

Chapter 2

Online Social Networks: Privacy Threats and Defenses

Shah Mahmood

Abstract With over 1 billion users connected through online social networks, user privacy is becoming ever more important and is widely discussed in the media and researched in academia. In this chapter we provide a brief overview of some threats to users' privacy. We classify these threats as: users' limitations, design flaws and limitations, implicit flows of information, and clash of incentives. We also discuss two defense mechanisms which deploy usable privacy through a visual and interactive flow of information and a rational privacy vulnerability scanner.

2.1 Introduction

The level of human connectivity has reached unprecedented levels with over 1 billion people using one or more online social networks including Facebook, Twitter, YouTube, and Google+.

The immense amount of data provided and shared on these social networks may include the following information about a user: date of birth, gender, sexual orientation, current address, hometown, email addresses, phone numbers, web sites, instant messenger usernames, activities, interests, favorite sports, favorite teams, favorite athletes, favorite music, television shows, games, languages, his religious views, political views, inspirations, favorite quotations, employment history, education history, relationship status, family members, and software applications. The user also provides updates in the form of status messages or Tweets, which could include: a thought, an act, a link they want to share, or a video. All these information reveal a lot about the user, which will be of interest to various groups including governments, advertisers, and criminals.

S. Mahmood (✉)

Department of Computer Science, University College London, London, UK

e-mail: shah.mahmood@cs.ucl.ac.uk

Employers have used these social networks to hire or fire employees on the basis of their behavior on social networks [33]. Universities can screen applicants or ensure discipline by monitoring their students on social networks [7]. According to a survey by Social Media Examiner, 92 % of marketers use Facebook as a tool [40]. Phishers have improved their techniques by personalizing their schemes based on the data they acquire from social networks and these have been shown to be more useful than traditional phishing schemes [21, 36]. A woman in Indiana (US) was robbed by a social network friend after she posted on her Facebook profile that she was going out for the night [32].

Social networks, due to many such unfavorable incidents, have been criticized for breaching the privacy of their users. Both in academia and in the media, the importance of a user's privacy has been repeatedly discussed. In addition to some proposed technical solutions, there have been a vast number of initiatives to educate users so that they do not provide an excessive amount of personal information online. Facebook privacy consciousness is displayed in this reply by President Obama when a school student who wanted to become the President of the United States asked him for advice. He replied [34],

Be careful about what you post on Facebook, because in the YouTube age, whatever you do will be pulled up again later somewhere in your life . . .

Nonetheless, despite increased user awareness, we still come across numerous stories in which a user's privacy has been breached, with unfortunate consequences.

Before further classifying the threat and its solutions, we think it is essential to define privacy.

2.2 Definitions of Privacy

There is no single agreed definition of privacy in academia or in government circles. Over the course of time several definitions have been proposed. In this section we look into some of those definitions.

One of the first definitions of privacy, by Aristotle, makes a distinction between political activity as public and family as private [37].

Implicit here are boundaries that might be suggested by the walls of a family house, an assumption which is made explicit, though also modified, in a far more recent definition, that of Associate Justice John Paul Steven of the US Supreme Court. For Steven [38]:

The 4th Amendment protects the individual's privacy in a variety of settings. In none is the zone of privacy more clearly defined than when bounded by the unambiguous physical dimensions of an individual's home - a zone that finds its roots in clear and specific constitutional terms: the right of the people to be secure in their . . . houses . . . shall not be violated.

Here, the home is not the exclusive locus of privacy, but is, rather, the informing image or motif in light of which privacy in other contexts may be construed.

This is an interesting definition. The Internet has managed to blur the boundaries that would have been suggested by the walls of a house. Is the right of security in one's house also assured when users are connected through the Internet? Would this extend to usage outside the home, such as in a car?

A definition of privacy less beholden to images of hearth and home and perhaps clearer in consequence is that adopted by the Calcutt Committee, headed by Sir David Calcutt, in 1990. The committee defined privacy as:

The right of the individual to be protected against intrusion into his personal life or affairs, or those of his family, by direct physical means or by publication of information.

However, privacy on the Internet is a more complex affair than physical metaphors of intrusion and exposure can capture alone. Defense against publication of private information can protect the exposure of that information, but what if it is used, rather to produce targeted advertisements, with no publication?

William Parent provides a definition of privacy which does not rest on an implicit physical dimension, as follows [35]:

Privacy is the condition of not having undocumented personal knowledge about one possessed by others. A person's privacy is diminished exactly to the degree that others possess this kind of knowledge about him.

This definition rests on the notion of "informed consent" as defined by Aristotle [37]. On this view, for any information to be held about another, there must be documentary proof of consent. An idea of privacy breach understood in these terms thus remains very valid in the era of cloud computing. But the importance of issues of control and freedom from interference of all kinds is suggested by definition of privacy that is most widely used in the computer science community, derived from Samuel Warren's and Louis Brandies' 1890 paper [42], "The Right to Privacy," in which they refer to Judge Cooley summarizing it as consisting of the right "to be let alone." The current formulation casts this as a right to:

Control over information about oneself.

It is in this tradition of thought that Alan Westin defined privacy as an individual right, held by all people, [45]:

To control, edit, manage, and delete information about them[selves] and decide when, how, and to what extent information is communicated to others.

In the light of the concerns of this thesis a survey by Alessandro Acquisti and Jens Grossklags is salient. More than 90 % of users agreed with this definition of privacy as ownership and control of their personal information. What is more, they found that while 61.2 % of users defined privacy as a matter of personal dignity, a significant minority, 26.4 %, nonetheless thought of it as the "ability to assign monetary value to each flow of personal information" [1], suggesting personal autonomy and control, rather than inviolability, is at issue.




