *privacy adjustment* to resolve conflicts. For example, in the third scenario depicted in Table 2, since the mutual friend is the owner of her/his own face and body, we allow the mutual friend to share her/his face and body with v. At the same time, since this decision compromises privacy of the user, we use *privacy adjustment* to automatically blur the user's *body*, in addition to

Table 2: Example of Relational PII Conflict Between Mutual Friend (MF) and User

MF's Own Policy		User's Policy				Conflict?	MF Conflict Resolution	User Conflict Resolution
Face	**Body**	**Own Face**	**Own Body**	**MF Face**	**MF Body**			
Blur	Blur	Blur	Blur	Blur	Blur	No	Blur Face and Body	Blur Face and Body
Share	Share	Blur	Blur	Share	Share	No	Share Face and Body	Blur Face and Body
Share	**Share**	**Blur**	**Share**	**Blur**	**Share**	**Yes**	**Share Face and Body**	**Blur Face and Body**

the face of the user so that there is a minimum compromise in the user's privacy. This ensures an optimum trade off between date sharing and privacy protection.

We use *privacy adjustment* for conflict resolution in case of multiple mutual friends and users in a multiparty photo. Since the PMAC model emphasizes the importance of privacy of individuals in multiparty photo sharing, it does not allow the majority decision to override the privacy concerns of users. Therefore, in case of conflict, PMAC model uses *privacy adjustment* to address the privacy concern of users by making the policy of users more restrictive.

We summarize the conflict resolution decision of mutual friend i and user j as follows.

$$Decision = \begin{cases} \text{Share} + Privacy\ Adjustment & \text{if } R_i < R_j \\ \text{Blur} & \text{if } R_i \geq R_j \end{cases} \quad (2)$$

Where R_i and R_j are the *policy restrictiveness* of the mutual friend and the user, respectively, for at least *minimum* R_j for user, which is to blur face.

Shareable PII Conflict Resolution Through Sharing Risk Measurement and Privacy Adjustment: Our basic premise for conflict resolution in case of stakeholder conflict is the following: a) a PII item conflicting stakeholder policies must be shared if the *majority* of stakeholder policies are in favor of sharing the PII item; and b) PMAC must use the *privacy adjustment* for the stakeholder policies, which are in favor of blurring the PII item. In order to facilitate effective conflict resolution for shareable PII conflicts, we define sensitivity of shared PII item and aggregate decision value as follows:

- **Sensitivity of shared PII item**: PII item sensitivity defines stakeholder's perception about the confidentiality of the PII item being shared. The sensitivity level of the shared PII item defined by a stakeholder j is denoted as sl_j. This factor depends on the stakeholder themselves, since certain PII items are more confidential for some stakeholders than some others; and
- **Aggregate decision value**: A *sharing risk based* scheme is used by PMAC to compute an aggregated decision value, in favor of sharing a PII item and in favor of blurring the PII item. The total number of stakeholders in favor of sharing a PII item k, is denoted by $N_{sh}(k)$ and the total number of stakeholders in favor of blurring the PII item k is denoted by $N_{Bl}(k)$.

In order to measure the aggregate value in favor of sharing and in favor of blurring a PII item k, denoted by $AV^{Sh}_{total}(k)$ and $AV^{Bl}_{total}(k)$, we can use following equations.

$$AV^{Sh}_{total}(k) = \sum_{i \in stakeholders_{Sh}(k)} sl_i \times N_{sh}(k) \quad (3)$$

and

$$AV^{Bl}_{total}(k) = \sum_{j \in stakeholders_{Bl}(k)} sl_j \times N_{bl}(k) \quad (4)$$

Where functions $stakeholders_{Sh}(k)$ and $stakeholders_{Bl}(k)$ return the number of stakeholders of k who wish to share and blur k, respectively.

