Guided Socialbots: Infiltrating the Social Networks of Specific Organizations' Employees

Aviad Elyashar, Michael Fire, Dima Kagan, and Yuval Elovici

Telekom Innovation Laboratories and Department of Information Systems Engineering
Ben-Gurion University of the Negev, Beer - Sheva, Israel

aviade@post.bgu.ac.il, mickyfi@post.bgu.ac.il, kagandi@post.bgu.ac.il, elovici@post.bgu.ac.il

A dimension of the Internet that has gained great popularity in recent years is the platform of online social networks (OSNs). Users all over the world write, share, and publish personal information about themselves, their friends, and their workplaces within this platform of communication. In this study we demonstrate the relative ease of creating malicious socialbots that act as social network "friends," resulting in OSN users unknowingly exposing potentially harmful information about themselves and their places of employment. We present an algorithm for infiltrating specific OSN users who are employees of targeted organizations, using the topologies of organizational social networks and utilizing socialbots to gain access to these networks. We focus on two well-known OSNs - Facebook and Xing - to evaluate our suggested method for infiltrating key-role employees in targeted organizations. The results obtained demonstrate how adversaries can infiltrate social networks to gain access to valuable, private information regarding employees and their organizations.

1 Introduction

Over the last decade, the Internet has been playing an increasingly central role in our lives. Today, there are billions of users all over the world who use the Internet on computers, tablets, and smartphones

for personal and business needs (*1*). Adolescents surf the web and play games (*2*); adults purchase goods (*3*), schedule appointments, download files and software, pay bills, conduct meetings, and read medical information (*4*) on a regular basis.

One sector of the Internet that has gained great popularity in recent years is the platform of online social networks (OSNs) (*5*), (*6*). OSNs are web-based services that provide individuals with an infrastructure to create a public or private profile within a bounded system. This system enables a user to establish a profile that includes a list of other users with whom he or she shares connections. Typically these other users are defined as the user's friends. Moreover, users can view and traverse their list of connections and those made by others within this system (*5*).

OSNs like Facebook,[1] LinkedIn,[2] MySpace,[3] and Xing[4] allow Internet users to create user accounts and maintain connections with others. The growth of OSN usage can be attributed to several reasons. Users want to stay connected with their surroundings, and by registering with OSNs, they can stay updated about their friends' whereabouts and maintain closer relationships with them (*6*). Others utilize OSNs to promote businesses. For years, self-employed individuals such as barbers, plumbers, or accountants have been using OSNs to publicize their businesses in order to increase their customer base. Still others rely upon OSNs to initiate romantic relationships (*7*).

There is a great diversity among OSNs; they offer a variety of networks to associate with different aspects of users' lives. The biggest OSN in the world is Facebook with more than one billion users (*8*), (*9*). However, there are many more OSNs which help users to connect based on shared interests, political views, or common activities. These include LinkedIn, one of the world's largest professional networks; Xing, a European social business network for business professionals; Academia.edu,[5] a social networking site for academics and researchers; Athlinks,[6] a social networking website aimed at competitive

[1] http://www.facebook.com
[2] http://www.linkedin.com
[3] http://www.myspace.com
[4] https://www.xing.com
[5] http://academia.edu
[6] http://athlinks.com

endurance athletes; and many more.

Alongside the numerous benefits OSNs provide, such as maintaining relationships, finding new colleagues, and promoting businesses, there are also many threats that may jeopardize OSN users as well as their places of work. These threats can be divided into three major areas.

First, threats exist to *individual security*. Today, many OSN users are unaware of the serious privacy issues that accompany the use of OSNs (*10*). Users often share personal data on OSNs without realizing the short-term or long-term consequences of such information flow (*11*), (*12*). Gathered data that disclose personal and sensitive information about users may cause security risks including identity theft (*13*); inference attacks (*14*); spreading spam (*10*), (*15*); privacy threats (*16*), (*17*); malware (*18*); fake profiles or sybils (*19*), (*20*); socialbots (*16*), (*21*), (*22*), (*23*); and sexual harassment (*24*), (*25*).

Second, there are threats to *business security*. Malicious users may engage in industrial espionage by creating fake profiles or bots in order to connect to users who are key employees in targeted organizations. By so doing, the hackers gain access to monitor the information users disclose (*22*), (*26*), (*27*). The exposed user's information regarding organizations may result in losses of intellectual assets and confidential business information as well as sensitive business data, which can include stock market manipulation and cybercrimes that may end up costing hundreds of millions of dollars every year (*28*). Moreover, malicious users may spread rumors regarding the targeted organization that could result in serious reputational damage without the ability to track the source of the rumors (*29*), (*30*).

Third, there are threats to *national security*. Soldiers may inadvertently disclose confidential operational information to their friends through OSNs (*31*). The enemy may collect these national secrets, like undisclosed locations, and use these against them in the future. Moreover, hackers may use virtual identities in order to spread propagandas by connecting to key users in the OSN to demoralize the opponent society. Furthermore, the enemy may use the exposed information to run astroturf campaigns for spreading propagandas or misinformation in important issues such as U.S. political elections (*16*), (*32*), (*33*).

In this study, we have specifically focused on businesses, that is, the threats to an organization's secu-

rity. We evaluated an algorithm for infiltrating employees on two OSNs, Facebook and Xing; measured its success in infiltrating these OSNs; and compared the success of the infiltration method between the two OSNs.

Our method was able to successfully infiltrate specific employees in targeted organizations. These results demonstrate that adversaries can infiltrate both low- and high-ranked employees and gain access to valuable information regarding their organizations, while the employees remain completely unaware of the infiltration. Both the employees and their organizations should keep in mind that they are vulnerable to information leakage. Attackers may even use identities of employees from other organizations in order to attack key-role employees in the targeted organizations.

1.1 Contributions

This study presents an approach for the infiltration of specific employees in organizations using socialbots, where infiltration occurs when a user accepts the friend request of a socialbot. This is an expansion of our previous study (*23*) which was the first study to present an algorithm for attacking specific employees in targeted organization within OSNs.

Specifically, this study offers the following contributions: In contrast to several studies that discuss attempts to infiltrate users through OSNs without making any distinction between low-ranked and high-ranked users (*13*), (*16*), we define only specific employees in an organization as targets to whom potential hackers would be interested in infiltrating, namely, those employees in key-roles. Second, this study, as well as the previous one (*23*), differs in that its focus is on organizations rather than users. This study is among the first that has changed the focus of attacks from user-oriented to organization-oriented in order to understand how vulnerable organizations are to cyber attacks within OSNs. Third, this is the first study that has evaluated a cyber attack on organizations within more than one OSN. Fourth, to the best of our knowledge this is the first study that offers a generic algorithm for infiltrating specific users within OSNs by means of socialbots.

1.2 Organization

The remainder of this paper is organized as follows: In Section 2 we provide a brief overview of the studies that focused on similar issues that we have discussed in this study. Section 3 describes the experimental framework and methods we used in order to infiltrate specific employees in targeted organizations within two different OSNs. Furthermore, this section includes the algorithm we developed to infiltrate specific employees, the obstacles we faced in this project, and the datasets we used in order to evaluate our method. Section 4 presents our numeric results. Section 5 includes ethical considerations that arose during this study. Section 6 presents a discussion regarding the results, and Section 7 presents our conclusions and future research directions.

2 Related Work

Over the past few years, the rapid development of OSNs worldwide has increased the number of studies regarding privacy issues of users within OSNs. In this section, we present several studies related to the issues we have focused on in our study to provide helpful background information and additional insights. Below we outline studies involving the crawling of OSNs, the use of socialbots, and the identification of socialbots.

2.1 Crawling Social Networks

In recent years, there have been several studies that have utilized crawling methods on OSNs to retrieve large amounts of information (*27*), (*34*), (*35*). In 2007, Chau et al. (*34*) implemented crawlers on an auction website. They were able to visit approximately 11 million online auction users, of which 66,000 were completely crawled. In 2008, Mislove et al. (*36*) examined growth data from the Flickr[7] OSN. They crawled Flickr once per day for three months and were able to identify 950,143 new users and over 9.7 million links. Also in 2008, Cheng et al. (*37*) crawled the YouTube[8] social network for four months

[7] https://www.flickr.com
[8] https://www.youtube.com

and collected more than three million YouTube videos. In 2009, Cha et al. (*38*) gathered and analyzed traces of information disseminated from the Flickr OSN. In their experiments they crawled markings of 2.5 million users on 11 million photos. In 2010, Kwak et al. (*35*) crawled Twitter[9] and gathered 41.7 million user profiles, 1.47 billion social relations, 4,262 trending topics, and 106 million tweets.

Most recently, Fire et al. (*27*) analyzed different types of organizations through data mining. The analysis was based on information that organizations' employees revealed on Facebook, LinkedIn, Google search results, their company's web page, and other publicly available sources. To accomplish their goal, Fire et al. designed and built a web crawler. The web crawler extracted the network of informal social relationships of employees of a given target organization by optimizing data collection from users associated with a specific group or organization. They collected publicly available data from six well-known high-tech companies on three different size scales using the new web crawler.

2.2 Utilizing Socialbots

Socialbots are defined as automatic or semi-automatic computer programs that take over OSN accounts and perform human behavior such as sending friend requests, messages, etc. (*16*), (*39*). Unlike a regular bot, such as a Twitter bot or spambot, a socialbot hides the fact that it is a robot. This robot typically is programmed to infiltrate communities within OSNs and pass itself off as a human being. Malicious users maintain socialbots to steal or attain sensitive information they could not otherwise access, or to reach an influential position in order to spread misinformation or propaganda (*32*), (*33*).