Fig. 2.1 Classification of the causes of users' privacy leaks

2.3 Privacy Threats

In this section we discuss four causes of privacy leaks along with real online social network examples, where possible. These four causes include: user limitations, design flaws or limitations, implicit flows of information, and clash of interests, as shown in Fig. 2.1. Let us discuss each of them in detail now.

2.3.1 Users' Limitations

Users of online social networks are (mostly) human beings. Humans have inherent limitations and flaws. It is due to these flaws that Alexander, in Malaysia, shared his atheistic views on Facebook resulting in him being sent to jail. He acted in accord with his free will, in spite of awareness of the possible risk. Similarly, Congressman Wiener shared his inappropriate pictures with a female correspondent online, resulting in a controversy leading to his resignation [3]. On many occasions, we share a lot of content on social networks without thinking about the short-term or long-term consequences of such information flow.

Human rationality in decision making is limited by the amount of time we have to make a decision; the amount of data involved; the arrangement of this data; and other cognitive limitations of the mind [4]. Two reasons for bad privacy decisions by human beings are: (1) bounded rationality and (2) limited working memory. Human decision makers should not be expected to make optimal decisions in complex scenarios, in a limited time. Similarly, when a user logs onto a social network, they are not expected to spend hours reading the privacy policies, understand the technical jargon, be aware of all risks and latest attacks, etc. With the lack of all this required information and the limited decision time, are we doing justice to the user by expecting them to make the best decision?

Human working or operant memory, the part of memory which is available for use at a particular instant, is limited [2]. A user's memory cannot hold too much information at the same time. So, whatever decision is made, it is made on the basis of whatever is present in the working part of the memory. When users log into an online social network or provide information on any other site, they are psychologically distracted in several different ways. These distractions make

the users forget about their privacy or cause them to make their privacy decision on the basis of very limited information. Spiekermann et al. [39] conducted an e-privacy experiment which suggested that people appreciate highly communicative environments and thus forget about their privacy concerns. One explanation of this could be that within this communicative environment, the working memory of the user is flooded with marketing information. Privacy consciousness is removed from operational memory, and the user ends up providing sensitive information without noticing it.

Humans are considered to be inherently wired to trust other human beings. This trusting factor enables an adversary to engage in social engineering attacks. Social engineering is the art of manipulating another person. Here, a victim allows the attacker into his sphere of trust and provides him with information or access that the attacker would not have otherwise been authorized to get. It is due to human trust and social engineering attacks that users have been found to add large numbers of strangers to their Facebook accounts, enabling these strangers to view their personal content, intended to be shared with friends only.

Boshmaf et al. [9] launched a social bot attack against users on Facebook. They created social bots¹ for this purpose, where each social bot sent 25 friend requests to randomly selected users per day. They limited the per day requests to avoid the social bots being detected by Facebook's Immune System (FIS) which protects Facebook from threats and prevents malicious activity in real time. In the first phase, they sent 5,053 friend requests over 2 days. 2,391 of the requests were sent from male social bots and 2,661 from female social bots. Over a period of 6 days, 976 of their requests were accepted, with an acceptance rate of 19.3 %. For the male social bots the acceptance rate was 15.3 % and for the female social bots it was 22.3 %. Then over the next 6 weeks they requested another 3,517 users from the extended neighborhood of those who already accepted their requests; 2,079 of their requests were accepted, putting the average acceptance rate at 59.1 %. They noted that the acceptance rate increased with an increased number of mutual friends. In total they added 3,055 Facebook users, using 102 social bots, over a period of eight weeks, to demonstrate that 35.6 % of their friend requests were accepted. Social bots have been previously used by criminals and are available online for a few dollars.

Mahmood and Desmedt [30] introduced the targeted friend attack. For this attack they created a pseudonymous profile. During the first phase, which lasted for 99, days, they sent 255 friend requests out of which 90 were accepted. They did not allow the receipt of any friend requests during this time. From Day 100 to Day 285 they started accepting friend requests. They sent a total of 595 friend requests in total, out of which 370 were accepted, resulting in an acceptance rate of 62 %, as shown in Fig. 2.2. The interesting part of their experiment was that they received another 3,969 friend requests, which is 6.67 times more than the number of users they requested. The received friend request distribution is shown in Fig. 2.3. In total, their single pseudonymous profile had access to the private data of a 4,339 users.

¹Social bots are bot nets on social networks.




Fig. 2.2 (a) Probability of acceptance of our friend requests on Facebook, (b) Total number of friends requests sent and total accepted [30]




Fig. 2.3 (a) Probability of acceptance of our friend requests on Facebook, (b) Total number of friends requests sent and total accepted [30]

Similarly, an attacker could create profiles for a famous scientist. For example, someone has created a Facebook profile for Claude Shannon. It has got 180 friends.² Moreover, there are Facebook profiles made for animals.³

²<http://www.facebook.com/claude.shannon>—Checked on September 12, 2012, Shannon’s profile may have been made by a fan to pay tribute but similar approaches can be used by attackers.

³The profile of a cat at the first author’s previous residential hall has 170 friends. This could again be used to spy on users—<http://www.facebook.com/CelebratingStarlight>.

Human limitations in privacy decision making is not limited to social networks. Despite the awareness that loyalty cards by supermarkets can be used to track the behavioral attributes of a user and launch correlation attacks against them, users sign up for them even when the benefits are minimal. For example, the author signed up for a Tesco Club Card in the United Kingdom and regularly shopped at the store for his groceries for a period of thirteen months only to get £7 worth of cash coupons. In this little experiment the author was paid only £7 to allow the supermarket track his eating habits for a period of over 1 year. Users have also been found selling their DNA samples for a McDonald's Big Mac burger or writing their PINs on their credit/debit cards. In the latter case, when they lose the card, they have already attached the secret to it, making the job of the adversary much simpler.

