Jiang J, Shan ZF, Wang X *et al*. JOURNAL OF COMPUTER SCIENCE AND TECHNOLOGY : 1–
Mon. Year

Understanding Sybil Groups in the Wild

Jing Jiang[1,2] (蒋 竞), *Member, CCF*, Zifei Shan[2] (单子非), Xiao Wang[2] (王 潇)
Li Zhang[1] (张 莉), *Senior Member, CCF*, Yafei Dai[2] (代亚非), *Distinguished Member, CCF*

[1]*State Key Laboratory of Software Development Environment Beihang University, Beijing, China*
[2]*Department of Computer Science and Technology, Peking University, Beijing, China*

E-mail: {jiangjing, lily}@buaa.pku.edu.cn, {shanzifei, wangxiao, dyf}@pku.edu.cn

Received April 25, 2014

Revised September 27, 2014

Abstract Sybil attack is one of the well-known and powerful attacks against online social networks (OSNs). In a sybil attack, a malicious attacker generates a sybil group consisting of multiple sybil users, and controls them to attack the system. However, data confidentiality policies at major social network providers have severely limited researchers' access to large-scale datasets of sybil groups. A deep understanding of sybil groups can provide important insights into questions of malicious behavior, as well as numerous practical implications on the design of security mechanisms. In this paper, we present an initial study to measure sybil groups in a large-scale OSN, Renren. We analyze sybil groups at different levels, including individual information, social relationships and malicious activities. Our main observations are: (1) User information in sybil groups are usually incomplete and in poor quality. (2) Sybil groups have special evolution patterns in connectivity structure, including bursty actions to add nodes, and monotonous merging pattern that lacks non-singleton mergings. (3) Several sybil groups have strong relationships with each other and compose sybil communities, and these communities cover a large number of users and pose great potential threat. (4) Sybil users are banned a long time after

This work is supported by National Science Foundation of China under the National Basic Research Program of China under Grant No.2011CB302305, National Natural Science Foundation of China under Grant No.61300006, and State Key Laboratory of Software Development Environment under Grant No.SKLSDE-2013ZX-26

registration in some sybil groups. The characteristics of sybil groups can be leveraged to improve the security mechanisms in OSNs, to defend against sybil attack. Specifically, we suggest that OSNs should (1) check information completeness and quality with our approach, (2) learn from dynamics of community connectivity structure to detect sybil groups, (3) monitor sybil communities and inspect them carefully to prevent collusion, and (4) inspect sybil groups that behave normally even for a long time to prevent potential malicious behaviors.

Keywords Online Social Networks, Measurement, Security, Sybil Groups, Sybil attack

1 Introduction

In recent years, online social networks (OSNs) become huge and they are still growing throughout the world [15]. Unfortunately, the openness and the tremendous growth of OSNs attract the interest of malicious parties. Sybil attack is one of the well-known and powerful attacks against OSNs. The malicious attacker generates a sybil group consisting of multiple accounts (called sybil users), and disguise them into different users. Sybil groups even collude together, build close relationships and generate communities. Sybil attacks are serious threats to OSNs: Multiple sybil users are utilized to unfairly increase influence and power of target users [24]. Moreover, sybil attackers target OSNs as media to propagate spam [4, 7, 9, 11, 12, 21, 23, 26, 27, 30]. Sybil attacks become increasingly dangerous as more people use OSNs as primary interfaces to the Internet [16].

Various techniques were applied to identify sybil users or spammers in OSNs, including rare social links between sybil users and normal users [8, 28, 32, 36, 37], honeypots [6, 12, 18, 26, 31], manual identification [2, 3, 24] and abnormal behavior [20, 27, 30, 34, 35]. Sybil users alone do not harm the system. What is really dangerous is that multiple sybil users are controlled together to form a sybil group [13]. Some research took initial steps and designed algorithms to detect sybil groups [7, 13].

Initial studies mainly designed methods to identify sybil users or sybil groups, but few works have measured characteristics of sybil groups in the wild. The difficulty of large-scale measurements is primarily due to the lack of datasets. The providers of some online social networks generally consider their data to be trade secrets, and have few incentives to make such data available for research. However, a deep understanding of sybil groups and sybil users can provide important insights into questions of malicious behavior, as well as numerous

practical implications on the design of security mechanisms for social networks. For example, understanding properties of sybil groups can help sybil detection and increase the cost for attackers to disguise sybil users as humans; lessons from how sybil groups are formed and connected can guide the design of new, more effective mechanisms for sybil attacks.

In this paper, we present a first-of-its-kind study to measure sybil groups in OSNs. Our work is based on datasets from Renren social network , one of the largest and oldest OSNs in China. In our previous work [13], we built a sybil group detector and designed automatic validation mechanisms. We successfully detected and confirmed 2440 sybil groups and 985,797 sybil users. These large-scale datasets in the wild are notable and provide us precious opportunities to study characteristics of sybil groups.

We analyzed the data to answer three key questions:

1. How do sybil groups provide individual information? Is the information complete and in high quality?

2. What is the connectivity inside sybil groups? How are relationships between different groups?

3. Do sybil groups take malicious actions? What is the characteristics of malicious activities?

To answer these essential and practical questions, we analyze the sybil groups with focus on three levels: individual information, social relationships and malicious activities. At each level, our data analysis reveals several facts about sybil groups.

Individual information. We study the completeness and quality of various user information. We discover that sybil groups have low completeness of optional user information. Moreover, some individual information has poor quality, and it is easily identified as fake by manual. These results show that the completeness and quality of user information can be utilized to improve the security mechanisms against sybil attack. Groups with low completeness and poor quality of user information can be identified as suspicious by the security mechanism.

Social relationships. we focus on the social aspect, and make measurement of relationships inside the sybil group, and between different sybil groups. We find that sybil groups have special evolution patterns in connectivity structure, including bursty actions to add nodes, and monotonous merging pattern that lacks

http://www.renren.com

non-singleton mergings. It implies that dynamics of community connectivity structure can be considered to modify the security mechanisms in future. We discover that several sybil groups build strong relationships among themselves, and they have a large number of users.These sybil groups have great potential threats to the system. If these sybil groups collude together, they all take malicious actions in a short time, which is dangerous in the system. Online social networks should pay special attention and monitor these sybil groups together.

Malicious activities. We study the action time and the type of malicious activities for sybil groups and sybil users. We discover that sybil users are banned a long time after registration in some sybil groups. Even if some sybil groups and sybil users behave normally for a long time, they should still be carefully monitored, in case of sudden attacks. We find sybil users perform different kinds of malicious activities, and thus a single mechanism is not enough to inspect them. It is important to identify sybil groups and sybil users beforehand, and then use various mechanisms to continuously monitor them.