Then, following equation can be utilized to make the decisions (sharing or blurring a PII item for a viewer request) for the stakeholder conflict resolution.

$$Decision = \begin{cases} \text{Share} + Privacy\ Adjustment & \text{if } AV^{Bl}_{total}(k) < AV^{Sh}_{total}(k) \\ \text{Blur} & \text{if } AV^{Bl}_{total}(k) > AV^{Sh}_{total}(k) \end{cases} \quad (5)$$

4 IMPLEMENTATION AND EVALUATION

4.1 System Implementation

We implemented a proof-of-concept Facebook application called *AppX*. *AppX* is a third-party Facebook application written in PHP and MySQL and hosted on Apache servers. The user interface also contains jQuery. External APIs were used extensively. Using Facebook's Graph API, the users' Facebook profile can be accessed to authenticate the user and import the friends list to the database. *AppX* is a social media application that implements the PMAC mechanism. The current implementation is restricted to photo sharing, but the system can be generalized to other forms of visual media sharing. We use Face++ for face recognition and Google Vision for logo and text detection and Microsoft Computer Vision for description tags. For object and body recognition, py-faster-rcnn was used with the VGG16 trainval on the Palmetto server, a high performance computing cluster with 12 GBs of RAM, 8 CPU cores, and a K40 GPU. A RESTful API was created for py-faster-rcnn using the Flask library in Python.

When first accessed, the user is asked to grant *AppX* Facebook permissions to view basic profile information (see Figure 4). After the user is authenticated, *AppX* will then download the user's profile picture and import the user's Facebook friends into the database. The profile picture is used for face detection. If *AppX* cannot detect a single face in the picture, it will ask the user to upload a portrait. After this, *AppX* presents the user the settings page. The settings page allows the user to access all of the PMAC policies. After initialization, the user is presented the picture feed, where they will see all of the picture uploaded to *AppX* by the user and the user's friends, similar to Facebook's home page. The user can change the settings or upload a picture.

When a picture is uploaded, it is saved locally on the server and processed on the fly when they are viewed. When the photo is viewed, it is first processed and edited in accordance to the PMAC policy before being shown to the user. Each face in the picture is recognized and their PMAC policies are loaded from the database

and applied. Any face, body, or item in the photo may be blurred if permission is not allowed.

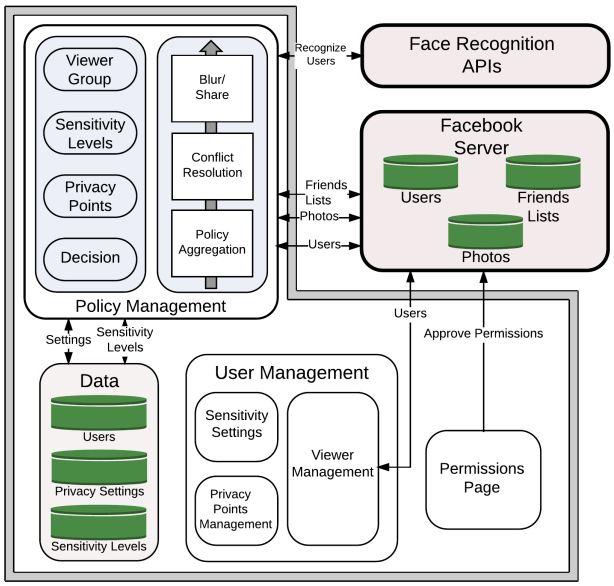

Figure 4: System Architecture of *AppX*.

The principal components of *AppX* are depicted in Figure 4. Since the sensitivity of PII items is specific to every user's privacy concerns, the *User Management* component provides options to adjust it. We have various PII items, such as location, belongings and affiliation. A user can add or remove the PII items that they are concerned about using the *User Management* component. In addition, due to the evolving nature of social media, *AppX* provides support for adding more PII items according to the PMAC model policy updates. Updates in the *policy settings* reflect in the *Data* component. As a result, *AppX* updates policy levels in the *Policy Management* component.