Lack of privacy and over-sharing on social networks has also affected criminals. Grasso, a drug fugitive from Italy was caught in London after he posted photos about his life in London. He posted photos with the wax figures of President Obama and Prime Minister Cameron, taken at London's Madame Tussauds [13].

These are certainly not examples of rational decision making. However, they may be classified as rational within the inherent bounds and limits of humans at the instant of that decision.

2.3.2 Design Flaws and Limitations

Having looked at the limitations of users, let us now turn to the design flaws and limitations in some social networks. The design limitations and flaws include the weak privacy controls by social networks and the possibility of explicit attacks including cloning attacks.

Facebook Privacy Setting Evolution and Devolution

Facebook was initially launched as a student social network, with its approach to privacy being initially network-centric, which meant all data shared by users was visible to all the members of the network [48]. A network could have been based on an academic institution or a city. In the former case the profile of a user would be visible to all his fellow students within that institution, while in the latter case, his profile would be visible to everyone who selected that city as their city of residence. As the network grew and its users increased to millions, the privacy settings were changed several times until they reached their present form, in which by default different levels of user information are visible to "Friends," "Friends of Friends," and "Everyone." Today, Facebook does allow a user the option of sharing information only with himself through an "Only Me" option in the privacy settings. It also allows making an exception for specific groups of people, called "lists." These lists are very similar to Google+ circles [26]. It has been shown in the literature that users rarely change the default settings [8].



Fig. 2.4 Reconstructing friend-list on Facebook from wall posts

As Facebook’s default settings evolved to a level where not all information was visible to a user’s network, attackers and researchers tried to glean information from what was publicly available. A new area of social graph extraction, where by the communities and friends of users were extracted, emerged, e.g. [7,22,43,46]. Although the information this provides is limited when compared to all the information provided by users on their profile, it still reveals a lot. In 2009, Facebook showed a list of eight friends of a user when he was searched for either within the network or through an external search engine by someone who was not their friend. These eight friends were selected on the basis of some internal algorithm and would change every time the searchers’ cache was reset. According to Facebook, revealing eight friends would help a searcher decide whether they have found the person they were looking for. If they were the person the searcher was looking for, the searcher could add them, and if they were not that person, then the searcher could look into the next user returned by the search. However, another consequence was that using this information, researchers were able to approximate the social graph of users on Facebook [8]. Facebook no longer displays the list of eight friends to a searcher.

For added privacy, Facebook users have the option to restrict who can view their friend-list, but, this does not mean a friend attacker⁴ cannot reconstruct that user’s friend-list [24]. For at least a partial reconstruction, a friend attacker can enumerate the names/user IDs of all the users who comment on posts visible to friends only. In Fig. 2.4, even though the user’s friend-list is not visible to the author, we are able to find the names of at least seven friends of the victim. These friends liked the post. Any users who make comments on these posts will also be visible. By analyzing more posts, over a longer duration of time, an attacker can find the names and user IDs of more friends of the victim.

Similarly, when a user is tagged in a photo, we can see the name of the person who tagged the user by rolling the mouse over their name. It displays “Tagged by” and the tagger’s name. As only a user’s friends are allowed to tag them on Facebook, this also helps in reconstructing the friend-list. Moreover, Facebook does not allow users to hide their mutual friends. These can also be added to the reconstruction of

⁴A friend attacker is an attacker who is a friend on Facebook.

the victim's friend-list. This way the attacker can reconstruct a very significant part of a user's friend-list.

Timeline, a new virtual space in which all the content of Facebook users is organized and shown, was introduced on December 15, 2011 [19]. In addition to re-organization of users' content, Timeline comes with some default and unchangeable privacy settings. First, it is no longer possible for a Facebook user to hide their mutual friends, which was possible before Timeline. Second, it is not possible to limit the public view of "cover photos." These cover photos could be a user's personal pictures or political slogans and their widespread sharing may have various short-term and long-term consequences for that user. Third, with Timeline, depending on the user's privacy settings, if the likes and friend-list of a user are shared with a list of users, then that list of users can also see the month and the year when those friends were added or when the user liked those pages. This will allow an attacker to analyze the sentiments and opinions of a user, e.g. when did a user start liking more violent political figures and un-liking the non-violent ones. Finally, with Timeline, if a user makes a comment on a page or a group, he does not have the option to disable the probability of being traced back to the profile. Before Timeline, a user could make themselves searchable by a specific group (e.g., "Friends" or "Friends of friends") and even if they commented on pages and groups, people outside those allowed groups would not be able to link back to the commenters profile. Facebook can solve these problems by allowing users to change the settings to share their content with their desired audience.

Fake Accounts and Cloning Attacks

Social networks do not currently offer any certification or verification of the authenticity of user accounts. Thus, it is very easy for an attacker to register accounts in the name of someone else, although it is prohibited by the privacy policies of most service providers. The act of creating fake accounts is also known as a sybil attack.

An attacker can use personal information, e.g. pictures and videos of the victim, on the fake profile to win the trust of his friends and let them allow the fake account into their circle of trust. This way the attacker will have access to the information of the friends of his victim, which his friends have agreed to share with the victim and not necessarily the attacker.

The process of creating fake accounts, is called a "cloning attack," when the attacker clones (creates almost exact copies) of real social network accounts and then adds the same and/or other contacts as their victim. Bilge et al. [5] showed the ease of launching an automated identity theft attack against some popular social networks by sending friend requests to friends of hundreds of their cloned victims.