In this paper, we present an initial study to study characteristics of sybil groups in the wild. Results provide new insights for OSNs to improve the security mechanisms, to ensure fairness and credibility in the system, to reduce users' burden of dealing with spam, and to positively impact the overall value of OSNs going forward. In particular, OSNs should utilize properties of whole group to detect sybil groups, rather than simply analyzing features of individual users. Moreover, OSNs should monitor sybil communities and inspect them carefully to prevent collusion. Tight-knit sybil groups can control a large number of sybil users to take malicious actions in a short time, which is dangerous in the system. Even if attackers know these modified mechanisms, attackers need to pay more cost and overhead, so as to change their behavior and adapt to the inspection. As the cost increases greatly, the benefit of sybil groups decreases significantly.

2 Individual information

Before diving into the measurement of sybil groups, we begin by describing the Renren social network and our datasets. We then study characteristics of personal information for sybil groups and sybil users. In Renren, each profile includes user information, such as name, age, phone number, email and profile picture. User information describes personal properties and contact information of the user, which is feasible for friends to know and contact the user. We ask two questions: Do users in sybil

groups fill in their information or just leave default values? If sybil groups fill in information, dose this information seems to be fake or true? Results provide deep understanding of characteristics of sybil groups and sybil users, and practical implications on the design of security mechanisms for social networks in future.

2.1 The Renren Social Network

Renren was set up in 2005, and it is one of the oldest and biggest OSNs in China [14]. Users maintain personal profiles, upload photos and write blogs. Users also establish bidirectional social relationships with friends, view friends' pages and exchange comments. User pages on Renren are similar to Facebook. A user profile includes a profile picture, personal information (name, age, hobbies, etc.), and contact information (phone number, email, etc.) The body of each user's page is a chronologically ordered "feed" of the user's actions: status updates, photos uploaded and tagged, blog entries written, etc.

As used by more and more users, Renren becomes an attractive platform for companies to promote products. Therefore, some attackers create sybil groups to unfairly increase the power of target users in social games, or spread advertisements for companies. Renren has deployed several orthogonal techniques to detect sybil users. In order to further improve security and defend against sybil attacks, Renren has built a collaboration with our research team since December 2010 [34]. To support the project, Renren provides anonymized user data on their servers, which is preprocessed to remove private information.

In our previous work, we built a sybil group detector to identify sybil groups and sybil users. Then we designed automatic validation mechanisms and successfully validated 2440 sybil groups and 985,797 sybil users [13]. These sybil groups constituted large-scale datasets in the wild and provided us precious research opportunities. In this work, we mainly study characteristics of sybil groups, rather than single sybil users. Our validation mechanisms confirmed that users in groups had high similarity of action time and whole groups were sybil. In order to make comparison, we also selected 1,480 normal groups composed by 179,319 normal users as the control group.

In order to protect trade secrets, Renren provides anonymized user data. Due to various detection algorithms and different original datasets provided by Renren, Our datasets are different from datasets in the work [34]. Furthermore, previous work [34] made contributions to sybil users, instead of sybil groups. Though our datasets may not cover all sybil

groups and sybil users, they are valuable for research and allow us to get the initial understanding of sybil groups in the wild.

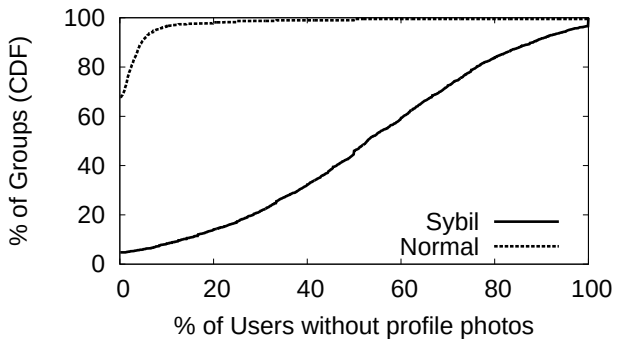

Fig. 1. Percentage of users who do not upload profile photos

2.2 Completeness of user information

User information describes personal properties and contact information. Complete user information allows friends to quickly find and recognize the user from a large number of people. It is also feasible for friends to know and contact the user. Therefore, normal users often fill in their information. However, sybil users are not created for social activities. It remains a question whether sybil users fill in complete information. In order to answer this question, we study several properties of user information. We discover that users in sybil groups always fill in basic user information, such as gender and birthday. This is because basic user information is required in registration process. Then we analyze optional user information and observe that sybil groups are inactive in filling optional information. We choose three typical attributes and introduce their results.

The first feature is the profile photo. The profile photo describes the user's appearance, therefore it is an important feature of a user. The profile photo is widely used to attract attention, and help friends quickly identify the user. The profile photo may be a personal photo or the other picture. For each sybil group as well as each normal group, we compute the percentage of users who do not upload profile photos. Figure 1 shows the distribution of sybil groups and normal groups. 45.9% of sybil groups have less than 50% of users without profile photos, while other 54.1% of sybil groups have more than 50% of users without profile photos. 27.5% of sybil groups have even more than 70% of users without profile photos. In contrast, In 67.7% of normal groups, all the users upload profile photos. 91% of normal groups have only less than 5% of users without profile photos. Normal users always upload photos, while the majority of sybil group-

s do not have profile photos. Uploading profile photos of sybil groups cost more time and bandwidth, therefore attackers are inactive in providing profile photos.

The Mann-Whitney-Wilcoxon (MWW) test is a non-parametric statistical test that assesses the statistical significance of the difference between two distributions [22]. We use the MWW test, because it does not assume any specific distribution, which is suitable for our datasets. In the MWW test, the p-value is used to make decision. Given a significance level $\alpha = 0.001$, if p-value is smaller than α, then the test rejects the null hypothesis, which implies that two datasets have different distributions at the significance level of $\alpha = 0.001$.

We use the MWW test to compare the distribution of sybil groups and normal groups for profile photos. The p-value is only $3.96E-30$, which is much smaller than 0.001. We find that at 0.1% significance level that the difference between sybil groups and normal groups are statistically significant.

The second feature is telephone number. Telephone number is feasible for friends to contact people. Moreover, telephone number can be bound to the account, so that if the password is lost or the account is compromised, the user can request the password through phone message. Telephone number is the important contact information of the user. Although we do not get users' phone numbers, we get the data whether users have filled their phone numbers in the system. For each sybil group as well as normal group, we compute the number of users who fill in telephone numbers, divided by the number of users in the sybil group. Figure 2 shows the percentage of users with telephone numbers. In 65.3% of sybil groups, none of users provide telephone information; In 94.5% of sybil groups, the percentage of users with telephone numbers is less than 2%. In contrast, only 38.4% of normal groups have no users with telephone numbers. 42.3% of normal groups have more than 10% of users who fill in telephone information. We have tested and confirmed that the two distributions of sybil groups and normal groups are significantly different using the MWW test at 0.001 significance level. Sybil groups have much lower percentage of telephone numbers than normal groups. To explain this phenomenon, sybil users are not created for social activities and telephone numbers are unnecessary.

