

On the Evolution of User Interaction in Facebook

Bimal Viswanath[†]

Alan Mislove[‡]

Meeyoung Cha[†]

Krishna P. Gummadi[†]

[†]Max Planck Institute for Software Systems (MPI-SWS)
Kaiserslautern/Saarbrücken, Germany

[‡]Rice University
Houston, Texas

ABSTRACT

Online social networks have become extremely popular; numerous sites allow users to interact and share content using social links. Users of these networks often establish hundreds to even thousands of social links with other users. Recently, researchers have suggested examining the *activity network*—a network that is based on the actual interaction between users, rather than mere friendship—to distinguish between strong and weak links. While initial studies have led to insights on how an activity network is structurally different from the social network itself, a natural and important aspect of the activity network has been disregarded: the fact that over time social links can grow stronger or weaker. In this paper, we study the evolution of activity between users in the Facebook social network to capture this notion. We find that links in the activity network tend to come and go rapidly over time, and the strength of ties exhibits a general decreasing trend of activity as the social network link ages. For example, only 30% of Facebook user pairs interact consistently from one month to the next. Interestingly, we also find that even though the links of the activity network change rapidly over time, many graph-theoretic properties of the activity network remain unchanged.

Categories and Subject Descriptors

H.3.5 [Information Storage and Retrieval]: Online Information Services—*Web-based services*; J.4 [Computer Applications]: Social and Behavioral Sciences—*Sociology*

General Terms

Human Factors, Measurement

1. INTRODUCTION

Online social networks have become a popular way for users to connect, express themselves, and share content. Popular social networks have hundreds of millions of registered users and are growing at a rapid pace. As these networks grow and mature, users have been observed to form hundreds to even

thousands of friendship links. For example, a user in Orkut has on average over 100 friends [1] and the average number of friends in Facebook is over 120 [2]. Certain individuals have much higher degrees than the average; in fact, in Flickr we found multiple users who have more than 25,000 friends!

While most social networking sites allow only a binary state of friendship, it has been unsurprisingly observed that not all links are created equal. A recent study [4] has demonstrated that the “strength of ties” varies widely, ranging from pairs of users who are best friends to pairs of users who even wished they weren’t friends. In order to distinguish between these strong and weak links, researchers have suggested examining the *activity network*, the network that is formed by users who actually interact using one or many of the methods provided by the social networking site. While the initial studies on activity networks [3, 5] have brought great insights into how an activity network is structurally different from the social network and how system designers can utilize such information, little attention has been paid to a rather natural and important aspect of user interaction: the fact that the level of interaction between two individuals can vary over time.

In this paper, we study the *evolution of activity* between users at both the microscopic and macroscopic levels. At the microscopic level, we investigate how pairs of users in a social network interact, and at the macroscopic level, we examine how the varying patterns of interaction affect the overall structure of the activity network. For this study, we collect data on both friendship relationship and interactions for a large subset of the Facebook social network. In total, we examine over 60,000 users and over 800,000 logged interactions between those users over a period of two years. We have made our data set available to the research community in an anonymized form.¹

We make a number of interesting observations. From a microscopic view, we find that interactions between pairs of users who interact infrequently are likely triggered by site mechanisms. For example, over 54% of the interactions between the infrequently interacting user pairs can be directly attributed to Facebook’s birthday reminder feature. Even when users do interact frequently, we find that the activity level of user pairs tends to decrease markedly over time, implying that most activity links die out. However, highly active user pairs exhibit this trend to a lesser degree, further emphasizing the strength of these links.

These findings together suggest that the underlying activity network evolves rapidly. Thus, we also examine how the macroscopic properties of the activity network vary over

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WOSN’09, August 17, 2009, Barcelona, Spain.

Copyright 2009 ACM 978-1-60558-445-4/09/08 ...\$5.00.

¹Data available at <http://socialnetworks.mpi-sws.org>

time. Surprisingly, we find that while the individual user pairs that compose the activity network changes rapidly over time (i.e., over the course of one month 70% of the links in the activity network disappear), many of the graph-theoretic properties (e.g., average node degree, average clustering coefficient, average path length) show remarkable stability over the course of two years.

