

Uncovering Facebook Side Channels and User Attitudes

Sai Lu
Rutgers University
sl914@cs.rutgers.edu

Janne Lindqvist
Rutgers University
janne@winlab.rutgers.edu

Rebecca N. Wright
Rutgers University
rebecca.wright@rutgers.edu

Abstract—Over the course of the last decade, Facebook has become an incredibly popular social networking site, reporting around a billion visitors monthly. Like any social networking site, Facebook’s design decisions have implications about what is shared, what is not shared, and how much control users have about such sharing. In this paper, we report a systematic analysis of side channels in Facebook—that is, channels that can reveal privacy-sensitive information to users through indirect mechanisms. While these side channels may not be particularly surprising to the security research community, they still represent potential threats to Facebook users, depending on user expectations and attitudes. We surveyed Facebook users to determine user expectations and attitudes, including whether users are aware of these channels, and whether there are privacy objections to the channels to those aware. We find that many users are unaware of the side channels and express surprise at finding out about the side channels. Among users who are aware of them, some users express concerns while others do not. Based on these results, we identify design implications for social network sites that wish to provide users with more control over such choices.

I. INTRODUCTION

Since mid-2012, Facebook has ranked as the most or second-most accessed web site globally [1], with over a billion users active monthly reported [2]—more people than the population of many countries, and nearly one in every seven people on the planet. Facebook and other online social networking sites (OSNs) have led to an unprecedented ability to share information distantly, bringing people closer online. People share photos, location, basic personal information, and interesting links with their Facebook friends. However, there are also concerns about such information sharing, with information divulged on Facebook sometimes resulting in unexpected or undesired outcomes such as dismissal from employment [3] or crime and fraud [4].

In order to allow users to make better decisions about information sharing, most OSNs—including Facebook—now allow users to make their own decisions about which groups of people they wish to share their posts, basic information, etc. with by adjusting their privacy settings. For example, a Facebook user Alice can choose who she would like to share her posts with: everyone, her friends, only herself, or “custom” settings where she can choose which subset of her friends can see each post.

In our work, we focus on information that is shared via “side channels.” Analogous to side channels in other computer systems and applications, such as storage channels [5], timing channels [6], and power consumption channels [7], a side channel is an information channel that is secondary or incidental to the intended communication channel but that can convey additional information. A trivial example of a side channel in

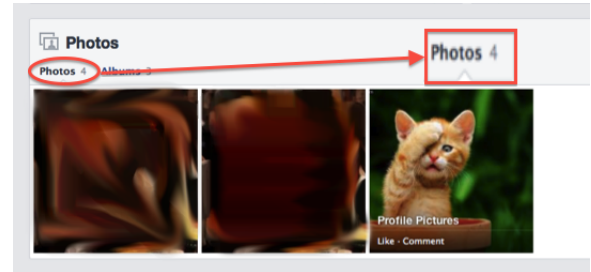


Fig. 1: An example of a side channel in Facebook. There are only three photos displayed. In contrast, Facebook lists that there are four photos in this photo album. Users viewing this can infer that one photo has been blocked from them by the user who posted these photos.

Facebook is depicted in Figure 1. Users are allowed to share their photos on Facebook, but also are able to limit who can view them. In this example, Facebook lists a total of four photos for this photo album but shows only three different photos. Therefore, the viewer could infer that he or she is blocked from viewing one photo.

It is not surprising that Facebook privacy controls have side channels. However, to determine whether these side channels should be viewed as potential privacy issues or have other social implications, we investigate: (1) whether they cause behaviors that violate users’ expectations because they are unaware of who can actually see what and (2) whether those users who are aware of the side channels and their implications view the side channels as problematic from a privacy perspective. Toward this end, we conducted a survey of 80 Facebook users over Amazon Mechanical Turk.

Our paper makes the following contributions. We systematically investigate the existence of side channels in Facebook and identify eleven such side channels of four major types. We conducted a survey of Facebook users to determine whether users are aware of these channels or not, and whether those who are aware have privacy objections to the channels. We find that many users are unaware of the side channels and express surprise at finding out about them. Among users who are aware of them, some users express concern while others do not. Additionally, based on these results, we suggest some design recommendations for social network sites that wish to provide users with more control over such choices.

II. RELATED WORK

Privacy protection is a difficult issue in OSNs because they are designed precisely to share and communicate information,

and indeed facilitating and encouraging such sharing drives the business model of the OSN itself [8]. Many researchers have addressed privacy concerns surrounding Facebook, including suggesting mechanisms for improving the security and privacy of Facebook [9]–[13].

Acquisti and Gross conducted early studies of privacy concerns of users on Facebook [14], [15]. They conducted a survey on users’ privacy concerns about Facebook and their awareness about the visibility of their online profiles. The study revealed that 40% of their participants had misperceptions about the visibility of their profiles and 30% reported that they did not know how to control correct privacy setting.

Several studies focus on audiences as a way to evaluate the effectiveness of privacy controls [10], [16]–[18]. These studies focus on the difference between users’ expected audience and the actual audience—for example, who gets to see a photo the user posted. Liu et al. [18] presented ten different scenarios to users and discovered that 37% of the time, users used incorrect privacy controls to post information online with respect to the target audience they were aiming for. They also found that half of the posted content on Facebook was published with then-default settings that exposed the content to all Facebook users. In order to mitigate privacy setting errors, Egelman et al. [19] found that simply alerting users to potential errors in their privacy settings is not sufficient to mitigate the errors; they designed a new privacy setting interface intended to reduce such errors.