Permanent Takeover of a Facebook Account

Facebook allows a user to recover their compromised account using several verification mechanisms, but they all fail if the attacker disassociates the victim's current




Fig. 2.5 Basic concept of the deactivated friend attack

login email address from the victim’s account and associates it with a new dummy account [24]. The attacker can associate a new email address with the victim’s original account and permanently take it over. As an example, let us suppose the victim is “Bob” and his login email address is “bob@victim.com.” The attacker will associate the original account of Bob with a new email address “attacker@hacked.com” and associate a new account named “Dummy” with “bob@victim.com.” Thus, when Bob recovers his account, he gets access to the Dummy account. Presently, there is no way to recover an account after such an attack.

Facebook should not allow the association of used email addresses with new accounts. This would prevent the “permanent takeover” attack.

Facebook’s Deactivated Friend Attack

*Deactivated friend attack*⁵ occurs when an attacker adds their victim on Facebook and then deactivates his own account. As deactivation is temporary in Facebook, the attacker can reactivate his account as he pleases and repeat the process of activating and deactivating for an unlimited number of times. While a friend is deactivated on Facebook, he becomes invisible. He cannot be unfriended (removed from a friend’s list) or added to any specific list. The only privacy changes that may apply to him are those applied to *all* friends, or to the particular list of which he is already a member.

This deactivated friend, i.e. the attacker may later reactivate the account and crawl his victims profiles for any updated information. Once the crawling has finished, the attacker will deactivate again. While activated, the attacker is visible on the victim’s friend list. The concept here is very similar to that of cloaking in Star Trek where Badass Blink or Jem’Hadar has to uncloak (be visible), even if only for a moment, to open fire. Figure 2.5 provides the abstract conceptual view of the attack. Facebook provides no notification of the activation or deactivation of friends to its users.

⁵Based on our paper [30].

The deactivated friend attack is only possible due to the presence of a cloaked channel, which is defined below.

Definition 1. *Cloaked Channel:* A channel is called a cloaked channel if, and only if, it is invisible when cloaked, and reappears when uncloaked [30].

This kind of attack is very serious for several reasons. First, it is very hard to detect this kind of attack. The attacker could activate his account at times when he is least likely to be detected and crawl his victims' profile for information, with which to update his records. Various groups of information aggregators could find this attractive as a permanent back door to the private information of Facebook users, including: marketers; background checking agencies; governments; hackers; spammers; stalkers; and criminals. Second, it continues to be useful even if a user becomes more security conscious. He may want to adjust his privacy settings but will not be able to affect his attackers access, unless he applies an update to all his friends. Third, by closely monitoring a few users on Facebook, an attacker can get a deeper insight into a large network. This possibility has been enhanced by Facebook's addition, last year, of a "browsing friendship" feature. This would help the attacker in analyzing the bond between two of his victims by browsing their friendship which provides information including: the month and year since which they have been Facebook friends; events they both attended; their mutual friends; things they both like; their photos; the messages they have written/write on each others' walls, etc. This would give a very deep insight into the level of their relationship, the intensity of their activity at a particular time, the degree of their interactivity, etc. This information could be used for several attacks including social engineering and social phishing attacks. This vulnerability was fixed by Facebook using one of our solutions [30]. Now the deactivated friends are visible to a user on their friend-list and a user can un-friend them or add them to another list, possibly with a limited view.

Google+ Photo Metadata

When a user uploads a photo on Google+, some metadata are made available to those with whom the photo is shared, including: the name of the photo owner; the date and time the photo was taken; the make and model of the camera, etc.⁶ This set of information, in particular the date and time, may at first seem relatively innocent and trivial, but could in reality lead to some serious privacy concerns. On August 10, 2007, in Pennsylvania (USA), a divorce lawyer proved his client's spouse to have been unfaithful, when electronic toll records showed him in New Jersey (USA) on a particular night and not in a business meeting in Pennsylvania as he had claimed [16]. With the metadata revealed by Google+ a user might leak enough information to be legally held liable on a similar basis.

⁶Based on our paper [26].




Fig. 2.6 Metadata from a photo by Larry Page on Google+ [26]

Similarly, the make of the camera with which a photo was taken could be another concern for privacy. Higher end cameras cost thousands of dollars. There have been past incidents where victims were killed for their cameras. In May 2011, a Greek citizen, 44, was killed for his camera when taking his wife to hospital for the birth of their child [44].

Just to give an example of the level of information a picture exposes about the camera, look at the metadata of the publicly shared pictures (from his Google+ profile) of Google co-founder Larry Page, shown in Fig. 2.6.⁷ It reveals that he used a Canon EOS 5D Mark II to shoot his vacation photographs. This camera is worth approximately USD 2400. This gives the robber incentives.

Zuckerberg's Photo Leak and Other Attacks on Facebook

In December, 2011, Facebook founder Mark Zuckerberg's photos were leaked by a relatively simple, presumably accidentally discovered, vulnerability in Facebook protocols [18]. The vulnerability was exploited when a user reported a victim's display photo as nude. Facebook responded by sending more photos of the victim, even from his private albums asking the reporting party, whether they were also nude. It took the leak of the personal photographs of its founder for Facebook, to fix

⁷Photo modified for didactic purposes.




Fig. 2.7 Helping Facebook identify information leakage in their protocols [29]

the bug, which would have been far easier to identify and rectify if the protocol had been visualized as shown in Fig. 2.7. The vulnerability was fixed after the leak.

Similarly, Felt [20] presented a cross-site scripting vulnerability in the Facebook Markup Language which allowed arbitrary JavaScript to be added to the profiles of users of an application, which led to session hijacking. Facebook by default uses HTTP which helps the attacker with traffic analysis. The attacker does not have to decrypt the packet sniffed in transit. Dhingra and Bonneau independently provided limited hacks into Facebook photos [6, 15].