The remainder of this paper is organized as follows. Section 2 describes our data collection methodology and introduces the data set. Section 3 presents an analysis of the evolution of pairwise user interactions, and discriminates between high and low rate interactions. Section 4 provides an analysis of the evolution of network-level properties of the activity network over time. Section 5 discusses implications of our findings. We describe related work in Section 6 and conclude in Section 7.

2. MEASUREMENT METHODOLOGY

In this section, we introduce our data set and relevant characteristics of the Facebook social network.

2.1 Data: Facebook New Orleans Network

We focus on the New Orleans regional network in Facebook for the study. We chose to crawl a regional network because we could easily join the network for crawling. In comparison, other networks require stronger authentication for joining; for example, university networks require email addresses from the university’s domain in order to create an account. We created several Facebook accounts and joined each to the New Orleans regional network, which allowed us to view most profiles of other users in the same regional network.

We collected our data set over two separate periods. First, we crawled the social network topology during December 29th, 2008 and January 3rd, 2009. In this period, we only collected information about friendship links (i.e., who is friends with whom). Our crawler started from a single user in the New Orleans network and visited all friends of the user and their friends in a breadth-first-search (BFS) fashion, consistent with previous social network crawls [1]. Our crawler could only view users who made their profiles visible to the network. We downloaded HTML documents of the public profiles of those visible users.

Second, we crawled the *wall* feature in Facebook as a form of user interaction in the social network. In Facebook, a user’s friends can post comments to the user’s wall; these comments appear on the user’s wall and can be seen by others who visit the user’s profile. Therefore, wall posts represent a broadcast-style messaging service within the site. We conducted our second crawl between January 20th, 2009 and January 22nd, 2009, and downloaded the entire wall history for all users whom we had previously discovered. Each wall post entry in our data set contains information about the wall owner, the user who made the post, the time at which the post was made, and the post content.

2.2 Data statistics

In total we gathered information about 90,269 users and 3,646,662 friendship links between those users. This accounts for 52% of the users in the New Orleans network based on the statistics provided by Facebook. However, not all users make their profiles public; we could view the wall postings of 60,290 (66.7%) users. These 60,290 users were connected together by 1,545,686 links in the social network with an average node degree of 25.3. In the remainder of the paper, we consider only this set of users.

Our wall post data spans from September 26th, 2006 to January 22nd, 2009. In total, we observed 838,092 wall posts, for an average of 13.9 wall posts per user. This covers communication between 188,892 distinct pairs of users, representing 12.2% of the links in the social network. The remaining 87.8% of the link in the social network did not exhibit any wall activity.

2.3 Limitations

Our measurement methodology has a few limitations. First, in Facebook, users can interact in many ways (e.g., messaging, applications, photo uploads, and chat). While wall posting is one of the most popular methods of user interactions [5], we do not know if it is representative of other forms of interaction. Second, our Facebook crawl was limited to the subset of people who made their profiles visible to the people in the same regional network, which is the default privacy setting in Facebook. Nonetheless, we expect that our results are representative of all other users in the same network, because our data set covers a majority of the network (66.7%). Third, we were only able to collect data for users in the giant connected component of the New Orleans networks; we were unable to crawl users who were not connected to this component. However, previous work has shown that the giant connected component tends to contain the majority of the users [1]. Finally, our data set is based on a single network. However, a recent study has demonstrated similarity across multiple Facebook networks [5] and gives us confidence that our findings will generalize to other regional networks in Facebook.

2.4 High-level characteristics of data

We present two high-level characteristics of the wall posting activity. First we use the time stamps of wall posts to examine the growth of user interactions over time. Figure 1 shows the number of wall posts over the entire trace period. We observe a steady increase from approximately the middle of 2008. While this sudden growth in mid-2008 seems abrupt, it is likely that Facebook’s launch of a new site design on July 20, 2008 allowed users to more easily view wall posts through friend feeds [6]. We speculate that this change led to the spike in wall posts towards the end of the trace.




Figure 1: Number of wall posts per month.

Second we compare the activity network with the social network. Using wall post data, we construct the *activity network* as an undirected graph, where a link exists between a pair of users if they interacted at least once. Previous studies [3, 5] have shown that node degree in the activity network is significantly lower than in the social network. We can confirm this in our data: only 12.2% of the social links showed any wall post activity. Figure 2 shows the correlation between the node degree of individuals in the social network

and the activity network. The plot clearly shows that users only interact with a small number of their social links, and that the maximal degree in the activity network tends to be bounded at around 100.




