To Add or not to Add: Privacy and Social Honeypots

Hamed Haddadi*, Pan Hui†
*University of London †Deutsche Telekom Laboratories

***Abstract*—Online Social Networks (OSNs) have become a mainstream cultural phenomenon in the past years, where million of people connect to each other and share memories, digital media and business relations. Many users also publish personal information about their activities, relationships, locations and interests on these sites, seemingly unaware of how these data can be used by other parties. Sites typically attempt to restrict data-sharing to members of a user's social network, but this is only effective if these social networks cannot be exploited by malicious users.**

In this paper we perform an experiment in order to assess the vulnerability and privacy awareness of users when engaging in online relations with random unknown users, or those pretending to be a famous character. We find that usually users do not accept random friendship requests, but some aggressively search for celebrities, making a perfect case for spammers to form *honeypots* using such fake profiles. We present a set of suggestions for enhancing privacy on social networks which could reduce the threats of identity theft in such environments.

I. Introduction

Recent years have seen the rapid growth of online communities of users who connect to each other, publish and share information such as photographs and articles, and track each other's personal and professional activities. These communities are commonly referred to as Online Social Networks (OSNs). They range from networks of business relationships to religious groups and classmates[1]. One common feature of these networks is the publish of large amount of personal data about a user and his acquaintances, such as notices of change of circumstance (e.g., marriage, divorce, new jobs), user profile information, travel photographs and personal statements. These data, which are accessible by almost everyone in the case of open social networks [1], are clearly extremely desirable for using beyond the, perhaps benign, original community setting. These include, but are not limited to, advertisements and marketing [2], police investigations, stalking, black mail [3], and social scam [4].

Joining OSNs, albeit a voluntary task, has become a norm for many adults, and teenagers, in many countries, specifically USA and UK where it has been reported that more than 2/3 of the online population is a member of at least one OSN. Once joined, it is also not easy to "completely" leave such networks as the personal information remains on the Internet, usually by the choice of the OSN designers who tend the take copyright over content uploads (e.g., Pictures loaded on Facebook will not be deleted even after a user leaves the network, although this is currently an issue under investigation in some courts and there are constant legal battles over such issues.). This is, of course, not a concern to the younger generation who has a strong desire for short term fame and attention. Later in career life though this choice may affect them. It is reported that in UK, one in five recruiters have used social networks for vetting candidates. News like this often raise the question of, how much privacy that users are willing to sacrifice in order to strengthen their social popularity or increase their circle of friends? Are users willing to add anyone as a friend? How sensitive are the users to the authenticity of a profile owner to prove who they claim to be? A growing concern in the age of digital identity is the validation of the authenticity of users in relation to their online existence.

An important concept about OSNs is that the users are on themselves as they want and need to socialize. But by doing so they naturally risk their privacy in a trade-off for publicity. Privacy advocates will not be able to prevent or discourage users form being on these networks, rather, it is important to provide information and guidelines for the users. The OSN designers should also provide mechanisms for transparent privacy controls, such as an opt-in system for information use rather than an opt-out mechanism [5]. On this basis, our paper has a number of contributions. We perform a privacy awareness experiment, which evaluate how users are aware of their privacy in a rather peculiar way. Usually they would not accept a link request from a celebrity, but they are more likely to accept it from a random, and indeed fake, user. On the other hand, some of them actively go *hunting* for celebrities and ad them without noticing their authenticity. In this way, such profiles can become a social honeypot, used by information collectors in order to harvest user's profile information for other uses.

Naturally, our experiment design leads to shortcomings. Just like any information gathering experiment on OSNs, a privacy-aware user can of course use pseudonym and publish faked profile information. In this way, his/her privacy can be well protected despite of how many unknown online friends the user has. Secondly, we performed experiments based on one OSN and in particular users in one country, which means that the conclusion may not be simply generalized, considering different cultural background and the implementation of such a OSN. However we still believe our work sheds light on an important risk model for data harvesting and identity theft on

[1]http://en.wikipedia.org/wiki/List_of_social_networking_websites

social networks.

In this paper we perform a preliminary analysis of user behavior in OSNs. In Section II we introduce the privacy concerns and issues in OSNs. We discuss examples of how private information can be used by advertisers, stalkers and recruiters. In Section III we look at some real profiles of ordinary, and famous, users and observe how their circle of friends grow in the course of one month in a society newly opened up to an OSN and in Section IV we provide the results of this experiment and give insight to the social impacts of these results over privacy concerns. Based on our findings, Section VI-A brings together a set of recommendations for improving the security on OSNs. We present the related work in Section V and finally Section VI concludes the paper with a discussion of results and discuss avenues for potential future works.