4.2 System Evaluation

4.2.1 Performance Evaluation. A single photo was used as a background image, and 1 to 10 people were added to the picture by substituting arguments to the API. Each person is box 72x153 pixes, and is 20 pixels apart from each other in a line. The processing and database time for the PMAC system were recorded. As more people are detected in the picture, more PMAC policies have to be retrieved and processed. Each person has a PMAC policy with only body blurring enabled. Each person's PMAC policy is then enforced. To prevent the test runs from interfering with each other, caching was turned off. The database calls increase linearly as the number of people increase because the database calls were made separately in sequential order for each person. Figure 5 depicts the performance of *AppX* as a plot of number of people in a photo against average time in milliseconds.

As depicted in figure 5, there is a slight and proportional increase in the processing time and the database time as the number of PII items in the photo increases. This is an expected observation, as there are more computations involved as the PII items in a photo

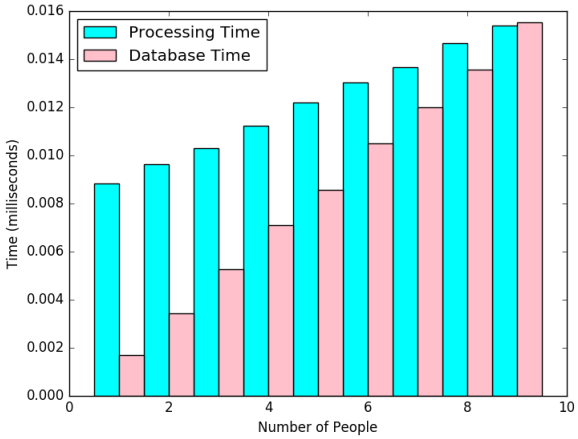

Figure 5: Performance Evaluation of *AppX*.

increases. For the maximum case, we only observed a difference of around 0.006 milliseconds. Hence, PMAC could only add little overhead to the current Facebook photo control mechanism.

4.2.2 Effectiveness Evaluation. To evaluate the effectiveness of our approach, we compared the privacy of a shared photo from the perspective of PII privacy for Facebook solution, Face/Off solution [12] and our PMAC solution for 30 randomly selected photos from Facebook. The metric we used for evaluation is the total Privacy of all controllers and all PII items in a photo, based on the assumption that a viewer is not authorized to see PII items of any of the controllers present in a photo.

The Facebook solution does not allow users to control viewer access to PII items. The Face/Off solution allows users to control access to their face only. Our model allows users to control access to their *unique*, *shareable* and *relational* PII items. In order to evaluate the impact of each type of PII in a photo, we used certain factors derived from [21] that are relevant to the context of online photo sharing. These factors are listed below:

- **Identifiability**: $Identifiability$ of a PII item is a measure of how easily the PII item can be used to identify a user.
- **Quantity of PII**: $Quantity\ of$ PII is the total number of PII items in a photo.
- **Data Field Sensitivity**: Some PII items, such as face of a user, are more sensitive than other PII items like location. The $data\ field\ sensitivity$ of a PII item is a measure of the sensitivity of the PII item.
- **Context of Use**: $Context\ of\ use$ reflects the purpose for which a PII item can be used for. For example, the location of a user can be used to learn about the whereabouts of the user.

In our evaluation, we allocated $identifiability$, $quantity\ of\ PII$, $data\ field\ sensitivity$ and $context\ of\ use$ scores for each PII item that we have considered in our model, in order to capture the impact of the privacy of these PII items for the 30 randomly selected Facebook photos. We then used the computed scores of each type of PII items to evaluate the effectiveness of all 30 cases of our randomly selected Facebook photos. In our effectiveness evaluation, a user's

unique PII items were allocated with the highest score, as they are most crucial to the user and uniquely identify a user. This is followed by the shareable PII items and relational PII items.

We evaluated the outcomes of 30 cases for Facebook, Face/Off and our PMAC solutions, as depicted in Figure 6. The Facebook privacy policy is owner centric. That is, the owner decides the privacy of the complete photo and hence there is no collaborative control. In addition, the Facebook solution is not *fine-grained* and hence, users cannot control their PII items. Since our evaluation is PII privacy centric, the Facebook solution does not address any of the above mentioned factors. Hence, in all 30 cases Facebook solution is evaluated as the lowest privacy solution for PII items among all three solutions.

