

Friendship and Privacy in Online Social Networks

Soudeh Ghorbani

Department of Computer Science
University of Illinois at Urbana Champaign

Yashar Ganjali

Department of Computer Science
University of Toronto

Abstract—Online social networks (OSNs) have become extremely popular in recent years. Users actively interact in these networks and share large amounts of personal information. This has led to emergence of a treasure trove of data for many entities, from marketers and spammers to employers and intelligence agencies, which has become a serious privacy concern. Previous works have addressed many aspects about privacy in OSNs such as characterizing potential privacy leakage [16], possible ways for inferring sensitive private information [9], [20], and appropriateness of default privacy settings [13]. In contrast, we focus on the entity who plays the main role in guarding privacy: the user. By sending out friend requests to unknown users in one of the largest OSNs, we provide evidence that a considerable portion of OSN users are willing to let a stranger, possibly an adversary, into their social network, thus granting her access to the users' personal information and to some extent to those of their friends. We study several factors that might foster such behavior, and measure the amount of information that will consequently become accessible. We find that for more than 95% of the users who accept our friend requests, we gained access to personal information that would not otherwise be accessible. We also show that the majority of the users who accept the requests have indeed changed their default privacy settings to restrict access to some parts of their personal information to their friends while making them publicly inaccessible. Our observations suggest that in addition to the technical attempts to protect user privacy, public awareness should be raised about the implications of accepting friendship requests from untrusted sources. Furthermore, the success rate of an adversary in establishing links in OSNs can help revisiting the key assumptions underpinning the design of some systems such as OSN-based Sybil defense systems.

Index Terms—Online Social Networks, Privacy, Personal Information, Privacy Settings

I. INTRODUCTION AND MOTIVATION

Online Social Networks (OSNs) have grown significantly in size and in terms of user activity over the past few years: Twitter has over 105 million registered users [7]; Facebook has more than 750 million active users [4], handles about a billion chat messages each day, and at peak times serves about 1.2 million photos every second [6].

Users usually maintain a connection between their online profiles and their real world identities, and actively provide considerable amounts of personal information on OSNs. Many factors cause OSN users to willingly share information. These include peer pressure and herding behavior, relaxed attitudes towards personal privacy, incomplete information about the possible privacy implications of information revelation, faith in the networking service or trust in its members [13]. They might also estimate the advantages of revealing data to strangers

as being larger than the perceived costs of possible privacy invasions.

OSN providers have incentives to encourage their users to become as visible and searchable as possible. OSNs are social goods closely aligned with the market demands. The more users an OSN has and the more actively they interact with the network, the more valuable that OSN will be.

Such a huge amount of information has made OSNs invaluable sources of information for marketers, employers, spammers, phishers, credit rating agencies, and intelligence agencies who might misuse the available information. Identity theft, online and physical stalking, embarrassment, price discrimination and blackmailing are only a few implications and privacy threats that may arise from the violation of user privacy in OSNs.

These possible risks have raised privacy concerns among privacy advocates and the research community. When talking about possible threats, the finger of blame is usually pointed towards the OSN providers who are believed to intentionally avoid or unintentionally fail to provide an acceptable level of privacy protection. In a report issued by Canadian Privacy Commissioner in July 2009, Facebook has been criticized for default privacy settings, collection and use of users' personal information for advertising purposes, and disclosure of users' personal information to third-party application developers [1]. Privacy advocates insist that OSN providers not only should avoid exposing their users' identities, but also they need to ensure that they do not make enough information public that may be combined with other publicly available data to expose users to identity theft. They also insist that there must be an efficient mechanism to control the ways OSN providers might share users' information with third parties.

Even under the strictest privacy standards, and with the most secure online environments and access control mechanisms personal information leakage may occur. In other words, no matter what the OSN providers do, there is still an entity with absolute power of nullifying all the privacy preserving attempts: that entity is the user herself. There have been many articles about misuse of OSNs for harassment that suggest that OSNs users are not cautious enough about what they post online and how they choose their online friends. In March 2010, a Facebook status update led to a house robbery, when an American woman posted the status update that she was heading out to a bar. While she was away, two men broke into her house. The incident was caught on camera, and she later realized that one of the men looked suspiciously like someone

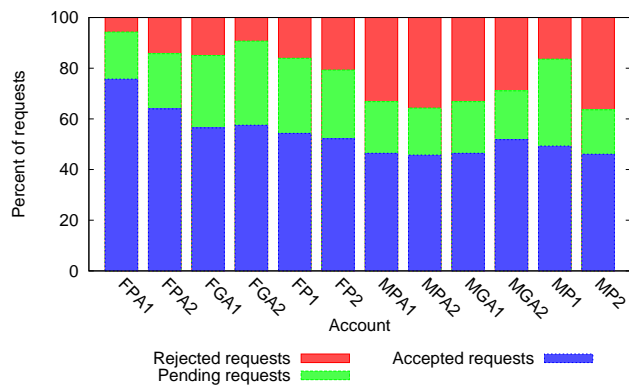


Fig. 1: Acceptance, pending, and rejection percentage for each agent.