Figure 2: Comparison between the node degree in the social network and the activity network.

3. PAIRWISE USER INTERACTIONS

In this section, we study the patterns of wall post activity between pairs of Facebook users. Our goal is to understand how relationships between pairs of individuals evolve over time, based on the frequency and the content of wall posts exchanged between them. Towards this goal, we conduct two sets of analyses. First, we investigate the distribution of the number of wall posts per link and examine to what extent certain user pairs generate more posts than others. Second, we characterize the patterns of communication of both the infrequently interacting user pairs and the frequently interacting user pairs.

3.1 Data used

In our trace, different user pairs formed their social links and initiated their first interaction at different times. To make our analysis fair across user pairs with differing periods of activity, in this section, we focus only on the first year of activity for each link, as measured from the date of link establishment. Thus, we only consider user pairs (a) for which we could determine when the link was established, based on the wall post, and (b) that created their social link before January 22nd, 2008, and thus had at least one year of activity. In total, there were 59,916 (31.7%) such user pairs, and we consider only these user pairs in the analysis that follows.

3.2 Distribution of the number of wall posts

We first investigate the overall distribution of the number of wall posts across each link, looking at the extent to which certain links possess more activity than others. For this, we calculate the cumulative distribution function (CDF) of the number of wall posts for each user pair in Figure 3. We observe a significant skew in the distribution. The median number of posts per user pair is 2 and 81% of the pairs exchange no more than 5 wall posts. However, some user pairs (<1%) exchanged more than 100 posts. In fact, the most active pair exchanged 767 total wall posts in the first year the link was established.

User pairs with low number of wall posts may be casual acquaintances in Facebook, while user pairs who generated high rates of wall posts may represent close friends. Likewise, different wall posting rates may indicate that the relationship between individuals evolve with different patterns.




Figure 3: Distribution of the number of wall posts across each link.

In the remainder of this section, we take a closer look at both sides of the distribution: user pairs with low and high interaction rates. We partition users into two groups based on a threshold of 5 wall posts in the first year of activity.

3.3 Patterns of infrequent interactions

We first examine the patterns of interaction for user pairs who sent no more than 5 wall posts across their link. This group represents a majority of the activity network, accounting for 48,689 (81%) of the user pairs. We are interested in understanding the reasons for the low level of interaction, and in particular, which events trigger a wall post to be sent across the link. Because this group shows an extremely low rate of wall posting, we initially expected that the creation of the social link would cause these users to interact a few times soon after link establishment and then discontinue their communication.

To test our hypothesis, we calculated the time between the establishment of the link and the time of the first wall posting. Figure 4 shows the CDF of this delay in the first wall post for all pairs in the low-rate interaction group. We find that 20% of the user pairs interacted the very same day they became friends in Facebook. However, the remaining pairs had an almost even distribution over the following year. In fact, more than half of the pairs exchanged their very first wall post over two months after they became friends! This demonstrates that, in most cases, the first wall post between user pairs with a low rate of interaction is not triggered by the creation of the social link.

We then investigated the cause of this unexpected delay in wall post interaction. For this, we examined the actual content of the wall posts in order to shed light on what triggered the interaction. We found that over 39% of the first wall posts contained either “birthday” or “bday”, suggesting that these wall posts are birthday messages. Because the birthdays of Facebook users are likely to be spread out across the year, we expect that this underlies the smooth curve in the time to first interaction in Figure 4. In fact, wall posts related to birthday greetings are even more prevalent if we consider wall posts other than the first: 54% of the user pairs in the low-rate interaction group exchanged at least one wall post related to birthday wishes. This phenomenon could be due to the fact that Facebook provides a birthday reminder feature which alerts users of their friends’ birthdays at login.

3.4 Patterns of frequent interactions

Now we turn our focus to user pairs who sent more than 5 wall posts in the first year of activity. In this group, the average number of wall posts exchanged is 16.2 and the median number of wall posts is 10. In contrast to the low-rate inter-




Figure 4: Time between link creation and the first wall post for user pairs with low levels of interactions.

action group, user pairs in this group exhibit very different interaction patterns. The median duration for time to first interaction is 6 days, compared to 74 days in the low-rate interaction group. This means that most users in the high-rate interaction group exchange their first wall post soon after social links are established. Put another way, these users probably did not need a specific event (e.g., birthday) to initiate their conversation.