In addition to privacy issues caused by the use and understanding of Facebook’s privacy control mechanisms, they can also arise through other kinds of leakage. Henne et al. [11] focused on privacy implications for users created by other users’ actions. Privacy issues have also been identified in other online social networks such as foursquare, Google+, and Twitter (e.g. [20]–[23]).

There are several examples of machine-learning and other inference techniques being used to identify sensitive information without being given direct access to it. A famous example is the MIT “Gaydar” project [24], which showed that it was possible to infer with high accuracy the sexual orientation of a user based on the percentage of the user’s friends that self-identify as gay male. Dey et al. [25], [26] show how to use such techniques to identify ages of Facebook users and, going further, to identify and profile high school students in a target high school, even though Facebook attempts to protect minors from excess disclosure.

To the best of our knowledge, we are the first to consider side channels in Facebook and OSNs. However, side channels and their close relatives, covert channels, have long been a known issue in computer systems, particularly in multi-user systems designed to support multi-level security [27]. Covert channels were first named in a seminal paper by Lampson in the early 1970s [28]. Since then, they have been studied in a variety of settings in computer systems [5], [29] and networking, including LANs [30], TCP/IP protocol suite [6], [31]–[35], anonymous routing [7], and HTTP (web traffic) [36]. It is a known standard security principle that because query denials (such as authorization denials) can themselves leak information, it is particularly important that such queries are run with the security privileges of the querier or that they

otherwise provide as little information as possible [27], [37], [38]. Several of the side channels we demonstrate could be avoided, if desired, by following this guideline.

III. SYSTEMATIC ANALYSIS OF FACEBOOK’S SIDE CHANNELS

We focus on side channels in Facebook. We deem a side channel to occur when a user can easily infer, through an indirect channel on Facebook, information that he or she does not directly have access to on Facebook. For example, a side channel can occur regarding relationship status. If Alice attempts to add Bob as a relationship, and Bob is already in a relationship, the action will fail (as described in more detail in Section III-D). In this case, Alice can then infer Bob is in a relationship with another person.

Systematic Analysis. To systematically explore side channels in Facebook, we created three Facebook accounts for users we named Denise, John, and Lance. In a series of experiments enumerated in in Appendix A, we focused on side channels that involve a maximum of three users. Our experiments, conducted in April 2013, were mainly based on the privacy control options and actions of Facebook that have associated privacy controls.

Each of our accounts plays a different role in our experiments: two accounts play as actors, one account plays as the audience. We studied how privacy controls work on 20 different features and actions on Facebook, for example, Photos, Taggings, Share, etc.. We also went through all privacy settings (Privacy, Timeline and tagging, Blocking) on Facebook for a total of 767 action combinations. We focused on leaks (i.e. apparent or potential violations of the intent of privacy controls) that an action or combination of actions enabled. Given an action or actions and a combination of privacy controls, we would like to know whether all actors’ intentions appear to be fulfilled.

For example, in one of our experiments, John (Actor) makes a post, tags Denise (Actor) in it, and enables everyone to see it—that is, sets the privacy setting as Public. Denise, on the other hand, sets her Tagging privacy setting as Friends-only, which allows only her friends to see posts she has been tagged in. We want to know that if Lance (Audience) who is not a friend of Denise can see this post on Facebook. In the experiments, we changed the friendship statuses and privacy controls systematically and for each case checked who actually gets to see what.

Based on this exploration, we identified 11 side channels that fall into four major types, as well as some additional privacy control concerns. In the rest of this section, we describe these side channels. (Examples of trials can be found in Appendix A.)

A. Side Channels Based on Counts

Among other actions, Facebook enables its users to post photos, post notes they have written, and “Like” “pages”. Users can also set privacy settings that determine which users these actions are visible to. In several cases, as already illustrated in Section I, Facebook will display both information related to an action for a user (such as the photos a user has posted) and

Fig. 2: The figure illustrates the side channel caused by the relationship feature. Facebook will produce an error message (“The relationship you chose isn’t supported by Facebook.”) if a user tries to publish a relationship with someone who is already in a relationship, even if the relationship is not visible to the user trying to publish the relationship.

the count of those actions (such as the number of photos the user has posted). In some cases, the privacy policy affects the content but not the count.

We discovered three side channels that involve counts: Likes, Photos, and Notes.

We found that the total number of photos that an owner has is displayed, is different from the number that a viewer is allowed to see, leading to an observable inconsistency. (See Figure 1.) For example, consider the scenario where Alice, Bob, and Carol are friends on Facebook with each other. Alice has three photos in total, and because of her privacy setting, Bob is only able to see two photos out of three. However, Bob can see that Alice has three photos and can infer that Alice must have hidden a photo from him.

The side channel here is that the viewer can figure out that the owner has hidden a photo or photos from the viewer from the fact that the number of photos listed does not match with the number of photos that can be seen. In the example above, Bob can determine that Alice has hidden a photo from him through the fact that Alice has three photos while he can only see two of them. In addition to the number of photos in the above example, this side channel also applies to notes and to pages a user likes.

B. Side Channels from Sharing

Thanks to the share button, news, links, posts can spread in users’ networks. Users can repost status, photos, links, and other Facebook content that their friends have already posted. Facebook offers sharing privacy settings for users to choose from in order to provide users some control over the visibility of content and actions. However, when sharing is involved, the privacy settings of one or more users may be relevant. For example, if a post was not “public”, then Facebook will only allow the post to be shared via private message.

We found two side channels related to the Share feature.

Again, assume that Alice, Bob, and Carol are Facebook friends. In this scenario, Alice posted a photo as “only-me” and tagged Bob, hiding the photo from Carol. Bob decided

to share it with Carol in a private message and added “Check out this photo Alice posted!”. However, Carol found that she could not see the photo, and from this she can infer that Alice did not share the photo with her.