In 2009, Bilge et al. (*13*) performed crawling and identity theft attacks against several OSNs, including Xing, StudiVZ,[10] MeinVZ,[11] Facebook, and LinkedIn in order to collect personal and sensitive information. The attacks were divided into two parts. The first attack was a classic identity theft in which several profiles of victims had been cloned; then, cloned profiles sent friend requests to their contacts. By being accepted to the friends' networks, they were able to access sensitive information of users within

[9] `https://twitter.com`
[10] `http://www.studivz.net`
[11] `http://www.meinvz.net`

the same network. The second attack was more complex. In this case they launched a cross-site profile cloning attack, recognizing users who had a profile account in one OSN and not in another. They cloned the profile account of the victim in the first OSN and forged a new one in another OSN where he or she was not registered. By this method, Bilge et al. rebuilt the OSN of the victim by contacting his or her friends in both OSNs.

Also in 2009, Bonneau et al. (*40*) described two techniques for working with false profiles: first, creating false profiles in several networks; second, sending friend requests to "highly connected users who are more likely than average to accept a friend request from a stranger."

In 2010, Ryan (*41*) conducted the Robin Sage Experiment in which false identities were created on various OSNs. Ryan succeeded in exploiting strangers' trust based on occupation, gender, education, and connections. Eventually, Robin gained hundreds of friends on various OSNs, including executives at government entities such as NSA, military intelligence groups, friends from Global 500 corporations, etc.

In 2011, Boshmaf et al. (*16*) built a network of socialbots and operated them on Facebook for eight weeks. Their results included three main conclusions. First, OSNs like Facebook can be infiltrated with a success rate of up to 80%. Boshmaf et al. were able to demonstrate that the more friends a user has, the more likely he or she is to respond positively to friend requests. Second, depending on users' privacy settings, successful infiltration can result in privacy breaches greater than if the users' data were exposed by purely public access. Third, for all practical purposes, the main security defense of Facebook - the Facebook Immune System (FIS) - was found to be poor for detecting or stopping large-scale infiltration campaigns.

Moreover, in 2012, Wagner et al. (*39*) performed an experiment in which three teams built several socialbots to influence user behavior on Twitter. Eventually, they were able to develop models for identifying gullible users and for predicting the users' susceptibility level.

Also in 2012, Elyashar et al. (*22*), concerned about organizational data leakage that had been exposed

by employees in OSNs, presented a method for the mining of data of a targeted organization by using OSNs and socialbots. Their method was based on accepting a socialbot's friend request, from which users unknowingly exposed information about themselves and their workplaces. They evaluated the described method within two organizations using Facebook. Elyashar et al. were successful in discovering up to 13.55% more employees and 18.29% more informal organizational links in the crawling process by the "friendly" socialbots in contrast to users without any friends.

2.3 Detecting Socialbots

Along with studies that tried to gather leaked information about users within OSNs and even in some cases to infiltrate them, there have been several studies that have attempted to suggest solutions to these privacy issues. The solutions were based on socialbot detection. Quick identification of malicious users within OSNs may help innocent users as well as OSN operators to defend themselves against these malicious profiles. The techniques for detecting these malicious users varies and includes identification using machine learning (*42*), (*43*), (*44*), (*45*) as well as creating honeypots for attracting spammers (*46*), (*47*).

In 2010, Benevenuto et al. (*42*) tried to detect spammers on Twitter by collecting a large dataset of the Twitter OSN and classifying users into spammers and non-spammers by machine learning techniques. Eventually, they succeeded in identifying approximately 70% of spammers and 96% of non-spammers.

Also in 2010, Chu et al. (*43*) showed interest in identifying malicious bots that spread spam or harmful content. They classified human, bot, and cyberblog accounts on Twitter in terms of tweeting behavior, tweet content, and account properties. Their classifications were based on a crawled collection of 50,000 Twitter accounts. Eventually, they designed an automated classification system.

In 2012, Fire et al. (*44*) presented a method for the detection of spammers and fake profiles in OSNs using OSN topological features. The proposed method was based on a combination of graph theory algorithms and machine learning, and it has been evaluated on datasets of three different OSNs.

In 2011, Stein et al. (*45*) described the FIS, which, like the human immune system, protects Facebook users from attacks. FIS is an adversarial learning system that performs real-time tests and classifications on every read and write action on Facebooks database. These tests exist in order to protect Facebook users and the entire social network from malicious activities.

Besides machine learning detection methods, honeypots were used when Lee et al. (*46*) proposed a honeypot-based approach in 2010 for discovering OSNs spammers. Their approach utilized social honeypots within MySpace and Twitter OSNs in order to attract spammers to attack. With these honeypots they developed statistical user models in order to distinguish between social spammers and legitimate users. Eventually, their honeypots succeeded in identifying social spammers with low false positive rates.

Likewise in 2010, Stringhini et al. (*47*) created several "honey-profiles" on three large OSNs: Facebook, MySpace, and Twitter. Later, they analyzed the collected data and were able to identify anomalous behavior of users. Eventually, they were able to detect and delete spam profiles.

3 Methods and Experiments

As opposed to similar studies that have been focused infiltrating OSN users without making any distinction between low-ranked and high-ranked users using socialbots, our current study offers a unique algorithm for infiltrating specific OSN users who are currently working or had worked in targeted organizations. We evaluate our algorithm using two OSNs: Facebook and Xing.

Most OSNs require users to create accounts in order to establish connections with other network users. In many cases, a user's account may include personal data such as photographs, birthday, hometown, ethnicity, and personal interests (*16*). In most undirected OSNs, users connect to one another by initiating friend requests. The recipient must accept the friend request in order to establish a friend link with the initiator of the request. When the process of link establishment has been completed, the two parties acquire the privilege of accessing each other's profile details whenever they please. Therefore,

we define accepting a friend request as an infiltration. Our study established socialbots on the OSNs, and after network users accepted these socialbots' friend requests, we received increased access to users' profiles and were able to gather additional information about the user and, in some cases, information about the user's friends and their friends.

To infiltrate specific employees of organizations, we had to take several actions (see Algorithm 1). First, in launching the infiltration process, we had to crawl in targeted organizations and gather public information about employees who had established user accounts and stated that they were working or had worked in the targeted organizations. For the crawling process, we created a public user account within the targeted OSN (Facebook and Xing). We used this account and a crawler similar to the crawler introduced by Fire et al. (*27*) in order to crawl the targeted network and identify an organization's employees (line 1).

Algorithm 1 Socialbot Organizational Infiltration

Input: Uids - a set of seed URLs to Facebook profile pages of an organization's employees,
S - socialbot,
O - targeted organization,
TU - targeted users,
OG - organization's graph

1: $OG \leftarrow Org - Crawler(S, Uids, O)$
2: $i \leftarrow 0$
3: **while** (i<10) **do**
4: $TU \leftarrow ChooseRandomTargetedUsers(O)$
5: $i \leftarrow i + 1$
6: **end while**
7: **while** ($NumOfFriends(S) \leq 50$) **do**
8: SendFriendRequestToRandomUsers(S)
9: **end while**
10: $TUFriends \leftarrow FindOrgFriends(O, TU, OG)$
11: **for** $f \in TUFriends$ **do**
12: SendFriendRequestInDescendingOrder(f)
13: **end for**
14: **for** $U \in TU$ **do**
15: SendFriendRequest(U)
16: **end for**

Second, by the end of the crawling process, we had gained sufficient information about users who

were working or had worked in the targeted organizations and their connections. In other words, we had intelligence on the targeted organizations' employees. We used this intelligence in order to randomly select ten users to serve as targets for infiltration (lines 2-6). We utilized their mutual friends in order to infiltrate the targeted users.

Third, in the analysis process we created a socialbot account for every organization that we planned to infiltrate. Prior to the infiltration of a targeted organization, we designed user profiles to look like reliable profiles of real users. Initially, we manually suggested friend requests to random users who had more than 1,000 friends regardless of any organizational affiliation (lines 7-9). We based this action on the observations of Boshmaf et al. (*16*): "the more friends they had, the higher the chance was that they accepted a friend request from a socialbot (i.e., a stranger)."

After a socialbot succeeded in gaining a positive response to its friend request from 50 random users, it automatically sent friend requests to targeted users' mutual friends who were employees in the same organization. Note that we waited for our socialbot to gain 50 friends from random users because we wanted our socialbots to look as much like real users as possible. We did not send friend requests when our socialbot had only a small number of friends because users tend to be suspicious when a user with few friends initiates a friend request; this lack of friends might cause other users to automatically reject our socialbot's friend requests (*22*).

Fourth, our goal was to use our socialbots' friend requests in order to become friends with specific employees in the targeted organizations. In this study, we specified the platform in which each socialbot would act; i.e., we provided a different targeted organization for each socialbot: A socialbot called S will try to infiltrate specific employees in an organization called O, in which S currently has no friends. The main goal was not to simply infiltrate a new organization, but rather to infiltrate specific, targeted employees inside an organization.

Fifth, for each chosen targeted user, we identified his or her friends inside the organization, and our socialbot sent friend requests to them (lines 10-13). The process of sending friend requests was

handled in descending order based on the targeted user's number of friends: at first the socialbot sent friend requests to the most "friendly" users, i.e., those with the largest number of friends, and at the end, requests were sent to the users with the fewest friends in the targeted organization. The idea was to gain as many mutual friends as possible, an accomplishment that would increase the probability that a socialbot's friend request would be accepted by the targeted user. It is important to mention that we did not send a friend request to the targeted user until we had finished sending friend requests to his or her mutual friends inside the organization.

Sixth, after we completed the process of sending friend requests to targeted users' mutual friends, we sent friend requests to the ten targeted users (lines 14-16) and observed how many of them accepted our socialbot's friend requests.