2.3.3 *Implicit Flows of Information*

Leakage of information on social networks does not always need to be explicit. There may be many more implicit flows of information, where one piece of information may leak information about another. For example, Dey et al. [14], in a study of 1.47 million user accounts on Facebook, found that only 1.5 % of users revealed their age. Using the high school graduation year and friend connections of these users, they were able to estimate the ages of 84 % of these users with a mean error of plus or minus 4 years.

Similarly, the “likes” of a user on Facebook or the videos that a user might watch on Facebook enable an attacker to deduce lots of implicit information about that user [10]. On Facebook we found 233 people liked Breivik’s Wikipedia page,⁸ and

⁸<http://www.facebook.com/pages/Anders-B-Breivik/265349460157301> Accessed: July 13, 2012.

according to Facebook 13 people were talking about him when we checked that page. We also found 954 people liked Al-Awlaki's Wikipedia page⁹ and according to Facebook 30 people were talking about him at that time [25]. These people may be much more likely to commit a violent act of terror than other Facebook users, on average. Moreover, "likes" of a particular brand of chocolate or medicine can also reveal a lot. For example, a user writing reviews about sweet chocolates is unlikely to be diabetic while someone who "likes" "Accu-Check," a brand of diabetes testing kit, is more likely to be either diabetic or closely involved with diabetic patients. Furthermore, a user who mostly Tweets in the evenings, and never in the morning, is likely to be either busy in the morning or sleeping. These are just a few examples. There is likely to be an enormous number of similar correlational flows of information. Unfortunately, this type of privacy leak is the hardest to defend against, because it adds an enormous number of not-clearly-known information flows to the threat model, which might otherwise be comprised of explicit flows of information to an adversary.

2.3.4 *Clash of Interests*

Most social networks, like other web services, are supported by the revenue generated from advertisements. This may create a conflict of interest on the part of the service provider when it comes to the rights of its users and the rights of advertisers. Without the advertiser's money, the service provider would not be able to function. Similarly, without a proper user base, it won't be able to generate lots of advertising money.

The users' interests and those of the advertisers do not always match. Users want their data to be inaccessible to anyone except the parties they have explicitly permitted under a well-documented, human-readable, and informed type of consent. Moreover, users want data to be used only for the exact purpose they have consented to. From a user's point of view, any deviation from these terms by the service provider is considered less than wholly honest at best, and actively dishonest at worst.

Advertisers, on the other hand, want to mine the maximum amount of data about a user over a long duration of time. They desire to keep the data for an indefinite time. More data will provide them with a competitive advantage and help advertisers in crafting better targeted advertisements. These well-crafted advertisements will increasingly influence users' choice of products resulting in making the advertisers happier. Thus, the advertisement networks will be able to

⁹<http://www.facebook.com/pages/Anwar-al-Awlaki/102248169830078?rf=134722733227304>
 Accessed: July 13, 2012.

charge their customers higher fees in the future. Zuckerberg confirmed the threat from advertising companies during an interview by saying [31]:

I think that these companies with those big ad networks are basically getting away with collecting huge amounts of information ...

Another issue of interest concerning advertisers is the use of behaviorally targeted advertising [47]. One may argue that personalization is useful for a person who is too busy to search through normal banner or textual ads, but it can have some seriously concerning implications. Suppose a neutral user stumbles across a right-wing article and opens it. Then, the next day the search engine recommends to her two such articles, and being curious the user opens them. After a few days all her search results may be those meant for extremely radicalized people, with advertisements pointing her to products that may potentially help her act on the radical views introduced to her by behavioral targeting. We wonder how many people have already been radicalized as a side effect of behavioral targeting, as the process is still ignored due to the monetary gains it makes for large corporations.

Considering the discussion above, it is hard to trust service providers when there is a conflict of interests. Computer scientist Lanier sums this up by saying:

Facebook says, 'Privacy is theft,' because they're selling your lack of privacy to the advertisers who might show up one day.

Social networks like Facebook charge 30 % of a third party application's revenue, for using its ecosystem [17]. These applications have their own privacy practices, which are not necessarily in line with those of Facebook. Facebook in their "terms of use" agree that they have no control over the data that is being stored by third party applications. Users have not necessarily given informed consent to such practices. Facebook will not risk losing some of the famous gaming applications, which provide a large chunk of Facebook's revenue, by asking them for better privacy practices. Such application providers may then further sell the data to potentially malicious parties without any control by the user, Facebook, or the application provider.

In a very similar way to advertisers, governments want more private information about people in order to gauge their sentiments and opinions. Facebook has recently announced that it will help advertisers by letting them target their adverts to users on the basis of their phone numbers. Advertisers using this mechanism can direct advertisements to customers who have previously used a service and provided that service with their phone numbers. Service providers, in several countries, are bound to provide governments with users' personal information, and in some cases are legally prohibited from informing users about such sharing of information.

This creates a situation where service providers either are forced by law to monitor users' actions on social networks or are economically incentivized to collect and store more information about users for advertisement and a government's profiling purposes.

2.4 Defense Mechanism

Anonymity, pseudonymity, and unlinkability [11, 12] have been proposed as some of the well-researched solutions for privacy protection; yet, there are dozens of attacks proposed against some of the best systems being researched and used. Another approach proposed and adopted in some quarters is to ensure users' privacy through regulations. Unfortunately the time taken for regulations to be drafted and become effective may exceed the valuable life of information [39]. There is limited research on how to quantitatively analyze one's vulnerability to different attacks. Moreover, there is the problem of what happens if a particular set of information about a user is leaked. A user cannot realistically be expected to be careful about their privacy protection at all times, considering a user's inherent limitations, including bounded rationality and limited working memory. If we cannot realistically expect a user to protect their privacy at all times, is there another mechanism that might provide users with more rational options in a limited time? And is there a model which could help a user identify whether it is possible for her to annul the effect of at least some of the leakage of information, so as to substantially minimize any adverse effects?