The Face/Off solution enables multiparty access control and users can control who can see their faces in a shared photo. From a PII privacy perspective, a user's privacy can be compromised by several other PII items. For example, a popular celebrity having a unique tattoo on her/his body can be easily identified by the tattoo on her/his body, even though her/his face is blurred. In addition to this, shareable PII items can be responsible for privacy compromise of more than one user in a shared photo, with the additional overhead of conflict resolution in a collaborative environment. We can also see from the example in Figure 2 that relational PII items also play an important role in privacy compromise is shared photos. Therefore, the Face/Off solution has been scored for only protecting the privacy of the face of a user in the 30 cases. The results are depicted in Figure 6. Obviously, the Face/Off solution is better than the Facebook solution in every case.

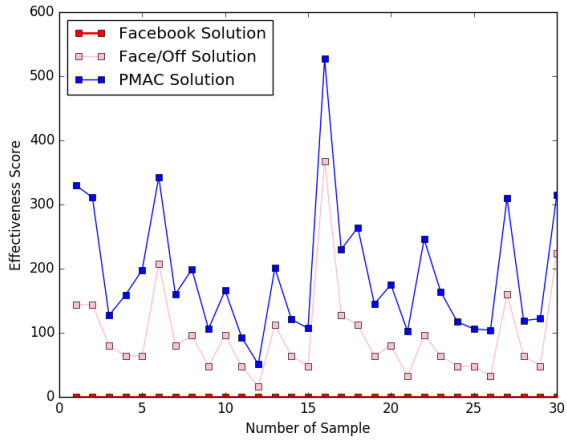

Figure 6: Effectiveness Evaluation of Facebook, Face/Off and PMAC Solutions.

As depicted in Figure 6, our solution performs the highest among all three solutions. This is because, from a PII privacy perspective, our solution provides the highest control to users over who can have viewer access to individual PII items. This observation is confirmed in all 30 cases of our experiment.

However, it can be observed that in some cases, such as the cases 3 and 12 the Face/Off solution comes close to our solution in terms of PII privacy protection, but in some cases such as the cases 1 and 27, there is a comparatively larger difference in effectiveness. This is due to the fact that the cases 3 and 12 had very few observed PII items, apart from face. Therefore, since both Face/Off and our solution are effectiveness in protection of face PII, the effectiveness is close. However, in the cases 1 and 27, a large number of PII items in addition to face are present. Since the Face/Off solution provides no control for these PII items, there is a large difference in the effectiveness between the Face/Off solution and our PMAC solution.

4.2.3 Preliminary Study of Privacy Filter. To determine the potential of our system, we conducted a preliminary survey for identifying the likability and adoption willingness of the privacy filter, *blurring*, used in our system. We conducted an online survey with 30 participants, and measured the users' likability towards as is (no filter) and body blurring filter with the question "I like the privacy filter" which derived the interface preference scale [23]. We also measured their general adoption willingness using the question "I want online social networks (Facebook etc.) to adopt the privacy filters so that I can be obscured in certain photos my friends upload (group photo etc.)". Both response scales are 7-point likert scale from 1 "Strongly disagree" to 7 "Strongly agree". Thirty participants were recruited from MTurk. Nineteen of them are female, ten are male, and one participant selected "I prefer not to answer".

We created two identical group photos with four people as foreground and a campus scene as background. In one photo, we applied body blurring filter on one target person. First, the participants saw one photo, and rated their likeability from 1 to 7; then they were shown another photo and rated the likeability. Afterwards, they rated their adoption willingness of this type of privacy filters.

The result shows the mean of likability of as is condition is 5.33 (Somewhat agree); and for body blurring filter, it is 4.1 (neither disagree nor agree). Seventy seven percent of the participants like the original photo without any filter (as is condition), while 53% like the body blurring filter. We expected this result as people may be preferential towards original photo as compared to the photo with body blurring privacy filter. However, the neutral mean of the body blurring filter also indicates that our privacy filter does not seriously degrade user experience.