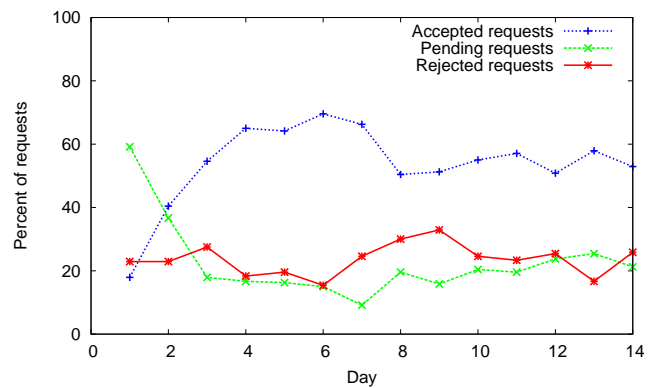


Fig. 2: Acceptance, pending, and rejection percentage vs. day.

who had recently befriended her on Facebook [3]. In a separate incident, a woman was stabbed repeatedly and killed by her ex-lover after he saw a picture of her with her new boyfriend on Facebook [5]. In 2006, an adult befriended a 13-year-old girl, sending her derogatory messages and cyberbullying via a fake Myspace account which caused the girl to eventually commit suicide [2].

In this paper, we study how people respond to friendship requests from strangers (and possibly adversaries) in an OSN, thus giving them access to their friend lists, and other personal information. We also try to quantify how a user's online friends might be affected, directly or indirectly, as a result of such behavior. To this end we use data gathered by 12 members of a popular OSN each sending 20 friend requests daily to completely unknown people. We measured how often these requests were accepted, whether the receiver of the request has changed her default privacy settings, and in case of acceptance of the request, to which previously inaccessible data fields we gain access. Our initial plan was to use a large number of robots to perform these experiments, but we were advised by the OSN provider that it is against their terms of service, and we were "highly discouraged" to perform such experiment. We understand the limitations of using real users for this task: they are a more limited number of agents that are not necessarily as accurate as a programmed robot. However, this was the only approach we could take without violating the terms of service. We believe increasing the number of agents will not have a qualitative impact on our results. At the same time, as explained in Section 5, we are working on a new set of experiments with a larger set of users, and hope to report the results in the near future.

We analyze the online behavior of more than 3,000 users of a large OSN when they receive a friend request from a stranger. More than half of that population accepted our friend requests. We evaluate the amount of information they disclose, and study the changes in accessibility of different information fields before and after the acceptance for those who accepted our requests. We found that almost always (for more than 95% of the friends of all our 12 agents) befriending reveals

information that would otherwise be inaccessible. We also study users' usage of the privacy settings, and show that more than 55% of people did alter the default setting, and more than half of those who have modified these settings still accept friend requests from a stranger. Of those who have changed the setting and accepted our requests 75% grant us access to all the information that they explicitly made inaccessible to public. This shows that, in contrast to the current consensus, modifying privacy settings might not be a good sign of privacy concern and awareness. This implies that persuading the OSN providers to set more restrictive settings as default setting might not be always adequate to guarantee the privacy of information. Our results provide a preliminary picture of lack of privacy concerns about personal information revelation.

We do not intend to underemphasize the significance of improving infrastructures and privacy setting. Our goal is to attract attention to the less-studied problem of users' uninformed behavior in OSNs and its consequences. Our work shows that users need to be educated about privacy and the implications of their behavior on OSNs. We believe even a simple warning sign that advises users against accepting friendship requests from strangers (especially when they are about to accept a request from a suspicious agent) can go a long way in protecting user privacy.

The remainder of the paper is structured as follows. We discuss related work in Section 2. We present our methodology in Section 3, follow by the results of our study in Section 4. Then we present a discussion about implications of our findings for designing more efficient Sybil defense systems. We conclude the paper in Section 6 with a summary of our contributions and a discussion of future work.

