

Understanding Sybil Groups in the Wild

Jing Jiang^{1,2} (蒋 竞), *Member, CCF*, Zi-Fei Shan² (单子非), Xiao Wang² (王 潇)
Li Zhang^{1,*} (张 莉), *Senior Member, CCF*, and Ya-Fei Dai² (代亚非), *Distinguished Member, CCF*

¹*State Key Laboratory of Software Development Environment, Beihang University, Beijing 100191, China*

²*Department of Computer Science and Technology, Peking University, Beijing 100871, China*

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

Received April 25, 2014; revised January 9, 2015.

Abstract Sybil attacks are one kind of 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 of 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 the characteristics 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 is usually incomplete and in poor quality; 2) sybil groups have special evolution patterns in connectivity structure, including bursty actions to add nodes, and a 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 threats; 4) some sybil users are not banned until a long time after registration in some sybil groups. The characteristics of sybil groups can be leveraged to improve the security mechanisms in OSNs to defend against sybil attacks. Specifically, we suggest that OSNs should 1) check information completeness and quality, 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 network, measurement, security, sybil group, sybil attack

1 Introduction

In recent years, online social networks (OSNs) has become huge and they are still growing throughout the world^[1]. Unfortunately, the openness and the tremendous growth of OSNs attract the interest of malicious parties. Sybil attacks are one kind of well-known and powerful attacks against OSNs. The malicious attacker generates a sybil group consisting of multiple accounts (called sybil users), and disguises 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 the influence and power of target users^[2]. Moreover, sybil attackers target OSNs as media to propagate spam^[3-12]. Sybil attacks become increasingly dangerous as more people use OSNs as primary interfaces to the Internet^①.

Various techniques were applied to identify sybil users or spammers in OSNs, including rare social links between sybil users and normal users^[13-17], honeypots^[7,10,18-20], manual identification^[2,21-22] and

Regular Paper

This work is supported by the National Basic Research 973 Program of China under Grant No. 2011CB302305, the National Natural Science Foundation of China under Grant No. 61300006, and the Project of the State Key Laboratory of Software Development Environment under Grant No. SKLSDE-2013ZX-26.

*Corresponding Author

① http://readwrite.com/2009/03/09/social_networking_now_more_popular_than_email, Dec. 2014.

©2015 Springer Science + Business Media, LLC & Science Press, China

abnormal behavior^[11-12,23-25]. 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^[26]. Some researches took initial steps and designed algorithms to detect sybil groups^[4,26].

Initial studies mainly designed methods to identify sybil users or sybil groups, but few studies measured the 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 the characteristics 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 a dataset from Renren social network^②, one of the largest and oldest OSNs in China. In our previous work^[26], we built a sybil group detector and designed automatic validation mechanisms. We successfully detected and confirmed 2 440 sybil groups and 985 797 sybil users. This large-scale dataset in the wild is notable and provides us precious opportunities to study the characteristics of sybil groups.

We analyze the data to answer the following three key questions.

- How do sybil groups provide individual information? Is the information complete and in high quality?
- What is the connectivity inside sybil groups? How are relationships between different groups?
- Do sybil groups take malicious actions? What are the characteristics of malicious activities?

To answer these essential and practical questions, we analyze the sybil groups with a 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 kinds of user information. We discover that sybil groups have low completeness of optional user information. Moreover, some individual in-

formation has poor quality, and it can be easily identified as fake by hand. These results show that the completeness and quality of user information can be utilized to improve the security mechanisms against sybil attacks. 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 a monotonous merging pattern that lacks 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 and all take malicious actions in a short time it is dangerous for the system. Online social networks should pay special attention to 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 some sybil users are not banned until 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. Furthermore, 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 on the 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 the properties of the whole group to detect sybil groups, rather than simply analyze the 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

②<http://www.renren.com>, Dec. 2014.

to take malicious actions in a short time, which is dangerous for the system. Even if attackers know these modified mechanisms, they need to pay more cost and overhead to change their behavior and adapt to the inspection. As the cost increases greatly, the benefit of sybil groups decreases significantly.

The remainder of the paper is organized as follows. Section 2 describes our dataset and studies personal information of sybil groups and sybil users. Section 3 measures relationships between sybil users and Section 4 analyzes malicious activities of sybil users and sybil groups. In Section 5, we present related studies. Finally, we conclude this paper in Section 6.