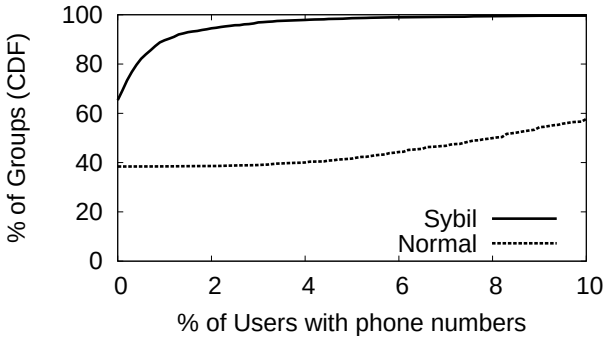

Fig. 2. Percentage of users who fill in their telephone numbers

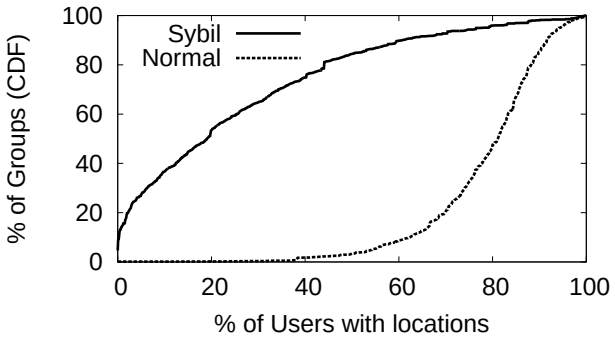

Fig. 3. Percentage of users who fill in their locations

Finally, the third feature is the location. The location describes the city or the province where the user lives. The location is a significant feature for users to find friends from many people with same names. Figure 3 shows the percentage of users with locations for sybil groups and normal groups. 74.8% of sybil groups have less than 40% of users who fill in their locations. In contrast, only 8.8% of normal groups have less than 60% of users with locations, and other 91.2% of normal groups have more than 60% of users with locations. Sybil groups have obviously lower completeness of locations than normal groups. MWW test is used to compare the distribution of sybil groups and normal groups for the completeness of locations. The p-value is only $3.44E-27$, which is much smaller than 0.001. We find that at 0.1% significance level that the difference between sybil groups and normal groups are statistically significant.

Results show that sybil groups have low completeness of user information. This is because that filling in information needs more computational resource, bandwidth and time. In order to save the overhead, attackers are inactive in providing optional information of sybil users. The completeness of user information can be utilized to improve the security system. People may argue that if attackers know the detection mechanisms, attackers can easily change their behavior to adapt to them. However, attackers need to pay additional cost to avoid detection. The inspection of information completeness increases overhead of sybil attacks, and the cost may be even heavier than the profit for attackers.

2.3 Quality of user information

In this subsection, we study the quality of user information for sybil groups. The quali-

ty measures whether the information seems to be fake or true. If the information is carefully fabricated and seems to be true, the quality is high. Obviously-fake information has low quality. Some simple attributes always have high quality, such as gender, location and birthday. No matter the gender is female or male, it is normal and seems to be true. Both sybil groups and normal groups have high quality of these attributes, and they are useless for security mechanism. Other complex attributes may have different quality between sybil groups and normal groups, and we mainly study these attributes in the following part.

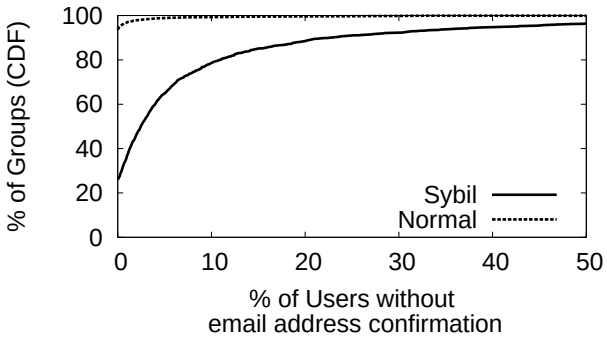

Fig. 4. Percentage of users who do not confirm their email addresses

When a user fills in an email address in the registration, the system sends a confirmation link to this email address. Then the user clicks the link and confirms the email address. Confirmed email addresses are likely to be true, and have better quality than email addresses without confirmation. For each group, we compute the percentage of users without email address confirmation. Figure 4 shows results of sybil groups and normal groups. In 93.6% of normal groups, all users have their email addresses confirmed. In contrast, only 26% of sybil groups have all users with email address confirmed. We again run the MWW test and find that the difference is statistically significant at a significance level of 0.001. These results indicate that sybil groups have low percentage of users who confirm email addresses, Therefore, sybil groups have lower quality of email addresses than normal groups. Since the confirmation of email address gives users special priority, normal users prefer to confirm the e-mail address. However, the confirmation of e-mail address needs additional cost and overhead, and attackers are unlikely to confirm e-mail addresses. Email confirmation is encouraged in OSNs, but it is not necessary. If people do not pass the email confirmation, their activities are limited, but they can still updates their status and publish blogs in Renren. We suggest that in future OSNs may not allow new users to join if they do not confirm their email addresses. This mechanism will guarantee high quality of email addresses and increase cost and price of the creation of sybil groups.

We take a further step and study the quality of the domain name of email address. The

email address is composed by two parts: the part before the @ sign is the username, and the part after the @ sign is the domain name. The domain name describes the organization of the email service. For example, if the email address is Alice@gmail.com, "Alice" is the username, and "gmail.com" is the domain name. It shows that the email service is provided by Google. We categorize domain names into three types: *real, temporary and nonexistent.* The *real* type means the email service really exists, and the email address might be real. This type has the best quality. The *temporary* type means the email address only exists for a short time and expires after a certain time period. Some companies provide the service and allow people to receive emails at temporary addresses. For example, yopmail provides several temporary domain names, such as yopmail.com and courriel.fr.nf. Temporary email addresses are created easier than real emails, and thus they are feasible for attackers to register accounts. The *nonexistent* type means the email service does not exist, and the email address is absolutely fake. The nonexistent type has the worst quality.

The email address was not necessary for the registration in the past, and some sybil users did not fill in email addresses. In total, 710,182 sybil users have email addresses. We extract domain names of email addresses. Many email addresses have the same domain name. After eliminating duplicate domain names, we obtain 2,025 unique domain names. 203 domain names are widely used and they cover 99.3% of email addresses. We manually check 203 domain names and categorize them into above types.