We first examine the rate of wall posting over time for the high interaction group. We are particularly interested in knowing whether wall posts are spread out evenly in time, or whether the overall interaction shows an increase or a decrease over time. Figure 5 shows the fraction of total wall posts that are sent each month, out of the total wall posts exchanged between each user pair. Because user pairs vary widely in terms of the number of wall posts exchanged, we show the average rate of interaction for different subsets of the active users: user pairs who exchanged more than 5, 10, 20, 30, and 100 wall posts.




Figure 5: Rate of interaction over time, based on the fraction of wall posts sent since link formation. Wall posts are grouped by month. The plot shows the average rate of interaction for user pairs with different levels of total interactions.

We make two observations. First, the level of the pairwise user interaction is highest right after link establishment and decays over time. In fact, a detailed look into the trace shows that within the first month interaction also drops substantially from week to week. This trend is contrary to what we observed in the low-level interaction group, where we found that the wall posts are more evenly spread out in time because these wall posts are often driven by external events (e.g., birthdays). Second, user pairs who display very high levels of activity (e.g., exchange more than 100 wall posts) show a more even distribution of wall posts over time, compared to the other groups. The characteristic drop in wall

post frequency after the first month (Figure 5) indicates a general trend: users interact less and less over time.

Finally, we examine how long interactions across each of these links last, focusing on the time of cessation of user activity (i.e., when the link is removed from the activity network). Figure 6 shows the fraction of user pairs that no longer interact after a given number of months. We show the trend for user pairs with different levels of activity. Across all levels of interaction, the largest fraction of user pairs discontinue their conversation after the first month of interaction, as expected. However, the dropout rate decreases over time across all activity levels. This indicates a reinforcing relationship between the duration of interaction and the probability of continued interaction: the longer one has engaged in an interaction, the more likely it is to continue. In total, 23% of the user pairs in the high-level interaction group stayed active during the entire one year period.




Figure 6: Fraction of user pairs with no further activity after a given number of months.

3.5 Summary

In this section we characterized the patterns of interaction between pairs of users in Facebook. We observed a strong skew in the total number of wall posts exchanged between user pairs; a majority of the user pairs display very little activity, but a minority are highly active. To better understand the patterns of both active and relatively inactive relationships, we separated user pairs into two groups and examined their interaction patterns in detail. The low-level interaction group typically takes more than a month to initiate their wall interaction since the time they become friends. Often a birthday greeting event serves to break the ice. In contrast, the high-level interaction group initiates conversation immediately after becoming friends, but their rate of conversation decays over time. In summary, our study captures the temporal dynamics of wall post activity for both the low- and the high-level interaction groups, and indicates that links in the activity network change rapidly. In the following section, we examine how such microscopic dynamics impact the structure of the activity network over time.

4. NETWORK EVOLUTION OVER TIME

In this section, we study the impact of the dynamics of user interactions on the global structure of the activity network. In particular, we are interested in capturing (a) how links in the activity network come and go and (b) to what extent overall properties of the activity network change over time.

We study how the activity network evolves over time by deriving multiple snapshots of the network. We generated snapshots of the activity network based on the wall posts

at 90 day intervals during the trace period. In total, we generated nine snapshots of the activity network between September 2006 to January 2009. Each snapshot contains links for all pairs of users who exchanged at least one wall post during the 90 day period. Below, we examine the evolving activity network based on these nine snapshots.

4.1 Refresh rate of the activity network

First, we examine what fraction of the network links persist from one snapshot to the next. To measure the overlap in network links in two consecutive snapshots, we use the notion of *resemblance*. Resemblance is used to measure the quantitative overlap between two sets. In our context, resemblance is defined as the proportion of the network links that remain unchanged over two network snapshots. Resemblance \mathcal{R}_t at time t is defined as

$$\mathcal{R}_t = \left| \frac{P_t \cap P_{t+1}}{P_t} \right|, \quad (1)$$

where P_t is the set of links that are active at time t . The value of \mathcal{R}_t varies between 0 and 1. If $\mathcal{R}_t=1$, then the entire set of active user pairs continued to interact at the next time step. Conversely, if $\mathcal{R}_t=0$, then this means that none of the user pairs who interacted in time t interacted in time $t+1$.