2 Individual Information

Before diving into the measurement of sybil groups, we begin by describing the Renren social network and our dataset. We then study the 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. The user information describes personal properties and contact information of the user, and it is feasible for friends to know and contact the user. We ask two questions. 1) Do users in sybil groups fill in their information or just leave default values? 2) If sybil groups fill in information, does this information seem to be fake or true? Results provide deep understanding of the characteristics of sybil groups and sybil users and practical implications on the design of security mechanisms for social networks in future.

2.1 Renren Social Network and the Dataset

Renren was set up in 2005, and it is one of the oldest and biggest OSNs in China^[27]. 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 those on 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^[24]. To support the project, Renren provides anonymized user data on its 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 2 440 sybil groups and 985 797 sybil users^[26]. These sybil groups constitute a large-scale dataset in the wild and provide us precious research opportunities. In this work, we mainly study the characteristics of sybil groups, rather than single sybil users. Our validation mechanisms confirm that users in groups have high similarity of action time and these groups were sybil. In order to make comparison, we also select 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 dataset is different from datasets in the work of [24]. Furthermore, previous work^[24] made contributions to sybil users, instead of sybil groups. Though our dataset may not cover all sybil groups and sybil users, it is valuable for research and allows us to get the initial understanding of sybil groups in the wild.

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 the 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 which is an important feature of a user. The profile photo is widely used to attract attention, and help friends identify the user quickly. The profile photo may be a personal photo or some other picture. For each sybil group as well as each normal group, we compute the percentage of users without profile photos, and plot the cumulative distribution functions (CDF) of results in Fig.1. Fig.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 have uploaded 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 groups do not have profile photos. Uploading profile photos costs more time and bandwidth; therefore attackers are inactive in providing profile photos.




Fig.1. Percentage of users without 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^[28]. We use the MWW test, because it does not assume any specific distribution, which is suitable for our dataset. 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 the 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 the difference between sybil

groups and normal groups is statistically significant at 0.1% significance level.

The second feature is the telephone number. The telephone number makes it feasible for people to contact others. Moreover, the telephone number can be bound to the account, and if the password is lost or the account is compromised, the user can request the password through phone message. The telephone number is 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 each normal group, we compute the number of users with telephone numbers divided by the number of users in the group. Fig.2 shows the percentage of users with telephone numbers. In 65.3% of sybil groups, none of users have 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 with telephone information. We have tested and confirmed that the distributions of sybil groups and normal groups are significantly different using the MWW test at the 0.001 significance level. Sybil groups have much lower percentage of users with telephone numbers than normal groups. Sybil users are not created for social activities and telephone numbers are unnecessary for them.




Fig.2. Percentage of users with telephone numbers.

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 the same names. Fig.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 with their locations. In contrast, only 8.8% of normal groups have less than 60% of users with locations, and the 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 $3.44\text{E}-27$, which is much smaller than 0.001. We find that the difference between sybil groups and normal groups is statistically significant at the 0.1% significance level.

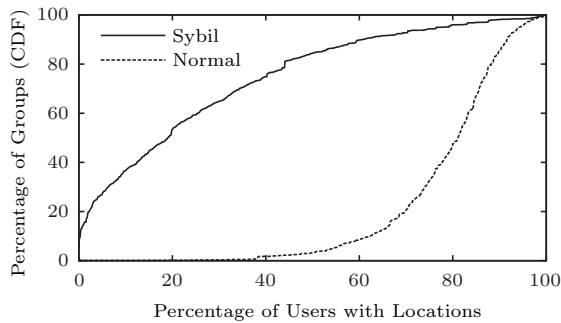


Fig.3. Percentage of users with locations.

Results show that sybil groups have low completeness of user information. This is because that filling in information needs additional computational resource, bandwidth, and time. In order to save the overhead, attackers are inactive in providing the optional information of sybil users. Therefore, the completeness of user information can be utilized to improve the security system. People may argue that if attackers know the detection mechanisms, they 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 the 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 quality 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 whether 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 qualities between sybil groups and normal groups, and we mainly study these attributes in the following part.