The side channel here is that a user is able to find out he has been intentionally hidden from a post through someone’s sharing of this post in a private message. It is still true when a post is shared within a group instead of a private message. In the example above, Carol will figure out that Alice has hidden a post from her through the private message Bob sent to her.

C. Side Channels from Tagging

Adding a tag is defined as creating a link on a user’s timeline. Facebook has two sets of settings involving tagging, one for users who make the tagging action (Public, Friends, Custom, Only me, as shown in Figure 3a) and another for users who have been tagged (Everyone, Friends of friends, Friends, Custom, Only, as shown in Figure 3b). The combinations lead to a number of “scenarios”, as shown in Tables I and II.

When there is a conflict between the audience of two relevant settings, Facebook will follow the setting of the original poster. In some cases, however, this will disobey the setting of the tagged users. Excluded users still can see the post, and they can figure out who has blocked them. Among all combinations of settings, we identified four side channels related to tagging in a post or other contents, and one side channel related to the restricted list.

For example, by comparing the timelines of two friends involved with a post, a mutual friend can figure out whether one of them has hidden a post from him. In this scenario, Alice tags Bob in a post and Facebook automatically creates a link on Bob’s timeline. Of note here is that: 1) Alice can choose the audience on her timeline, and 2) Bob can select the audience of posts he has been tagged with. As a result, there will be possible conflicts between Alice’s audience and Bob’s audience. For example, Alice’s setting allows their common friend Carol to see her post, in which Bob was tagged. At the same time, Bob customizes to hide posts that he has been tagged in from Carol. Facebook allows Carol to see this post from Alice’s timeline, but Carol cannot see this post on Bob’s timeline. Carol will then figure out that Bob has hidden this post from her if she can see this post on Alice’s timeline but cannot see it on Bob’s timeline.

The side channel here is that a user can figure out he has been hidden from a post when he cannot see a post on another user’s timeline while being able to see it somewhere else. This side channel also applies to the Restricted List, which can be used to block friends from viewing any protected contents. Adding a user to Restricted List will not prevent the user from seeing tagged contents on others’ timeline.

D. Side Channels from Relationship Status

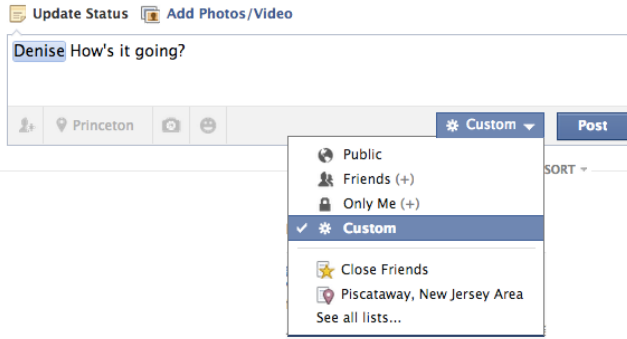
Facebook has a strong connection with people’s real lives, since it forces people to use their real names and encourage them to add friends from their real lives. Sometimes, it affects people’s lives and romantic relationships [39]. After several times changing the setting panel, Facebook has replaced its former design on relationships. Representatively, a user can

Tagging / tagged	Everyone	Friends of friends	Friends	Only me	Custom
Public	Everyone	Everyone 1	Everyone 2	Everyone 3	Everyone 4
Friends	A's friends and B's friends	B's friends	A's friends 5 B's friends	A's friends 6 B's friends	A's friends 7 B's friends
A's Only me	A+ tagged	A+ tagged	A+ tagged	A+ tagged 8	A+ tagged
Custom	Custom	Custom	Custom	Custom 9	Custom

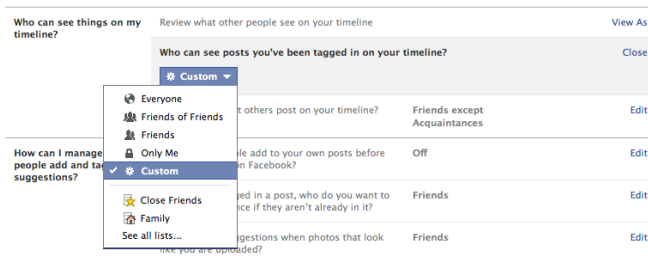
TABLE I: Audience of a post on Alice's timeline. A stands for Alice, B stands for Bob, and the numbers in the cells label scenarios representing different combinations of privacy settings.

Tagging/ tagged	Everyone	Friends of friends	Friends	Only me	Custom
Public	Everyone	Everyone 1	Everyone 2	Everyone 3	Everyone 4
Friends	A's friends and B's friends	A's friends B's friends	A's friends 5 B's friends	A's friends 6 B's friends	A's friends 7 A+ friends
Only me	A+ tagged	A+ tagged	A+ tagged	Only me 8	A+ tagged
Custom	Custom	Custom	Custom	Only me 9	Custom

TABLE II: Audience of a post on Bob's timeline. A stands for Alice, and B stands for Bob. If we can make a comparison between audience of this post on Alice's timeline (Table I) but not on Bob's timeline (Table II), scenario No.3, No.6, No.7, and No.9 can be considered as side channels. Because in those scenarios, they are audience that Bob has limited from this post.



(a) A user can choose different privacy settings before he or she publishes a post.



(b) A user can choose privacy settings for posts he or she has been tagged in.

Fig. 3: Users can choose audience of their posts as well as posts they have been tagged in. Figure 3a and 3b show different choices that Facebook users have. We consider scenarios based on different combinations of these choices, and determine which of them lead to side channels.

no longer have multiple relationships at the same time, though it is allowed in the life event part.