3.1 Overcoming Obstacles

The process of infiltrating the OSNs and infiltrating specific users employed in targeted organizations revealed several obstacles. The complicated infiltration process contains several sub-processes including crawling the targeted organization's OSN, creating infiltration profiles for our socialbots, and finally infiltrating specific employees of organizations. Each of these sub-processes faces challenges. The operators or administrators of OSNs should take into account the dangers that can be caused by these processes and actively try to prevent them.

3.1.1 Adjusting the Crawling Process

The crawling phase is a fundamental step in the general infiltration process. This phase includes several sequential actions.

First, we had to select an organization to target. The selected organization could be any organization or company that has employees who use Facebook or Xing OSNs. In Xing, it is much easier to find employees by a given organization because this OSN is designated for professionals, who in most cases expose their workplaces. We could also make use of existing pages that were created and operated

by organizations in order to present themselves and display their activities publicly. In many cases, these pages included minimal information (for example, a small list of employees) that were information sources for the crawling process.

After we chose our targeted organizations, we had to create public user accounts in order to initiate crawling on the OSNs. With these public accounts, we did not try to infiltrate, but rather to crawl on the OSNs and find user profiles, which fit the given criteria. Because of their different goals, we did not define these accounts as socialbots. These accounts did not include social properties except for a name and an image of an animal of some kind. Moreover, they did not send any friend requests to anyone.

When we finished creating a public user account, and after we had already chosen a targeted organization, we had to connect to Facebook or Xing with the created public user account and operate the crawler. As the crawler ran, it provided us with the profiles as text or HTML files, and it also gave a list of connections between the users. The full implementation is described by Fire et al. (*27*).

In the described procedure, the crawler was able to collect hundreds of users in just a few hours. The crawler ran for days and actually downloaded thousands of OSN user profiles. We did, however, face the obstacle of having user profiles blocked. The OSN providers identified the behavior of our crawler as an anomaly because it is uncommon for a user to surf hundreds of profiles so rapidly. Therefore, a few public users were blocked by the OSN providers. We subsequently added time-outs and delays in order to slow down or even stop the crawler's running for several hours, thus overcoming this obstacle.

3.1.2 Creating Realistic Infiltration Profiles

The process of creating infiltration socialbot accounts at Facebook and Xing was largely done manually. The challenge faced at this stage was to avoid suspicion and look like a real OSN user; otherwise our socialbots' friend requests would have failed. In order to prevent failure, we chose common names for our socialbot accounts with the intent of looking familiar to other users. Then we had to select images for socialbots' profiles.

In contrast to Boshmaf et al. (*16*), for female users on Facebook we selected obscure images of real women, such as images without a face, in order to make recognition unfeasible. For male users, we chose images of cute puppies, fancy cars, etc.

In order to look like an authentic user, we added interests and other properties for each socialbot. The properties were "likes" to popular singers and movies, posting beautiful nature images, adding posts to the user's timeline, etc.

We based the selection of profile images for our Xing socialbots on a unique profile image which included a suit and a tie in order to look like a professional employee. The image was based on profile images of 2-3 real people, which we then combined to create a new image of a person who does not exist in reality.

3.1.3 Infiltrating Specific Employees

The first phase of becoming friends with 50 random users, who each have more than 1,000 friends, passed without any special problems in Facebook. Users with more than a thousand friends tend to accept strangers' friendship requests, and the high percentage of friends we gained helped our Facebook socialbots look like real users to other potential friends. However, because of the small size of the Xing OSN compared to Facebook, it was hard for us to find 50 random users who have more than 1,000 connections. In order to overcome this obstacle, we reduced the threshold of the number of friends a potential user needed from 1,000 to 400 in Xing. Moreover, one of the parameters that the OSN providers use to detect suspicious activity is the community structure of users. Therefore, socialbots are allowed sending friend requests to users only in a limited number of such communities. OSN providers use algorithms for detecting anomalies in suspicious account's network topology. According to Bosmaf et al. (*16*), users who connect randomly with strangers may be fake profiles. Hence, most of the legitimate users are connected only to a small number of communities. Fake profiles, on the other hand, tend to establish friendships with users from a large number of different communities (*44*). In order to overcome

these obstacles, we sent few friend requests in the beginning of the process to users who had a large number of friends. After a "friendly user" accepted the socialbot's friend request, our next friend requests were sent to mutual friends of the "friendly user"- who also had a large number of friends.

Although we did not face significant obstacles in the first phase, the next phase presented numerous challenges. This second phase involved sending friend requests to mutual friends of the targeted users in the targeted organization. In this phase we needed to avoid being blocked or disabled by the OSN, a process that could be activated based on how many users accepted (or declined) our socialbots' friend requests. A low acceptance percentage on Facebook, as well as on Xing, can trigger the OSN's anomaly detection mechanism. This can prompt a warning and a decline of our friend requests. For example, if Facebook suspects misuse, i.e., that we really do not know the specific user to whom we sent a friend request, Facebook may send us a message indicating that they have decided to halt the friend request in order to prevent misuse. When Facebook realizes that the warnings do not interfere with the socialbots and a low rate of acceptance persists, Facebook may then block the socialbot. Once blocked, Facebook requires the socialbot to provide a real phone number to verify its identity in order to avoid being disabled or losing access to Facebook.

On Xing, we did not receive a warning prior to being blocked. In this case, a user logs on to Xing and simply receives a message that the account has been deactivated. Additionally, as a measure of protection, Xing counts the number of unconfirmed users, i.e., users who did not confirm a specific user. When the count reaches 100 unconfirmed users, Xing prevents the user from sending any additional friend requests. In order to overcome this obstacle and send additional requests, we moved unconfirmed contacts to a bookmarked list. This procedure enabled sending additional requests to users.

3.2 Datasets

We utilized our presented algorithm on two different OSNs: Facebook and Xing. Facebook is the most popular OSN in the world and Xing is a well-known European OSN that is designed for use by business

professionals.

3.2.1 Facebook

With more than 1.28 billion monthly active users as of March 31, 2014, Facebook stands out among all other popular OSNs in the world. According to Facebook, there are 1.01 billion active monthly users of its mobile products. On average, Facebook has 802 million active daily users, and 81.2% of them are outside the U.S. and Canada (*8*). The average Facebook user has around 190 friends (*48*). Additionally, Facebook users have made 140.3 billion friend connections and used over 1.13 trillion "likes" (*49*).

According to Facebook's estimations (*50*), 8.7% of its accounts are defined as fake. This means that Facebook includes tens of millions of fake accounts. Moreover, 4.8% of Facebook accounts are defined as duplicate accounts - ones that a user maintains in addition to his or her principal account. Furthermore, 2.4% of Facebook users are defined as user-misclassified accounts, referring to users who have created personal profiles for a business, organization, or non-human entity such as a pet. Of Facebook's fake accounts, 1.5% are defined as undesirable accounts, i.e., belonging to malicious users. Undesirable accounts are defined as fake accounts which were created with the intent of being used for purposes that violate Facebook's terms of service, such as spamming or distributing other malicious links and content.

Regarding security issues, Facebook uses FIS (*45*) and provides an additional privacy settings tool which enables users to edit their profile and decide which information will be accessible to others. However, this tool matched Facebook users' expectations only 63% of the time. Furthermore, the Facebook privacy settings tended to be more open than users' expectations (*51*).

Targeted Organizations. We decided to infiltrate specific employees who stated on their Facebook profile that they were employees of one of four targeted organizations. Eventually, we succeeded in infiltrating three of these organizations: O_{F1}, O_{F2}, and O_{F3}.

O_{F1} organization is an international software company that develops, licenses, implements, and supports software applications for its customers. According to public sources, O_{F1} organization employs

thousands of employees and operates offices in North America, Europe, and the Middle East.

In the crawling process we identified 2,199 informal links of 330 Facebook users who, according to their Facebook profiles, work or have worked in this organization (see Table 1 and Figure 1).

Table 1: Organizational datasets statistics

Organizations	Social Networks	Nodes	Edges
O_{F1}	Facebook	330	2,199
O_{F2}	Facebook	469	3,831
O_{F3}	Facebook	918	10,986
O_{X1}	Xing	107	369
O_{X2}	Xing	1,237	14,408
O_{X3}	Xing	416	4,153

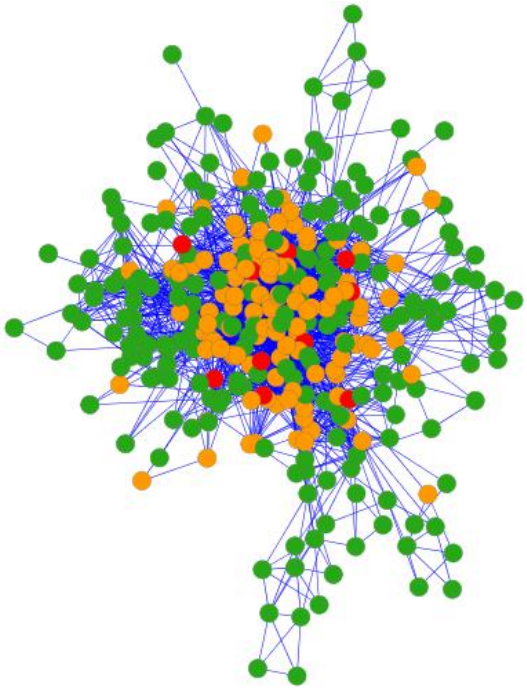

Figure 1: O_{F1} organizational social network - Red nodes represent targeted users and orange nodes represent users who received friend requests.

The O_{F2} organization is a leading information technology company that specializes in the integration, development, and application of technologies, solutions and software products, hardware, infrastructure, etc. The company is located in Eastern Europe and the Middle East and has thousands of

employees around the world.

In the crawling process we discovered 3,831 informal links of 469 Facebook users who, according to their Facebook profiles, work or have worked in this organization (see Table 1 and Figure 2).

