Making social interactions accessible in online social networks

Fredrik Erlandsson [a,*], Roozbeh Nia [a], Henric Johnson [b] and Felix S. Wu [a]

[a] *University of California, Davis, CA, USA*

[b] *Blekinge Institute of Technology, Karlskrona, Sweden*

Abstract. Online Social Networks (OSNs) have changed the way people use the internet. Over the past few years these platforms have helped societies to organize riots and revolutions such as the Arab Spring or the Occupying Movements. One key fact in particular is how such events and organizations spread through out the world with social interactions, though, not much research has been focused on how to efficiently access such data and furthermore, make it available to researchers. While everyone in the field of OSN research are using tools to crawl this type of networks our approach differs significantly from the other tools out there since we are getting all interactions related to every single post. In this paper we show means of developing an efficient crawler that is able to capture all social interactions on public communities on OSNs such as Facebook.

Keywords: Online social networks, crawling, social graph

1. Introduction

Recently, online social networks, OSNs, have gained significant popularity and are among the most popular ways to use the Internet. Additionally, researchers have become more interested in using the social interaction networks, SINs [9], in order to further enhance and personalize their services [11]. OSNs are also redefining roles within the publishing industry, allowing publishers and authors to reach and engage with readers directly [8]. However, SINs are not very easily available as of today through the current APIs provided by most OSNs. Such applications would therefore spend tremendous amount of time trying to gather the required SINs for their services. Therefore, our research problem is how we can design a system that makes social interactions in OSNs accessible. This also refers to the problem of how to crawl OSNs in a structured way, which is the focus of this short paper.

The nature of OSNs and the amount of information available makes the problem of what to crawl interesting. To narrow down the scope of the proposed research, we are focusing on the interactions in OSNs. By doing this, we noticed that there exist a gap and segregation between content and social graph. To simply provide social informatics for social computing applications, we have developed a crawler that serves as a bridge between the content and social graph in the online world, by not only providing which users have interacted with each other but around exactly which content these interactions have occurred.

Privacy of the user is a major concern when it comes to all online social interactions and crawling as discussed in [5]. We are treating the crawled data with high respect to the integrity of the people behind the users.

---

[*]Corresponding author. E-mail: fredrik.erlandsson@bth.se.

2. Related work

Despite the huge number of social network publications, few have been dedicated to the data collection process. Chau et al. [3] briefly describe using a parallel crawler running breadth-first search, BFS to crawl eBay profiles quickly. The measurement conducted by Mislove et al. [7] is, to the best of our knowledge the largest OSN crawling study ever published. From four popular OSNs, Flickr, Youtube, LiveJournal and Orkut, 11.3 M users and 328 M links are collected. Their analysis confirms known properties of OSNs, such as a power-law degree distribution, a densely connected core, strongly correlated in-degree and out-degree, and small average path length.

Other studies on OSN crawlers include [2,6]. Gjoka et al. [6] proposed two new unbiased strategies: Metropolis–Hasting random walk (MHRW) and a re-wighted random walk (RWRW). Where Catanese et al. [2] described the detailed implementation of a social network crawler. It used the BFS and uniform sampling as the crawling strategies to run the crawler on Facebook, and then compared the two strategies.

3. A platform to make interactions accessible

We have designed and developed a system that is able to crawl open data. Initially, we will focus on the Facebook Graph API to crawl all content that is viewable to users; such as posts, comments, likes on posts and comments, and shares of posts.

3.1. Design

We have designed our crawler to operate in two stages. *Stage one* uses the Facebook's unique identifier of a public community (page or a group) to find the id of all posts, messages, photos, and links posted on the given community by admins and members. For readability, a post will refer to anything shared on a community on Facebook in this paper. *Stage two* is a bit more complicated; for each post gathered in *stage one* we send at least three to four separate requests (assuming that there are no "likes" on comments), one for the post itself, one for the "likes" on the post (if there exist any), one to get information on who have shared the post and finally one to get all comments (if there exist any). If one of the responses is paginated we have to make consecutive requests to gather the complete view. This also means that for posts with a lot of interactions we have to make multiple requests to the graph. For instance, we have crawled posts with hundreds or thousands of comments each with a few likes, where we have to make a request for each comment to get its likes. To scope the huge number of requests and the requirement to be efficient, our crawler is built as a distributed service much like discussed in [3]. Figure 1 shows a basic sketch how the controller and the crawling agents are connected.

3.2. Statistics

Over the last eight months our tool have gathered a bit over 150 GB of structured data, including: 93 million unique Facebook users, 14 million posts, 126 million comments and over 800 million likes.

4. Challenges and requirements

Our crawler highly depends on Facebook's API, and therefore, bugs in Facebook's API will cause problems that we have no control over. Also, resource limitations has forced us to be picky about which
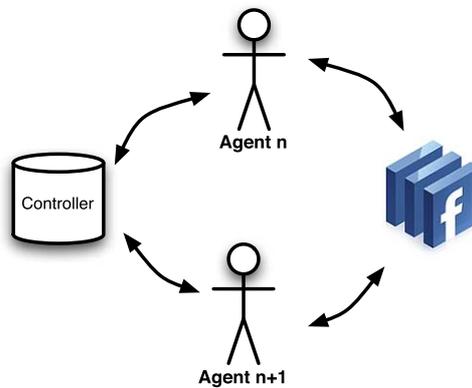
Fig. 1. Our distributed crawling mechanism. (Colors are visible in the online version of the article; http://dx.doi.org/10.3233/ISU-130702.)

communities to crawl. Given enough resources, our crawler can be modified to automatically crawl all public communities on Facebook and other OSNs given an initial set of seeds.