To announce a relationship on Facebook one requires the following steps: selecting a proper privacy setting, tagging a partner's name, then waiting for the partner's response. Once the partner approves this relationship, it is officially posted

on Facebook. However, this kind of publicity can cause some problems to users' lives. Romantic partners may have different views on the disclosure of this relationship on Facebook [39], and, further, a user can find out whether his or her partner is in an existing relationship on Facebook.

We found one side channel related to the Relationship feature, which can leak a user's relationship status, illustrated by the following 4: Alice and Bob are in a relationship and Bob is also in another relationship with Carol. Because of Bob's privacy setting, Alice cannot see Bob's relationship status. When Alice adds Bob in her relationship on Facebook, she will receive a notice saying that the relationship chosen is not supported by Facebook. (Previously, the error message said "That user is already in a relationship." Although the error message has changed, the circumstances that produce the message have not.) Alice can then figure out that Bob is already in a relationship with someone else, and Bob is not only in a relationship with her.

The side channel described is that a user can find out the status of another user's relationship status by adding this user in a relationship on Facebook, even if the relationship status had been blocked from the user. This side channel discloses the relationship status of those who wish to keep their status secret.

IV. USER STUDY METHOD

In the previous section, we described side channels that we have identified on Facebook. In order to understand whether these side channels should be considered privacy-invasive, we conducted a user study to determine whether users know about the side channels as well as whether they consider them to be problematic or reasonable information disclosures. To this end, we conducted a survey on Amazon Mechanical Turk (AMT). Both the study and the associated consent form were approved by the Rutgers Institutional Review Board (IRB).

A. Participants

We surveyed a total of 80 participants. The participants included 47 females (56.25%) and 33 males (43.75%), age

19 to 61 (mean = 35) with an average time on Facebook of 4.7 years (max = 9.1 years, min = 3 days). To mitigate confounding factors due to cultural or language issues, we restricted our participants to people located within the United States. To help guarantee the quality of responses, we also required participants to have received their AMT “Master’s” and to have a lifetime approval rate of at least 95% (i.e., the rate of successfully completing previous tasks). The survey was open for a total 26 days during August 2013 to October 2013. Each participant with a successful submission received a payment of \$2.00. We checked that all participants spent at least a minimum threshold of time on the survey that we estimated would be needed to reflect on the questions on the survey rather than simply clicking through without paying attention. The mean finishing time for participants was 26 minutes. We refer to the participants in this study as P1-P80.

B. Experimental Design

Our survey included 50 to 70 questions, depending on participants’ choices within the survey. Some questions were not shown if the answer to a previous question was incorrect or showed that a participant had no experience related to a previous question. For example, in the tagging scenarios, if a participant answered the first question incorrectly, this participant would not be shown the next four questions based on that scenario.

C. Procedure

Our survey had four parts: a consent form, a set of Facebook usage questions, a set of questions exploring side channels, and demographic questions.

a) Facebook Usage: We asked our participants about their frequency of Facebook use and time spent on each Facebook session. Participants were also asked five open questions about their opinions on Facebook, which included their views on Facebook’s drawbacks and benefits. They were also asked about any potential concerns they may have about Facebook and how (or whether) they work around or address their concerns in practice.

b) Privacy and Side Channels: Most of our survey questions were based on the side channels identified in Section III, and others were used to evaluate users’ awareness of privacy controls on Facebook. We designed several scenarios to illustrate the identified side channels. For each scenario, the survey described the scenario and then asked a few questions related to the scenario. Participants were also asked whether they have experienced similar situations. To provide richer qualitative data, participants were asked to elaborate whenever they reported having experienced a similar situation.

Relationship Scenarios. Two scenarios were presented about relationships. One is based on the side channel described in Section III-D, another is about the privacy settings of relationship status.

Tagging Scenarios. Participants were asked about actions related to tagging on Facebook. The survey included four scenarios about tagging, one for each of the four side channels related to tagging. Participants

were given a scenario first and then asked questions about what they think will happen in that scenario.

Edit History For Comments. Participants were asked if they know how to view the edited history of a comment and their opinions about this function.

Number of Photos, Notes, and Likes. Participants were presented with a scenario about Facebook’s Photos feature to explore their behavior and experience related to this side channel.

Share Button. The survey included two questions about whether the participant could recall situations in which they could not use the Share button on Facebook when it would normally be available. A third question asked whether participants’ friends have ever complained to them about not being able to view their shares. A fourth question asked whether participants have ever been aware of being unable to view photos that had been shared with them on Facebook.

Apps. We asked participants about their knowledge of Facebook apps’ privacy settings.

Friends List. Five survey questions related to privacy and friend lists on Facebook.

c) Demographics: We asked participants typical demographics questions including gender, age, education, and occupation. To control for confounding factors, participants were also asked about their nationality, native language, state and city of residency, and how long have they lived in the United States. The participants were also asked questions to estimate their technical expertise and what online social networks they use. Finally, to quantify participants’ attitudes towards online privacy, we used an online privacy scale developed by Tsai et al. [40] with a 7-point Likert scale.

V. USER STUDY RESULTS

In this section, we describe the results of our user study. We start by describing our participants’ privacy attitudes and awareness of Facebook privacy controls, and continue with their understanding of the side channels we discovered and their experiences with side channels, if any. The quotations included below are representative of particular findings from the survey.

A. Privacy Attitudes and Awareness of Facebook Privacy Controls

We probed our participants’ general awareness of Facebook’s privacy controls and their attitudes towards privacy.

Our participants reported that they were very highly concerned about their online privacy as indicated by the results of using the online privacy scale developed by Tsai et al. [40]. The mean and median values were 6, and the mode was 7 on a 7-point Likert scale.