II. Social Networks and Privacy

In this section, we briefly highlight some status quo of the privacy issue in social network and how can this private information be used in advertisement.

Evidently some OSN users are indeed interested in privacy. In September 2006, Facebook introduced a "News Feeds" feature, which broadcasted all of a user's activities to anyone in that user's social network. Even though this information was already available to members of the social network, the "sense of exposure and invasion" [6] led to protest and the feature was eventually removed. Similarly, an advertising feature named "Beacons" and a change to the Facebook privacy policy, which attempted to give the company ownership of users' data were also removed after protest from users.

On the other hand, some OSN users perhaps appear to be less concerned about privacy. The security company Sophos created a fake Facebook profile for "Freddi Staur" (an anagram of "ID Fraudster") and sent friend requests to 200 Facebook users.[2] 87 of these responded, with 82 leaking personal information useful for identity theft, such as e-mail and physical addresses, dates of birth, employment data and mothers' maiden names. Moreover, "Freddi" was also able to access personal information about friends of these victims, such as photographs. Although this experiment was just to demonstrate the potential threat of OSNs, more recently fraudsters have used this to steal money[3]. Five Facebook-specific security threats including malware designed to steal personal data were disclosed in March 2009.[4]

There are also a range of "applications" developed by third parties on these OSNs. These applications, once added to the user's profile, give unrestricted access of the user's data to the developer. In the case of Facebook, Applications can be tried out or even seen only if the user gives full access. Users may believe that some of the OSNs have strong privacy restrictions, for example Facebook only allow authorized person to review profiles. However, not many people use this privacy setting, and some features of the system setting make privacy information collection even easier. Wilson *et. al* [7] were able to download information from 10 million users from Facebook, the data contains user profile, Wall and photo data. Mini-Feed of the users can also be completely collected, which shows the complete activities of the users. Facebook provides a feature to show 10 randomly selected users from a given regional network; the researchers performed repeated queries to this service to gather 50 user IDs to "seed" their breadth-first searches of social links on each network. They were able to complete each crawl in under 24 hours, while averaging roughly 10 MB/s of download traffic. Their completed data set is approximately 500 GB in size, and includes full profiles of more than 10 million Facebook users.

Some OSNs have also been ordered to shut down due to privacy concerns. For example the operators of imbee.com[5], a social networking site specifically targeting kids and teenagers were penalized as their data-collection practices violated federal laws.[6]

III. Privacy Awareness Experiment

In order to understand how much the OSN users aware of their privacy and to give them an alarm about this issue, we designed a real experiment targeting users on a popular OSN, and in particular users in a regional network. We created 40 fake identities (20 men and 20 women) on the network. 10 of each group adopted the identities of well-known film stars from that country: these were our *celebrities*, and the other 10 were just ordinary people with random names but no other personal information. Each week, for each fake identify, we selected 10 people at random, who had self-declared to be from the same country, to befriend with. At the same time, we also accept any incoming friendship requests. This is carried out for five consecutive weeks. We do not publish or respond to any private or public messages with these fake characters, hence it was completely passive, and provided no clues as to our true identities. [7]

Previous privacy studies of OSNs [8], [9] have collected data by crawling hundreds of profiles from a given social network (e.g., a particular university). This approach leads to a large quantity of data but only provides a snapshot of data useful for an observational study. In this study, we selected a small number of users whose profiles were monitored, and combined this with intervention and observed the resulting changes in the social network. Our aim was to compare the user's behavior when they faced different types of users and friendship requests. Using our method, one could see the potential methods for an attacker to *passively* collect user information with interests in sports, politics, cinema or music. Such a rich and filtered dataset can then be used for marketing purposes.

[2]http://www.sophos.com/pressoffice/news/articles/2007/08/facebook.html
[3]http://www.cnn.com/2009/TECH/02/05/facebook.impostors/index.html
[4]http://news.bbc.co.uk/1/hi/technology/7918839.stm
[5]http://www.imbee.com/
[6]http://www.ftc.gov/opa/2008/01/imbee.shtm
[7]It is practically impossible to verify the true identity of a person on almost all OSNs.

