Adaptive Conflict Resolution Mechanism for Multi-party Privacy Management in Social Media

Jose M. Such
School of Computing and Communications
Lancaster University
Lancaster, UK
j.such@lancaster.ac.uk

Natalia Criado
School of Computing & Mathematical Sciences
Liverpool John Moores University
Liverpool, UK
n.criado@ljmu.ac.uk

ABSTRACT

The lack of multi-party privacy management support in current mainstream Social Media infrastructures makes users unable to appropriately control to whom co-owned items are shared. Computational mechanisms that are able to merge the privacy preferences of multiple users into a single policy for these kind of items can help solve this problem. As privacy preferences may conflict, these mechanisms need to consider how users' would actually reach an agreement in order to propose acceptable solutions to the conflicts. We propose the first computational mechanism to resolve conflicts for multi-party privacy management in Social Media that adapts to different situations that may motivate different users' concessions and agreements.

Categories and Subject Descriptors

K.4 [**Computers and Society**]: Privacy

Keywords

Online Social Networks; Multi-party Privacy; Co-owned items

1. INTRODUCTION

There are massive amounts of items shared through Social Media that involve multiple users [7], e.g., photos that depict multiple people, comments that mention multiple users, events in which multiple users are invited, etc. The problem is that users involved in one particular item (e.g., users depicted in a photo) may have different privacy preferences for that item. For instance, Alice and Bob are depicted together in a photo in which Bob appears drunk, and Alice would like to share the photo with her friend Charlie, but Bob would not like to share the photo with Charlie because Bob feels embarrassed about his looks in the photo and Charlie is just a distant acquaintance of him.

Computational mechanisms to aid the negotiation process in these cases are one of the biggest gaps in privacy management in social media [4, 9]. The use of these mechanisms does not mean that users would lose control in any way. These mechanisms would only *suggest* a possible solution to the conflicts that users will need to accept to be finally applied, so that if users do not accept the suggestion they will need to enter into a manual negotiation by other means. Thus, the main challenge is to suggest solutions that are accepted most of the times by all the users involved, so that users are forced to negotiate manually as little as possible.

Mechanisms to resolve conflicts have already been proposed [6, 3, 7]. Some like [6] still need a great deal of human intervention during the negotiation process. Other mechanisms such as [7] are more automated, but only consider one fixed way of aggregating user's privacy preferences without considering how users would actually achieve compromise and the *concessions* they might be willing to make to achieve it depending on the specific situation. One mechanism considers more than one way of aggregating users' privacy preferences [3], but the user that uploads the item chooses the aggregation method to apply, which becomes a unilateral decision on a multi-party setting. All of this causes these mechanisms to have difficulties to adapt to different situations that may motivate different users' concessions.

In this paper, we present the first *adaptive* computational mechanism for social media that, given the individual privacy preferences of each user involved in a shared item, is able to find and resolve conflicts by applying a different aggregation method based on concessions users' may be willing to make in different situations.

2. BACKGROUND

We assume a set of users $U = N \cup T$, where a set of *negotiating* users $N = \{a_1, \ldots, a_n\}$ negotiate whether they grant a finite set of *target* users $T = \{i_1, \ldots, i_m\}$ access to a particular co-owned item. For simplicity and without loss of generality, we will consider only a negotiation for one item over the course of this paper — e.g., a photo that depicts the negotiating users together — and hence, we do not include any additional notation for the item in question.

Negotiating users can specify their individual privacy preferences about the item using *any* of the access control models already proposed for Social Media. In this paper, we used relationship-based access control [1]

Relationship-based privacy policies usually consider different relationship types $R = \{r_1, \ldots, r_l\}$ — e.g., family, friends, colleagues, etc. These policies usually define a mapping $r : U \times U \rightarrow R$ so that $r(a, b)$ is the relationship type between users a and b. Privacy policies in most relationship-based access control approaches [1] also consider the strength of the relationships that a user has to other users. Intimacy (also called relationship or tie strength) has already been defined in previous works, and there are tools to obtain the

WPES '14 November 03 2014, Scottsdale, AZ, USA.
ACM 978-1-4503-3148-7/14/11 ...$15.00.
http://dx.doi.org/10.1145/2665943.2665964.

intimacy values for all the user's friends for particular Social Media infrastructures such as Facebook [2] with minimal user intervention, e.g., few refinements to the output of the tool mainly to consider that not all users' interactions happen through Social Media. Even if these tools are not used, users could be asked to self-report their intimacies to their friends, which would obviously mean more burden on the users but would still be possible. In this paper, we only use the final intimacy value assigned to each friend whatever the method is used to do so as follows:

DEFINITION 1. *Given two users $a, b \in U$, and a maximum positive integer intimacy value $\mathcal{Y} \in \mathbb{N}$, the* intimacy *between a and b is given as $int(a, b)$, where $int : U \times U \rightarrow \{0, \ldots, \mathcal{Y}\}$.*

Based on this, the following definition for relationship-based privacy policies is used:

DEFINITION 2. *A privacy policy P is a tuple so that $P = \langle \theta_1, \ldots, \theta_{|R|}, E \rangle$, where $\theta_j \in \{0, \ldots, \mathcal{Y}\}$ is the intimacy threshold for the relationship type $r_j \in R$, and $E \subseteq U$ is the set of exceptions to the policy.*

P_a denotes the individual privacy policy of negotiating user $a \in N$ for the item. The role of exceptions E in Definition 2 is to allow an individualised treatment of friends if need be. Finally, we denote $\mathcal{P} = \{P_{a_1}, \ldots, P_{a_n}\}$ as the set containing the individual privacy policy of every negotiating user in N for the item.

The problem we tackle is the following: *how the set of negotiating users N who co-own an item and have individual (possibly conflicting) privacy policies $\mathcal{P}$ for the item can agree on to whom, from the set of the target users T, the item is shared?* This problem can be decomposed into:

1. Given the set of the individual privacy policies $\mathcal{P}$ for the item, how can we identify if policies from $\mathcal{P}$ have contradictory decisions — or *conflicts* — about whether or not granting target users T access to the item.
2. If conflicts are detected, how can we propose a solution to the conflicts found that respects as much as possible the preferences of negotiating users N.

We propose the use of a mediator that helps negotiating users detect conflicts and find an agreement. In the following sections, we show how this mediator can detect conflicts and propose a solution for each conflict found.

3. DETECTING CONFLICTS

To compare privacy policies from different negotiating users for the item, we consider the *effects* that each particular privacy policy has on the set of target users T. Privacy policies dictate a particular action to be performed when a user in T tries to access the item. In particular, we assume that the available actions are either 0 (denying access) or 1 (granting access). The action to perform according to a given privacy policy is determined as follows:

DEFINITION 3. *Given an user $a \in N$, her privacy policy $P_a = \langle \theta_1, \ldots, \theta_{|R|}, E \rangle$, and a user $i \in T$, we define the function:*

$$act(P_a, i) = \begin{cases} 1 & \textit{iff } int(a,i) \geq \theta_{r(a,i)} \wedge i \notin E \\ 1 & \textit{iff } int(a,i) < \theta_{r(a,i)} \wedge i \in E \\ 0 & \textit{otherwise} \end{cases}$$

We also consider so-called *action vectors* $\vec{v} \in \{0, 1\}^m$, i.e., complete assignments of actions to all users in T, such that $v[i]$ denotes the action for user $i \in T$. When a privacy policy is applied to the set of users T, it produces such an action vector:

DEFINITION 4. *The* action vector *induced by privacy policy P in T is $\vec{v} = (v_1, \ldots, v_m)$, where $m =\mid T \mid$ and $v[i] = act(P, i)$.*

Given the action vectors induced by the privacy policies of the negotiating users, we are now in a position to be able to detect whether there are any conflicting actions suggested for the same target user. That is, if all the action vectors assign the same action for all target users, then there is no conflict. Otherwise, there are at least two action vectors that assign different actions to the same target user, and there is a conflict. In other words, a conflict arises when some negotiating users would like to grant access to one target user while the others would not. Formally:

DEFINITION 5. *Given a set of negotiating users N, their privacy policies $\mathcal{P}$ for the item, and a set of target users T, $i \in T$ is said to be* in conflict *if the privacy policies of at least two users $a, b \in N$ are P_a and P_b, so that $v_a[i] \neq v_b[i]$.*

Further, we say that the set of *users in conflict* is the set $\mathcal{C} = \{i \in T \mid \exists a, b \in N, v_a[i] \neq v_b[i]\}$. If the mediator does not detect any conflict — i.e., $\mathcal{C} = \emptyset$, it will return to the users without changes to their preferred privacy policies. If it detects conflicts, it will then run the conflict resolution module, which is described in the following section.

	i_1	i_2	i_3	i_4
a	9	6	4	1
b	8	6	9	4

Table 1: Intimacies for Example 1.