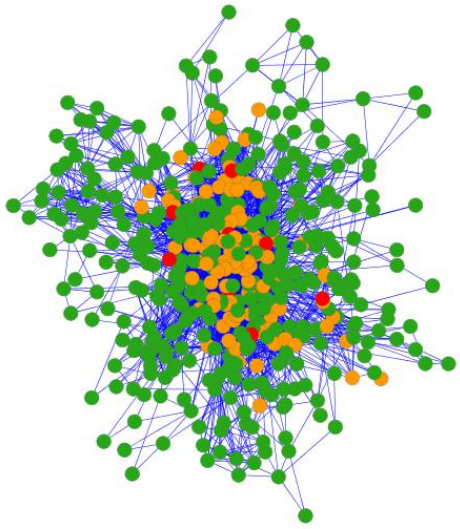

Figure 2: O_{F2} organizational social network - Red nodes represent targeted users and orange nodes represent users who received friend requests.

The O_{F3} organization is a telecommunications networking product provider that provides communication products and develops services for carriers, cable/multiple system operators, wireless/cellular service providers, etc. O_{F3} organization is located in the Middle East and the Far East, and according to public sources, employs thousands of employees.

In the crawling process we identified 10,986 informal links of 918 Facebook users who, according to their Facebook profiles, work or have worked in this organization (see Table 1 and Figure 3).

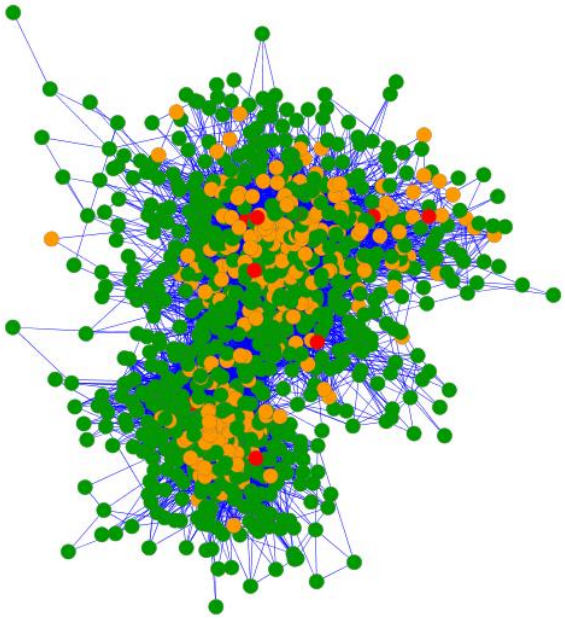

Figure 3: O_{F3} organizational social network - Red nodes represent targeted users and orange nodes represent users who received friend requests.

3.2.2 Xing

Xing is a social network for business professionals. It was founded in Hamburg, Germany, in 2003 and has been publicly listed since 2006. Xing has around 14 million members worldwide, 7 million of whom are based in German-speaking countries (*52*). Most Xing users use this OSN to promote their businesses, boost their career, or find a job.

Furthermore, Xing provides a suitable platform for professionals to meet, find jobs, connect with colleagues, collaborate, share new ideas, etc. (*52*). Regarding security issues, Xing provides information about how active a given user is (based on the frequency of a user's visits), and this information is referred to as activity. With this data we can assess the activity level of Xing users.

Moreover, Xing keeps track of the number of unconfirmed users, i.e., users who did not confirm a specific user. When a user is unconfirmed 100 times, Xing prevents the user from sending additional friend requests (*53*). To the best of our knowledge Xing does not publish statistics regarding the number

of fake users within the network.

Targeted Organizations. On Xing, we used five socialbots for specific user infiltration of five different organizations. Our socialbots sent friend requests to the targeted users' mutual friends in order to gain as many mutual friends as possible, with the goal to eventually be accepted by the targeted users themselves. Eventually, we succeeded in infiltrating two of the organizations. We present the results from the attempts of S_{X1}, S_{X2}, S_{X3}, S_{X4}, and S_{X5} socialbots that tried to infiltrate specific employees in targeted organizations O_{X1}, O_{X2}, and O_{X3} respectively (see Section 4). S_{X4}, and S_{X5} had been blocked in the first stage so they did not attack an organization.

The O_{X1} organization is a large crude oil and natural gas producer. The activities of this organization include exploration and production of oil and natural gas as well as the natural gas trade, including transport and storage. This organization is based in Eastern and Western Europe, North Africa, and South America. O_{X1} employs thousands of employees worldwide.

In the crawling process, we identified 369 informal links of 107 Xing users who, according to their profiles, work or have worked in this organization (see Table 1 and Figure 4).

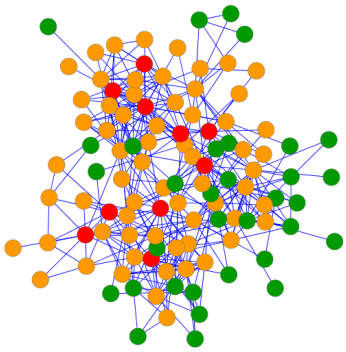

Figure 4: O_{X1} organizational social network - Red nodes represent targeted users and orange nodes represent users who received friend requests.

The O_{X2} organization is a food company that serves as a manufacturer, retailer, and marketer of beverage concentrates. The organization is based in the U.S but has branches worldwide. It employs

hundreds of thousands of employees worldwide, and in this study we focused on the European branch of O_{X2} organization.

In the crawling process we identified 14,408 informal links of 1,237 Xing users who, according to their profiles, work or have worked in this organization (see Table 1 and Figure 5).

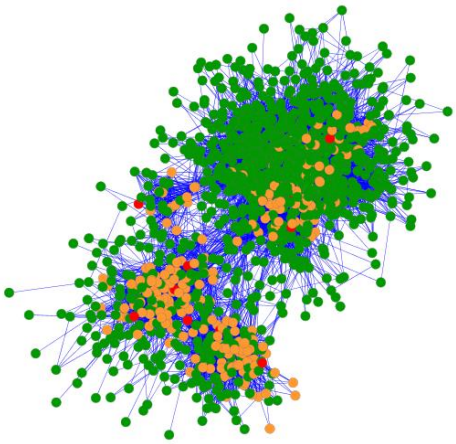

Figure 5: O_{X2} organizational social network - Red nodes represent targeted users and orange nodes represent users who received friend requests.

The O_{X3} organization is a corporation that develops, manufactures, and sells computer software, electronics, etc. This organization is American, but has lots of branches worldwide. O_{X3} has hundreds of thousands of employees worldwide, and for this study we focused on the European branch.

In the crawling process we identified 4,153 informal links of 416 Xing users who, according to their Facebook profiles, work or have worked in this organization (see Table 1 and Figure 6).

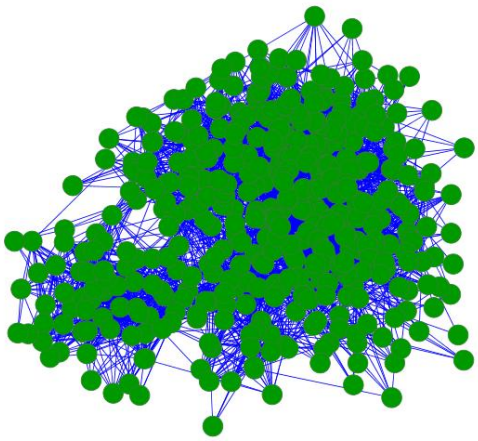

Figure 6: O_{X3} organizational social network.

4 Results

In this section we present the results of our study based on the utilization of our suggested algorithm (see Algorithm 1). The following results are categorized by the OSNs used and by the organizations selected to infiltrate.

4.1 Facebook Infiltration

We used four socialbots on Facebook for specific user infiltration of four different organizations. Our socialbots sent friend requests to targeted users' mutual friends in order to gain as many mutual friends as possible to help be accepted by the targeted users. We present the results we gained from the S_{F1}, S_{F2}, and S_{F3} socialbots that attempted to infiltrate specific employees in targeted organizations O_{F1}, O_{F2}, and O_{F3} respectively (The FIS disabled the S_{F4} socialbot, which was going to infiltrate organization O_{F4}).

4.1.1 S_{F1} Socialbot

By the end of the infiltration process, S_{F1} sent 124 friend requests to 124 users in the O_{F1} organization (including the ten targeted users). Among them, 46 users accepted, and 78 users rejected S_{F1}'s requests

(see Figure 1).

First, we randomly chose ten users who stated on their Facebook profiles that they work or had worked in the O_{F1} organization. We then collected the friends of the ten targeted users who also work in the O_{F1} organization and sent them friend requests. Next, socialbot S_{F1} sent friend requests to the ten targeted users (TU_1- TU_{10}). In total, socialbot S_{F1} sent 124 friend requests and was successful in connecting with 46 different users (see Table 2).