In this section we discuss two such proposed tools for users' privacy protection.

2.4.1 *Usable Privacy Through Visual and Interactive Flow of Information*

The world of computer technology has evolved from stroking keys to point and click, and recently to touch and tap. Users need a more graphical and interactive view of developments in respect to privacy, so that they may better understand the currently obscure world of communications, up and down the stack.

Facebook, in January 2011, offered users an opt-in to secure browsing. Prior research has shown that users are reluctant to make changes to the default settings [23]. Normally it is too cumbersome for the user to search into all the options and visualize the impact of switching between them. It will be much easier for the user if they are provided with a view as shown in Fig. 2.8. With this user-friendly view, users will be encouraged to change from their current settings to another without making too much extra effort.¹⁰

Facebook, by default, emails notifications about almost any activity relevant to the user including any friend requests received, friends commenting on the users' wall, any photos in which the user has been tagged, any new message for the user on Facebook, etc. These email notifications contain links to the activity in question. Phishers also send users emails with links seemingly coming from legitimate

¹⁰This section is based on our paper [29].




Fig. 2.8 Easier, graphical and interactive switch to secure browsing [29]




Fig. 2.9 Identify possible flows of information that may aid in social engineering attacks [29]

sources, but actually directing the user to the attackers' pages. These attacker pages steal their victims' credentials or launch a man-in-the-middle attack. When users regularly receive similar emails from Facebook it is easier for them to fall victim to phishers claiming to be Facebook. It is possible to turn off these email notification and thus reduce the risk of phishing attacks, but currently the user needs to search into several levels of options by first clicking the downward pointing arrow in the top right corner of the Facebook page, then clicking on "Account Settings" from the pull down menu, followed by clicking the third option in the left side labelled "Notifications," then clicking on "Facebook" under the "All Notifications" in the center of the page and, finally unchecking the desired notifications. It would be much more user-friendly if the flow were represented as shown in Fig. 2.9, and if deactivation were possible with a click on the graphic.

When Google+ was launched, one aspect that was widely appreciated, and marketed as a privacy feature, was the concept of circles. A user could divide his social network contacts into different circles and share content on a circle




Fig. 2.10 Visualizing Facebook updates and providing users with interactive control [29]

by circle basis. Though Facebook termed all the user's social network contacts as "Friends," there was a possibility of classifying users into lists [26]. Google+ circles were appreciated more than Facebook lists (which were widely unknown to users) due to the graphical and interactive un-usability of the latter. In the real world, relationships with people change with time, i.e., some people get closer and others get emotionally distant. It is hard for users to keep on reflecting such relationship changes in our social network presence. A better way would be to visualize information flow on the basis of each update, as shown in Fig. 2.10. The picture on the left is that of the user who is updating his Facebook profile. The information is flowing, shown by the directed arrows, to the list of those friends of the user shown on the right, with whom the information is currently being shared, each represented by a thumbnail picture. Since some of the pictures may have been uploaded recently and not be readily recognizable by the user, the user can move the pointer over any of the arrows or pictures to get additional information, such as their name. Clicking on any of the arrows will disconnect the flow of that particular information to the specific person and save the user from regretting the sharing of private and embarrassing information.

These are only a few examples where the graphical and interactive view of information flow can be considered as a more usable means of ensuring users' privacy. In practice, the visual and interactive view may identify many currently obscure sources of leaks of information.

Advantages of Graphical and Interactive View of Information Flow

The use of graphical and interactive view of information has several advantages including the following:

Countering Limitations of the Working Memory. When a user logs into a social network or provides information on any other website, he is psychologically distracted in many ways. These distractions make users less conscious about privacy. With a graphical flow of information available, a user can be reminded of what information is being shared and with whom. This brings a concern about privacy back into the working memory.

Countering Bounded Rationality Problem. Human rationality in decision making is limited by the amount of time there is to make a decision, the amount of data available, the arrangement of this data, and the cognitive limitations of the mind. Decision makers should not be expected to make optimal decisions in complex scenarios, in limited time. When a user visits a web site, he should not have to spend hours reading the privacy policy, understanding the technical jargon, getting aware of all the parties with whom the information will be shared, etc. A more graphical representation of his information flow can put him in a better position for quick decisions.

As a Protocol Verification Technique. The graphical and interactive view of information flow can help privacy experts, allowing service providers testing a unit to identify leaks of information.

Transparency and User's Increased Trust in the System. Using graphical representation of information flow, users will feel more in command of their private data. This will build their trust and confidence in the system giving them better control when sharing information.

2.4.2 Rational Privacy Vulnerability Scanner

Users can be provided with a rational privacy vulnerability scanner, which is based on stochastic almost combinatorial games. Stochastic almost combinatorial games are stochastic games in a combinatorial setup. Here a user is shown the threat level to his privacy by assigning cost and probability to the possibility of an attacker accessing a user's information that they may share with a limited set of users on a social network. As a motivating example, let us suppose an attacker wants to make a hotel reservation with the credit card details of a victim. There are normally three units of information required for this attack to be successful, namely, the name of the cardholder (1), the address to which it is registered (2), and the card number (3) (occasionally the card expiry date and the three digit security code may be required, but these are not considered in this example). The attacker may acquire this information in various orders and "units".