II. RELATED WORK

Studies about people's behavior and their interactions in OSNs thus far have relied on surveys or interviews of OSN users. These studies suggest that people are highly concerned that a stranger would know the information that they usually provide in their online profiles (e.g., the place they live [8]), they very rarely use the OSNs to initiate new connections [14],

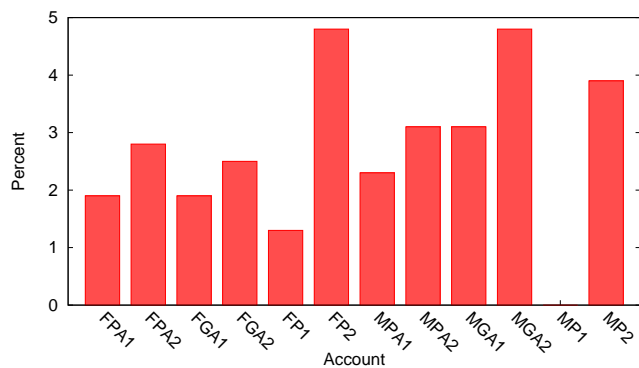


Fig. 3: Percentage of friends whose account visibility did not change before and after accepting friend request.




Fig. 4: Percentage of friends with modified privacy settings, who made *all* the information they restricted public access to available to our agents.

[17], [18], and most of them never befriend a person they have never met in person [15].

While our sample size does not make us eligible to firmly reject these claims, our preliminary findings show a significantly different behavior in the majority of users. For all our 12 agents, even those without any online activity or information, more than 45% of the unknown people to whom the friendship requests were sent accepted our requests. For one of the accounts with an attractive profile picture and some OSN activity this number was as high as 76%. This illustrates that the studies relying solely on people's claimed behavior might not provide accurate and reliable results to make general conclusions about their actual attitude.

Also, it is generally known that a very small number of people change their privacy settings in OSNs[13], [21]. Our study shows that at least 55% of the people to whom we sent requests have changed from the default setting. It should also be noted, as we will explain later, that this practice is not a reliable signal of privacy awareness.

III. METHODOLOGY

Previous studies have investigated the elements that play a role in acquiring trust in computer mediated environments, such as profile's photo [10], [23], type of language used [19], and sex of the profile [23]. Studying the effect of such elements required creating and operating different accounts with different characteristics, as well as aggregating the results automatically. Although that is technically feasible, it is against the statement of rights and responsibilities of the OSN we considered. The OSN, much to our regret, rejected our request for temporarily bypassing those restrictions during the limited period of our study. Consequently, the current study presents the results of a limited experiment with only 12 agents.

More precisely, we used 6 female and 6 male accounts whose owners did not previously have profiles on the OSN, and kindly agreed to act as our agents in the network: they created the accounts, sent 20 friend requests to completely

unknown people each day for 2 weeks and gathered and reported to us the specific statistics after that period of time. Of those agents, 8 (4 female and 4 male) were active, meaning they imitate the behavior of an average user as that OSN reported in its general statistics, such as leaving 12 comments on the site's contents and becoming a member of 12 groups¹. Half of these active accounts, 2 female and 2 male, had attractive portraits as profile pictures, while the profile pictures of the other 2 active accounts were a photo of a large number of people whose faces could not be recognized. The rest of the accounts, 2 female and 2 male, were passive, where we define a passive account as an account that is not involved in any type of activity except for sending 20 friend requests each day, only the mandatory information fields are completed, and the account has no profile picture.

We will call active female accounts with portrait pictures FPA1 and FPA2, active female accounts with group pictures FGA1 and FGA2, and passive female accounts FP1 and FP2. Similarly, active male accounts with portrait pictures are called MPA1 and MPA2, active male accounts with group pictures MGA1 and MGA2, and passive male accounts MP1 and MP2.

IV. RESULTS

We categorize friend requests into 3 groups: accepted, pending reply, and rejected. We waited up to 14 days after the request was made for the response. Since the OSN we study claims that 50% of its active users log on to it in any given day, it makes sense to consider the condition of the request after that time period as final. From the total number of 3,360 requests that our agents sent to unknown people over a period of 2 weeks, 54% of them were accepted. Figure 1 shows the percent of accepted, pending and rejected requests for each of our agents. For all 12 accounts, more than 45% of the requests were accepted, while this number is as high as 76% for FPA1 (active female account with an attractive picture). It

¹For the sake of privacy of our agents, we keep their names and the name of that OSN anonymous. Therefore, we are not able to mention the citation for these statistics.