Table 1. Number of email domain names in each type

Type	# of Domain names	# of Email addresses
Real	35	647,212
Temporary	21	19,210
Nonexistent	147	38,532

Table 1 shows the number of domain names and the number of email addresses in each type. Though real type makes up 91.8% of email addresses, it only accounts for 17.2% of domain names. Nonexistent type accounts for 72.4% of domain names. Temporary type still accounts for a non-negligible fraction. We also check domain names of normal users. We find that temporary or nonexistent types of domain names are never used by normal users, and only sybil groups use temporary or nonexistent types. These results also explains why

some sybil groups do not confirm their email addresses. When attackers use nonexistent domain names, these domains are fake and impossible to receive confirmation emails; When attackers use temporary domain names, email addresses may expire and they are also unable to receive confirmation emails. The type of domain name is useful to improve the security system. If the user fills in temporary or nonexistent domain name, the user is suspicious and needs further checks.

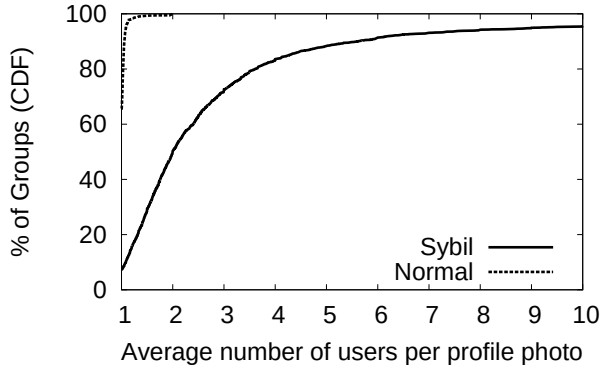

Fig. 5. Average number of users per profile photo

Next, we study the quality of profile photo. It is not easy to automatically judge the quality of a single profile photo. Luckily, we can analyze profile photos from the perspective of a group. The real photo is an authentic personal picture of a person. If a photo is used by multiple users inside a group, this photo is likely to be fake and has poor quality. We compute the number of users divided by the number of unique profile photos in a group. The result is the average number of users per profile photo, and it reflects the similarity of profile photos in a group. Figure 5 shows results of sybil groups and normal groups. In only 7% of sybil groups, a profile photo is used by one user. In 50% sybil groups, a profile photo is used by more than 2 users on average. These profile photos are used by multiple users, and they are likely to be fake and have poor quality. In 94% of normal groups, average number of users per profile photo is less than 1.08. Profile photos in normal groups have low similarity, since most of normal users upload different photos. We make the MWW test and find that the difference is statistically significant at the significance level of 0.001. These results show that sybil groups have obviously worse quality of profile photos than normal groups, which is demonstrated by a much higher profile photo similarity inside groups.

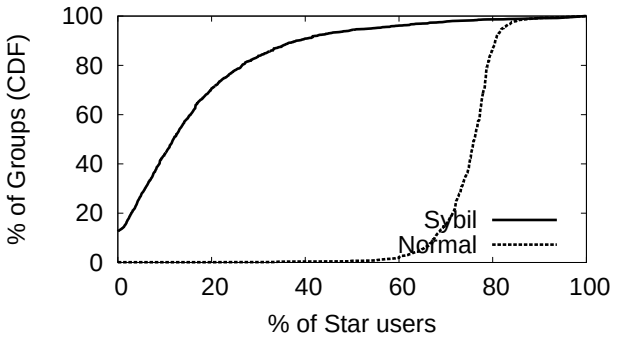

Fig. 6. Percentage of star users

It is not easy to automatically evaluate the quality of some other kinds of user information.

For example, some attackers fill in celebrities' information as personal information. People easily identified that the information belongs to celebrities, instead of sybil users. But it is hard for computers to identify the information as fake. Renren administrators manually check the authenticity of user information. If the user fills in a real Chinese name, provides an authentic personal photo and fill in other real information, then the user is given the title as star user. The star user has some privileges, such as unlimited album space and free gifts. Therefore, normal user prefers to achieve the title of star user. For sybil groups and normal groups, we compute the percentage of star users and show results in Figure 6. In 70.6% of sybil groups, the percentage of star users is less than 20%. The majority of users in sybil groups do not pass the manual check, and their user information has low quality. In contrast, 1.9% of normal groups have the percentage of star users less than 60%, and other 98.1% of normal groups have the percentage of star users more than 60%. Results show that information quality of sybil groups is much worse than that of normal groups. Most of users in normal groups provide real information, and they pass the manual check. In the majority of sybil groups, sybil users seldom get the title of star users, because it requires high quality of user information, which increases the cost and overhead for attackers. MWW test is used to compare the distribution of sybil groups and normal groups for the percentage of stars. The p-value is only $1.93E-31$, which is much smaller than 0.001. We find that at 0.1% significance level that the difference between sybil groups and normal groups are statistically significant.

2.4 Summary of Observations

Our efforts on analyzing user information of sybil groups lead to the following key findings:

- Most of sybil groups have low completeness of optional user information.
- The majority of sybil groups have poor quality of user information: rare confirmed email addresses, high similarity of profile photos, and rare star users.

The completeness and quality of user information can be utilized to improve the security system. Providing complete information with good quality need much more resources than simply registering accounts. Therefore, some attackers do not carefully fabricate information and their sybil groups are obviously different from normal groups. Groups with low completeness and poor quality of user information are suspicious. Even if attackers know the

inspection of information completeness and authenticity, attackers need to pay more cost and overhead, so as to change their behavior and adapt to the inspection. As the cost increases greatly, the benefit of sybil groups decreases significantly.

3 Social Relationships

In this section, we focus on the social aspect, and make measurement of relationships between sybil users. We firstly research on friend relationships inside the sybil group, and study the connectivity structure of the sybil group. Secondly, we analyze friend relationships between different sybil groups, and identify several sybil groups with close connections. Results provide deep understanding of connectivity inside the sybil group, and among several sybil groups. Relationships between sybil users and normal users are already analyzed in our previous work [13].

3.1 Relationships inside the sybil group

For a sybil group, sybil users and friend relationships between them construct a sybil graph, which describes connections inside the sybil group. We study the component structure of the sybil graph in details. Our goal is to understand connectivity structure of the sybil graph as it evolves over time. In particular, we ask: what are the dynamics of component formation and evolution inside sybil groups?

In order to answer this question, We apply a connected-component algorithm on the sybil graph. A connected component is a subgraph in which any two nodes are connected to each other by paths, and which is connected to no additional nodes in the supergraph. The connected component effectively captures the connectivity in the sybil graph. Therefore, it is an effective abstraction with which to measure the dynamics of component formation and evolution.

According to initial study [25], evolution of the connected component has three types: First of all, *growth event* means a new edge is built inside a connected component. The scale of the component increases, but the number of the component remains the same. Secondly, *merging event* means a new edge is built between two connected components, causing them to merge together into one larger component. In this event, two components merge into one component, and the number of components decreases by 1. Finally, *birth event* means that a node is created and a new component is built, and the number of components increases by 1. The growth event is the most common type, but it does not modify the connectivity structure. No component is formed or removed.