All but one of our female participants had set privacy settings of their posts to more restrictive than Public (43 Friends-only, Only Me 1, Custom 2, 1 had not set it). Six of our male participants had Public profiles and 25 were Friends-only and one Custom. According to our participants, if you do

Name	Participants who were able to discover the side channel presented	Percentage
Tagging 1	24	30%
Tagging 2	25	31.25%
Tagging 3	24	30%
Tagging 4	13	16%
Relationship	24	31%
Photos	7	8.25%
Edit History	19	23.75%
Sharing	11	13.75%

(a)

Name	Participants who answered correctly	Percentage
Generic awareness 1	31	38.75%
Generic awareness 2	29	36.25%
Both 1 and 2	5	6.25%

(b)

TABLE III: Numerical Survey Results: (a) and (b) shows some numerical results of participants' answers. (a) shows the number and percentage of participant who are aware of each side channel scenarios, (b) shows the number and percentage of participants who answered correctly about generic awareness questions.

not understand the privacy settings, this will lead your actions to having unexpected audience, which can cause unpleasant experiences with friends and friends of friends.

Most of our participants did not publicly share their posts or posts they have been tagged in. 68 participants (85%) chose "Friends" as privacy setting of posts. For posts they have been tagged in, 20 participants (25%) chose "friends of friends" as their privacy setting, 36 participants (45%) chose "friends".

We gave our participants two questions based on two different scenarios to test their knowledge on privacy settings of tagging. 31 people (38.75%) answered the questions based on scenario A correctly, and 29 participants (36.25%) answered the questions based on scenario B correctly. Only five participants (6.25%) gave correct answer for both questions. The result shows our participants lacked knowledge on what each privacy setting changes audience of a post.

Participant P17 shared: *"I have noticed on some posted that people liked them that were not my friends. It ended up being if I tagged a friend then all their friends could not see the post. It is annoying sometimes."* Similarly, participant P35 mentioned, *"I was new to Facebook and wasn't totally well versed with the privacy settings. I published a status and unknowingly allowed people who are not my friends yet share mutual friends with me be able to see it too."* Participant P46 discussed not knowing how to restrict access to some of his friends: *"Yes, some of my friends have different religious views than I do. I did not mean to start a conflict and would have changed the settings if I knew."*

Edited history of a comment on Facebook can be viewed by anyone who has access to the post. Our study indicates that Facebook users do not have a good sense of this function. 61 participants (76.25%) did not know how to view the edited history of comments. 65 participants (81.25%) have never viewed the edited history of comments before. After participants became aware of the function with our survey, some of them became concerned about some comments they have made in the past.

Participant P7 commented, *"Yes, this doesn't seem fair you should be able to delete or edit things as you see fit. It may show that something has been edited but it should not show the edited."*

Participant P47 became aware of this function from our survey and became worried if her friend could see this edited history, *"I deleted a comment I made about a mutual friend in a post....I am now concerned that he actually did see it or can see it...? Yikes! :-"*

B. Understanding Side Channels

Relationship. 24 participants (31%) claimed they have knowledge of the relationship side channel, but none of our participants reported having experienced it. 49 participants (63%) reported that they did not have enough knowledge about what happens if they receive a relationship request. 32 participants (40%) would not publish a relationship status on Facebook unless they know the privacy setting of it. On the other hand, participant P36 was concerned about who can see relationship status, *"This is why I don't accept relationship requests, because I don't know who will see them."*

Sharing. 11 participants (13.75%) reported having an experience that they could not find a Share button when they expected to find one. Participant P67 told us: *"Sometimes I see a graphic I want to share but the button isn't there."*, while Participant P29: *"I've noticed some photos posted do not have a share button."* Participant P24 thought there were some issues with Facebook's user interface: *"I don't really find Facebook all that userfriendly sometimes. It is almost counterintuitive, I want to close a box after reading a comment and I realize I am reporting it as spam instead."*

Photos. Only seven participants (8.75%) recalled noticing that Facebook sometimes showed the total number of photos instead of the exact number a user can see.

Participant P19 mentioned, *"The picture count listed did not match the number of photos shown."* If a user sets a photo in an album not to be accessed by all friends, his or her friends can notice this when they browse this album, as participant P17 had noticed *"I could tell b/c when I was looking through a list of photos one came up that said either this had been removed by user or you do not have access to view this item."* Participant P65 had noticed a photo been blocked from him via his wife's Timeline: *"Because it showed up fine on my wife's Facebook page."*

Tagging. In order to learn how much Facebook users know about tagging related side channels, how they deal with

them, and what their experiences with the side channel, we designed four scenarios based on our four discussed tagging side channels. In these scenarios, we described that Facebook users John, Denise, Lance are friends. John tagged Denise in his post with a privacy setting that allowed Lance to see his posts, and Denise blocked Lance from posts she had been tagged in. We asked five questions in each scenario to understand if participants understood how tagging works on Facebook. In our scenario, Lance will see this post in all scenarios from John's timeline, which enables him to figure out Denise hid this post from him. Please see again Section 3.3. for further details.

Our participants tended to think that when Denise blocks Lance from seeing this post, then Denise's setting will override John's original setting to protect Denise's setting. However, Facebook will follow John's privacy setting instead of Denise's. 55 participants (68.75%) answered incorrectly in at least one of the four scenarios. They thought Lance will not see this post. 16 participants (20%) mentioned uncertainty words like "I assume..." and "sure" and subjectivity words like "should".