EXAMPLE 1. *Assume a set of users $U = \{a, b, i_1, i_2, i_3, i_4\}$ and only one relationship type among them $R = \{r_1\}$ for the sake of simplicity. Negotiating users $N = \{a, b\}$ are in the process of deciding to which target users $T = \{i_1, i_2, i_3, i_4\}$ they grant access to a photo in which both of them are depicted, and the intimacy values of users a and b toward others are as shown in Table 1, with $\mathcal{Y} = 10$. Suppose now that user a would prefer the policy $P_a =< 5, \emptyset >$, so that $\vec{v}_a = (1, 1, 0, 0)$ — i.e., user a wants to grant access to the photo to users i_1 and i_2, towards whom user a has an intimacy greater or equal to 5, but not to users i_3 and i_4 who are less intimate to user a than that. However, user b would prefer the policy $P_b =< 4, \emptyset >$, so that $\vec{v}_b = (1, 1, 1, 1)$ — i.e., user b wants to grant access to users i_1, i_2, i_3, and i_4. As $v_a[i_3] \neq v_b[i_3]$ and $v_a[i_4] \neq v_b[i_4]$, the set of users in conflict is $\mathcal{C} = \{i_3, i_4\}$.*

4. RESOLVING CONFLICTS

After conflicts are detected, the mediator runs an automated conflict resolution mechanism and proposes a solution. We shall firstly explain the method by which the mediator estimates the willingness of each user to accept applying each possible action for the conflicts detected. Secondly, we shall explain how the mediator estimates whether each negotiating user would concede or not based on this willingness and on the implications for the other negotiating users. Finally, we shall describe how the mediator solves the conflicts based on the concessions users would do.

4.1 Willingness to change an action

We need a measure of how *willing* a user would be to change the action most preferred by her/him to find a solution to the conflict that can be acceptable by all negotiating users. We call this measure *willingness* and it is based on:

1) Sensitivity of the item to be shared. If a user feels that an item is very sensitive for her/him, she/he will be less willing to accept sharing it than if the item is not sensitive for her/him. In our proposal in this paper, users do not need to specify how sensitive an item is for them as we estimate this by considering the intimacy thresholds assigned to the relationship types in the privacy policy for the item. In particular, the higher the intimacy thresholds the higher the sensitivity of the item. For instance, suppose that Alice would only like to share photos about partying with her close friends, but she would not mind sharing photos about her travels around the world with her distant friends, colleagues, and family. Thus, partying photos would be more sensitive for Alice than travelling photos.

2) Intimacy distance to the target user. Intimacy (or relationship strength) is also a factor that plays a role when sharing items [8]. We account for how willing a negotiating user would be to accept an action for a target user based on the difference between two intimacy measures: (i) the intimacy between the negotiating user and the conflicting target user—i.e., $int(a, i)$; and (ii) the intimacy threshold set by the negotiating user in her/his privacy policy for the relationship type of the conflicting target user—i.e., $P.\theta_{r(a,i)}$. The higher this distance, the less willing the user may be to accept the action for the target user — being it granting or denying. For instance, if Alice would only like to share an item with close friends, she would probably be unwilling to concede and share the item with a distant acquaintance. However, if Alice would like to share with friends but not with Charlie, who is close friend of her, this means that she would probably be unwilling to concede and share the item with Charlie (e.g., Alice could be creating an event in which she invites all her friends except Charlie because the event is a surprise party for Charlie's birthday).

We formalise the willingness function as follows:

DEFINITION 6. *Given user $a \in N$, its preferred privacy policy P_a, the maximum possible intimacy value $\mathcal{Y}$, a conflicting target user $i \in \mathcal{C}$, the willingness of user a to accept an action $d \in \{0, 1\}$ to be assigned to target user i is a function $\mathcal{W}_a : \mathcal{C} \times \{0, 1\} \to [0, 1]$ so that:*

$$\mathcal{W}_a(i, d) = \begin{cases} 1 & \text{if } d = v_a[i] \\ \frac{\mathcal{Y} - |P.\theta_{r(a,i)} - int(a,i)|}{\mathcal{Y} + P.\theta_{r(a,i)}} & \text{otherwise} \end{cases}$$

4.2 Modelling Concessions

The mediator models users' decision making during negotiations based on the willingness to accept an action defined above and existing evidence about *manual* negotiations in this domain [4, 9]. Users' decision making on continuous variables like the willingness to accept an action is commonly modelled using fuzzy sets that characterize subranges of the domain of the continuous variables under consideration [10]. Figure 1 depicts the subranges for the willingness to accept an action: LOW and HIGH. Based on this, the mediator considers fuzzy IF-THEN rules to model *concessions* in different situations as described below.