Figure 7 shows the resemblance values over the consecutive snapshots. The average resemblance across all snapshots is 0.45, indicating that, on average, 45% of the links remain active over time. Conversely, the remaining 55% of the links no longer exist in the consecutive snapshot. Moreover, there is a general downward trend, whereby the final pair of snapshots have a resemblance of only 0.30, indicating that the activity network has become even more volatile over time. This may be related to the growth of the Facebook network and the increase in the wall post participation of new users as seen in Figure 1.




Figure 7: Resemblance of the activity network.

The resemblance value is sensitive to the size of the window. When we generate the activity network using a smaller window of 30 days, the average resemblance is 0.29; i.e. 71% of the links refresh over time. However, even when we use a much larger window of 180 days, resemblance value only increased marginally to 0.53. This demonstrates that the links in the activity network changes dynamically even over a large time scale.

Finally, we examine the extent to which a *core* exists in the activity network, where a core is defined as the set of links that exist in every snapshot. Note that the definition of the core is dependent on the window size. For example, for a window size of 180 days, the core contains of 12% of links in the first snapshot. Smaller snapshot windows have a correspondingly smaller core; this fraction goes down to 3% when we consider a window of 90 days. While the

core clearly represents a minority of the links in the activity network, the presence of such a set even over a period of over two years indicates that a certain fraction of the network is stable over time.

4.2 Structural properties over time

We now examine to what extent this rapid fluctuation in the activity network links affects the overall network structure. To do so, we calculated various network metrics for each of our nine network snapshots. We present three popular metrics here: average node degree, clustering coefficient, and average path length. In each of the snapshots, the largest connected component typically accounted for 86% of the entire activity network. Our analysis results are similar for other metrics such as density and betweenness, but we omit these due to lack of space.

Figure 8 shows how each of the three metrics vary over the consecutive snapshots.² Surprisingly, the different network measures are all relatively stable over time. There is some fluctuation towards the end of the trace, which may be influenced by Facebook’s launch of new web design in July 2008 (month 22). Stable trends were also observed for alternative time windows of 30 days and 180 days.




Figure 8: Structural properties of the evolving activity network based on three popular metrics.

Combined with our finding about the fast refresh rate of the activity network, our results on the network properties of the evolving activity network indicate a rather surprising pattern. The set of active links change rapidly over both a short and a large time scale, but over any given period of time, the activity network exhibits strikingly stable structural properties!

4.3 Summary

In this section, we analyzed the patterns of the evolving activity network from two perspectives: the refresh rate of links (or the persistence of user interactions over time) and

²Because it is computationally expensive to find the shortest paths between all user pairs, consistent with previous work [1], we randomly picked 1,000 user pairs that belonged to the largest component of the activity network and computed their average path length.

the resulting global structural properties. We found that while the individual links that constitute the activity network change rapidly over time, the average network properties remained relatively stable.

5. DISCUSSION

One of the primary concerns about online social networks is the extent to which links in the social network represent actual trust relationships between users. The fact that many social networks contain users with hundreds or thousands of friends only makes this concern more worrisome. Since it is difficult to quantify trust on a link, we rely on the patterns of activity to distinguish the strong links (representing users who interact often) from the weak links (representing users who become friends out of courtesy). Our work sheds more light on exactly how users are interacting on these sites. One of the interesting findings of our study is that the mechanisms present on the online social networking sites can affect the activity network in unexpected ways. For example, we found that Facebook's birthday reminder feature caused much of the activity that we observed, implying that this activity may not have occurred had that feature not been in place. This suggests that, going forward, one may have to examine the content and the cause of interaction in order to determine the nature of the link strength between users.

Recently many systems have been proposed that leverage the properties of social networks in various ways. For example, SybilGuard [7] and SybilLimit [8] use the difficulty of forming social network links to prevent Sybil attacks, and Ostra [9] uses the same property to block users who send large amounts of unwanted communication. However, one of the primary concerns with these systems is whether a social network exists which has the requisite difficulty in establishing and maintaining links.