The mean of general adoption willingness is 4.7 (somewhat agree). Sixty percent of the participants have a positive attitude (rating from 5-7), suggesting the majority of the participants may be willing to use the blurring filter.

4.2.4 Survey (User Study) of PMAC Model. We conducted a survey ("user study") to evaluate users' naive perceptions about the policy settings as implemented in a Facebook application (*AppX*). First, we presented users with an *AppX* policy setting user interface. Next, we presented five scenarios focused on privacy of each of the *PII items* supported by our model. All scenarios had sample photos so that participants can understand the privacy concerns in the scenarios. Figure 7 illustrates five scenarios in *AppX*. Next, we asked them questions about their perceptions of the app and the settings it enables. We used two criteria for evaluation: *adoption willingness* and *control*. *Adoption willingness* is a measure of a user's willingness to adopt a particular feature in current OSNs. It can also help us identify if users perceive a feature as useful. *Control* is a measure of the user's perceived control of their private information.

Table 3: Adoption Willingness and Perceived Privacy Control for PMAC Model. Response scales are 7-point likert scale from 1 "Strongly disagree" to 7 "Strongly agree".

Metric	PII items														
	Body			Mutual Friend			Affiliation			Belonging			Location		
	Mean	SEM	% PR	Mean	SEM	% PR	Mean	SEM	% PR	Mean	SEM	% PR	Mean	SEM	% PR
Adoption Willingness	4.32	0.16	90.32	3.87	0.20	70.97	4.17	0.19	90.00	3.96	0.18	78.57	4.14	0.17	85.72
Control	4.52	0.14	90.32	4.23	0.16	83.87	4.14	0.17	82.76	4.00	0.18	78.58	4.32	0.16	85.71

We measured users' *adoption willingness* of *AppX* by asking "If Facebook implemented this feature in its privacy policy, I would use it". We measured *control* by asking "I believe that I can better control the visibility of my *PII* to only users that I want to share it with in *AppX* when compared to Facebook", where *PII* was each of body, mutual friend, affiliation, belonging and location. We asked these two questions for each of the five *PII items*. Both response scales are 5-point likert scale from 1 "Strongly disagree" to 5 "Strongly agree".

We conducted the online survey using the Qualtrics platform. There were a total of thirty nine participants, with 40.74% of participants in the age group of 18-24 years, 48.15% in the age group of 25-34 years and 11.11% in the age group of 35-44 years. The participants consisted of students as well as working professionals. Among the participants, a majority (96.30%) claimed to use Facebook and 85.19% of the participants claimed to use Instagram and Google+.

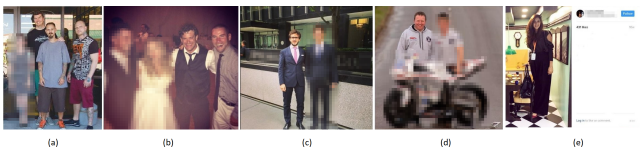

Figure 7: Illustration of PII Scenarios in *AppX*: (a) Body, (b) Mutual Friends, (c) Affiliation, (d) Object, and (e) Location.

Table 3 depicts the results from the survey, where we record mean, standard error mean (SEM) and percentage positive response (PR). The means for all the scenarios for *adoption willingness* are above 3, which indicates that naive users, without having used the system, may be willing to consider using *AppX*. The means for *control* are above three for all scenarios, indicating that users may perceive that *AppX* provides them some control over the five *PIIs* in photos. Participants rated their potential willingness to adopt *body* privacy the highest as compared to the other *PII items*. One possible reason is that participants may consider their body as the most sensitive *PII item* in the photos shared in OSNs. However, it could be that they simply find body blurring the least offensive in terms of visual affect to the photos. A similar speculation can be applied around the degree of perceived control of *body* in a multiparty photo.