Table 2: O_{F1}, O_{F2}, and O_{F3} targeted users summary results

Organization	Targeted Users	Accepted/ All Friends	Acceptance Percentage	Accepted?
O_{F1}	TU_1	5/13	38.46%	**Yes**
	TU_2	4/13	30.76%	**No**
	TU_3	5/16	31.25%	**No**
	TU_4	6/16	37.50%	**No**
	TU_5	6/17	35.29%	**Yes**
	TU_6	21/42	50%	**Yes**
	TU_7	7/21	33.33%	**No**
	TU_8	4/14	28.57%	**Yes**
	TU_9	7/13	53.84%	**No**
	TU_{10}	13/32	40.62%	**Yes**
	Total	**46/124**	**37.09%**	**50%**
O_{F2}	TU_1	5/12	41.16%	**Yes**
	TU_2	6/11	54.54%	**Yes**
	TU_3	7/17	41.17%	**Yes**
	TU_4	5/16	31.25%	**No**
	TU_5	11/25	44%	**Yes**
	TU_6	6/12	50%	**No**
	TU_7	8/19	42.10%	**Yes**
	TU_8	6/22	27.27%	**No**
	TU_9	8/21	38.10%	**Yes**
	TU_{10}	5/15	33.33%	**Yes**
	Total	**38/114**	**33.33%**	**70%**
O_{F3}	TU_1	11/24	45.83%	**No**
	TU_2	17/34	50%	**Yes**
	TU_3	6/22	27.27%	**Yes**
	TU_4	8/18	44.44%	**No**
	TU_5	16/45	35.55%	**No**
	TU_6	19/44	43.18%	**No**
	TU_7	18/42	42.85%	**No**
	TU_8	11/24	45.83%	**Yes**
	TU_9	4/14	28.57%	**No**
	TU_{10}	19/36	52.77%	**Yes**
	Total	**88/219**	**40.18%**	**40%**

With regard to the targeted users, S_{F1} was able to become a friend of five targeted users (TU_1, TU_5, TU_6, TU_8, and TU_{10}), giving a success rate of 50% (see Table 3 and Figure 7). Moreover, S_{F1} was able to become a friend of 37.09% of all users who received friend requests (see Table 2).

Table 3: Facebook socialbots total results

Facebook Socialbots	Mutual Friends	Targeted Users
	Accepted/Total	Accepted/Total
S_{F1}	46/124	5/10
S_{F2}	38/114	7/10
S_{F3}	88/219	4/10

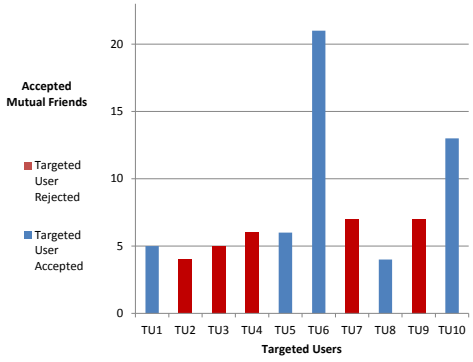

Figure 7: O_{F1}'s targeted users - the numbers represent how many mutual friend S_{F1} had before it sent a friend request to the targeted users. A blue column represents a targeted user who accepted S_{F1} whereas a red column represents a targeted user who rejected S_{F1} friend request.

4.1.2 S_{F2} Socialbot

S_{F2} sent 114 friend requests to 114 users (including the ten targeted users). Among them, 38 users accepted, and 76 users rejected S_{F2}'s requests (see Table 3, and Figure 2).

Regarding the targeted users, socialbot S_{F2} was able to become a friend of seven targeted users (TU_1, TU_2, TU_3, TU_5, TU_7, TU_9, and TU_{10}), with a success rate of 70% (see Table 2 and Figure 8). Moreover, S_{F2} was able to become a friend of 33.33% of all the users who received friend requests (see Table 2).

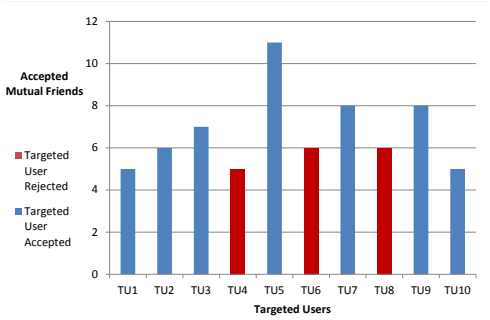

Figure 8: O_{F2}'s targeted users - the numbers represent how many mutual friend S_{F2} had before it sent a friend request to the targeted users. A blue column represents a targeted user who accepted S_{F2} whereas a red column represents a targeted user who rejected S_{F2} friend request.

4.1.3 S_{F3} Socialbot

First, in order to look like a real user S_{F3} sent 58 friend requests to random users with more than 1,000 friends. Among them, 33 users accepted S_{F3}'s friend requests, and 19 users asked S_{F3} to be their friend. This means that S_{F3} reached the threshold of 50 users in the first three days (see Figure 9[12]). Overall, S_{F3} gained 129 random users and achieved an acceptance rate of 56.9% (see Table 4).

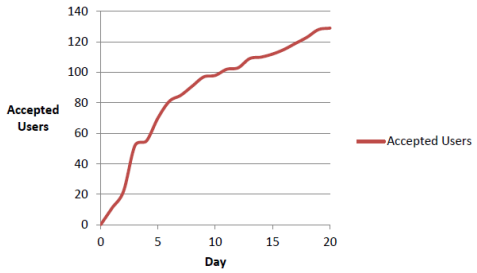

Figure 9: S_{F3}'s Random Users Accumulation.

[12]Due to a failure of stored information we could not build similar graphs for S_{F1} and S_{F2}

Table 4: S_{F3}'s random users accumulation summary results

Day	Accepted Users			Total Sent Requests	Percent
	Users Who Got Request	Users Who Sent Request	Total		
1	11	0	11	17	65%
2	9	2	11	20	45%
3	13	17	30	21	62%
4	0	3	3	0	
5	0	15	15	0	
6	0	11	11	0	
7	0	4	4	0	
8	0	6	6	0	
9	0	6	6	0	
10	0	1	1	0	
11	0	4	4	0	
12	0	1	1	0	
13	0	6	6	0	
14	0	1	1	0	
15	0	2	2	0	
16	0	3	3	0	
17	0	4	4	0	
18	0	4	4	0	
19	0	5	5	0	
20	0	1	1	0	
Total	**33**	**96**	**129**	**58**	**56.90%**

In the infiltration process S_{F3} sent 219 friend requests to 219 users in the O_{F3} organization (including the ten targeted users). Among them, 88 users accepted, and 131 users rejected S_{F3}'s requests (see Table 2, and Figure 3).

Regarding the targeted users, S_{F3} was able to become a friend of 4 targeted users (TU_2, TU_3, TU_8, and TU_{10}), with a success rate of 40% (see Table 2 and Figure 10). Moreover, S_{F3} was able to become a friend of 40.18% of all the users who received friend requests in the infiltration process (see Table 2).

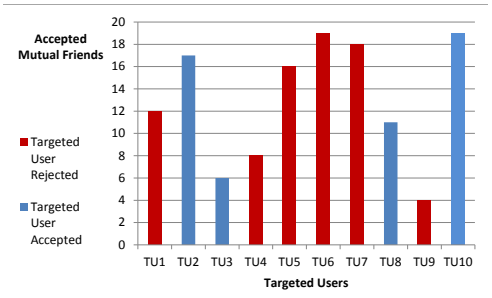

Figure 10: O_{F3}'s targeted users - the numbers represent how many mutual friend S_{F3} had before it sent a friend request to the targeted users. A blue column represents a targeted user who accepted S_{F3} whereas a red column represents a targeted user who rejected S_{F3} friend request.

4.2 Xing Infiltration

We used five socialbots on Xing for specific user infiltration on five different organizations. To achieve these goals, our socialbots sent friend requests to targeted users' mutual friends in order to gain as many mutual friends as possible to help be accepted by the targeted users. In total, our Xing socialbots sent 850 friend requests to 850 Xing users. Among them, 439 accepted our socialbots' friend requests (51.64% acceptance rate). We present the results we gained from the Xing socialbots S_{X1}, S_{X2}, and S_{X3}, which attempted to infiltrate specific employees in the targeted organizations O_{X1}, O_{X2}, and O_{X3} respectively. Socialbots S_{X4} and S_{X5} were blocked by Xing.

4.2.1 S_{X1} Socialbot

First, S_{X1} sent 101 friend requests to random users with more than 400 friends. Among them, 68 users accepted S_{X1}'s friend requests, creating an acceptance rate of 67.33% in six days (see Table 5 and Figure 11).

Table 5: S_{X1}, S_{X2}, S_{X3}, S_{X4}, and S_{X5} random users summary results

Socialbot	Day	Accepted Users			Total Sent Requests	Percent
		Users Who Got Requests	Users Who Sent Request	Total		
S_{X1}	1	13	0	13	20	65%
	2	11	0	11	20	55%
	3	13	0	13	20	65%
	4	13	0	13	20	65%
	5	17	0	17	20	85%
	6	1	0	1	1	100%
	Total	**68**	**0**	**68**	**101**	**67.33%**
S_{X2}	1	13	0	13	20	65%
	2	18	0	18	20	90%
	3	18	0	18	19	95%
	4	13	0	13	19	68%
	5	9	0	9	9	100%
	6	0	1	1	0	
	Total	**71**	**1**	**72**	**87**	**81.61%**
S_{X3}	1	7	0	7	15	47%
	2	12	0	12	20	60%
	3	12	0	12	20	60%
	4	17	0	17	20	85%
	5	11	1	12	15	73%
	Total	**59**	**1**	**60**	**90**	**65.56%**
S_{X4}	1	12	0	12	20	60%
	2	12	0	12	20	60%
	3	9	0	9	20	45%
	4	8	0	8	15	53%
	Total	**41**	**0**	**41**	**75**	**54.67%**
S_{X5}	1	6	0	6	15	40%
	2	5	0	5	10	50%
	3	10	0	10	15	67%
	4	5	0	5	20	25%
	5	11	0	11	20	55%
	6	8	0	8	20	40%
	Total	**45**	**0**	**45**	**100**	**45%**

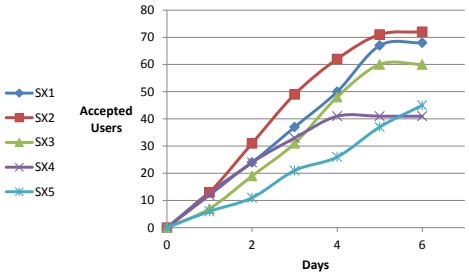

Figure 11: Xing socialbots random users accumulation.