One possible method for attacker is to buy all three units of information from an underground market [41] for a certain cost, $c_{\{\} \rightarrow \{123\}}$ (cost of finding all three units of information from the initial state). This case is shown in Fig. 2.11j. The details bought may not always be valid, and thus there is a probability of success attached, represented by $p_{\{\} \rightarrow \{123\}}$. The costs are incurred even if the details are invalid, as shown by the loop back arrow. The probability of the information being incorrect would be $1 - p_{\{\} \rightarrow \{123\}}$. In this case, this is because of the illegal nature of underground markets, but even information trading companies may not guarantee accurate and up-to-date information.

Another possibility is that the attacker finds a card, with name and number, at an ATM machine. The cost $c_{\{\} \rightarrow \{13\}}$ here is zero and the probability that the card is not yet cancelled is $p_{\{\} \rightarrow \{13\}}$. This case of getting all 3 units of information in one go is shown in Fig. 2.11h. The attacker now needs to find the correct address, using the information she has. The address from a name could be found using social networks like Facebook, people-search services like www.123people.com or by other means. This information again will have a cost (time is money) and have a certain probability of being valid. There are several other possible combinations, as shown by the rest of the cases in Fig. 2.11. The attacker needs to find the most feasible path, where the determination of feasibility depends on the criteria of the attacker.

It is also possible that the victim comes to know about the compromise of her information. Suppose the victim knows that all three units of her information are compromised. To render an attack useless, the victim will need to annul only one bit of information, although changing name or address might cost more than ordering a new card with a new number. So the victim can push the attacker from state $s_{\{1,2,3\}}$ to a state $s_{\{1,2\}}$, thus diminishing the probability of an attack.


In this example, there were only three relevant units of information, and we were able to find the most practical transitions without any external aid. But when the number of relevant information units increases, the number of possibilities of transitions between the states also increases, creating the need for a framework to aid the user in such decisions.

The graph presented from the viewpoint of the attacker is called the *attack graph* while the graph presented from the viewpoint of the defender is called the *defense graph*. For three units of information Fig. 2.12 shows an attack graph and Fig. 2.13 shows a defense graph, with edges labelled with cost and probability of success. Both the views combined will create provide us with a stochastic almost combinatorial game scenario.

We envision a tool, which provides users with an estimated cost that an attacker will have to incur in order to acquire his information and launch an attack if he uploads that information on social networks. The tool can use the economic models in [28].

The aims of the scanner can be to mine all possible units of relevant information, gather it in one place, create links between the different units of information and provide a threat report. Using the provided data it can show the cost and success probabilities of any possible remedies, in case of a privacy breach. The tool can be

Fig. 2.11 All possible ways to find 3 units of data






Fig. 2.12 Attack Graph for 3 units of information




Fig. 2.13 Defense Graph for 3 units of information

used as a guide on what information to avoid giving in cases where provision of some information is optional. This will afford a ready answer to the question: How much do we lose by providing such non-compulsory units of information?

The dichotomy between privacy valuation and experimental observation is widely noted in literature. This dichotomy is observed under constraints of limited

working memory and the bounded rationality of users. An interesting direction to take will be to use the framework outlined here and repeat some of the tests performed by Acquisti and Grossklags in [1]. Our hypothesis is that with the increased working memory and extended bounds of rationality due to the use of these tools, a user's dichotomous behavior may be reduced.

2.5 Summary

In a nutshell, in this chapter we discussed four different causes of privacy leaks in online social networks. These four causes include: users' limitations; design flaws and limitations; implicit flow of information; and clash of interests. Then we discussed two defense mechanisms involving the visual and interactive control of flow information, and economic modeling of the privacy threat, thus providing users with the means to make a more rational choice. Interested readers may follow the given references for a detailed review of the provided threats and defenses.

Acknowledgements The author would like to thank Professor Yvo Desmedt with whom he co-authored some of the work cited in this chapter. The author would also like to thank University College London for providing him financial support through the University College London PhD Studentship Program. This chapter is based on our work in [24–30].

References

1. Acquisti, A., Grossklags, J.: Uncertainty, ambiguity and privacy. In: WEIS, 2005
2. Baddeley, A.: Working memory. *Science* **255**(31), 556–559 (1992)
3. Barret, D., Saul, M.H.: “weiner now says he sent photos”. *Wall St. J.* (2011)
4. Berger, P.L.: *Models of Bounded Rationality*, Vol. I–III. MIT Press, Cambridge, MA (1982)
5. Bilge, L., Strufe, T., Balzarotti, D., Kirda, E.: All your contacts are belong to us: automated identity theft attacks on social networks. In: WWW, pp. 551–560, 2009
6. Bonneau, J.: New facebook photo hacks. <http://www.lightbluetouchpaper.org/2009/02/11/new-facebook-photo-hacks/>, (2009). Accessed 15 July 2011
7. Bonneau, J., Anderson, J., Danezis, G.: Prying data out of a social network. In: ASONAM, pp. 249–254, 2009
8. Bonneau, J., Anderson, J., Stajano, F., Anderson, R.: Eight friends are enough: Social graph approximation via public listings. In: SNS, 2009
9. Boshmaf, Y., Musluhkov, I., Beznosov, K., Ripeanu, M.: The socialbot network: When bots socialize for fame and money. ACSAC, Sept 2011
10. Chaabane, A., Acs, G., Kaafar, M.: You are what you like! information leakage through users' interests. In: Proc. Annual Network and Distributed System Security Symposium, 2012
11. Chaum, D.: Untraceable electronic mail, return addresses, and digital pseudonyms. *Comm. ACM* **24**(2), 84–88 (1981)
12. Chaum, D.: Blind signatures for untraceable payments. In: CRYPTO, pp. 199–203, 1982
13. Cooper, B.: Italian drugs fugitive jailed after posting pictures of himself with Barack Obama waxwork in London on Facebook. *Mail Online* February 14, 2012
14. Dey, R., Tang, C., Ross, K.W., Saxena, N.: Estimating age privacy leakage in online social networks. In: INFOCOM, pp. 2836–2840, 2012