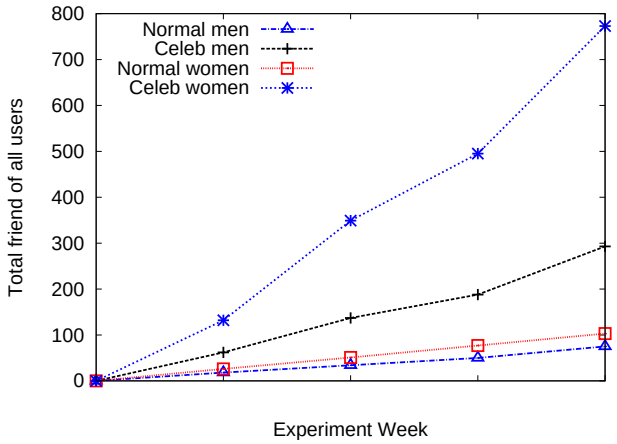

Fig. 1. Evolution of the online friend population.

As this study was intended to have a true results of users' awareness toward online privacy, it was impossible to provide informed consent. Instead, participants were debriefed after the study. Data were anonymized and the fake OSN user profiles were deleted at the end of the study.

IV. Results

In this section we present the results of our experiment. Figure 1 displays the cumulative sum of the number of friends for the four distinct groups of normal and celebrity male and female profiles. Note that in our experiment we only ever added 40 random friends for each user, with equal number of male and female contacts added on each attempt. One can immediately spot the fact that women are more successful than men in attracting friends[8]. Celebrity women, ideally should have had 40 friends each, bringing their total sum to 400. But we notice that they acquired just less than 800 friends. On the other hand celebrity men have been less successful in friend hunting. Similar pattern can be seen for normal users.

Table I displays the number of public and private messages that our users received. Inevitably, celebrity women attracted a large number of private and public messages. An interesting observation is the fact that the ratio of private to public messages is much larger in celebrity profiles. This indicates that users like to communicate openly with celebrities, but may prefer to keep a private conversation with those that they are not sure of their identity. Needless to mention that nearly all of the normal user's private messages were literally an authentication request, in most cases after the connection request had been accepted.

The table also displays the number of phone numbers collected by our accounts. These details, alongside addresses obtained form other social networks can be used in a variety of malicious ways for identity theft.

[8]We do not comment on the sociological factors behind the results in this short paper

Group	Public Messages	Private Messages	Phone Numbers
Normal men	4	13	18
Celebrity men	57	52	38
Normal women	6	42	18
Celebrity women	127	174	133

TABLE I
PHONE NUMBERS AND MESSAGES IN THE EXPERIMENTAL DATA SET

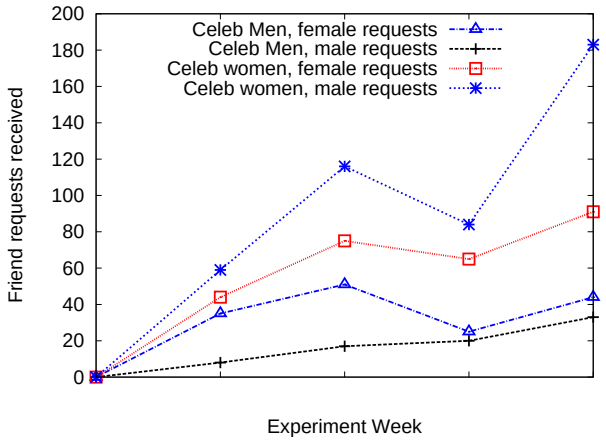

Fig. 2. Friendship requests received by celebrities.

A. Normal friendships trends

Our normal user profiles also received interesting responses. In one occasion, a user accepted our friendship request, indicating that she does not recognize us, however she finds our "chosen" name appealing. Overall, the more "traditional" the fake name is, the less likely it is for it to collect friends. In the OSN that we conducted our experiment in, there is an easy method to report suspicious users. To our surprise, only on a single occasion this facility was used against one of our profiles, resulting in a warning.

B. Celebrity friendships trends

Figure 2 displays the number of incoming friendship requests for our celebrity profiles in each week over the experiment period. It can be seen that the highest figures are for male users seeking friendship female celebrities. It can also be observed that the lowest inter-action group is between male users and male celebrities. There is a fall across nearly all categories in the third week of the experiment which coincided a public holiday. Unlike in Europe and America, where the peak Internet usage is on Sundays and holidays, the users in our study region do not usually have Internet access at home.

There were also some interesting observations about our celebrity profiles. After the 4 weeks of experiment, many of the celebrities had a large number of friends in common. These were the users who were acting as celebrity aggregators. Some of our celebrity profiles even received birthday invitations and in one case one of them was approached by a known amateur producer, inviting them for collaboration.