Merging and birth events greatly influence the connectivity structure, and change the number of components. In order to study the dynamics of component formation and evolution, we mainly study merging and birth events.

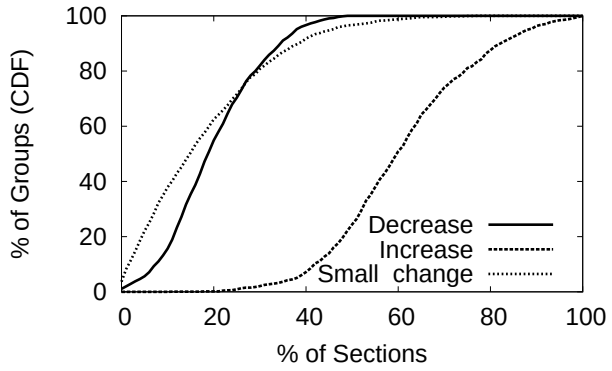

Fig. 7. Evolution of the number of connected components

For each sybil graph, we sort birth events and merging events in time order. In order to visualize the statistics, we divide these events into 100 sections evenly by time, and each section has k events. If all events are birth events, then number of connected components increases by k; If all events are merging vents, the number of connected components decreases by k. The variation range of the number of connected components is $[-k,k]$. We equally divide the variation range into three parts, namely $[-k,-k/3]$,$[-k/3,k/3]$,$[k/3,k]$. According to the change of the number of connected comments, we mark the section as 'decrease', 'small change' or 'increase'. More specifically, If the ratio of birth events is more than 2/3 and the number of connected components increase by at least $k/3$, the section is marked as 'increase'. If the ratio of merging events is more than 2/3 and the number of connected components decrease by at least $k/3$, the section is marked as 'decrease'. Otherwise, the section is marked as 'small change'.

Figure 7 shows statistical results of sections, which can be concluded as following points: (1) Only 24.4% of sybil groups have less than 50% of increase sections, while other 75.6% of sybil groups have more than 50% of increase sections. This shows that in the majority of sections, new nodes are created and the number of components increases. (2) Over 70% of sybil groups have less than 25% of decrease sections. This shows that the majority of groups do not often merge their components into big ones. Note that decrease sections can at most account for about 50% of sections in a group, since the number of merges cannot be larger than that of creating nodes. (3) Small changes account less than 30% in over 80% of groups, indicating that most groups take *bursty* actions to create nodes or add links, rather than doing them at the same time which will cause more "small change" sections.

In sybil groups, birth events are much more than merging events. More birth events create more accounts and strengthen the at-

tack power. Merging events increase connectivity between sybil users. However, initial study [29] observes that good connectivity between sybil users can be leveraged to defend against sybil attacks. Therefore, attackers prefer birth events.

We now investigate how components merge with other components as nodes and edges arrive in sybil groups. Table 2 shows component sizes in merging events. The (i, j)-th entry is the percentage of merging events that a component of size i merges with a component of size j. 95.11% of merging events are in the left column. The main type of component merges is that a singleton merges with another component. In Flickr and Yahoo! 360, there are two most-common types of merging events: (1) a singleton merges with another component; (2) a giant component merges with another component[17]. Compared to their results, our work shows that most of the merges in sybil groups is that a singleton merges with other components (Type 1 in previous work [17]); it is very rare that two non-singletons merge. We conclude that sybil groups has monotonous merging pattern: most attackers merge singletons to existing components, but they rarely merge two non-singletons. It is usual for a normal user from a giant component to know a user in another component, thus connecting the two components. But in the artificial and manipulated network of sybil groups, sybil users lack the diversity in merging patterns. In addition, security mechanism can leverage good connectivity between sybil users to defend against sybil attacks [29], so sybil attackers prefer not to merge large components. In summary, we discover that sybil groups and normal groups have different distribution of component sizes in merging events. This dynamic connectivity structure can be leveraged to improve the security system.

3.2 Relationships between sybil groups

Friend relationships exist between sybil users in different sybil groups, and connections are built between sybil groups. Few links between sybil groups may be created randomly. However, if some sybil groups have strong relationships, they are likely to be controlled by same attackers, or attackers in the same organization. These sybil users are created by IP addresses in different regions， and thus they are divided into several groups by our detection algorithm [13]. If these sybil groups collude together to attack the system, they control more sybil users and have more power than a single sybil group, and they are extremely dangerous for online social networks. Therefore, it is im-

Table 2. Distribution of component sizes in merging events (%)

	1	2	3-5	6-10	11-20	21-100	> 100
1	23.13						
2	16.62	0.84					
3-5	15.71	1.06	0.27				
6-10	10.78	0.58	0.28	0.07			
11-20	8.14	0.34	0.17	0.08	0.02		
21-100	10.91	0.41	0.2	0.08	0.04	0.01	
> 100	9.82	0.24	0.1	0.05	0.02	0.02	0.01

portant to detect sybil groups with close relationships. In this subsection, we firstly build a new graph to describe close relationships betweens sybil groups. Then we detect communities in this graph to identify several sybil groups with strong connections, and analyze their characteristics.

The original graph describes relationships between users. We build a new graph to analyze relationships between sybil groups. In the new graph, each sybil group A is represented as a node A'. In the original graph, almost any two sybil groups have some connections [19]. However, it does not mean that any two sybil groups have strong ties. Actually, most of sybil groups have weak ties, namely one or two links between them. Few links can hardly indicate close relationships between sybil groups. Therefore, we ignore these weak ties and only keep strong ties. More specifically, two groups are considered as having strong ties, if they have the number of links between them more than the threshold $T_{A,B} = \sqrt{n_A * n_B}$. n_A, n_B are the number of users in sybil group A and B, respectively. For any two sybil groups A and B, $S_{A,B}$ is the number of edges between sybil group A and sybil group B in the original graph. If $S_{A,B}$ is larger than $T_{A,B}$, then we create an edge between nodes A' and B' in the new graph. Two sybil groups may have some links between them by incident, but the possibility is very low if the number of links goes over

the threshold. $n_A * n_B/2$ is the maximum value of the number of edges between sybil group A and sybil group B. The threshold depends on possible edges between two sybil groups, without considering other edges of sybil groups. This is because we mainly analyze relationships between two sybil groups. It does not matter whether sybil groups have many internal edges or external edges with other sybil groups. In the new graph, each node stands for a sybil group, and each edge stands for the strong relationship between two sybil groups.

Communities are groups of nodes which are densely connected with each other because of similar backgrounds [10]. Communities effectively capture "neighborhoods" in the graph. Therefore, we believe communities represent the best abstraction with which to measure close relationships between sybil groups. Community detection is a well-studied area, and there are many different algorithms. Our goal is not to evaluate different community detection algorithms or propose new ones. Instead, We use the community detection algorithm proposed in [5]. Because it has been shown to work well, and has been applied in open graph software Gephi [1]. Details of this algorithm are described in the work [5].