Account Code	Personal Info					Contact Info			
	Photo	Birthday	Current City	Home Town	Relationship	Email	Mobile Number	Address	IM IDs
FPA1	6.1	26.9	10.8	19.3	20.8	82.1	34.4	25.0	42.5
FPA2	59.2	17.9	6.1	12.8	15.1	92.7	50.8	36.3	15.6
FGA1	12.7	48.7	6.3	31.6	36.1	80.4	28.5	4.4	50.6
FGA2	9.3	31.7	14.9	21.1	24.2	75.2	34.8	23.6	20.5
FP1	9.2	13.2	8.6	9.2	6.6	95.4	17.1	17.1	16.4
FP2	18.5	55.5	7.5	54.1	29.5	84.2	19.9	10.3	58.2
MPA1	8.5	44.6	6.9	37.7	43.1	91.5	35.4	10.0	37.7
MPA2	19.5	43.0	12.5	17.2	29.7	67.2	5.5	3.9	18.8
MGA1	14.6	50.0	10.8	28.5	31.5	80.8	15.4	5.4	35.4
MGA2	17.2	53.1	8.3	43.4	46.9	90.3	25.5	5.5	22.1
MP1	9.4	15.2	5.1	13.0	15.2	90.6	61.6	39.1	7.2
MP2	21.7	58.1	9.3	68.2	46.5	87.6	24.8	9.3	29.5

TABLE II: Percentage of people who revealed information *after* accepting the requests.

Account Code	%males who accepted req	%females who accepted req	no. received req
FPA1	88.1	69.5	396
FPA2	80.0	61.7	104
FGA1	67.8	61.1	6
FGA2	63.9	60.0	16
FP1	59.7	54.5	16
FP2	73.6	49.0	6
MPA1	49.2	48.4	4
MPA2	46.2	45.8	5
MGA1	61.9	50.0	3
MGA2	66.2	56.8	5
MP1	51.1	55.7	3
MP2	55.0	49.0	4

TABLE I: People's reaction to friend requests they received and the requests they sent.

is interesting to note that even when the information fields are left blank and the accounts have no activity or photo, approximately half of the people accept the requests (54%, 52%, 49%, and 46% acceptance for FP1, FP2, MP1, and MP2, respectively).

When considering users to whom we made requests that had gender information, 50.19% of them were male and 49.81% were female. We observe that 63% of the male users and 55% of the female users accepted our requests. For all 12 agents except MP1, as summarized in Table 1, acceptance rates were higher among males. In other words, a male user is more likely to accept a friend request from a stranger. The difference was particularly noticeable for the requests sent by FPA1, FPA2, and FP2 (in which cases percent of males who accepted the requests were, respectively, 19, 18, and 25% more than the females).

Surprisingly, during the project execution, unknown people made friend requests to our 12 agents. The minimum number of friend requests that female accounts received, 6, was greater than the maximum number of the friend requests that male

accounts received, 5. While active female profiles with attractive photos received a significantly larger number of friend requests (396 and 194 friend requests for FPA1 and FPA2) than active female profiles with group images and passive female profiles (which received 6 or 16 friend requests each), having account activity or a portrait picture do not seem to have such considerable effects on the number of friend requests that male accounts received (all between 3 to 5, inclusive). Also, for female accounts, friend requests from active accounts with an attractive profile picture were accepted more frequently than those of the active accounts with group photos or passive accounts. But we do not observe this effect for male agents.

Also, the number of friends someone has effects whether others accept her request or not. Since measuring the exact number of friends our agents had at the time someone received their request was not feasible, we used a rough approximation to that: we plotted percentage of accepted, pending, and rejected requests versus the day the requests were made. We expected that on day 2, our agents would have more friends than day 1 of the experiment; on day 3, the number of friends is higher than day 1 and day 2; and so on. As Figure 2 suggests, in the very first days of the experiment, people seemed reluctant to accept the requests. In other words, an agent with an acceptable number of friends would be more successful in establishing new friendships in OSNs than someone with no or very few friends.

A. Visibility

Previous survey-based studies suggest that people will not normally accept requests from strangers. One might argue that such behavior is characteristic of more open people. That is, people whose public view of their profiles contain equal amount of information as their profiles as seen by their friends, and consequently they will not reveal any previously hidden information by accepting the requests. Figure 3, however, shows that for the friend lists of all the 12 accounts under 5% of the friends had identical public and private profiles.

While some information like email addresses, mobile phone numbers, and addresses were visible in only 2%, 1%, and 1%

of public profiles of the users to whom we sent the requests, 85%, 30%, and 17% of those who accepted our requests made these information visible to us.

Tables 2 and 3 show what fields of what percentage of the friends became visible *after* accepting the friend requests. While some information such as address, birthday, hometown, or mobile phone puts the users at the risk of being identified or stalked, other information, such as email, activities, interests, and favorites, can be used for marketing and targeted advertising.