Another interesting result we noticed was the fact that celebrities were much less successful in hunting for friends.

On average, less than 20% of the users they added confirmed the friendship. While this figure was much higher for the 20 ordinary identities. However, the ordinary identities literally received no *incoming* friendship requests. This shows that in order to attract users, all that an attacker needs to do is to create a variety of celebrity profiles in different categories (cinema, sport, music to name a few) and literally act idle and wait for incoming requests, acting as a *honeypot* for users' information.

V. Related Work

Our results indicate that although OSNs may purport to build personal relationships, the inability of verifying the identity of a user may lead to data leaking. But even if a user is really who he or she claims to be, it is not clear what information should they really have access to. Krishnamurthy and Wills [8] look into the extent to which users share their private data in an OSN. They compare the level if information accessible to third party advertisement agencies on OSNs when compared to traditional websites. They find that tracking and advertising agents have a strong presence in all popular OSNs, taking advantage of user's unawareness of privacy setting.

Privacy concerns are not limited to Facebook. Hinduja and Patchin [10] study MySpace and find that many adolescent users provide personal information to the public, with 57% providing photos. Dwyer *et al.* [11] perform an online survey of two popular social networking sites, Facebook and MySpace, comparing perceptions of trust and privacy concern, along with willingness to share information and develop new relationships. Out of 55% of Facebook members and 60% of MySpace members who access the site every day, many indicated their preference to use these sites for personal communications with groups of their friends rather than emails. Their results suggest that in online interaction, trust is not as necessary in the building of new relationships as it is in face to face encounters. These findings indicate the low bar for advertisers and also information collectors for gaining trust within the network.

Guha *et al.* [9] propose a privacy extension for OSNs that enables the user's personal details to be partitioned into individual atoms of information, and then encrypted. using an external channel for key exchange, it is not possible for the social network or related applications to correlate the atoms in order to get the user's full profile. This approach prevents targeted advertisements based on user's personal information and gives back control to the users over whom they share their private information with. This work is similar to the effort of Zhou and Pei [12] on anonymization of user data in social networks in order to preserve privacy. Lucas and Borisov [13] also use an encryption for protecting information transmitted by the user through different Facebook applications. However in this paper we show that some users are rather careless about adding friends in the first place, which breaks down the anonymization and security chain. In addition, while such ideas may appear interesting from a research perspective, they disable the commercial income from advertising which is the dominant source of income of such portals.

Gross *et al.* [14] analyze the online behavior of more than 4,000 Carnegie Mellon University students who have joined the Facebook network site, catered to colleges at the time of the study. They show that only a minimal percentage of users changes the highly permeable privacy preferences. Our study improves on this work vastly, as now Facebook, like many other OSNs, is open to any one. This increases the risk of collecting information for fraud and criminal activity vastly. We also perform our experiment on a much wider network (of a whole nation), where users can be thousands of miles away from each other. While at a certain college, the risks involved with leaving a profile accessible is much less.

VI. Discussions

A. Can we improve privacy in OSNs?

The OSN designers are in a comfortable position. Users want to, and need to, socialize using the new digital media. However OSNs also should be required to present a trustworthy and comfortable environment for their users. On the first level, social networks should become directed graphs in a way that they give the users complete control over the level of interactions they have with their contacts. For example one can only interact partially with a newly added friend, while having access to the required information, until he/she can prove authenticity challenges. There can be an embedded challenge-response protocol when adding people (request connection with a hint of how one knows another).

Another useful idea is to have a recommendation-style "authenticity factor", a measure of user's authenticity, based on a weighted factor of user's friends, the ratio of invites sent/accepted and the number of invites received. Such a measure will avoid Sybil attacks [15] by reducing authenticity of a group of fake users adding each other. In this way, even if the attackers formed virtual circles of trust, the negative feedback outside the circle will still affect them. As they start adding other authentic users, their accounts can be partially disabled if many approved user's reject their requests. In this way, a group of Sybil attackers can't harvest authentic users. There are also OSNs which are by-invitation-only.[9]

In a similar manner, third party applications should only access information they really need, and data must be stored on the OSN provider's servers, or there should at least be a mechanism for an independent party to be able to verify the appropriate storage and management of collected information. There has been propositions for decentralized OSNs where users virtually host the OSN services in a peer-to-peer manner [16]. if such techniques are used, the advertisements can also be fed in a decentralized manner which would avoid privacy issues with the OSN hosts having access to all user data [17].