We use the community detection approach and identify 1091 communities in the new graph. Sybil groups in the same community have close relationships. With the results of community detection, we visualize the new graph in Figure 9 using Gephi[1]. Node size stands for the size of the sybil group, and link width stands for the number of links between two groups it connected. Different colors of nodes stand for different communities they belong to. From the graph we see that there is a giant connected component that connects most of the sybil groups, and we can see several communities with tight connections among themselves. Moreover, the majority of sybil groups have weak relationships with other sybil groups, and they are distributed in the out layer of the figure.

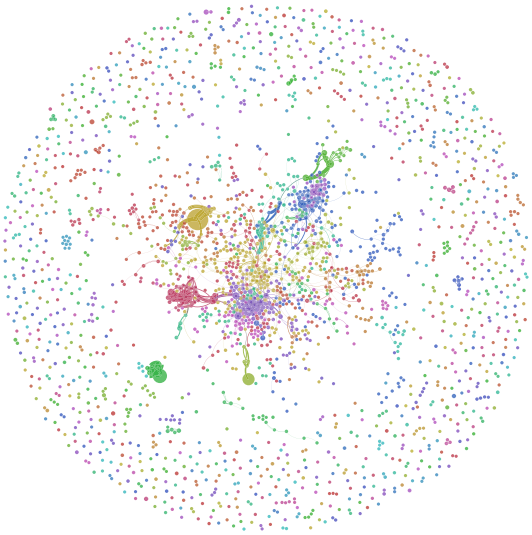

Fig. 9. Communities of Sybil Groups

Next, we study the size of communities. We compute the number of nodes in each com-

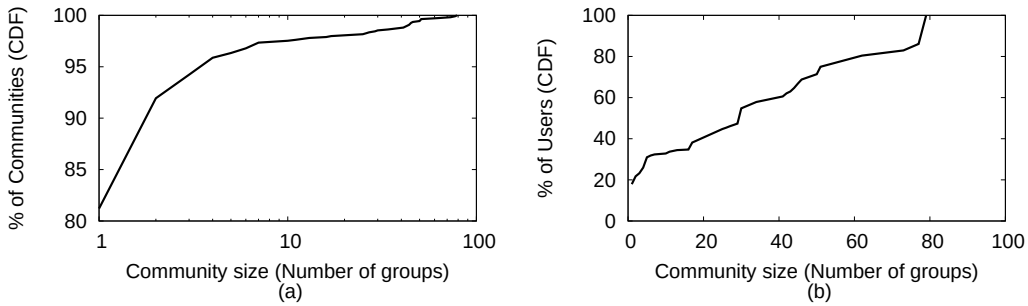

Fig. 8. The distribution of community size

munity, namely the number of sybil groups. We plot the distribution of the community size in Figure 8 (a). 81.2% of communities only have 1 sybil group, and only 3.7% of communities have more than 5 sybil groups. It shows that the majority of sybil groups do not build strong ties with others. Next, we consider the weight of each node, namely the number of users in a sybil group. For each community, we compute the total number of users in sybil groups of this community. In Figure 8 (b), the x-axis is the community size, and the y-axis is the percentage of users whose sybil groups are in communities smaller than x. Only 17.9% of users are in communities which have 1 sybil group. In contrast, 69.1% of users are in communities which have more than 5 sybil groups. Comparing results in Figure 8 (a) and (b), we find that only a few communities contains multiple sybil groups, but they cover a large number of users. The largest community includes 79 sybil groups and 137,004 sybil users. If these sybil groups collude together, they have strong attack power. They can control a large number of sybil users to take malicious actions in a short time. For example, when 79 sybil groups are joined up to propagate rumors, 137,004 sybil users together publish fake information in a short time and leave malicious wall posts in a large number of profiles. Lies when repeated a thousand times appear to be truth. Thousands of people receive rumors and some of them even believe rumors. Even if OSNs quickly detect rumors and delete fake content, rumors may already disseminate to many people and attack the system. These sybil groups have great potential threats to the system, and they should be monitored carefully.

3.3 Summary of results

Our analysis of social relationships of sybil groups produces several conclusions:

- The merging pattern in sybil groups is monotonous: it is very rare that two non-singletons merge.

- Strong relationships are built among several sybil groups, and they control a large number of sybil users.

Our results provide insights into the fight against sybil attack. First of all, sybil groups and normal groups have various dynamic connectivity structure, which can be leveraged to improve the security system in future. Secondly, some sybil groups have close relationships, and they have great potential threats to the system. This is because these sybil groups have a large number of sybil users. If these sybil groups collude together, they have strong attack power. Online social networks should pay special attention and carefully monitor these sybil groups together.

4 Malicious activities

In our previous work, we identified sybil groups and reported them to Renren [13]. Renren carefully monitor these sybil groups. Once users in sybil groups take abnormal actions, they are detected and banned by Renren. We ask several questions: Have these sybil groups and sybil users already performed malicious activities to attack the system? What is characteristics of malicious activities? In order to answer these questions, we contact Renren and know status of these sybil groups and sybil users. We study characteristics of their abusive behavior in this section.

4.1 Measurement and analysis

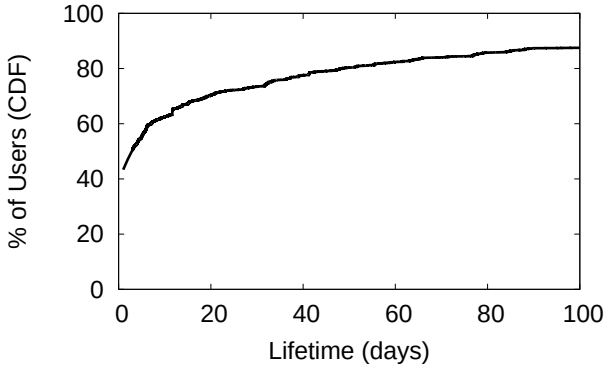

Fig. 10. Duration of lifetime for sybil users

Due to malicious activities, 147,388 sybil users have already been banned until now. It shows that a part of sybil groups and sybil users have already performed abusive behavior and attacked the system. We take a further step and study the time interval between registration and ban. Figure 10 shows the duration of lifetime for sybil users banned by Renren. 43.4% of users are immediately banned within 1 day of their registration, while 26.5% of users survive more than 1 month. 13% of users are even active for more than 100 days before they are banned. It demonstrates that a part of sybil users are banned a long period after their registration. These covert sybil users may behave normally for a long time and suddenly attack the system. Even if some sybil groups

and sybil users take actions normally, they are still potentially dangerous. Therefore, it is important to detect sybil groups and sybil users in advance, and continuously monitor their behavior. Once they perform abusive behavior, they are quickly banned to prevent serious attacks.