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 a better quality than email addresses without confirmation. For each group, we compute the percentage of users without email address confirmation. Fig.4 shows the results of sybil groups and normal groups. In 93.6% of normal groups, all users have their email addresses confirmed. In contrast, only in 26% of sybil groups, all users have confirmed email addresses. We again run the MWW test and find that the difference is statistically significant at the significance level of 0.001. These results indicate that sybil groups have low percentage of users with confirmed 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 email address. However, the confirmation of email address needs additional cost and overhead, and attackers may not to confirm email addresses. Email confirmation is encouraged in OSNs, but it is not necessary to some extent. If people do not pass the email confirmation, their activities are limited, but they can still update 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 the high quality of email addresses and increase the cost of the creation of sybil groups.




Fig.4. Percentage of users without confirmed email addresses.

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 will expire 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, 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 the email addresses. Many email addresses have the same domain names. After eliminating duplicate domain names, we obtain 2025 unique domain names. Two hundred and three domain names are widely used and they cover 99.3% of email addresses. We manually check the 203 domain names and categorize them into above types.

Table 1 shows the number of domain names and the number of email addresses in each type. Though the real type makes up 91.8% of email addresses, it only accounts for 17.2% of domain names. The nonexistent type accounts for 72.4% of domain names. The 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 explain 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, they are also unable to receive confirmation emails after the expiration of these email addresses. The type of domain name is useful to improve the security system. If the user fills in a temporary or nonexistent domain name, the user is suspicious and needs further checks.

Table 1. Number of Email Domain Names in Each Type

Type	Number of Domain Names	Number of Email Addresses
Real	35	647 212
Temporary	21	19 210
Nonexistent	147	38 532

Next, we study the quality of profile photo. It is not easy to automatically judge the quality of a single profile photo. Fortunately, 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. Fig.5 shows the 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 two 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, the 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 sybil groups.

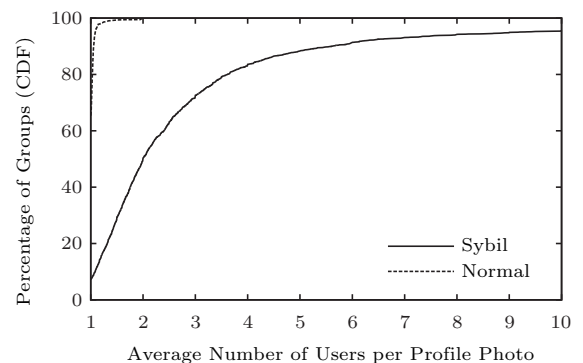


Fig.5. Average number of users per profile photo.

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

ple, some attackers fill in celebrities' information as personal information. People can easily identify that the information belongs to celebrities, instead of ordinary persons. 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 fills 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 users prefer to achieve the title of star user. For sybil groups and normal groups, we compute the percentage of star users and show results in Fig.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 the 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 the high quality of user information, which increases the cost and overhead for attackers. The 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 the difference between sybil groups and normal groups is statistically significant at the 0.1% significance level.

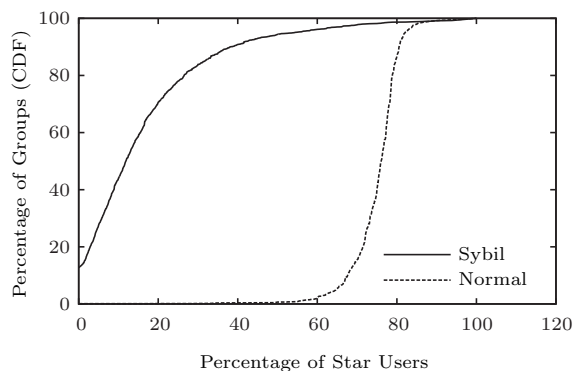


Fig.6. Percentage of star users.

2.4 Summary of Observations

Our efforts on analyzing the 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 needs 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 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 multiple sybil groups with close connections. Results provide deep understanding of connectivity inside the sybil group, and among multiple sybil groups. Relationships between sybil users and normal users were already analyzed in our previous work^[26].

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 the connectivity structure of the sybil graph as it evolves over time. In particular, we study: 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 not connected to additional nodes in the supergraph. A subgraph, H , of a graph, G , is a graph whose vertices are a subset of the vertex set of G , and whose edges are a subset of the edge set of G . The connected component effectively captures the connectivity in the sybil graph. Therefore, it is an effective abstraction to measure the dynamics of component formation and evolution.