The activity network would seem like a natural fit for such systems as it represents some form of real interaction between users. In [5], the authors present findings that suggest that social network-based systems should be based on the activity network, rather than the social network. Our study informs this argument, showing that the activity network exhibits remarkable stability in high-level structural properties. However, designers of such systems should take into account that a majority of the individual links in the activity network fluctuate rapidly, even though there is a small fraction of links which remain consistently active over time.

6. RELATED WORK

Recently, much work has focused on understanding the structure and evolution of large-scale online social networks [1, 10, 11]. While these studies shed light on the high level growth and topological characteristics of social network, they do not examine how the social network links are being used by users to interact. This has led to a new direction of research focusing on the activity network.

Hyunwoo et al. studied the activity network from the guest book logs of the Cyworld online social network [3]. They constructed the activity network from the comments posted by users in each others' guest books. They observed that the structure of the activity network was similar to the social network and that user interactions tended to be bidirectional. However, Christo et al. [5] used user interactions to demonstrate the opposite trend in Facebook: the structure of the activity network differed significantly from

the social network. For example, they showed that the activity network exhibits fewer small world properties compared to the social network. They also looked at the evolution of number of interactions per user. In contrast, our study mainly focusses on the evolution of pairwise user interactions. Leskovec et al. [12] analyzed an activity network based on user interaction in a large instant messaging network. They studied the activity network structure, online interaction characteristics and its dependency on user demographics. The authors discovered that the activity network displayed strong influence of homophily in interaction, where similar users interact more often.

7. CONCLUSION

In this paper, we studied the interaction between users in the Facebook online social network. We collected data on friendship links and the activity data based on wall posts for a large subset of the Facebook New Orleans network. Our analysis demonstrate that there is a significant skew in the distribution of activity across links: a minority of the user pairs generate a majority of the activity. We observed a general decay in the amount of interaction between pairs of users, suggesting that the activity network is rapidly changing over time. Analysis of the activity network as a whole revealed that even though there is high churn in the user pairs that interact over time, many of the global structural properties remained relatively constant over time.

8. REFERENCES

- [1] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel and B. Bhattacharjee. Measurement and Analysis of Online Social Networks. In *Proc. of IMC*, 2007.
- [2] Facebook Factsheet. <http://www.facebook.com/press/info.php?statistics>.
- [3] H. Chun, H. Kwak, Y-H. Eom, Y-Y. Ahn, S. Moon and H. Jeong. Online Social Networks: Sheer Volume vs Social Interaction. In *Proc. of IMC*, 2008.
- [4] E. Gilbert and K. Karahalios. Predicting Tie Strength With Social Media. In *Proc. of CHI*, 2009.
- [5] C. Wilson, B. Boe, A. Sala, K. P. N. Puttaswamy and B. Y. Zhao. User Interactions in Social Networks and their Implications. In *Proc. of Eurosys*, 2009.
- [6] Facebook Unveils Next Evolution of Site Design. <http://www.facebook.com/press/releases.php?p=47448>.
- [7] H. Yu, M. Kaminsky, P. B. Gibbons and A. Flaxman. SybilGuard: Defending Against Sybil Attacks via Social Networks. In *Proc. of SIGCOMM*, 2006.
- [8] H. Yu, P. B. Gibbons, M. Kaminsky and F. Xiao. SybilLimit: A Near-Optimal Social Network Defense against Sybil Attacks. In *Proc. of IEEE S&P*, 2008.
- [9] A. Mislove, A. Post, K. P. Gummadi and P. Druschel. Ostra: Leveraging Trust to Thwart Unwanted Communication. In *Proc. of NSDI*, 2008.
- [10] Y-Y. Ahn, S. Han, H. Kwak, S. Moon and H. Jeong. Analysis of Topological Characteristics of Huge Online Social Networking Services. In *Proc. of WWW*, 2007.
- [11] R. Kumar, J. Novak and A. Tomkins. Structure and Evolution of Online Social Networks. In *Proc. of SIGKDD*, 2006.
- [12] J. Leskovec and E. Horvitz. Planetary-Scale Views on a Large Instant-Messaging Network. In *Proc. of WWW*, 2008.