Participant P9 thought Denise's setting *should* prevent Lance seeing the post, *"Her settings should prevent him from being able to see."* (Q7.1), similarly, P39 shared *"Lance should be block from viewing this picture."* Participant P23 assumed that *"I assume Only-me means Lance can't see the post anywhere."* Participant P9 told us he was *not sure*, in his words *"I'm not sure about this one, but with the settings I would assume no."*

17 participants (21.25%) answered the question correctly in all four scenarios and elaborated correctly as well. For example, participant P45 told us that Lance could see it on John's page but not on Denise's, *"I think that Lance would be able to see it on John's page, but if Lance went to Denise's page he wouldn't be able to see it."*

If participants answered correctly that Lance will see the post in the scenario, we then asked if they are aware the side channel in these scenarios. From scenario one to scenario four, the percentage of participants who said Lance can figure out Denise has hidden from him were 30% (24), 31.25% (25), 30% (24), 16.25% (13). Participant P24 shared to us that *"If he can see it on John's and he can't see it on Denise's then he will know he is blocked."*

Even though users have noticed the fact that Lance can still see the post even though Denise has hid it from him, some of them would not consider them a privacy problem. For example, P47 answered *"He could assume that he can't see it on her wall because she's hidden it from him, but maybe she's hidden it from everyone except whomever she is having a conversation with - so he'd just be assuming that it's all about him."*

In the scenarios, John would like Lance to see this post, but Denise does not want it to be visible to Lance. Which user's privacy setting to follow becomes a debatable question. John publishes this post, and can be considered as the first owner, and Lance can be viewed as the second. But when it comes to privacy, we need to figure out whose privacy to protect. Our participants' answers divided into two aspects about it. Some of them hold that John's privacy setting as a higher priority because he is the one who posts it. Like P14 and P43 told

us, *"It was made public by the author."* *"Yes, with hesitation. Because the person who originally posted it intended for it to be viewable on his page, I think he should have the right to post it. This is why I hate the myraid of permissions. You need to expect that people might post things to/about you from time to time."*

However, some of them believe that since Denise is tagged in the post, her settings need to be executed, too. In Participant P24's words, *"Denise has a right to set her account any way she wants. No one has a right to her private information."* Participant P34 contributed her idea to protect Denise's privacy, she wrote *"Since John's post was public and he is a friend, then Lance should be allowed to see the post, but not that Denise was tagged."*

C. Experiences with Side Channels

Two of our participants shared with us additional specific experiences with side channels.

Participant P47 shared: *"I have posted on a mutual friends' wall, and I assume the person I unfriended, clicked to add me when she saw that I had posted on our friends' wall. I did not specifically exclude her from seeing my posts though."*

In some cases, users cannot use privacy controls to limit audience, which causes unpleasant experience between users and their friends. For example, participant P42 shared with us, *"I have tried to make private events, but they show up on my timeline. People not invited can't see the actual event, but they can see that I posted an event."*

VI. DISCUSSION AND CONCLUSIONS

A. Discussion

Our study revealed eleven side channels on Facebook, in four major types. Three of these functions (Likes, Photos, Notes) show a total number of items to users even though the users do not have access to all of them. Four combinations of privacy settings made by posts' owners and users who get tagged in them, will lead a mutual friend to discover that he or she is blocked from seeing a post. Sharing a non-public post in a message and in a group causes side channels, too. A restricted-access user is able to find out that he or she is blocked by the poster based on the different contents that show on two timelines.

Some of the side channels we have identified are the result of design choices that Facebook has made, possibly in some cases with full knowledge of the fact that the channels would be created. For example, the choice to show Bob the total number of Alice's photos rather than show him the number of photos viewable to him could be viewed as a way of ensuring some transparency and protecting Bob from being misled. Alice can hide a photo from him, but she cannot hide the existence of the hidden photo. Similarly, one can view the handling of relationship status as being designed to ensure that it used only for the kinds of relationships that are intended to exclude the possibility of additional relationships.

We found that improper privacy controls will lead to information leaks to unexpected audiences. Confirming previous work (e.g. [15]), we have learned that many users do not

have enough knowledge of privacy controls on Facebook. Their understanding of privacy settings is limited and different from the real meaning of those settings. Their unawareness of side-channels may cause unpleasant experience among families and friends. We also found that female users are more careful with the privacy of their profiles.

Facebook has removed some side channels that existed before. For example, Facebook has removed the side channel from the Poke feature. Poke is a way of showing interest or intent to speak to another user. When one user is poked by another, a notification appears on the user's page saying "Person *X* poked you", the other user can poke back, ignore the poke, or delete it. Facebook previously used to notify the user who pokes whether the other user ignored the poke or whether the other user has not received it yet by showing messages "This user received your poke" and "The user has not yet received your last poke. He or she will get it the next time he or she logs in". However, Facebook now only shows "The user has not responded yet" no matter the other user chooses to ignore the poke or does not see the poke.

Our study shows that Facebook will leak information to the social network about privacy settings that have been enabled. The privacy setting and some related information of a post is visible to its audience, and in some cases, it is also visible to a non-audience. The privacy of contents on Facebook is protected by privacy settings that users made, but the privacy of privacy settings has not been preserved on Facebook. For example, Alice hides a post from Bob, when Bob found out that Alice blocks him from seeing her post, Bob will know that Alice's setting is to hide from him. We suggest that both contents and privacy settings should be invisible to non-audience.

Facebook only enables privacy settings of users' timeline. However, users tend to believe that Facebook could fulfill the privacy settings of themselves on the whole site. There is a misunderstanding between what privacy controls can provide and what privacy controls that users think Facebook enable. To implement a thorough protection of users' privacy, we suggested to enable an extra privacy control for users' contents on others' timelines.

Relationship on Facebook is exclusive, which means each user can only be in one relationship. If one user is secretly in a relationship on Facebook, anyone who tries to add this user in a relationship will be warned that this relationship cannot be added, which can be inferred that this user is already in a relationship.