According to initial study^[29], the evolution of the connected component has three types. First of all, the growth event means a new edge is built inside a connected component. The scale of the component increases, but the number of components remains the same. Secondly, the 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, the 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.

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 the number of connected components increases by k ; if all events are merging events, 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 components, 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 increases 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 decreases by at least $k/3$, the section is marked as “decrease”. Otherwise, the section is marked as “small change”.

Fig.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 the 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 for less than 30% in over 80% of sybil groups, indicating that most sybil groups take bursty actions to create nodes or add links, instead of doing them at the same time, which will cause more “small change” sections.

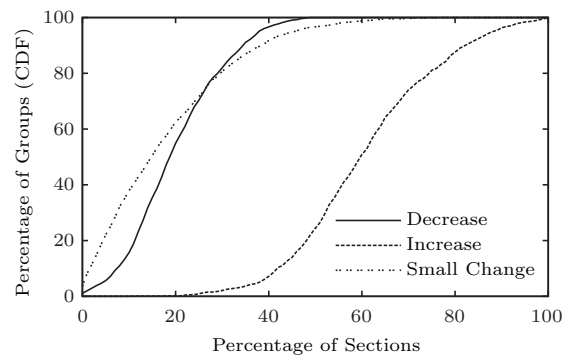


Fig.7. Evolution of the number of connected components.

In sybil groups, birth events are much more than merging events. More birth events create more accounts and strengthen the attack power. Merging events increase the connectivity between sybil users. However, initial study^[30] observed that the 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 second 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^[31]. Compared with their results, our work shows that the most common type of merges in sybil groups is that a singleton merges with other components (type 1 in previous work^[31]), and it is very

rare that two non-singletons merge. We conclude that sybil groups have a monotonous merging pattern: most attackers merge singletons to existing components, but they rarely merge two non-singletons. It is common 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^[30], thereby sybil attackers do not prefer to merge large components. In summary, we discover that sybil groups and normal groups have different distributions of component sizes in merging events. This dynamic connectivity structure can be leveraged to improve the security system.

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	0.20	0.08	0.04	0.01	
> 100	9.82	0.24	0.10	0.05	0.02	0.02	0.01

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^[26]. 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 important to detect sybil groups with close relationships. In this subsection, we firstly build a new graph to describe close relationships between 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^[32]. 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 to have strong ties, if they have the number of links between them more than the threshold $T_{A,B} = \sqrt{n_A \times n_B}$. n_A, n_B are the number of users in sybil groups 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. $n_A \times n_B / 2$ 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^[33]. Communities effectively capture “neighborhoods” in the graph. Therefore, we believe communities represent the best abstraction 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 [34]. It has been shown to work well, and has been applied in open graph software Gephi^[35]. Details of this algorithm are described in the work of [34].

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 Fig.8 using Gephi^[35]. Node size stands for the size of the sybil group, and link width stands

for the number of links between two groups. 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.8. Communities of sybil groups.

Next, we study the size of communities. We compute the number of nodes in each community, namely the number of sybil groups. We plot the distribution of the community size in Fig.9(a). 81.2% of communities only have one sybil group, and only 3.7% of communities have more than five 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 Fig.9(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 one sybil group. In contrast, 69.1% of users are in communities which have more than five sybil groups. Comparing results in Figs.9(a) and 9(b), we find that only a few communities contain multiple sybil groups, but they cover a large number of users. The largest community includes 79 sybil groups and 137004 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, 137004 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 will 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.