15. Dhingra, A.: Where you did sleep last night? ...thank you, i already know! *iSChannel* **3**(1) (2008)
16. Donald, A.M., Cranor, L.F.: How technology drives vehicular privacy. *J. Law Pol. Inform. Soc.* **2**, (2006)
17. Ebersman, D.A.: Facebook Inc., Form S-1 registration statement. United States Securites and Exchange Commission, February 1, 2012
18. Facebook bug sees Zuckerberg pictures posted online. *BBC*, December 7, 2011
19. Facebook Timeline: <http://www.facebook.com/about/timeline>. Accessed 16 May 2012
20. Felt, A.: Defacing Facebook: A security case study. 2007
21. Jagatic, T.N., Johnson, N.A., Jakobsson, M., Menczer, F.: Social phishing. *Comm. ACM* **50**(10), 94–100 (2007)
22. Lindamood, J., Heatherly, R., Kantarcioglu, M., Thuraisingham, B.M.: Inferring private information using social network data. In: *WWW*, pp. 1145–1146, 2009
23. Mackay, W.E.: Triggers and barriers to customizing software. In: *CHI*, pp. 153–160, 1991
24. Mahmood, S.: New privacy threats for Facebook and Twitter users. In: *IEEE 3PGCIC*, 2012
25. Mahmood, S.: Online social networks: The overt and covert communication channels for terrorists and beyond. In: *IEEE HST*, 2012
26. Mahmood, S., Desmedt, Y.: Poster: preliminary analysis of Google+’s privacy. In: *ACM Conference on Computer and Communications Security*, pp. 809–812, 2011
27. Mahmood, S., Desmedt, Y.: Online social networks, a criminals multipurpose toolbox (poster abstract). In: Balzarotti, D., Stolfo, S.J., Cova, M. (eds.) *Research in Attacks, Intrusions, and Defenses*, vol. 7462 of *Lecture Notes in Computer Science*, pp. 374–375. Springer, New York (2012)
28. Mahmood, S., Desmedt, Y.: Two new economic models for privacy. In: *ACM SIGMETRIC-S/Performance Workshops, PER*, 2012
29. Mahmood, S., Desmedt, Y.: Usable privacy by visual and interactive control of information flow. In: *Twentieth International Security Protocols Workshop*, 2012
30. Mahmood, S., Desmedt, Y.: Your Facebook deactivated friend or a cloaked spy. In: *IEEE PerCom Workshops*, pp. 367–373, 2012
31. MailOnline: Zuckerberg defends Facebook...by saying Microsoft, Google and Yahoo! are even worse at ignoring user privacy. *Daily Mail*, November 8, 2011
32. Henderson, M., de Zwart, M., Lindsay, D., Phillips, M.: Will u friend me? Legal risks of social networking sites. *Monash University*, 2011
33. Monkovic, T.: Eagles employee fired for Facebook post. *New York Times*, March 10, 2009
34. Obama advises caution in use of Facebook. *Associated Press*, September 8, 2009
35. Parent, W.: Privacy, morality and the law. *Philos. Publ. Aff.* **12**, 269–288 (1983)
36. Polakis, I., Kontaxis, G., Antonatos, S., Gessiou, E., Petsas, T., Markatos, E.P.: Using social networks to harvest email addresses. In: *WPES*, pp. 11–20, 2010
37. Privacy: *Stanford Encyclopedia of Philosophy*, 2002
38. Samaha, J.: *Criminal Justice*. Thomson Wadsworth, Belmont, CA (2006)
39. Spiekermann, S., Grossklags, J., Berendt, B.: E-privacy in 2nd generation e-commerce: privacy preferences versus actual behavior. In: *ACM Conference on Electronic Commerce*, pp. 38–47, 2001
40. Stelzner, M.: Social media marketing industry report. <http://www.socialmediaexaminer.com/SocialMediaMarketingReport2011.pdf>, 2011
41. The underground credit card blackmarket. <http://www.stopthehacker.com/2010/03/03/the-underground-credit-card-blackmarket/>, 2010
42. Warren, S.D., Brandeis, L.D.: The right to privacy. *Harv. Law Rev.* **4**(5), 193–220 (1890)
43. Wasserman, S., Faust, K.: *Social Network Analysis*. Cambridge University Press, Cambridge (1994)
44. Weeks, N.: Greek police detain 24 in athens immigrant clash after murder. <http://www.bloomberg.com/news/2011-05-11/greek-police-detain-24-in-athens-immigrant-clash-after-murder.html>, 2011
45. Westin, A., Blom-Cooper, L.: *Privacy and Freedom*. Bodley Head, London (1970)

46. Xu, W., Zhou, X., Li, L.: Inferring privacy information via social relations. In: International Conference on Data Engineering, 2008
47. Yan, J., Liu, N., Wang, G., Zhang, W., Jiang, Y., Chen, Z.: How much can behavioral targeting help online advertising? In: WWW, pp. 261–270, 2009
48. Yardi, S., Romero, D.M., Schoenebeck, G., Boyd, D.: Detecting spam in a Twitter network. *First Monday* **15**(1) (2010)

Security and Privacy Preserving in Social Networks

Chbeir, R.; Al Bouna, B. (Eds.)

2013, XVI, 367 p. 107 illus., 66 illus. in color., Hardcover

ISBN: 978-3-7091-0893-2