Finally, only two participants shared specific experiences of side channels. Given that most participants were unaware of most side channels they were presented, we conclude that this further emphasizes how Facebook's privacy controls should be designed to be more user-centered.

B. Further Work

In this paper, we studied only side channels on Facebook. Even though Facebook has the largest user space so far, it is only one possible OSN of interest. Further work should look at other OSNs and how they deal with side channels. We also discussed above that some of the side channels on Facebook

might be intentional, or just artifacts of how Facebook wants some privacy controls work. Clear further work would be to explore the space for alternative approaches, and formally model the side channel space.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant Numbers 1018557 and 1211079. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

- [1] "Alexa Facebook siteinfo," <http://www.alexa.com/siteinfo/facebook.com>, accessed December 2013.
- [2] "Facebook Keyfacts," <http://newsroom.fb.com/Key-Facts>.
- [3] "Social Media: 10 Employment Cases Involving Facebook," <http://www.xperthr.co.uk/blogs/employment-tribunal-watch/2013/04/social-media-10-employment-cases-involving-facebook/%20long%20url>.
- [4] "Social media-related crime reports up 780% in four years," <http://www.theguardian.com/media/2012/dec/27/social-media-crime-facebook-twitter>.
- [5] C.-R. Tsai, V. D. Gligor, and C. S. Chandrasekaran, "A formal method for the identification of covert storage channels in source code," in *IEEE Symposium on Security and Privacy*, vol. 74, 1987.
- [6] S. Cabuk, C. E. Brodley, and C. Shields, "Ip covert timing channels: design and detection," in *Proceedings of the 11th ACM conference on Computer and communications security*. ACM, 2004, pp. 178–187.
- [7] S. J. Murdoch, "Hot or not: Revealing hidden services by their clock skew," in *Proceedings of the 13th ACM conference on Computer and communications security*. ACM, 2006, pp. 27–36.
- [8] J. Anderson and F. Stajano, "Must social networking conflict with privacy?" *IEEE Security & Privacy*, vol. 11, no. 3, pp. 51–60, 2013.
- [9] A. L. Young and A. Quan-Haase, "Privacy protection strategies on facebook: The internet privacy paradox revisited," *Information, Communication & Society*, vol. 16, no. 4, pp. 479–500, 2013.
- [10] C. M. Hoadley, H. Xu, J. J. Lee, and M. B. Rosson, "Privacy as information access and illusory control: The case of the facebook news feed privacy outcry," *Electronic commerce research and applications*, vol. 9, no. 1, pp. 50–60, 2010.
- [11] B. Henne, C. Szongott, and M. Smith, "Snapme if you can: privacy threats of other peoples' geo-tagged media and what we can do about it," in *Proc. WiSec '12*. ACM, 2013, pp. 95–106.
- [12] M. Johnson, S. Egelman, and S. M. Bellovin, "Facebook and privacy: it's complicated," in *Proc. SOUPS'12*. ACM, 2012, p. 9.
- [13] P. Mittal, C. Papamanthou, and D. Song, "Preserving link privacy in social network based systems," *arXiv preprint arXiv:1208.6189*, 2012.
- [14] R. Gross and A. Acquisti, "Information revelation and privacy in online social networks," in *Proc. WPES'05*. ACM, 2005, pp. 71–80.
- [15] A. Acquisti and R. Gross, "Imagined communities: Awareness, information sharing, and privacy on the facebook," in *Privacy enhancing technologies*. Springer, 2006, pp. 36–58.
- [16] M. S. Bernstein, E. Bakshy, M. Burke, and B. Karrer, "Quantifying the invisible audience in social networks," in *Proc. CHI'13*. ACM, 2013, pp. 21–30.
- [17] J. Bonneau, J. Anderson, R. Anderson, and F. Stajano, "Eight friends are enough: social graph approximation via public listings," in *Proc. EuroSys '09*. ACM, 2009, pp. 13–18.
- [18] Y. Liu, K. P. Gummadi, B. Krishnamurthy, and A. Mislove, "Analyzing facebook privacy settings: User expectations vs. reality," in *Proc. IMC '11*. ACM, 2011, pp. 61–70.
- [19] S. Egelman, A. Oates, and S. Krishnamurthy, "Oops, i did it again: Mitigating repeated access control errors on facebook," in *Proc. CHI '11*. ACM, 2011, pp. 2295–2304.