Users are generally willing to accommodate others' privacy preferences [9, 4], so if they do not mind much about which action is finally applied they will be willing to concede and accept applying the action that is not the most preferred for them. Hence, if the willingness to accept changing an action is high, then this may mean that the user would not mind much about which action is finally taken. Assuming

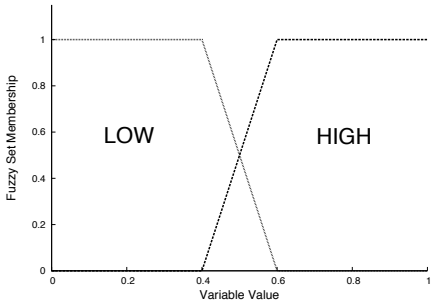

Figure 1: Fuzzy sets low and high.

a negotiating user $a \in N$, a conflicting target user $i \in \mathcal{C}$, and an action $d \in \{0, 1\}$ so that $v_a[i] \neq d$ — i.e., d is not a's most preferred action for i, this can be formalised as the following fuzzy IF-THEN rule:

$$\texttt{IF } \mathcal{W}_a(i, d) \texttt{ IS high THEN concede} \quad (1)$$

Note that `concede` means that user a would accept changing its initial most preferred action to reach an agreement. Thus, users that would initially prefer granting/denying the particular conflicting target user access to the corresponding item would accept denying/granting access. For instance, Alice and Bob could be depicted in a photo with very low sensitivity — e.g., a photo in which both Alice and Bob are depicted with a view to a monument — and both of them could have defined privacy policies for the photo so that all their friends can see it. Suppose that Charlie is friend of Alice but is distant acquaintance of Bob, so according to Alice's privacy policy Charlie should be granted access to the photo but according to Bob's privacy policy Charlie should not be granted access to the item. However, given that the photo is not sensitive for Bob, Bob would probably accept sharing also with Charlie and solve the conflict.

Even when the willingness to accept an action could be low for some of the negotiating users, users do not want to cause any deliberate harm to their friends and will normally listen to their objections [9]. That is, if some of the negotiating users prefers denying a conflicting target user access and her/his willingness to accept granting access is low, then users whose most preferred action for the target user is granting access and the willingness to accept denying is also low would *concede* and accept denying access to the conflicting target user. Indeed, considering others self-presentation online has been reported as a way of reaffirming and reciprocating user's relationships [9]. Assuming a negotiating user $a \in N$, a conflicting target user $i \in \mathcal{C}$, and $v_a[i] = 1$ — i.e., a would prefer granting i access to the item, this can be formalised as the following fuzzy IF-THEN rule:

$$\texttt{IF } \mathcal{W}_a(i, 0) \texttt{ IS low } \wedge \exists b \in N,\ \mathcal{W}_b(i, 1) \texttt{ IS low THEN concede} \quad (2)$$

For instance, Alice, Bob, and Charlie are depicted together in a photo in which Bob is clearly inebriated. Initially, Alice and Charlie might very much like to share the photo with friends because Alice, Bob and Charlie could agree they had a very good time together that day in which the photo was taken. However, Alice and Charlie would probably understand the privacy implications this may have to Bob. Thus, if Bob opposes to share the photo, Alice and Charlie would probably accept not sharing the photo.

Finally, when the willingness to accept granting grating access to the item is low, users very much seek to avoid

sharing the item [5], because it can cause them a privacy breach, i.e., a sensitive item ends up shared with someone they would not like. Assuming a negotiating user a, a conflicting target user $i \in \mathcal{C}$, and $v_a[i] = 0$ — i.e., a would prefer denying i access to the item, this can be formalised as the following fuzzy IF-THEN rule:

$$\texttt{IF} \quad \mathcal{W}_a(i,1) \ \texttt{IS low THEN do not concede} \quad (3)$$

For instance, in the example above, Bob would most probably not accept sharing the photo in which he appears inebriated with Alice and Charlie's friends because he may feel embarrassed about the photo.

4.3 Computing Conflict Resolution

The mediator computes the solution for each conflict found by applying the concession rules defined above. The solution is encoded into an action vector $\vec{o}$, so that $o[i]$ contains the action for target user i. If i is not conflicting, the mediator assigns to this target user the action shared by all negotiation users. If i is conflicting, the mediator assigns to $o[i]$ its proposal to solve the conflict by executing Algorithm 1.

If for all negotiating users, their willingness to accept the action that is not the one they prefer is high, then, according to concession rule (1), the mediator assumes that all users are willing to concede if need be, so that the final action to be applied for target user i can be both grating and denying. In order to select one of these two actions, the mediator runs a modified majority voting rule (Lines 3-6). In particular, this function selects the action that is most preferred by the majority of users. In case that there is a tie, then the one that uploaded the item is given an extra vote. Note that this function is only used if all the users have a high willingness to accept the action that is not the most preferred for them, i.e., it does not really make much of a difference for them, and all of them are willing to concede to reach an agreement.