We see that the lowest mean of *adoption willingness* among all the *PII items* is *mutual friends*, although it is above three (indicating a positive perception). We speculate that the comparatively low score for *mutual friends* could be attributed to the lack of familiarity of the people in the photo. The mutual friends PII is based on relationship between the viewer and the people in the photo. Since participants were not related to the people in the photo, we could argue that they might not recognize or appreciate the importance of mutual friends privacy presented in this scenario.

5 RELATED WORK

The need of collaborative management for data sharing, especially photo sharing, in OSNs has been addressed by some recent research [5, 9–11, 15, 29, 31]. For example, Hu et al. [10] formulated a MultiParty Access Control (MPAC) model to capture the essence of multiparty authorization requirements. They also investigated a collaborative data sharing mechanism to support the specification and enforcement of multiple privacy concerns, along with a conflict detection and resolution mechanism. In addition, they proposed an approach with the support of both theoretical and empirical analyses on privacy control in OSNs through analyzing the strategic behaviors of rational users using a game-theoretic model [11]. However, all those work can only facilitate collaborative control for an *entire* photo and lacks the support for the control of *individual PII items* in a shared photo.

The importance of PII items for privacy control has been discussed in several prior work [13, 14, 26]. For example, the ReCon system [26] discusses PII privacy in mobile networks and provides machine learning solutions to reveal leakage of PII items and also provides tools to control such a kind of privacy leakage. Recently, Face/Off solution [12] was proposed to model and express access control policies to control the view of users' faces in shared photos in OSNs. Face/Off adopts a face blurring technique to hide a user's identity in multiparty photos. In addition, extensive user studies conducted in [12] provide us with valuable inferences about user opinions regarding more fine-grained privacy control for user faces in shared photos in OSNs. However, Face/Off solution cannot support the specification of fine-grained access control policies for other important PII items, such as body, affiliation, and belonging in shared photos in OSNs. Besides, an access control scheme for the control of individual parts of a shared object in OSNs is provided by the CooPeD system [8]. The CooPeD system presents a model for the co-management of decomposable parts in shared objects in OSNs. However, the CooPeD model does not provide a solution to address the issue of the same shared parts belonging to multiple users. In ideal scenarios, there may exist many shared parts in an image that could be co-owned by multiple users. In contrast, our PMAC model can address such an issue and also use its conflict resolution and privacy adjustment mechanisms to effectively resolve PII-based privacy conflicts.

A solution for the video privacy protection in OSNs was provided by the BEPS system [24]. BEPS can detect and separate *intentionally captured persons* (ICP) from *non-intentionally captured persons* (non-ICPs), following which the non-ICPs are removed from the image using in-painting techniques. However, this work fails to recognize

the importance of PIIs with respect to the privacy and identification of subjects in shared videos. Although a non-ICP subject could be removed from a video, the non-ICP subject's PIIs can still cause privacy compromise of the subject. In our approach, in addition to blocking access to directly identifiable parts such as face and body of a subject, we further offer an access control mechanism that allows users to control visibility of their other PIIs.

6 CONCLUSION AND FUTURE WORK

In this paper, we have proposed a new mechanism for collaboratively controlling PII items in multiparty photos in OSNs. A PII-based multiparty access control model has been formulated, along with a policy specification and corresponding policy evaluation mechanism. In addition, our conflict resolution strategy leverages the flexible control of individual PIIs for effective conflict resolution. We have also described a proof-of-concept implementation of our solution called *AppX*, and provided system evaluation and usability study of our approach.

As part of our future work, we plan to conduct more comprehensive user studies to evaluate the user needs with respect to PII privacy in multiparty photo sharing in OSNs. In addition, we would extend our work in multiparty policy specifications to use machine learning techniques so that intuitive policy models that need minimal user interaction can be formulated.

ACKNOWLEDGMENT

This work was partially supported by grants from National Science Foundation (NSF-IIS-1527421, NSF-IIS-1527268, and NSF-CNS-1537924).