- [20] H. Mao, X. Shuai, and A. Kapadia, "Loose tweets: an analysis of privacy leaks on twitter," in *Proc WPES'11*. ACM, 2011, pp. 1–12.
- [21] J. Lindqvist, J. Cranshaw, J. Wiese, J. Hong, and J. Zimmerman, "I'm the mayor of my house: examining why people use foursquare-a social-driven location sharing application," in *Proc. CHI '11*. ACM, 2011, pp. 2409–2418.
- [22] S. Mahmood and Y. Desmedt, "Poster: preliminary analysis of google+'s privacy," in *Proc. CCS '12*. ACM, 2011, pp. 809–812.
- [23] B. Meeder, J. Tam, P. G. Kelley, and L. F. Cranor, "RT@ IWantPrivacy: Widespread violation of privacy settings in the Twitter social network," in *Proc. of the W2SP'10*, vol. 2, 2010.
- [24] C. Jernigan and B. Mistree, "Gaydar: Facebook friendships expose sexual orientation," *First Monday*, vol. 14, no. 10, 2009.
- [25] R. Dey, C. Tang, K. Ross, and N. Saxena, "Estimating age privacy leakage in online social networks," in *Proc. of INFOCOM '12*, 2012.
- [26] R. Dey, Y. Ding, and K. Ross, "The high-school profiling attack: How online privacy laws can actually increase minors' risk," in *Proc. of Internet Measurement Conference '13*, 2013.
- [27] R. Anderson, *Security Engineering*, 2nd edition. Wiley, 2008.
- [28] B. W. Lampson, "A note on the confinement problem," *Communications of the ACM*, vol. 16, no. 10, pp. 613–615, 1973.
- [29] D. P. Company, *A Guide to Understanding Covert Channel Analysis of Trusted Systems*. DIANE Publishing Company, 1994. [Online]. Available: <http://books.google.com/books?id=sAo7NN1lr5sC>
- [30] M. Wolf, "Covert channels in lan protocols," in *Local Area Network Security*. Springer, 1989, pp. 89–101.
- [31] S. M. Bellovin, "Security problems in the tcp/ip protocol suite," *ACM SIGCOMM Computer Communication Review*, vol. 19, no. 2, pp. 32–48, 1989.
- [32] C. H. Rowland, "Covert channels in the tcp/ip protocol suite," *First Monday*, vol. 2, no. 5, 1997.
- [33] K. Ahsan and D. Kundur, "Practical data hiding in tcp/ip," in *Proc. ACM Workshop on Multimedia Security*, vol. 2002, 2002.
- [34] S. J. Murdoch and S. Lewis, "Embedding covert channels into tcp/ip," in *Information Hiding*. Springer, 2005, pp. 247–261.
- [35] N. B. Lucena, G. Lewandowski, and S. J. Chapin, "Covert channels in ipv6," in *Privacy Enhancing Technologies*. Springer, 2006, pp. 147–166.
- [36] K. Borders and A. Prakash, "Web tap: detecting covert web traffic," in *Proceedings of the 11th ACM conference on Computer and communications security*. ACM, 2004, pp. 110–120.
- [37] K. Kenthapadi, N. Mishra, and K. Nissim, "Simulatable auditing," in *Proc. of PODS'05*, 2005.
- [38] S. Gurses and C. Diaz, "Two tales of privacy in online social networks," *IEEE Security & Privacy*, vol. 11, no. 3, pp. 29–37, 2013.
- [39] X. Zhao, V. Schwanda Sosik, and D. Cosley, "It's complicated: how romantic partners use facebook," in *Proc. CHI'12*. ACM, 2012, pp. 771–780.
- [40] J. Y. Tsai, P. Kelley, P. Drielsma, L. F. Cranor, J. Hong, and N. Sadeh, "Who's viewed you?: the impact of feedback in a mobile location-sharing application," in *Proc. of CHI '09*, 2009.

APPENDIX

In this section, we show some of the experiments we ran in our study by showing the different settings and features that we adjusted, ordered by side channel. All side channels are marked with an asterisk (*). Each table includes the following columns:

- Feature: to show which feature we focus on.
- Actions: to show what actions will implement the feature. (This column is not applicable to the "counts" feature.)
- Operations: to show what actions have been done on this feature.
- Comments: useful comments.

The following lists show the full list of features we explored, the ones shown first are included below:

Counts. Figure 4 shows the features we checked to determine whether the displayed total is different from the number of actually displayed items.

Tagging. Figure 5 shows the settings for experiments associated with tagging.

Share. We tested the share feature to determine whether the original privacy settings are obeyed in the share chains. See Figure 6.

Due to space limitations, the following tables are not included:

Friends. The settings for experiments associated with friends.

Comment, Like. The combinations for experiments on the "comment" and "like" features.

App. There are some privacy controls associated with Facebook apps. We tested these via several settings.

Relationship. We test features related to relationship as well as other personal information to determine whether it might be leaked to someone to whom it would not normally be available.

Feature	Operation	Comments
Games	change different settings on a user's timeline, check if the number matches with the viewer's side.	A user cannot edit privacy setting of game pages individually
Places	similar to Games	N/A
Music	similar to Games	same as Games
Movies	similar to Games	same as Games
TV Shows	similar to Games	same as Games
Books	similar to Games	same as Games
Events	N/A (no number displays)	N/A
Groups	similar to Games	N/A
Notes*	similar to Games	N/A
Apps, instagram etc.	similar to Games	N/A
Photos*	similar to Games	N/A
Likes*	similar to Games	N/A

Fig. 4: Features associated with our experiments on counts.

Feature	Action	Operations	Comments
post+tagging*	status / photo / place	privacy setting→future post or set doing a post	we set one user's setting as everyone the other's setting as onlyme
post+tagging*	status / photo / place	privacy setting→post or set when doing a post	we set one user's setting as everyone the other's setting as custom
post+tagging*	status / photo / place	privacy setting→future post or set when doing a post	we set one user's setting as friends the other's setting as custom
post+tagging*	status / photo / place	privacy setting→future post or set when doing a post	we set one user's setting as custom the other's setting as onlyme
post+tagging	status / photo / place	privacy setting→future post or set when doing a post	other combinations
post*	life event	privacy setting→future post or set when doing a post	compared with status / photo / place tagging
restricted list*	privacy setting → blocking	put a friend in the restricted list	the restricted users still can view non-public contents

Fig. 5: Features associated with our experiments on tagging.

Feature	Action/feature	Operations	Comments
share* on timeline	share a post on your own timeline	post → share	cannot share a post if privacy setting of this post is custom
share* on timeline	share a post on your own timeline	post → share	if Alice shared a post to Bob which Bob can't see, Bob may figure out he has been hidden from this post.
share on timeline	share a post on your own timeline	post → share	post privacy setting + share privacy setting
share on timeline	sharing history	share → show share history	none

Fig. 6: Features associated with our experiments on sharing.