By the end of the infiltration process, S_{X1} sent 71 friend requests to 71 users in the O_{X1} organization (including the ten targeted users). Among them, only ten users accepted (14.08%), and 61 users rejected S_{X1}'s requests (see Table 6).

Table 6: O_{X1}, O_{X2}, and O_{X3} targeted users summary results

Organization	Targeted Users	Accepted/ All Friends	Acceptance Percentage	Accepted?
O_{X1}	TU 1	1/8	12.50%	**No**
	TU 2	1/11	9.09%	**Yes**
	TU 3	3/14	21.42%	**No**
	TU 4	0/11	0.00%	**No**
	TU 5	2/9	22.22%	**Yes**
	TU 6	1/11	9.09%	**No**
	TU 7	1/8	12.50%	**No**
	TU 8	2/15	13.33%	**No**
	TU 9	4/24	16.67%	**No**
	TU 10	0/9	0.00%	**No**
	Total	**10/71**	**14.08%**	**20%**
O_{X2}	TU 1	9/24	37.5%	**No**
	TU 2	20/41	48.78%	**No**
	TU 3	6/12	50%	**Yes**
	TU 4	10/29	34.48%	**Yes**
	TU 5	8/19	42.10%	**No**
	TU 6	18/27	66.67%	**Yes**
	TU 7	18/26	69.23%	**Yes**
	TU 8	20/41	48.78%	**Yes**
	TU 9	15/30	50%	**No**
	TU 10	19/34	55.88%	**Yes**
	Total	**123/241**	**51.04%**	**60%**
O_{X3}	TU 1	5/11	45.5%	**No**
	TU 2	3/10	30%	
	TU 3	7/29	24.13%	
	TU 4	6/19	31.57%	
	TU 5	4/15	26.67%	
	TU 6	2/11	18.18%	
	TU 7	2/12	16.67%	
	TU 8	2/21	9.52%	
	TU 9	1/12	8.33%	
	TU 10	3/9	33.33%	
	Total	**22/85**	**25.88%**	

With regard to targeted users, S_{X1} was able to become a friend of two targeted users (TU_2, and TU_5), with a success rate of 20% (see Tables 6, 7, and Figure 12). Moreover, S_{X1} was able to become a friend of 14.08% of all users who received friend requests (see Table 6).

Table 7: Xing socialbots total results

Xing Socialbots	Mutual Friends	Targeted Users
	Accepted/Total	Accepted/Total
S_{X1}	10/71	2/10
S_{X2}	123/241	6/10
S_{X3}	22/85	

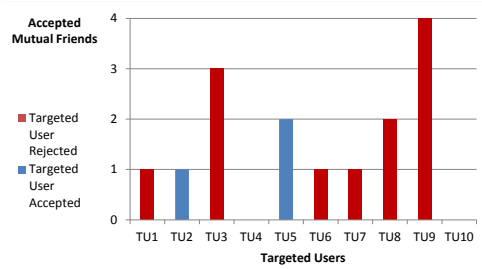

Figure 12: O_{X1}'s targeted users - the numbers represent how many mutual friend S_{X1} had before it sent a friend request to the targeted users. A blue column represents a targeted user who accepted S_{X1} whereas a red column represents a targeted user who rejected S_{X1} friend request.

4.2.2 S_{X2} Socialbot

S_{X2} sent 87 friend requests to random users with more than 400 friends. Among them, 71 users accepted S_{X2}'s friend requests, generating an acceptance rate of 81.61% in five days (see Table 5 and Figure 11).

By the end of the infiltration process, S_{X2} sent 241 friend requests to 241 users in the O_{X2} organization (including the ten targeted users). Among them, 123 users accepted, and 118 users rejected S_{X2}'s requests (see Table 6 and Figure 5).

With regard to targeted users, S_{X2} was able to become a friend of six targeted users (TU_3, TU_4, TU_6, TU_7, TU_8, and TU_{10}), with a success rate of 60% (see Tables 6 and 7, and Figure 13). Moreover, S_{X2} was able to become a friend of 51.04% of all users who received friend requests (see Table 6).

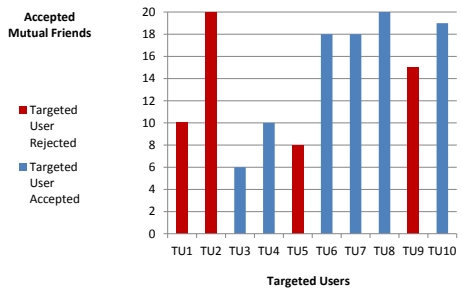

Figure 13: O_{X2}'s targeted users - the numbers represent how many mutual friend S_{X2} had before it sent a friend request to the targeted users. A blue column represents a targeted user who accepted S_{X2} whereas a red column represents a targeted user who rejected S_{X2} friend request.

4.2.3 S_{X3} Socialbot

In order to look like a real user, S_{X3} sent 90 friend requests to random users with more than 400 friends. Among them, 59 users accepted S_{X3}'s friend requests, creating an acceptance rate of 65.56% in five days (see Table 5).

By the end of the infiltration process, S_{X3} sent 85 friend requests to 85 users in the O_{X3} organization before it was blocked by Xing. Among the users who received friend requests, 22 users accepted and 63 users rejected S_{X3}'s requests - an acceptance rate of 25.88% in five days (see Table 6 and Figure 11).

4.2.4 S_{X4} Socialbot

S_{X4} did not manage to reach to the threshold of 50 friends when sending friend requests to random users with more than 400 friends (see Figure 11); in the middle of this process it was blocked by Xing. S_{X4} did manage to send 75 friend requests to random users with more than 400 friends before it was blocked. Among them, 41 users accepted S_{X4}'s friend requests. producing an acceptance rate of 54.67% in four days (see Table 5).

4.2.5 S_{X5} Socialbot

S_{X5} was not able to reach the threshold of 50 friends when sending requests to users with more than 400 friends (see Figure 11). As was the case with S_{X4}, in the middle of the infiltration process, S_{X5} was blocked by Xing. S_{X5} managed to send 100 friend requests to random users before it was blocked, and 45 users accepted S_{X5}'s friend requests. This created an acceptance rate of 45% in six days (see Table 5).

5 Ethical Consideration

Today, most OSNs do not allow free access to personal information due to the privacy concerns of network users and the OSN's terms of use (*54*). As a result, much of the research using OSNs involves

various techniques of collecting sensitive data by circumventing OSN limitations. Elovici et al. (*54*) performed a comprehensive review of research involving OSNs. They described two kinds of OSN research: "Whitehat" research is defined as legitimate academic and industrial investigation. "Blackhat/greyhat" research is defined as studying and exploiting vulnerabilities of OSNs in order to extract sensitive information, actively connect to users, create fake identities, and even perform malicious activities. The actions that researchers have to do in order to get this kind of data are controversial and raise many ethical questions.

In order to perform accurate blackhat/greyhat research, a researcher must base his or her study on actual OSN data, such as real connections between users, correct textual content, authentic files, etc. Researchers must monitor many real-life OSN users in order to study the diffusion of data in OSNs. An effective and widely employed technique to obtain data from OSNs and their users is based on establishing connections with users, typically by creating a large number of fake OSN user accounts, which are then used to connect other users.

Moreover, Elovici et al. claimed that the main goal of academic blackhat/greyhat research is precise purpose: to study vulnerabilities in order to create improved defenses for OSNs and their users.

Given the very nature of OSNs, we should ask the following question: Is it ethical to perform research such as ours? We believe that the answer to this question is positive for several reasons.

First, over the period we have been conducting this study, Ben-Gurion University of the Negev, did not question our work or require approval by the research ethics committee in order to conduct the study. Furthermore, we have made great efforts in order to increase the standard of ethics in the domain of OSNs (*54*), (*10*), (*55*).

Second, in the initial crawling process of our study on Facebook and Xing, we collected only publicly available data that is accessible to every registered user.

Third, we avoided using profile images of real users when creating the identities of our socialbots. Instead, we selected profile images that either did not include users' faces or that presented the faces in

such a way that it was impossible to identify the person, or we photoshoped images of fictional entities.

Fourth although this study included real OSN users, and the results may inadvertently provide knowledge of concern to OSN users and operators, ignoring the problem does not provide a solution. We can rid ourselves of responsibility for these challenging situations and choose not to perform such research; however, the problem will continue to exist and in fact, increase. Performing this type of research aids the development of new forms of protection by OSNs.

Lastly, given the enormous large number of OSN users and the extensive opportunities to exploit the personal information of each and every one of them, it is crucial to study the dangers and privacy issues that exist for users of OSNs. As aforementioned, Facebook itself has estimated that 8.7% of its accounts are defined as fake (*50*), and certainly some percentage of these are malicious. OSN users need to be aware of the online dangers that exist and modify their actions accordingly. Online security represents an acute problem that must be studied by legitimate researchers in order to be effectively addressed by industry and the academia.

6 Discussion

Using the methods described in Section 3, our socialbots were able to infiltrate specific employees in three different targeted organizations within the Facebook OSN and two targeted organizations within the Xing OSN. However, we should extend our discussion beyond these successful infiltrations.