B. Default Privacy Settings

Previous work suggests that very few people ever change the default privacy setting. Hence, privacy advocates constantly urge that the default setting should be set in such a way as to maximize users' privacy. Our findings show that majority of people have indeed changed their privacy settings to make some of their information publicly inaccessible. In the OSN we studied, there are information fields that are required for creating profiles, and are publicly available by default. To see whether someone has changed her default setting or not, we checked the availability of these fields on her public profile. While the absence of these fields implies that she has surely changed the default privacy setting, the vice versa is not necessarily true. Since the OSN provides fine-grained setting, the user might have changed the setting of other fields. So the availability of these fields should not be interpreted as *not* having changed the default setting. Consequently, what we measure and report, in Table 4, as the percent of people who have surely changed the privacy setting is actually the lower bound on the percentage of people who might have actually done so. Yet, more than 55% of those to whom we sent requests (52% of those who accepted the requests, 56% of those who did not reply, and 64% of those who rejected) have changed their default settings. Around 75% of those who had changed the privacy settings and accepted our requests made exactly the same information, that they made publicly inaccessible, accessible to us. Figure 4 reports the numbers for each agent. It suggests that changing the privacy setting might not be such a good signal of privacy awareness.

C. Direct and Indirect Effect on Friends' Privacy

The default setting of the OSN we studied makes some fields visible to friends of friends. Since in the OSN users can have several thousand friends, checking the visibility of profiles of all the friends of friends of our agents, without using automatic means, was challenging. As a rough approximation, we randomly selected 10 friends of friends of each of our agents who had not befriended our agents, and compared their public profiles with their profiles as seen by friends of friends. We found that for only 7% of them friends of friends do not have access to any publicly unavailable information.

Apart from that direct implication, accepting the friend request from a stranger might have an indirect effect on the privacy of the friends of someone: people with no mutual

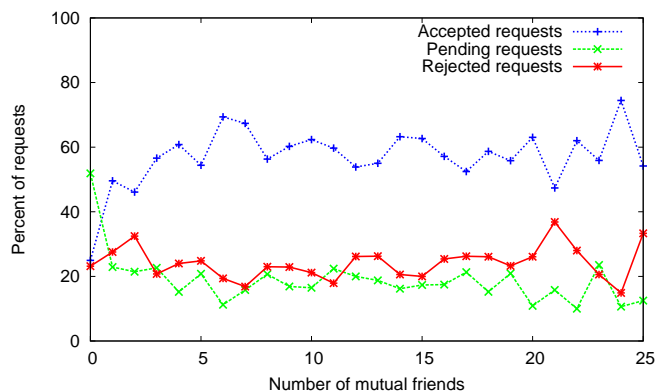


Fig. 5: Acceptance, pending, rejection vs. number of mutual friends.

Account Code	People who have accepted the friend requests	People who have not replied to the friend requests	People who have rejected the friend requests
FPA1	41.5	55.8	62.5
FPA2	32.1	59.1	56.7
FGA1	65.2	74.1	91.2
FGA2	47.8	44.9	59.1
FP1	12.5	20.7	53.3
FP2	76.7	76.1	77.6
MPA1	59.0	65.4	68.8
MPA2	48.4	38.3	47.0
MGA1	65.4	51.8	53.8
MGA2	65.5	80.0	72.8
MP1	50.0	50.7	67.4
MP2	75.2	61.7	67.6

TABLE IV: Lower bound on percentage of people who have modified privacy settings.

friends with our agents accepted the requests with much lower rates than the average rate.

Occasionally, people to whom the friend requests were sent, send messages to our agents. Sometimes these messages provide clues about the incentive that tempt people to accept the request. It seems, from messages like “Hello, I added u since u know some of my friends but ur name doesn’t ring a bell. Can u please refresh my memory?”, or “hey hi... know it’s odd... i know u seem to have a number of mutual friends...But i can’t remember you” that after receiving a request from a stranger with several mutual friends, people sometimes doubt that she is an acquaintance who they have forgotten. Intuitively, these requests should have higher chance of being accepted. Figure 5 shows the percent of accepted friend requests as a function of the number of mutual friends our agent has. We confirm that acceptance rates are higher among those with 1 or more mutual friends. Surprisingly we note that there is a diminishing return for our agents after the first few mutual friends.

Account Code	Views		Education			Employment	
	Political Views	Religious Views	College	Degree	High School	Employer	Position
FPA1	14.6	22.2	9.0	29.7	11.8	13.2	2.8
FPA2	15.1	17.9	5.0	9.5	8.4	2.8	0.0
FGA1	24.1	26.6	38.6	5.1	34.8	15.8	2.5
FGA2	19.9	25.5	11.2	7.5	11.2	7.5	0.6
FP1	2.0	3.3	8.6	0.7	7.9	2.0	0.0
FP2	34.9	27.4	6.2	0.7	5.5	1.4	0.0
MPA1	13.1	18.5	26.2	9.2	41.5	25.4	6.2
MPA2	14.8	18.8	10.2	0.0	7.8	13.3	7.8
MGA1	14.6	22.3	30.0	0.0	26.9	13.1	3.1
MGA2	20.0	29.7	41.4	0.7	49.0	28.3	6.9
MP1	10.1	13.8	2.9	0.0	10.9	3.6	2.2
MP2	27.9	23.3	42.6	1.6	46.5	35.7	6.2

TABLE III: Percentage of people who revealed information *after* accepting the requests.