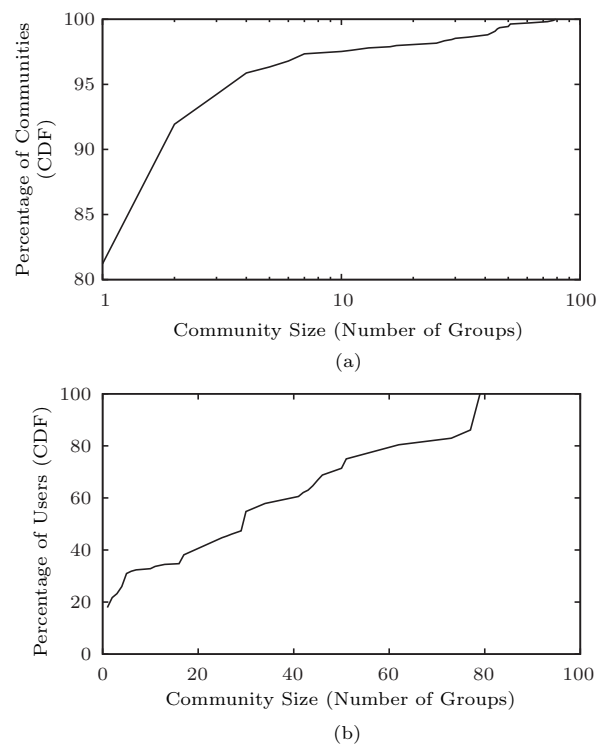


Fig.9. Distribution of community size.

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 attacks. First of all, sybil groups and normal groups have various dynamic connectivity structures, 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 to and carefully monitor these sybil groups together.

4 Malicious Activities

In our previous work, we identified sybil groups and reported them to Renren^[26]. Renren carefully monitors these sybil groups. Once users in sybil groups take abnormal actions, they are detected and banned by Renren. We have several questions. 1) Have these sybil groups and sybil users already performed malicious activities to attack the system? 2) What are characteristics of malicious activities? In order to answer these questions, we contact Renren and get the status of these sybil groups and sybil users. We study the characteristics of their abusive behavior in this section.

4.1 Measurement and Analysis

Due to malicious activities, 147388 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. Fig.10 shows the duration of lifetime for sybil users banned by Renren. 43.4% of users are immediately banned within one day of their registration, while 26.5% of users survive for more than one 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 for a long period after their registration. These covert sybil users may behave normally




Fig.10. Duration of lifetime for sybil users.

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.

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^[5,11]. 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 the 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 the Network Address Translation (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 the activity difference between sybil users and normal users who are behind an NAT. Thirdly, frequent requests to make friends with users in a short period also cause suspension^[24]. When normal users receive many friend requests from 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^[2]. They disrupt the fairness of online social networks. Finally, 14.9% of malicious activities belong to the other type. For example, administrators receive the reports of malicious actions from normal people, and decide to stop relevant sybil users.

Table 3. Types of Malicious Activities

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

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 not banned until a long time after registration.
- Sybil users perform different kinds of malicious activities.

The long incubation period makes us 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 have been applied to study sybil users or spammers in OSNs. First of all, several sybil defense schemes^[13-17] are based on the assumption that sybil users can hardly make friends with normal users^[30]. Secondly, honeypots were deployed to trap spammers who attempt to make friends with them in Twitter^[10,18-19] and MySpace^[7,19-20]. Thirdly, researchers manually identified spam tweets in Twitter^[22], phantom profiles in Facebook^[2] and spammers in Youtube^[21]. Fourthly, some studies designed algorithms to detect anomalies by their clustering characteristics^[4,9]. Finally, Thomas *et al.*^[11] identi-

fied accounts suspended by Twitter for disruptive activities; Yang *et al.*^[24,36] analyzed friend requests to detect sybil users; Yardi *et al.*^[25] examined spam around the Twitter meme to detect spammers; Wang *et al.*^[12] observed that some spammers are real users working in a crowd-sourcing system.

Our studies are much different from these studies. Previous studies 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 Conclusions

In this paper, we presented the first attempt to understand sybil groups in a large online social network, Renren, using a dataset that covers 2440 sybil groups and 985797 sybil users. More specifically, we focused on the analysis of sybil groups at different levels, including the completeness and quality of individual information, the 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 found 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 hand. These results showed the completeness and quality of user information can be utilized to improve the security mechanisms against sybil attacks. 2) At the sociality level, we discovered that compared with normal groups, sybil groups have a monotonous merging pattern: most of the mergings are the merging of a singleton with a component; they rarely merge two non-singletons. This connectivity structure can also be considered to modify the security mechanisms in future. We further discovered 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 discovered that in some sybil groups, some sybil users are banned long 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 found

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 attacks.

While we may not generalize our results from Renren to all social networks, our analysis provides a template for understanding sybil groups in the wild. Our work suggests that security mechanisms should 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 the detection of sybil attacks.