[9]http://www.asmallworld.net/privacy/

B. Conclusions and Future Work

In today's OSNs, the formation of trust and the willingness to share information is not bound to forming a social interaction. There are numerous communities, groups and virtual clubs where users may reveal information about themselves. The critical part of this information availability is to have control over *what* information are out there and *who* is able to gain access to them. Personal information at disposal of criminals can lead to identity theft crimes and also affect user's friends.

In this paper we performed a preliminary privacy-aware experiment in a popular OSN, trying to form links to randomly selected users, from a group of 40 fake identities, divided into male and female and acting as ordinary people or celebrities. Our findings indicated that users show mixed behavior with regards to *strangers*. They are happy to connect to a random, unknown profile, but they also show awareness when it comes to celebrity profiles and act with care. On the other hand, some users go actively hunting for celebrities in the wild, and become a target for information harvesters.

We believe simple techniques can increase user awareness in improving the ability of users to manage their privacy [18]. We also have suggested simple ways for improving security and trustworthiness of OSNs. Using these simple techniques and increasing user awareness will prevent OSN users from large scale spam attacks and identity theft cases in the future. We are planning to conduct larger scale of experiments on different OSNs with different cultural background. We hope this will give us a more macroscopic picture of the privacy awareness of general OSN users, and we want to use this to raise the awareness of privacy not only from the user sides but also the OSN designers. Together with our research on detection and prevention of stalkers, we hope to be able to contribute in making OSNs a safer place for the ordinary users.

REFERENCES

[1] D. Rosenblum, "What anyone can know: The privacy risks of social networking sites," *IEEE Security and Privacy*, vol. 5, no. 3, 2007.

[2] B. Krishnamurthy and C. E. Wills, "On the leakage of personally identifiable information via online social networks," in *WOSN '09*, Barcelona, Spain.

[3] D. Randall and V. Richards, "Facebook can ruin your life. and so can myspace, bebo..." The Independent, 10 Feb 2008.

[4] D. Frommer, "What a nigerian facebook scam looks like," The Business Insider, Jan 2009.

[5] J. Grimmelmann, "Facebook and the social dynamics of privacy," *Iowa Law Review*, vol. 95, no. 4, August 2009.

[6] D. Boyd, "Facebook's privacy trainwreck: Exposure, invasion, and social convergence," *Convergence*, vol. 14, no. 1, February 2008.

[7] C. Wilson, B. Boe, A. Sala, K. P. Puttaswamy, and B. Y. Zhao, "User interactions in social networks and their implications," New York, NY, USA, pp. 205–218, 2009.

[8] B. Krishnamurthy and C. E. Wills, "Characterizing privacy in online social networks," in *WOSN '08, Seattle, WA, USA*.

[9] S. Guha, K. Tang, and P. Francis, "NOYB: privacy in online social networks," in *WOSN '08, Seattle, WA, USA*.

[10] S. Hinduja and J. Patchin, "Personal information of adolescents on the internet: A quantitative content analysis of myspace," *Journal of Adolescence*, vol. 31, no. 1, pp. 125–146, February 2008. [Online]. Available: http://dx.doi.org/10.1016/j.adolescence.2007.05.004

[11] C. Dwyer, S. R. Hiltz, and K. Passerini, "Trust and privacy concern within social networking sites: A comparison of facebook and myspace," in *Proceedings of the Thirteenth Americas Conference on Information Systems*, August 2007.

[12] B. Zhou and J. Pei, "Preserving privacy in social networks against neighborhood attacks," *ICDE 2008*.

[13] M. M. Lucas and N. Borisov, "Flybynight: mitigating the privacy risks of social networking," in *WPES '08*, New York, NY, USA, October 2008.

[14] R. Gross, A. Acquisti, and H. J. Heinz, "Information revelation and privacy in online social networks," in *WPES '05*, 2005.

[15] J. R. Douceur, "The Sybil attack," in *International Workshop on Peer-to-Peer Systems*, Mar. 2002.

[16] A. Shakimov, A. Varshavsky, L. P. Cox, and R. Cáceres, "Privacy, cost, and availability tradeoffs in decentralized osns," in *WOSN '09*, Barcelona, Spain.

[17] H. Haddadi, S. Guha, and P. Francis, "Not all adware is badware: Towards privacy-aware advertising," in *IFIP Conference on e-Business, e-Services, and e-Society, I3E 2009, september 2009*, Nancy, France.

[18] H. R. Lipford, A. Besmer, and J. Watson, "Understanding privacy settings in facebook with an audience view," in *Proceedings of UPSEC '08 (Usability, Psychology and Security)*, April 2008.