V. DISCUSSION: DEFENDING AGAINST SYBIL ATTACKS BY LEVERAGING OSNs

We believe that our results also indicates the need to revisit the underpinning assumption of an important class of applications and algorithms that try to protect the social networks and users from Sybil attacks (as it will be defined in the next section). Therefore, in this section, we first define Sybil attacks and then present the results of another (larger) set of experiments to justify our claim.

A. Sybil Attacks

Decentralized distributed systems are known to be vulnerable to Sybil attacks [25], in which a single faulty entity, called *Sybil*, can obtain and control multiple identities [12]. A recent ongoing topic of research has been leveraging OSNs to identify Sybils [11], [24], [25]. These studies try to tell apart the benign parts of a network from the malicious parts, by looking for subgraphs with special properties within the graphs of social networks. They generally embody some assumptions about the characteristics of malicious and benign parts of the network. Recent studies has unveiled a major similarity among these schemes: they all work by detecting local communities around a trusted node, i.e., groups of nodes more tightly connected to each other than the rest of the graph [22]. They, therefore, reckon on the limited ability of Sybils or malicious nodes to establish links with benign users [22], [24], [25]. To test the efficiency of their schemes, they typically simulate simple attack scenarios in which links between malicious and benign users are limited in number and much less than the average degree of nodes.

In absence of any social engineering technique, we show that acceptance rate of our requests are still relatively high, which implies that an attacker can reach the average degree in the OSNs quite easily. This observation would help in ameliorating the designs of Sybil defenses, and testing them with more realistic attack scenarios.

B. Experimental Results

We repeat the experiment, as described before, with two major difference: (a) larger data set: this time we create

20 female and 20 male accounts, (b) avoiding any social engineering practice (like using profile pictures, having any activity, etc.) for all the accounts. More precisely, our 20 female and 20 male accounts are created using names of the objects (rather than human names) and without any profile pictures. We only provide the information that are obligatory for creating the accounts and set the privacy settings such that these information are not accessible. Thus, the only visible parts of our accounts are their names and the default OSN photo thumbnails. Each of these accounts send out 200 friend requests to completely unknown people.

It should be noted that such constraints dictates a worst case scenario. Put differently, our results provide a lower bound on the ability of an attacker to establish links with users and collect their information, as the attacker can leverage social engineering techniques to tempt people to accept the requests or to boost the trust in them by pretending that the requests originate from benign users. These techniques could vary from designing attractive profiles to targeting those people with whom the attacker has mutual friends i.e., sending requests to the friend lists of the users who have already accepted her requests.

As before, we categorize friend requests into 3 groups: accepted, pending reply, and rejected. From the total number of 8,000 requests that our accounts send to unknown people, 26% of them are accepted, where requests send out by female accounts have slightly better chance of acceptance (28% vs. 24%). Figure 6 shows the distribution of accepted, pending and rejected requests for each of our female and male accounts.

Quite similar to the result of our first set of experiments, when considering users to whom we make requests that had gender information, we observe that men are more likely to accept requests from unknown people: around 30% of the male users (vs. 21% of female users) accept our requests. It is interesting to note that the difference mainly stems from the sex of the accounts that send out the requests, despite the fact that the OSN default photos are the only bit of information that reveal the sex of the accounts. In other words, when receiving requests from female accounts, 20% of women and 35% of men accept the requests, whereas these values are respectively




Fig. 6: Acceptance, pending, and rejection percentage for each agent (2nd experiment).

around 23% and 26% when the requests originate from male accounts. This observation implies that very simple social engineering techniques can ameliorate the acceptance rates, and strengthen our argument that the assumptions about limitations of adversaries in establishing links in OSNs warrants further investigation.

Regarding visibility, as Figure 7 shows, for the friend lists of all our accounts under 6% of the friends have identical public and private profiles. While some information like email addresses, mobile phone numbers, and addresses were visible in less than 4%, 1%, and 1% of public profiles of the users to whom we send the requests, 85%, 11%, and 7% of those who accept our requests make these information visible to us.