REFERENCES

[1] 2011. The State of Social Media 2011: Social is the new normal. (2011). http://www.briansolis.com/2011/10/state-of-social-media-2011/.
[2] 2017. Facebook Privacy Policy. (2017). http://www.facebook.com/policy.php/.
[3] 2017. Facebook Sharing Settings. (2017). www.facebook.com/help/459934584025324/.
[4] F. Adu-Oppong, C.K. Gardiner, A. Kapadia, and P.P. Tsang. 2008. Social circles: Tackling privacy in social networks. In *Symposium on Usable Privacy and Security (SOUPS)*. Citeseer.
[5] A. Besmer and H. Richter Lipford. 2010. Moving beyond untagging: Photo privacy in a tagged world. In *Proceedings of the 28th international conference on Human factors in computing systems*. ACM, 1563–1572.
[6] J.Y. Choi, W. De Neve, K.N. Plataniotis, Y.M. Ro, S. Lee, H. Sohn, H. Yoo, WD Neve, CS Kim, YM Ro, and others. 2010. Collaborative Face Recognition for Improved Face Annotation in Personal Photo Collections Shared on Online Social Networks. *IEEE Transactions on Multimedia* (2010), 1–14.
[7] L. Fang and K. LeFevre. 2010. Privacy wizards for social networking sites. In *Proceedings of the 19th international conference on World wide web*. ACM, 351–360.
[8] Lorena González-Manzano, Ana I González-Tablas, José M de Fuentes, and Arturo Ribagorda. 2014. Cooped: Co-owned personal data management. *Computers & Security* 47 (2014), 41–65.
[9] Hongxin Hu, Gail-Joon Ahn, and Jan Jorgensen. 2011. Detecting and Resolving Privacy Conflicts for Collaborative Data Sharing in Online Social Networks. In *Proceedings of the 27th Annual Computer Security Applications Conference (ACSAC'11)*. ACM.
[10] Hongxin Hu, Gail-Joon Ahn, and Jan Jorgensen. 2013. Multiparty access control for online social networks: model and mechanisms. *IEEE Transactions on Knowledge and Data Engineering* 25, 7 (2013), 1614–1627.
[11] Hongxin Hu, Gail-Joon Ahn, Ziming Zhao, and Dejun Yang. 2014. Game theoretic analysis of multiparty access control in online social networks. In *Proceedings of the 19th ACM symposium on Access control models and technologies*. ACM, 93–102.
[12] Panagiotis Ilia, Iasonas Polakis, Elias Athanasopoulos, Federico Maggi, and Sotiris Ioannidis. 2015. Face/off: Preventing privacy leakage from photos in social networks. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 781–792.
[13] B. Krishnamurthy and C.E. Wills. 2010. On the leakage of personally identifiable information via online social networks. *ACM SIGCOMM Computer Communication Review* 40, 1 (2010), 112–117.
[14] Balachander Krishnamurthy and Craig E Wills. 2009. On the leakage of personally identifiable information via online social networks. In *Proceedings of the 2nd ACM workshop on Online social networks*. ACM, 7–12.
[15] A. Lampinen, V. Lehtinen, A. Lehmuskallio, and S. Tamminen. 2011. We're in it together: interpersonal management of disclosure in social network services. In *Proceedings of the 2011 annual conference on Human factors in computing systems*. ACM, 3217–3226.
[16] Yair Levy and Michelle M Ramim. 2016. Towards an Evaluation of Cyber Risks and Identity Information Sharing Practices in e-Learning, Social Networking, and Mobile Texting Apps. (2016).
[17] H.R. Lipford, A. Besmer, and J. Watson. 2008. Understanding privacy settings in facebook with an audience view. In *Proceedings of the 1st Conference on Usability, Psychology, and Security*. USENIX Association Berkeley, CA, USA, 1–8.
[18] Eden Litt and Eszter Hargittai. 2014. Smile, snap, and share? A nuanced approach to privacy and online photo-sharing. *Poetics* 42 (2014), 1–21.