Table 3. Types of malicious activities

Type	# of Users
Advertisement	48,318 (32.8%)
Abnormal IP address	35,520 (24.1%)
Aggressive making friend	24,117 (16.4%)
Malicious gamer	17,408 (11.8%)
Other	22,025 (14.9%)

Next, we study malicious activities causing the suspension of sybil users. we classify malicious activities into various types, and analyze their distribution in Table 3. First of all, we see that spreading advertisement is the most common type of abusive behavior. Attackers target OSNs as media to propagate advertisements [9, 27]. They post advertisements through different ways, such as publishing diaries, posting comments and forwarding status. Secondly, 24.1% of sybil users are forbidden for abnormal IP addresses. For example, many users in the same sybil group login in to the system through exactly the same IP address in a short time. This large-scale suspicious logins are likely to be preparation of attacks, and these sybil users are banned. Note that many login requests in a short time may be sent by a large number of users behind a NAT, who are seen as coming from the same IP address. User activities should be further analyzed to confirm anomalies in the inter-arrival time of login requests. In future work, we will contact Renren, apply for additional dataset, and analyze activity difference between sybil users and normal users who are behind a NAT. Thirdly, frequent requests to befriend users in a short period also cause suspension [34]. When normal users receive many friend requests form sybil users, they are disturbed and even permanently leave the system. Fourthly, 17,408 malicious gamers are controlled to achieve higher status in social games [24]. They disrupt the fairness of online social networks. Finally, 14.9% of malicious activities belong to the other type. For example, administrator receive reports of malicious actions from normal people, and decide to stop relevant sybil users.

Table 3 shows that sybil users perform various kinds of malicious activities. Therefore, a single mechanism is not enough to inspect all evil actions. It is significant to identify sybil

groups and sybil users beforehand, and continuously monitor their behavior. Furthermore, the identification of sybil groups is useful to detect collective and malicious behaviors. For example, if users in the same sybil group post similar content, they are likely to send advertisements. if users in the same sybil group participate in the same game, they are likely to cheat in the game. When these sybil groups are detected in advance, security mechanisms can utilize user lists of sybil groups and easily discover their abnormal activities.

4.2 Summary of results

Our efforts on analyzing malicious activities lead to the following key findings:

- Some sybil users are banned a long time after registration.
- Sybil users perform different kinds of malicious activities.

The long incubation period makes us to understand great potential threats of sybil groups and sybil users, which seems to be normal. Attackers never use them until a later date hint, at the possibility of stockpiling accounts and seriously attacking the system. Even if some sybil groups and sybil users behave normally for a long time, they should still be carefully monitored. Furthermore, malicious activities have different types, and thus a single mechanism is not enough to inspect them. Therefore, it is important to identify sybil groups and sybil users beforehand, and then use various mechanisms to continuously monitor them and prevent serious attacks.

5 Related Work

Various techniques are applied to study sybil users or spammers in OSNs. First of all, several sybil defense schemes [8, 28, 32, 36, 37] are based on the assumption that sybil users can hardly make friends with normal users [29]. Secondly, honeypots are deployed to trap spammers who attempt to make friends with them in Twitter [26, 18, 6] and Myspace [12, 18, 31]. Thirdly, researchers manually identify spam tweets in Twitter [3], phantom profiles in Facebook [24] and spammers in Youtube [2]. Fourthly, some works design algorithms to detect anomalies by their clustering characteristic [7, 23]. Finally, Thomas et al. identify accounts suspended by Twitter for disruptive activities [27]; Yang et al. analyze friend requests to detect sybil users [33, 34]; Yardi et al. examine spam around the Twitter meme to detect spammers [35]. Wang et al. observe that some spammers are real users working in a crowd-sourcing system [30].

Our works are much different from these s-

tudies. Previous works focus on detecting spam messages and malicious behaviors, or identifying sybil users and spammers. In this paper, we present a first of its kind study to measure sybil groups in the wild. Our results provide deep understanding of sybil groups and insights into the improvement of security system.

6 Conclusion

In this paper, we present the first attempt to understand sybil groups in a large online social network, using a dataset that covers 2440 sybil groups and 985,797 sybil users. More specifically, we focus on analysis of sybil groups at different levels, including the completeness and quality of individual information, social relationships inside the sybil group and between different sybil groups, and the action time and type of malicious activities.

Our analysis produced a number of interesting findings of sybil groups. (1) At the level of individual information, we find that sybil groups have low completeness of optional user information. Moreover, some individual information has poor quality, including rare confirmed email addresses, high similarity of profile photos, and rare star users. They are easily identified as fake by manual. These results show the completeness and quality of user information can be utilized to improve the security mechanisms against sybil attack. (2) At the sociality level, we discover that compared with normal groups, sybil groups have monotonous merging pattern: most of the mergings are a singleton merging with another component; they rarely merge two non-singletons. This connectivity structure can also be considered to modify the security mechanisms in future. We further discover that several sybil groups have strong relationships, and they control a large number of sybil users and have great potential threats. These sybil groups should be monitored carefully and continuously. (3) At the activity level, we discover that in some sybil groups, sybil users are banned a long time after registration. Even if some sybil groups and sybil users behave normally for a long time, they should still be carefully monitored in case of sudden attacks. We further find sybil users perform different kinds of malicious activities. Therefore, it is important to identify sybil users beforehand, and then use various mechanisms to continuously monitor them. All these results have important implications on the improvement of mechanisms against sybil attack.

While our results from Renren may not generalize to all social networks, our analysis provides a template for understanding sybil groups in the wild. A significant take-away from our work is that security mechanisms not

only utilize properties of users, but also leverage features of groups to fight against sybil attacks. Collective behavior of groups can also provide insights into detection of sybil attacks.

References

[1] Mathieu Bastian, Sebastien Heymann, and Mathieu Jacomy. Gephi: An open source software for exploring and manipulating networks. In *International AAAI conference on weblogs and social media*, volume 2. AAAI Press Menlo Park, CA, 2009.

[2] Fabrício Benevenuto, Tiago Rodrigues, Virgílio Almeida, Jussara Almeida, and Marcos Gonglves. Detecting spammers and content promoters in online video social networks. In *Proc. of SIGIR*, Boston, USA, July 2009.

[3] Fabricio Benevenuto, Gabriel Magno, Tiago Rodrigues, and Virgilio Almeida. Detecting spammers on twitter. In *Proc. of CEAS*, Washington, USA, July 2010.

[4] Sajid Yousuf Bhat and Muhammad Abulaish. Community-based features for identifying spammers in online social networks. In *Proc. of ASONAM*, Niagara Falls, Canada, August 2013.

[5] Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, P10008, October 2008.

[6] Zi Chu, Steven Gianvecchio, Haining Wang, and Sushil Jajodia. Who is tweeting on twitter: Human, bot, or cyborg? In *Proc. of Annual Computer Security Applications Conference*, pages 21–30, December 2010.