Also, our findings for this second experiment confirms that majority of people have indeed changed their privacy settings to make some of their information publicly inaccessible: more than 53% of those to whom we send requests (53% of those who accept the requests, 57% of those who do not reply, and 65% of those who reject) have changed their default settings. Around 83% of those who have changed the privacy settings and accept our requests make exactly the same information, that they previously made publicly inaccessible, accessible to us.

Overall, the results suggest that any user in the network (including the attackers) can easily establish large number of links to other, normal users that are not under the control of the attackers.

VI. CONCLUSIONS AND FUTURE WORK

While the consequences of excessive disclosure of personal information on OSNs have been examined from various angles in previous work [13], [16], we tried to test peoples' privacy concerns by running an experiment in one of the most popular OSNs. We find that the majority of people freely accept friend requests from unknown users. We observed that 95% of the users who accepted our requests revealed some previously unknown information. Based on the assumption that few people change the default settings, having impermeable default privacy settings has long been thought of as a mean for maximizing privacy. However, we found that more than

52% of the users who accepted our requests have changed their default privacy settings. This might be interpreted as an indication that they have at least some concerns about privacy. But more than 75% of them reveal all the information they have explicitly made publicly invisible to us by accepting our requests. By providing evidence of peoples' propensity for adding unknown accounts as friends, we show the need for raising public awareness about potential risks of sharing personal information with strangers. Our results could also be regarded as an initial step in revisiting the underlying assumptions of Sybil defense designs.

We understand the limited scope of this work, and in order to ensure statistical significance of the results larger sample sizes are needed. Also, as mentioned in Section 3, other factors can effect others willingness to trust someone and initiate friendship with her in OSNs. Some related work that studies the effect of these factors are survey-based experiments [10], [23]. We believe than since peoples' claimed attitude in surveys might have meaningful differences with their actual behavior. The role of such elements can be more reliably investigated by running similar experiments in actual networks. Given the concerns from the OSN providers, using automated agents is not possible. To address this problem, we are working on two different solutions. First, we are organizing a contest – hopefully with a large number of participants with diverse demographic characteristics – in which contestants will act as agents. We will gather data through automated means but as participants will be asked to give us permission for collecting this data we believe it will not be against terms of service of the OSN. As the second attempt for addressing the aforementioned problems, we are working on assignments for a undergraduate course in which students will perform various experiments designed to answer questions we are interested in this paper.

REFERENCES

- [1] Commissioner's findings - PIPEDA case summary #2009-008: Report of findings: CIPPIC v. Facebook Inc. - July 16, 2009. <http://www.priv.gc.ca/cf-dc/2009/2009Date: 29 May, 2012>.




Fig. 7: Percentage of friends whose accounts' visibility do not change before and after accepting friend request (2nd experiment).