[19] Y. Liu, K.P. Gummadi, B. Krishnamurthy, and A. Mislove. 2011. Analyzing Facebook Privacy Settings: User Expectations vs. Reality. In *Proceedings of the 2011 annual conference on Internet measurement (IMC'11)*. ACM.
[20] M. Madejski, M. Johnson, and S.M. Bellovin. 2011. The Failure of Online Social Network Privacy Settings. Technical Report CUCS-010-11, Columbia University, NY, USA. (2011).
[21] Erika McCallister, Timothy Grance, and Karen A Scarfone. 2010. Sp 800-122. guide to protecting the confidentiality of personally identifiable information (pii). (2010).
[22] Mainack Mondal, Yabing Liu, Bimal Viswanath, Krishna P Gummadi, and Alan Mislove. 2014. Understanding and specifying social access control lists. In *Symposium on Usable Privacy and Security (SOUPS)*. 11.
[23] Kyle B Murray and Gerald Häubl. 2010. Freedom of choice, ease of use, and the formation of interface preferences. (2010).
[24] Yuta Nakashima, Noboru Babaguchi, and FAN Jianping. 2016. Privacy Protection for Social Video via Background Estimation and CRF-Based Videographer's Intention Modeling. *IEICE Transactions on Information and Systems* 99, 4 (2016), 1221–1233.
[25] F.K. Ozenc and S.D. Farnham. 2011. Life "Modes" in Social Media. In *Proceedings of the 2011 annual conference on Human factors in computing systems*. ACM, 561–570.
[26] Jingjing Ren, Ashwin Rao, Martina Lindorfer, Arnaud Legout, and David Choffnes. 2016. Recon: Revealing and controlling pii leaks in mobile network traffic. In *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services*. ACM, 361–374.
[27] Cooper Smith. 2013. Facebook users are uploading 350 million new photos each day. *Business insider* 18 (2013).
[28] Craig Smith. 2016. By the Numbers: 200+ Amazing Facebook Statistics. (2016).
[29] A.C. Squicciarini, M. Shehab, and F. Paci. 2009. Collective privacy management in social networks. In *Proceedings of the 18th international conference on World wide web*. ACM, 521–530.
[30] K. Strater and H.R. Lipford. 2008. Strategies and struggles with privacy in an online social networking community. In *Proceedings of the 22nd British HCI Group Annual Conference on People and Computers: Culture, Creativity, Interaction-Volume 1*. British Computer Society, 111–119.
[31] K. Thomas, C. Grier, and D. Nicol. 2010. unFriendly: Multi-party Privacy Risks in Social Networks. In *Privacy Enhancing Technologies*. Springer, 236–252.
[32] Y. Wang, S. Komanduri, P. Leon, G. Norcie, A. Acquisti, and L. Cranor. 2011. I regretted the minute I pressed share": A qualitative study of regrets on Facebook. In *Symposium on Usable Privacy and Security*.
[33] R. Wishart, D. Corapi, S. Marinovic, and M. Sloman. 2010. Collaborative Privacy Policy Authoring in a Social Networking Context. In *2010 IEEE International Symposium on Policies for Distributed Systems and Networks*. IEEE, 1–8.
[34] Li Yifang, Vishwamitra Nishant, Knijnenburg Bart, Hu Hongxin, and Caine Kelly. (2017). Blur vs. Block: Investigating the Effectiveness of Privacy-Enhancing Obfuscation for Images. In *The First International Workshop on The Bright and Dark Sides of Computer Vision: Challenges and Opportunities for Privacy and Security (CV-COPS 2017)*.
[35] Li Yifang, Vishwamitra Nishant, Hu Hongxin, Knijnenburg Bart, and Caine Kelly. 2017. Effectiveness and Users' Experience of Face Blurring as a Privacy Protection for Sharing Photos via Online Social Networks. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Vol. 61. SAGE Publications.
[36] Y. Zhu, Z. Hu, H. Wang, H. Hu, and G.J. Ahn. 2010. A Collaborative Framework for Privacy Protection in Online Social Networks. In *Proceedings of the 6th International Conference on Collaborative Computing (CollaborateCom)*.