[7] Hanbo Dai, Feida Zhu, Ee-Peng Lim, and HweeHwa Pang. Mining coherent anomaly collections on web data. In *Proc. of CIKM*, Maui, USA, October 2012.

[8] George Danezis and Prateek Mittal. Sybilinfer: Detecting sybil nodes using social networks. In *Proc. of NDSS*, February 2009.

[9] Hongyu Gao, Jun Hu, Christo Wilson, Zhichun Li, Yan Chen, and Ben Y. Zhao. Detecting and characterizing social spam campaigns. In *Proc. of ACM Internet Measurement Conference*, pages 35–47, November 2010.

[10] Mao-Guo Gong, Ling-Jun Zhang, Jing-Jing Ma, and Li-Cheng Jiao. Communi-

ty detection in dynamic social networks based on multiobjective immune algorithm. *Journal of Computer Science and Technology*, 27(3):455–467, May 2012.

[11] Xia Hu, Jiliang Tang, Yanchao Zhang, and Huan Liu. Social spammer detection in microblogging. In *Proc. of IJCAI*, Beijing, China, August 2013.

[12] Danesh Irani, SteveWebb, and Calton Pu. Study of static classification of social spam profiles in myspace. In *Proc. of ICWSM*, May 2010.

[13] Jing Jiang, Zifei Shan, Wenpeng Sha, Xiao Wang, and Yafei Dai. Detecting and validating sybil groups in the wild. In *ICDCS workshops*, Macau, China, June 2012.

[14] Jing Jiang, Christo Wilson, Xiao Wang, Peng Huang, Wenpeng Sha, Yafei Dai, and Ben Y. Zhao. Understanding latent interactions in online social networks. In *Proc. of ACM Internet Measurement Conference*, pages 369–382, November 2010.

[15] Long Jin, Yang Chen, Tianyi Wang, Pan Hui, and Athanasios V. Vasilakos. Understanding user behavior in online social networks: A survey. *IEEE Communications Magazine*, 51(9):144–150, September 2013.

[16] Marshall Kirkpatrick. Social networking now more popular than email, report finds. ReadWriteWeb, March 2009.

[17] Ravi Kumar, Jasmine Novak, and Andrew Tomkins. Structure and evolution of online social networks. In *Proc. of KDD*, Pennsylvania, USA, August 2006.

[18] Kyumin Lee, James Caverlee, and Steve Webb. Uncovering social spammers: Social honeypots + machine learning. In *Proc. of SIGIR*, pages 435–442, July 2010.

[19] Zhenhua Li, Guihai Chen, and Tongqing Qiu. Partition node: the topologically critical nodes of unstructured p2p network. *Chinese Journal of Software*, 19(9):2376–2388, September 2008.

[20] Jian-Yun Liu, Yu-Hang Zhao, Zhao-Xiang Zhang, Yun-Hong Wang, Xue-Mei Yuan, Lei Hu, and Zhen-Jiang Dong. Spam short messages detection via mining social networks. *Journal of Computer Science and Technology*, 27(3):506 - 514, May 2012.

[21] Cristian Lumezanu and Nick Feamster. Observing common spam in tweets and email. In *Proc. of IMC*, Boston, USA, November 2012.

[22] H. B. Mann and D. R. Whitney. On a test of whether one of two random vari-

ables is stochastically larger than the other. *The annals of mathematical statistics*, 18(1):50–60, 1947.

[23] Zachary Miller, Brian Dickinson, William Deitrick, Wei Hua, and Alex Hai Wang. Twitter spammer detection using data stream clustering. *Information Sciences*, 260:64 – 73, March 2014.

[24] Atif Nazir, Saqib Raza, Chen-Nee Chuah, and Burkhard Schipper. Ghostbusting facebook: Detecting and characterizing phantom profiles in online social gaming applications. In *Proc. of The 3rd Workshop on Online Social Networks*, June 2010.

[25] Gergely Palla, Albert-László Barabási, and Tamás Vicsek. Quantifying social group evolution. *Nature*, 446(664-667), 2007.

[26] Gianluca Stringhini, Christopher Kruegel, and Giovanni Vigna. Detecting spammers on social networks. In *Proc. of Annual Computer Security Applications Conference*, pages 1–9, December 2010.

[27] Kurt Thomas, Chris Grier, Vern Paxson, and Dawn Song. Suspended accounts in retrospect: An analysis of twitter spam. In *Proc. of ACM Internet Measurement Conference*, November 2011.

[28] Nguyen Tran, Bonan Min, Jinyang Li, and Lakshminarayanan Subramanian. Sybil-resilient online content voting. In *Proc. of NSDI*, pages 15–28, April 2009.

[29] Bimal Viswanath, Ansley Post, Krishna P. Gummadi, and Alan Mislove. An analysis of social network-based sybil defenses. In *Proc. of SIGCOMM*, pages 363–374, August 2010.

[30] Gang Wang, Christo Wilson, Xiaohan Zhao, Yibo Zhu, Manish Mohanlal, Haitao Zheng, and Ben Y. Zhao. Serf and turf: Crowdturfing for fun and profits. In *Proc. of WWW*, Lyon, France, April 2012.

[31] Steve Webb, James Caverlee, and Calton Pu. Social honeypots: Making friends with a spammer near you. In *Proc. of CEAS*, Mountain View, USA, August 2008.

[32] Wei Wei, Fengyuan Xu, Chiu C. Tan, and Qun Li. Sybildefender: A defense mechanism for sybil attacks in large social networks. *IEEE Transactions on Parallel and Distributed Systems*, 24(12):2492–2502, December 2013.

[33] Jilong Xue, Zhi Yang, Xiaoyong Yang, Xiao Wang, Lijiang Chen, and Yafei Dai. Votetrust: Leveraging friend invitation graph to defend against social network sybils. In *Proc. of INFOCOM*, Turin, Italy, April 2013.

[34] Zhi Yang, Christo Wilson, Xiao Wang, Tingting Gao, Ben Y. Zhao, and Yafei Dai. Uncovering social network sybils in the wild. In *Proc. of ACM Internet Measurement Conference*, November 2011.

[35] Sarita Yardi, Daniel Romero, Grant Schoenebeck, and Danah boyd. Detecting spam in a twitter network. *First Monday*, 15(1-4), January 2010.

[36] Haifeng Yu, Phillip B. Gibbons, Michael Kaminsky, and Feng Xiao. Sybillimit: A near-optimal social network defense against sybil attacks. In *Proc. of the IEEE Symposium on Security and Privacy*, pages 3–17, May 2008.

[37] Haifeng Yu, Michael Kaminsky, Phillip B. Gibbons, and Abraham D. Flaxman. Sybilguard: Defending against sybil attacks via social networks. *IEEE/ACM Transactions on Networking*, 16(3):576–589, June 2008.