- [2] Conviction in MySpace suicide case tentatively overturned. <http://www.cnn.com/2009/CRIME/07/02/myspace-suicide/index.html?ref=allsearch> (Access Date: 29 May, 2012).
- [3] Facebook posting allegedly led to house robbery. <http://amfix.blogs.cnn.com/2010/03/26/facebook-posting-allegedly-led-to-house-robbery/?iref=allsearch> (Access Date: 29 May, 2012).
- [4] Facebook statistics. <http://www.facebook.com/press/info.php?statistics> (Access Date: 29 May, 2012).
- [5] Man killed ex-lover over Facebook photo with new man. <http://news.bbc.co.uk/2/hi/uk> (Access Date: 29 May, 2012).
- [6] Technology review: How Facebook copes with 300 million users. <http://www.technologyreview.com/web/23508/page1/> (Access Date: 29 May, 2012).
- [7] Twitter claims 105 million registered users. <http://scitech.blogs.cnn.com/2010/04/14/twitter-claims-105-million-registered-users/?iref=allsearch> (Access Date: 29 May, 2012).
- [8] A. Acquisti and R. Gross. Imagined communities: Awareness, information sharing, and privacy on the Facebook. In *Proceedings of the 6th Workshop on Privacy Enhancing Technologies*, 2006. <http://www.springerlink.com/content/gx00n8nh88252822/> (Access Date: 29 May, 2012).
- [9] Alessandro Acquisti and Ralph Gross. Predicting Social Security numbers from public data. *Proceedings of the National Academy of Sciences*, 106(27), 2009. <http://www.pnas.org/content/106/27/10975.short> (Access Date: 29 May, 2012).
- [10] C. Aguiton, D. Cardon, A. Castelain, P. Fremaux, H. Girard, F. Granjon, C. Nepote, Z. Smoreda, D. Trupia, and C. Ziemlicki. Does showing off help to make friends? In *Proceedings of the 3rd International Conference Conference on Weblogs and Social Media (ICWSM)*, 2009. <http://www.aaai.org/ocs/index.php/ICWSM/09/paper/download/178/562> (Access Date: 29 May, 2012).
- [11] G. Danezis and P. Mittal. SybilInfer: Detecting Sybil Nodes using Social Networks. In *Proceedings of NDSS*, 2009. <http://research.microsoft.com/en-us/um/people/gdane/papers/sybilinfer.pdf> (Access Date: 19 September, 2012).
- [12] J. Douceur. The Sybil attack. In *IPTPS*, 2002. <http://www.few.vu.nl/~mconti/teaching/ATCNS2010/ATCS/Sybil/Sybil.pdf> (Access Date: 19 September, 2012).
- [13] R. Gross, A. Acquisti, and H. J Heinz III. Information revelation and privacy in online social networks. In *Proceedings of the 2005 ACM Workshop on Privacy in the Electronic Society*, 2005. <http://dl.acm.org/citation.cfm?id=1102199.1102214> (Access Date: 29 May, 2012).
- [14] A. N Joinson. Looking at, looking up or keeping up with people?: motives and use of Facebook. In *Proceeding of the SIGCHI Conference on Human Factors in Computing Systems*, 2008. <http://dl.acm.org/citation.cfm?id=1357054.1357213> (Access Date: 29 May, 2012).
- [15] H. Jones and J. H Soltren. Facebook: Threats to privacy. *Project MAC: MIT Project on Mathematics and Computing*, 2005. <http://ocw.ispros.com/bd/courses/electrical-engineering-and-computer-science/6-805-ethics-and-the-law-on-the-electronic-frontier-fall-2005/projects/facebook.pdf> (Access Date: 29 May, 2012).
- [16] B. Krishnamurthy and C. E Wills. Characterizing privacy in online social networks. In *Proceedings of the 1st Workshop on Online Social Networks*, 2008. <http://dl.acm.org/citation.cfm?id=1397744> (Access Date: 29 May, 2012).
- [17] C. Lampe, N. Ellison, and C. Steinfield. A Face(book) in the crowd: Social searching vs. social browsing. In *Proceedings of the 2006 Conference on Computer Supported Cooperative Work*, 2006. <http://dl.acm.org/citation.cfm?id=1180901> (Access Date: 29 May, 2012).
- [18] C. Lampe, N. B Ellison, and C. Steinfield. Changes in use and perception of Facebook. In *Proceedings of the ACM 2008 Conference on Computer Supported Cooperative Work*, 2008. <http://dl.acm.org/citation.cfm?id=1460675> (Access Date: 29 May, 2012).
- [19] E. J Lee. Categorical person perception in computer-mediated communication: Effects of character representation and knowledge bias on sex inference and informational social influence. *Media Psychology*, 9(2), 2007. <http://www.tandfonline.com/doi/abs/10.1080/15213260701286007> (Access Date: 29 May, 2012).
- [20] J. Lindamood, R. Heatherly, M. Kantarcioglu, and B. Thuraisingham. Inferring private information using social network data. In *Proceedings of the 18th International Conference on World Wide Web*, 2009. <http://dl.acm.org/citation.cfm?id=1526899> (Access Date: 29 May, 2012).
- [21] W. E Mackay. Triggers and barriers to customizing software. In *Proceedings of the ACM CHI '91 Human Factors in Computing Systems*, 1991. <http://dl.acm.org/citation.cfm?id=108867> (Access Date: 29 May, 2012).
- [22] B. Viswanath, A. Post, K. P. Gummadi, and A. Mislove. An Analysis of Social Network-Based Sybil Defenses. In *Proceedings of SIGCOMM*, 2010. <http://dl.acm.org/citation.cfm?id=2043164.1851226> (Access Date: 19 September, 2012).
- [23] Shaojung Sharon Wang, Shin-Il Moon, Kyounghee Hazel Kwon, Carolyn A. Evans, and Michael A. Stefanone. Face off: Implications of visual cues on initiating friendship on Facebook. *Computers in Human Behavior*, 26(2), 2010. <http://www.sciencedirect.com/science/article/pii/S0747563209001502> (Access Date: 29 May, 2012).
- [24] H. Yu, P. B. Gibbons, M. Kaminsky, and F. Xiao. SybilLimit: A near-optimal Social Network defense against Sybil attacks. In *Proceedings of IEEE Symposium on Security and Privacy*, 2008. <http://dx.doi.org/10.1109/TNET.2009.2034047> (Access Date: 19 September, 2012).
- [25] H. Yu, M. Kaminsky, P. B. Gibbons, and A. Flaxman. SybilGuard: Defending Against Sybil Attacks via Social Networks. In *Proceedings of SIGCOMM*, 2006. <http://dx.doi.org/10.1109/TNET.2008.923723> (Access Date: 19 September, 2012).