First, there are points of comparison between the socialbots that completed the infiltration process successfully and the ones that were exposed and eventually blocked by the OSNs. In total, we operated four socialbots within Facebook. Three of four socialbots completed the infiltration process. S_{F1}, S_{F2}, and S_{F3} gained success rates of 50%, 70%, and 40% respectively (see Table 3). Socialbot S_{F4}, on the other hand, was blocked by Facebook operators in the middle of the infiltration process. We believe that the failure of S_{F4} lies in the *location feature*: namely, the fact that there was a great geographic distance between the targeted organization O_{F4} and S_{F4} socialbot's current location attribute. The air

travel distance between the location of most of O_{F4} organization employees and S_{F4}'s current location was more than 4,000 kilometers, spanning several countries. It is important to note that the difference between the targeted organization's employees and the socialbot's current location existed only in the case of O_{F4}; in all other cases, the location of the socialbot was similar to the location of the targeted organization. We can assume from this incident that likely there are cultural differences, which logically correspond to geographic distance, between OSN users, and such differences can raise suspicions when accepting friend requests. We hope to verify this assumption in a future study. Therefore, we decided to fit the socialbot's identity to its target. Moreover, this incident reinforces the conclusion of Liben-Nowell et al. (*56*) regarding their finding of a strong correlation between friendship and geographic location in LiveJournal.[13]

With regard to our infiltration of Xing, we operated five socialbots. Among them, two of the five completed the infiltration process. S_{X1} and S_{X2} succeeded in infiltrating specific employees; however, S_{X3}, S_{X4}, and S_{X5} failed in their mission. We believe that the failure in these three cases resulted from the socialbot's *organizational affiliation*. Our successful socialbots, S_{X1} and S_{X2}, were defined as users who were not connected directly to the targeted organizations: S_{X1} was defined as a freelancer coach, whereas S_{X2} was defined as a recruiter in a technology-oriented organization. They were able to infiltrate specific employees in targeted organization with the success rate of 20% and 60% (see Tables 6 and 7). In contrast to these two socialbots, we defined S_{X3} as an employee within the O_{X3} organization. S_{X3} succeeded in passing the first stage, the accumulation of 50 random users (see Table 5), but once S_{X3} began initiating friend requests to employees of O_{X3}, suspicions arose. Several suspicious users used their organization's human resource software to verify S_{X3}'s false identity. Some of them even notified us that they knew for sure that he or she was not an employee in their organization. A few days later, the S_{X3} socialbot was blocked by Xing. This described activity by the employees using organizational software tools for recognition prior to accepting a friend request demonstrates how a well-formed policy

[13] `http://www.livejournal.com`

and clear instructions to employees can benefit the security of employees and their organizations on OSNs.

As far as S_{X4} and S_{X5}, they failed in the middle phase of accumulating of 50 random users. We defined these two socialbots as employees in two random organizations. It is important to mention that these two random organizations were not the targeted organizations that we wanted to infiltrate, but separate organizations we had chosen for the identity of the socialbots. While trying to evade detection by Xing operators based on the communities structure of OSN users (see Section 3.1.3), our socialbot's friend requests became too embedded within the random organizations: Among the 75 friend requests that were sent by S_{X4}, 28 users were employees in the randomly chosen organization (37%). Similarly, among 100 friend requests that were sent by S_{X5}, 58 users were employees in the randomly chosen organization (58%). S_{X4} as well as S_{X5} were blocked by Xing operators when employees from these organizations verified the identities of these two socialbots.

Another issue we wish to discuss is the low infiltration of Xing socialbot S_{X1}, which was able to infiltrate only 20% of the targeted users (see Tables 6 and 7). We believe that this poor performance was related to the fact that among 61 users who did not respond S_{X1}'s friend requests were 25 users who had 0% activity. This means that 41% of the users who did not respond to S_{X1}'s friend requests were inactive, i.e, users who did not get friend requests in reality. It is important to understand that we made a point to not avoid these zero-activity users in our study; we sent friend requests to users who had 0% activity regardless in order to fairly evaluate their role.

Lastly, we want to consider the recommended threshold of what number of mutual friends would influence a specific user to accept our socialbot's friend request. Our previous study (*23*) found that a socialbot will typically be accepted as a friend of specific employees when it has gained seven or more mutual friends of the targeted user. This number of mutual friends corresponds to the results illustrated in Figures 7 and 8. However, when expanding the study to include one more organization on Facebook (see Figure 10) and two organizations on Xing (see Figures 12 and 13), we can suggest that the threshold is

increased to 17 or more mutual friends, and the probability that a targeted user will accept our socialbot is 70%. There are 7 targeted users with 17 or more mutual friends who accepted our socialbots' friend requests: one targeted user on O_{F1}, two targeted users on O_{F3}, and four targeted users on O_{X2}. There are three targeted users with 17 or more mutual friends who rejected our socialbots' friend requests: two targeted users on O_{F3}, and one targeted user on O_{X2}. These suggestions are further reinforced by the conclusions of Boshmaf et al. (*16*) that the more a user's mutual friends accept the socialbot's requests, the more likely the targeted user is to accept the socialbot's friend request.

7 Conclusion

In this study, we demonstrated an attack of socialbots which were able to infiltrate specific employees in targeted organizations within two different OSNs. This further emphasizes the persistent privacy issues surrounding OSNs. Based on our results, we can draw the following conclusions and recommendations regarding the infiltration of specific employees.

First, we were able to show that OSN users still tend to accept friend requests from complete strangers. Most of our socialbots in both OSNs were able to reach the threshold of 50 random users within 5-6 days (see Figures 9 and 11). OSN users should realize how easy it is for an attacker to create socialbots, and how insignificant is the price that the attacker must pay for having socialbots blocked. In case a socialbot is blocked, the attacker can quickly create a new, improved fake profile and continue with the infiltration process. OSN users must understand the risks of accepting friend request from those they do not know and should ignore a friend request received from strangers.

Second, our experimental results indicate that there is a link between having mutual friends and the acceptance of friend requests. The step of first being accepted as a friend of a mutual friend of the targeted users in targeted organizations was significant to our socialbots' ability to infiltrate. Without this step, the acceptance rate would have been much lower (see Tables 2 and 6). Moreover, we found that if a socialbot has 17 or more mutual friends of the targeted user, the probability that the targeted user will

accept the socialbot's friend request is 70% (see Section 6).

Third, malicious socialbots operate within OSNs and can function at a high level of sophistication. As we learned from executing our suggested algorithm, most of our socialbots were able to infiltrate specific employees in targeted organizations on both Facebook and Xing, despite the differences between these two OSNs. Please note that we intend to further explore cases of socialbot blocking in a future study. Furthermore, our results of the S_{F4} socialbot may indicate that users tend to trust strangers on the basis of their mutual attributes like current location, mutual friends, etc. We recommend users not to rely upon these mutual features when a stranger asks for their friendship.

Fourth, organizations should understand the risks of organizational information leakage that might occur due to their employees using OSNs. Moreover, we strongly recommend that organizations should take responsibility for raising the level of awareness of employees to this problematic phenomenon and for underscoring the risks of employees accepting unfamiliar users as friends. Organizations should explain to OSN users the risks of data leakage and provide them with tools to verify users who declare themselves to be employees in the organizations. This kind of software can help employees verify whether or not the stranger who sends them a friend request is a real employee. This recommendation also endorses a recommendation by Fire et al. (*10*) that suggested performing a short security check on a stranger.

This study has several future research directions. One possible direction is more thorough testing of the conclusions we found regarding the cultural differences between users and organizational affiliation of socialbots when we define their identity. Another possible direction is to use the algorithm presented in this paper to investigate whether our results are consistent over time and to assess whether there are changes in users' awareness and responses to privacy issues. Moreover, we can use the algorithm on other OSNs and observe the differences between them. Furthermore, we could differentiate between female and male profiles when infiltrating OSNs to investigate any gender differences that exist.

In both the present and the future, individuals and organizations need to be aware that harmful so-

cialbots exist on OSNs, and consequently they must use social networks wisely and should establish effective security and privacy measures.

Acknowledgements

The authors would like to thank Robin Levy-Stevenson for proofreading this article. Especially, we want to thank Carol Teegarden for her editing expertise and endless helpful advice, which guided this article to completion. We also want to thank the anonymous reviewer for his or her valuable comments and suggestions to improve the manuscript.

References and Notes

1. A. Ocass and T. Fenech, "Web retailing adoption: exploring the nature of internet users Web retailing behaviour," *Journal of Retailing and Consumer services*, vol. 10, no. 2, pp. 81–94, 2003.

2. E. F. Gross, "Adolescent Internet use: What we expect, what teens report," *Journal of Applied Developmental Psychology*, vol. 25, no. 6, pp. 633–649, 2004.

3. A. Lenhart, K. Purcell, A. Smith, and K. Zickuhr, "Social Media & Mobile Internet Use among Teens and Young Adults. Millennials.," *Pew Internet & American Life Project*, 2010.

4. J. A. Diaz, R. A. Griffith, J. J. Ng, S. E. Reinert, P. D. Friedmann, and A. W. Moulton, "Patients' Use of the Internet for Medical Information," *Journal of general internal medicine*, vol. 17, no. 3, pp. 180–185, 2002.

5. D. M. Boyd and N. B. Ellison, "Social Network Sites: Definition, History, and Scholarship," *Journal of Computer-Mediated Communication*, vol. 13, no. 1, pp. 210–230, 2007.

6. R. E. Wilson, S. D. Gosling, and L. T. Graham, "A Review of Facebook Research in the Social Sciences," *Perspectives on Psychological Science*, vol. 7, no. 3, pp. 203–220, 2012.

7. K. Manning, "The impacts of Online Social Networking and Internet Use on Human Communication and Relationships,"

8. "Facebook Newsroom." http://newsroom.fb.com/Key-Facts/, Sept. 2013. (last accessed on May 2nd, 2014).

9. "Our List Of The World's Largest Social Networks Shows How Video, Messages, And China Are Taking Over The Social Web." `http://www.businessinsider.com/the-worlds-largest-social-networks-2013-12`, Dec. 2013. (last accessed on June 3rd, 2014).

10. M. Fire, R. Goldschmidt, and Y. Elovici, "Online Social Networks: Threats and Solutions Survey," *arXiv preprint arXiv:1303.3764*, 2013.