If there are users whose willingness to accept the action that is not their preferred one is low (Lines 8-14), then the mediator considers two cases: (i) if there are at least two users with low willingness and different preferred actions, according to concession rules (2) and (3), the action to be taken should be denying the conflicting target user access to the item in question; (ii) otherwise, rule (1) applies so that the users that have high willingness will concede and the user/users who has/have low willingness will determine the action that is finally chosen as the solution.

Algorithm 1 Conflict Resolution

Input: N, $\mathcal{P}$, C
Output: $\vec{o}$

```
 1: for all i ∈ C do
 2:
 3:     if ∀a ∈ N, W_a(i, ¬v_a[i]) is HIGH then
 4:         o[i] ← modified_majority_voting(P, i)
 5:         continue
 6:     end if
 7:
 8:     if ∃a ∈ N, W_a(i, ¬v_a[i]) is LOW then
 9:         if ∃b ∈ N, W_b(i, ¬v_b[i]) is LOW ∧ v_a[i] ≠ v_b[i]
    then
10:             o[i] ← 0
11:         else
12:             o[i] ← v_a[i]
13:         end if
14:     end if
15: end for
```

	$\mathcal{W}(i_3,0)$	$\mathcal{W}(i_3,1)$	$\mathcal{W}(i_4,0)$	$\mathcal{W}(i_4,1)$
a	-	HIGH	-	LOW
b	LOW	-	HIGH	-

Table 2: Fuzzy Memberships for Example 3.

EXAMPLE 2. *Suppose again Example 1. Table 2 shows the fuzzy set membership for negotiating users a and b in case they would accept changing their most preferred action for the conflicting target users $\mathcal{C} = \{i_3, i_4\}$. We can see that for negotiating user a and target user i_3 rule 1 applies, so that the mediator assumes that user a is willing to concede (in this case accept granting i_3 access to the item). As there is only one negotiating user (b) with willingness LOW, then the action suggested by this user would be taken to solve the conflict, and the computed solution would be to grant i_3 access to the item. Regarding target user i_4, we have a similar situation. In this case, the willingness is HIGH for b, so that rule 1 applies and b would concede. Moreover, there is only one negotiating user (a) with willingness LOW, so the action suggested by this user is taken to solve the conflict. Therefore, the solution to the conflict would be to deny i_4 access to the item. The resulting action vector for the item would be $\vec{o} = \{1, 1, 1, 0\}$.*

5. REFERENCES

[1] B. Carminati, E. Ferrari, and A. Perego. Enforcing access control in web-based social networks. *ACM TISSEC*, 13(1):6, 2009.

[2] R. L. Fogués, J. M. Such, A. Espinosa, and A. Garcia-Fornes. Bff: A tool for eliciting tie strength and user communities in social networking services. *Information Systems Frontiers*, 16(2):225–237, 2014.

[3] H. Hu, G. Ahn, and J. Jorgensen. Multiparty access control for online social networks: model and mechanisms. *IEEE TKDE*, 2012.

[4] A. Lampinen, V. Lehtinen, A. Lehmuskallio, and S. Tamminen. We're in it together: interpersonal management of disclosure in social network services. In *Proc. CHI*, pages 3217–3226. ACM, 2011.

[5] M. Sleeper, R. Balebako, S. Das, A. L. McConahy, J. Wiese, and L. F. Cranor. The post that wasn't: exploring self-censorship on facebook. In *CSCW*, pages 793–802, 2013.

[6] A. Squicciarini, M. Shehab, and F. Paci. Collective privacy management in social networks. In *WWW*, pages 521–530. ACM, 2009.

[7] K. Thomas, C. Grier, and D. M. Nicol. unfriendly: Multi-party privacy risks in social networks. In *Privacy Enhancing Technologies*, pages 236–252. Springer, 2010.

[8] J. Wiese, P. Kelley, L. Cranor, L. Dabbish, J. Hong, and J. Zimmerman. Are you close with me? are you nearby? investigating social groups, closeness, and willingness to share. In *UbiComp*, pages 197–206. ACM, 2011.

[9] P. Wisniewski, H. Lipford, and D. Wilson. Fighting for my space: Coping mechanisms for sns boundary regulation. In *Proc. CHI*, pages 609–618. ACM, 2012.

[10] L. A. Zadeh. The calculus of fuzzy if/then rules. In *Proceedings of the Theorie und Praxis, Fuzzy Logik*, pages 84–94. Springer-Verlag, 1992.