References

- [1] Jin L, Chen Y, Wang T, Hui P, Vasilakos A V. Understanding user behavior in online social networks: A survey. *IEEE Communications Magazine*, 2013, 51(9): 144-150.
- [2] Nazir A, Raza S, Chuah C N, Schipper B. Ghostbusting facebook: Detecting and characterizing phantom profiles in online social gaming applications. In *Proc. the 3rd Workshop on Online Social Networks*, June 2010.
- [3] Bhat S Y, Abulaish M. Community-based features for identifying spammers in online social networks. In *Proc. ASONAM*, August 2013, pp.100-107.
- [4] Dai H, Zhu F, Lim E P, Pang H. Mining coherent anomaly collections on web data. In *Proc. the 21st CIKM*, October 29-November 2, 2012, pp.1557-1561.
- [5] Gao H, Hu J, Wilson C, Li Z, Chen Y, Zhao B Y. Detecting and characterizing social spam campaigns. In *Proc. the 10th ACM SIGCOMM Internet Measurement Conference*, November 2010, pp.35-47.
- [6] Hu X, Tang J, Zhang Y, Liu H. Social spammer detection in microblogging. In *Proc. the 23rd IJCAI*, August 2013, pp.2633-2639.
- [7] Irani D, SteveWebb, Pu C. Study of static classification of social spam profiles in MySpace. In *Proc. the 4th ICWSM*, May 2010, pp.82-89.
- [8] Lumezanu C, Feamster N. Observing common spam in Tweets and email. In *Proc. the 12th ACM SIGCOMM IMC*, November 2012, pp.461-466.
- [9] Miller Z, Dickinson B, Deitrick W, Hua W, Wang A H. Twitter spammer detection using data stream clustering. *Information Sciences*, 2014, 260: 64-73.
- [10] Stringhini G, Kruegel C, Vigna G. Detecting spammers on social networks. In *Proc. the 26th Annual Computer Security Applications Conference*, December 2010, pp.1-9.
- [11] Thomas K, Grier C, Paxson V, Song D. Suspended accounts in retrospect: An analysis of Twitter spam. In *Proc. the 11th ACM SIGCOMM Internet Measurement Conference*, November 2011, pp.243-258.
- [12] Wang G, Wilson C, Zhao X, Zhu Y, Mohanlal M, Zheng H, Zhao B Y. Serf and turf: Crowdturfing for fun and profits. In *Proc. the 21st WWW*, April 2012, pp.679-688.
- [13] Danezis G, Mittal P. SybilInfer: Detecting sybil nodes using social networks. In *Proc. NDSS*, February 2009.
- [14] Tran N, Min B, Li J, Subramanian L. Sybil-resilient online content voting. In *Proc. the 6th NSDI*, April 2009, pp.15-28.
- [15] Wei W, Xu F, Tan C C, Li Q. SybilDefender: A defense mechanism for sybil attacks in large social networks. *IEEE Transactions on Parallel and Distributed Systems*, 2013, 24(12): 2492-2502.
- [16] Yu H, Gibbons P B, Kaminsky M, Xiao F. SybilLimit: A near-optimal social network defense against sybil attacks. In *Proc. IEEE Symposium on Security and Privacy*, May 2008, pp.3-17.
- [17] Yu H, Kaminsky M, Gibbons P B, Flaxman A D. SybilGuard: Defending against sybil attacks via social networks. *IEEE/ACM Transactions on Networking*, 2008, 16(3): 576-589.
- [18] Chu Z, Gianvecchio S, Wang H, Jajodia S. Who is tweeting on Twitter: Human, bot, or cyborg? In *Proc. the 26th Annual Computer Security Applications Conference*, December 2010, pp.21-30.
- [19] Lee K, Caverlee J, Webb S. Uncovering social spammers: Social honeypots + machine learning. In *Proc. the 33rd SIGIR*, July 2010, pp.435-442.
- [20] Webb S, Caverlee J, Pu C. Social honeypots: Making friends with a spammer near you. In *Proc. the 5th CEAS*, August 2008.
- [21] Benevenuto F, Rodrigues T, Almeida V, Almeida J, Gonglves M. Detecting spammers and content promoters in online video social networks. In *Proc. the 32nd SIGIR*, July 2009, pp.620-627.
- [22] Benevenuto F, Magno G, Rodrigues T, Almeida V. Detecting spammers on Twitter. In *Proc. CEAS*, July 2010.
- [23] Liu J Y, Zhao Y H, Zhang Z X, Wang Y H, Yuan X M, Hu L, Dong Z J. Spam short messages detection via mining social networks. *Journal of Computer Science and Technology*, 2012, 27(3): 506-514.
- [24] Yang Z, Wilson C, Wang X, Gao T, Zhao B Y, Dai Y. Uncovering social network sybils in the wild. In *Proc. the 11th ACM SIGCOMM Internet Measurement Conference*, November 2011, pp.259-268.
- [25] Yardi S, Romero D, Schoenebeck G, Boyd D. Detecting spam in a Twitter network. *First Monday*, 2010, 15(1).
- [26] Jiang J, Shan Z, Sha W, Wang X, Dai Y. Detecting and validating sybil groups in the wild. In *Proc. the 32nd ICDCS Workshops*, June 2012, pp.127-132.
- [27] Jiang J, Wilson C, Wang X, Huang P, Sha W, Dai Y, Zhao B Y. Understanding latent interactions in online social networks. In *Proc. the 10th ACM Internet Measurement Conference*, November 2010, pp.369-382.
- [28] Mann H B, Whitney D R. On a test of whether one of two random variables is stochastically larger than the other. *The Annals of Mathematical Statistics*, 1947, 18(1): 50-60.
- [29] Palla G, Barabási A L, Vicsek T. Quantifying social group evolution. *Nature*, 2007, 446(7136): 664-667.
- [30] Viswanath B, Post A, Gummadi K P, Mislove A. An analysis of social network-based sybil defenses. In *Proc. SIGCOMM*, August 30-September 3, 2010, pp.363-374.
- [31] Kumar R, Novak J, Tomkins A. Structure and evolution of online social networks. In *Proc. the 12th KDD*, August 2006, pp.611-617.