11. S. B. Barnes, "A privacy paradox: Social networking in the United States," *First Monday*, vol. 11, no. 9, 2006.

12. S. Mahmood, "Online Social Networks: Privacy Threats and Defenses," in *Security and Privacy Preserving in Social Networks*, pp. 47–71, Springer, 2013.

13. L. Bilge, T. Strufe, D. Balzarotti, and E. Kirda, "All Your Contacts Are Belong to Us: Automated Identity Theft Attacks on Social Networks," in *Proceedings of the 18th international conference on World wide web*, pp. 551–560, ACM, 2009.

14. 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*, pp. 1145–1146, ACM, 2009.

15. S. Ghosh, B. Viswanath, F. Kooti, N. K. Sharma, G. Korlam, F. Benevenuto, N. Ganguly, and K. P. Gummadi, "Understanding and Combating Link Farming in the Twitter Social Network," in *Proceedings of the 21st international conference on World Wide Web*, pp. 61–70, ACM, 2012.

16. Y. Boshmaf, I. Muslukhov, K. Beznosov, and M. Ripeanu, "The Socialbot Network: When Bots Socialize for Fame and Money," in *Proceedings of the 27th Annual Computer Security Applications Conference*, pp. 93–102, ACM, 2011.

17. A. Mislove, B. Viswanath, K. P. Gummadi, and P. Druschel, "You Are Who You Know: Inferring User Profiles in Online Social Networks," in *Proceedings of the third ACM international conference on Web search and data mining*, pp. 251–260, ACM, 2010.

18. J. Baltazar, J. Costoya, and R. Flores, "The Real Face of KOOBFACE: The Largest Web 2.0 Botnet Explained," *Trend Micro Research*, vol. 5, no. 9, p. 10, 2009.

19. Q. Cao, M. Sirivianos, X. Yang, and T. Pregueiro, "Aiding the Detection of Fake Accounts in Large Scale Social Online Services," in *Proc. of NSDI*, 2012.

20. G. Stringhini, G. Wang, M. Egele, C. Kruegel, G. Vigna, H. Zheng, and B. Y. Zhao, "Follow the Green: Growth and Dynamics in Twitter Follower Markets," in *Proceedings of the 2013 conference on Internet measurement conference*, pp. 163–176, ACM, 2013.

21. Y. Boshmaf, I. Muslukhov, K. Beznosov, and M. Ripeanu, "Design and Analysis of a Social Botnet," *Computer Networks*, 2012.

22. A. Elyashar, M. Fire, D. Kagan, and Y. Elovici, "Organizational Intrusion: Organization Mining using Socialbots," 2012 International Conference on Social Informatics (SocialInformatics), (Lausanne, Switzerland).

23. A. Elyashar, M. Fire, D. Kagan, and Y. Elovici, "Homing Socialbots: Intrusion on a specific organizations employee using Socialbots," SNAA - Social Network Analysis in Applications 2013, (Niagara Falls, Ontario, Canada), 2013.

24. J. Wolak, D. Finkelhor, K. J. Mitchell, and M. L. Ybarra, "Online 'Predators' and Their Victims," *Psychology of violence*, vol. 1, pp. 13–35, 2010.

25. M. L. Ybarra and K. J. Mitchell, "How Risky Are Social Networking Sites? A Comparison of Places Online Where Youth Sexual Solicitation and Harassment Occurs," *Pediatrics*, vol. 121, no. 2, pp. e350–e357, 2008.

26. M. Balduzzi, C. Platzer, T. Holz, E. Kirda, D. Balzarotti, and C. Kruegel, "Abusing Social Networks for Automated User Profiling," in *Recent Advances in Intrusion Detection*, pp. 422–441, Springer, 2010.

27. M. Fire, R. Puzis, and Y. Elovici, "Organization Mining Using Online Social Networks," *arXiv preprint arXiv:1303.3741*, 2013.

28. "The Economic Impact of Cybercrime and Cyber Espionage," McAfee, jul 2013.

29. J. Kostka, Y. A. Oswald, and R. Wattenhofer, "Word of Mouth: Rumor Dissemination in Social Networks," in *Structural Information and Communication Complexity*, pp. 185–196, Springer, 2008.

30. M. Nekovee, Y. Moreno, G. Bianconi, and M. Marsili, "Theory of rumour spreading in complex social networks," *Physica A: Statistical Mechanics and its Applications*, vol. 374, no. 1, pp. 457–470, 2007.

31. K. Peterson and K. A. Siek, "Analysis of Information Disclosure on a Social Networking Site," in *Online Communities and Social Computing*, pp. 256–264, Springer, 2009.

32. J. Ratkiewicz, M. Conover, M. Meiss, B. Gonçalves, A. Flammini, and F. Menczer, "Detecting and Tracking Political Abuse in Social Media," 2011.

33. J. Ratkiewicz, M. Conover, M. Meiss, B. Gonçalves, S. Patil, A. Flammini, and F. Menczer, "Truthy: Mapping the Spread of Astroturf in Microblog Streams," in *Proceedings of the 20th international conference companion on World wide web*, pp. 249–252, ACM, 2011.

34. D. H. Chau, S. Pandit, S. Wang, and C. Faloutsos, "Parallel Crawling for Online Social Networks," in *Proceedings of the 16th international conference on World Wide Web*, pp. 1283–1284, ACM, 2007.

35. H. Kwak, C. Lee, H. Park, and S. Moon, "What is Twitter, a Social Network or a News Media?," in *Proceedings of the 19th international conference on World wide web*, pp. 591–600, ACM, 2010.

36. A. Mislove, H. S. Koppula, K. P. Gummadi, P. Druschel, and B. Bhattacharjee, "Growth of the Flickr Social Network," in *Proceedings of the first workshop on Online social networks*, pp. 25–30, ACM, 2008.

37. X. Cheng, C. Dale, and J. Liu, "Statistics and Social Network of YouTube Videos," in *Quality of Service, 2008. IWQoS 2008. 16th International Workshop on*, pp. 229–238, IEEE, 2008.

38. M. Cha, A. Mislove, and K. P. Gummadi, "A Measurement-driven Analysis of Information Propagation in the Flickr Social Network," in *Proceedings of the 18th international conference on World wide web*, pp. 721–730, ACM, 2009.

39. C. Wagner, S. Mitter, C. Körner, and M. Strohmaier, "When social bots attack: Modeling susceptibility of users in online social networks," 2012.

40. J. Bonneau, J. Anderson, and G. Danezis, "Prying Data out of a Social Network," in *Social Network Analysis and Mining, 2009. ASONAM'09. International Conference on Advances in*, pp. 249–254, IEEE, 2009.

41. T. Ryan, "Getting In bed with Robin Sage,"

42. F. Benevenuto, G. Magno, T. Rodrigues, and V. Almeida, "Detecting Spammers on Twitter," in *Collaboration, electronic messaging, anti-abuse and spam conference (CEAS)*, vol. 6, p. 12, 2010.

43. Z. Chu, S. Gianvecchio, H. Wang, and S. Jajodia, "Who is Tweeting on Twitter: Human, Bot, or Cyborg?," in *Proceedings of the 26th annual computer security applications conference*, pp. 21–30, ACM, 2010.

44. M. Fire, G. Katz, and Y. Elovici, "Strangers Intrusion Detection - Detecting Spammers and Fake Profiles in Social Networks Based on Topology Anomalies," *HUMAN*, vol. 1, no. 1, pp. pp–26, 2012.

45. T. Stein, E. Chen, and K. Mangla, "Facebook Immune System," in *Proceedings of the 4th Workshop on Social Network Systems*, p. 8, ACM, 2011.

46. K. Lee, J. Caverlee, and S. Webb, "Uncovering Social Spammers: Social Honeypots + Machine Learning," in *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval*, pp. 435–442, ACM, 2010.

47. G. Stringhini, C. Kruegel, and G. Vigna, "Detecting Spammers on Social Networks," in *Proceedings of the 26th Annual Computer Security Applications Conference*, pp. 1–9, ACM, 2010.

48. J. Ugander, B. Karrer, L. Backstrom, and C. Marlow, "The Anatomy of the Facebook Social Graph," *arXiv preprint arXiv:1111.4503*, 2011.

49. "Facebook Newsroom - One billion key metrics." `http://newsroom.fb.com/imagelibrary/downloadmedia.ashx?MediaDetailsID=4227&SizeId=-1`, Sept. 2013. (last accessed on Jan 4th, 2014).

50. "Facebook Company Filings." `http://www.sec.gov/Archives/edgar/data/1326801/000119312512325997/d371464d10q.htm`, June 2013.

51. Y. Liu, K. P. Gummadi, B. Krishnamurthy, and A. Mislove, "Analyzing Facebook Privacy Settings: User Expectations vs. Reality," in *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*, pp. 61–70, ACM, 2011.

52. "Xing Corporate Pages." https://corporate.xing.com/no_cache/english/company/xing-ag/, Aug. 2013.

53. "Tweaking Social Networks." http://tweakingsocialnetworks.wordpress.com/tag/bookmark-unconfirmed-contacts-in-xing/, Aug. 2010.

54. Y. Elovici, M. Fire, A. Herzberg, and H. Shulman, "Ethical Considerations when Employing Fake Identities in Online Social Networks for Research," *Science and engineering ethics*, pp. 1–17, 2013.

55. M. Fire, D. Kagan, A. Elyashar, and Y. Elovici, "Friend or Foe? Fake Profile Identification in Online Social Networks," *arXiv preprint arXiv:1303.3751*, 2013.

56. D. Liben-Nowell, J. Novak, R. Kumar, P. Raghavan, and A. Tomkins, "Geographic routing in social networks," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 102, no. 33, pp. 11623–11628, 2005.