- [32] Li Z, Chen G, Qiu T. Partition nodes: Topologically-critical nodes of unstructured peer-to-peer networks. *Journal of Software*, 2008, 19(9): 2376-2388. (in Chinese)
- [33] Gong M G, Zhang L J, Ma J J, Jiao L C. Community detection in dynamic social networks based on multiobjective immune algorithm. *Journal of Computer Science and Technology*, 2012, 27(3): 455-467.
- [34] Blondel V D, Guillaume J L, Lambiotte R, Lefebvre E. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008, 2008(10): P10008.
- [35] Bastian M, Heymann S, Jacomy M. Gephi: An open source software for exploring and manipulating networks. In *Proc. the 3rd International AAAI Conference on Weblogs and Social Media*, May 2009, pp.361-362.
- [36] Xue J, Yang Z, Yang X, Wang X, Chen L, Dai Y. VoteTrust: Leveraging friend invitation graph to defend against social network sybils. In *Proc. INFOCOM*, April 2013, pp.2400-2408.



Jing Jiang received her B.S. and Ph.D. degrees in computer science from Peking University in 2007 and 2012, respectively. She is now an assistant professor in the State Key Laboratory of Software Development Environment of Beihang University, Beijing. Her research interests include social network, data mining, human factors and social aspects of software engineering.



Zi-Fei Shan received his Bachelor's degree in computer science from Peking University in 2013. His recent research focuses on building and using knowledge bases to help people better understand and exploit data.



Xiao Wang received his B.S. and M.S. degrees in computer science from Peking University in 2009 and 2012, respectively. His research interests include social network and data mining.



Li Zhang received her Ph.D. degree in computer science from Beihang University, Beijing, in 1996. She is a professor and Ph.D. advisor in the State Key Laboratory of Software Development Environment of Beihang University. She is also the associate dean of the School of Software and the vice director of Software Engineering Institute. Her research interests include software engineering, system modeling and simulation, software architecture, empirical software engineering, requirements engineering and enterprise modeling.



Ya-Fei Dai received her Ph.D. degree in computer science from Harbin Institute of Technology in 1993. She is a professor in the Department of Computer Science and Technology of Peking University. Her research interests include distributed system, cloud storage, social network and graph computing.