4.1. Requirements

Our crawler tool, from a high level, is simply a black box that takes the identifier of a Facebook community as input and outputs a stream of documents. In addition to capturing the response of API requests, our crawler has to satisfy the following requirements:

Real-time. The information and interactions on Facebook public communities is extremely time-sensitive. In most cases, it is very important to crawl and parse a given post in a community on Facebook online. A few important questions that arise due to the nature of how the interactions around posts evolve are (1) "Which posts do we have to re-crawl to get the most updated information" and (2) "When would be the best time to re-crawl these posts".

Coverage. It is important and desirable to be able to crawl each and every post thoroughly and completely. However, if resources do not allow this, it is more desirable to get all the data from a limited set of posts, rather than less data from a larger set of posts.

Scale. As of today there are over a billion users and millions of public communities on Facebook [10]. There are over 2.7 billion likes and comments posted on Facebook on a daily basis as of February 1st [4].

Data quality. The crawler should output good quality and uncorrupted data. Therefore, it needs to be able to detect failures in Facebook's current API and be able to restart from exactly where it stops when a failure occurs.

5. Applications of SINs

There are a vast number of applications where SINs can be used, here we give a brief description of two we have used to evaluate our dataset.

Dynamic news feed: People spend hours on Facebook every day. However, they are only bound to see the posts shared by their immediate friends and pages they have liked. Using social interactions, it is possible to identify the type of posts the user has been interacting with and find similar posts based on the SIN formed around it that the user has not interacted with. This will create a more dynamic newsfeed

rather than the current one where users see the same posts over and over again throughout the day. We can identify which posts the user would be interested in using social interactions but has not interacted with yet. Therefore, the user will only see posts that he/she has not seen before and the content is socially related to what he/she likes.

Social search: Social search [1] is one of the hottest areas in the market and companies like Google, Facebook and Microsoft are spending billions of dollars in the race of building the best social search experience. We believe that the SINs formed around the content shared on these pages and groups give better results when combined with a search engine than the friendship networks currently used. While a group of users have very similar and close interactions around the content shared on Facebook, we can use this information when a person from this group queries something. We know the group's interests and that will help us serve the user with better social search results. Since there is a cap on how many friends users can have on Facebook, the social search will be limited to the number of direct friends. In addition to the limited social network, there are no guarantees that users immediate friends will share the same taste, thought process, or needs. In our approach we can link users with many interactions on related content to provide better search results. Based on the query we can identify the context and use the matching SIN to find related content.

6. Conclusion

We have shown means of building an extensive tool to gather data from public communities on OSNs. Our distributed crawler satisfies all of our requirements in order to retrieve the complete set of non-corrupted data, including all the content shared and all the user interactions around them. We discuss various applications and how they can benefit from leveraging SINs in order to further personalize their services. Finally, we have given a short description of how to design a data-mining tool for OSNs that can be used to gather data.

References

[1] P. Bhattacharyya, J. Rowe, S.F. Wu, K. Haigh, N. Lavesson and H. Johnson, Your best might not be good enough: Ranking in collaborative social search engines, in: *7th International Conference On Networking, Applications and Worksharing*, 2011, pp. 87–94.

[2] S.A. Catanese, P. De Meo, E. Ferrara, G. Fiumara and A. Provetti, Crawling Facebook for social network analysis purposes, in: *Proceedings of the International Conference on Web Intelligence, Mining and Semantics, WIMS'11*, ACM, New York, NY, USA, 2011, pp. 52:1–52:8.

[3] D.H. Chau, S. Pandit, S. Wang and C. Faloutsos, Parallel crawling for online social networks, in: *Proceedings of the 16th International Conference on World Wide Web, WWW'07*, ACM, New York, NY, USA, 2007, pp. 1283–1284.

[4] T. Cheredar, *Facebook User Data*, February 2012.

[5] F. Erlandsson, M. Boldt and H. Johnson, Privacy threats related to user profiling in online social networks, in: *Privacy, Security, Risk and Trust (PASSAT), 2012 International Conference on and 2012 International Conference on Social Computing (SocialCom)*, September 2012, pp. 838–842.

[6] M. Gjoka, M. Kurant, C.T. Butts and A. Markopoulou, Walking in Facebook: a case study of unbiased sampling of OSNs, in: *Proceedings of the 29th Conference on Information Communications, INFOCOM'10*, IEEE Press, Piscataway, NJ, USA, 2010, pp. 2498–2506.

[7] A. Mislove, M. Marcon, K.P. Gummadi, P. Druschel and B. Bhattacharjee, Measurement and analysis of online social networks, in: *Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement, IMC'07*, ACM, New York, NY, USA, 2007, pp. 29–42.

[8] A. Mrva-Montoya, Social media: New editing tools or weapons of mass distraction?, *The Journal of Electronic Publishing* **15** (2012).

[9] R. Nia, F. Erlandsson, P. Bhattacharyya, R. Rahman, H. Johnson and F. Wu, Sin: A platform to make interactions in social networks accessible, in: *ASE International Conference on Social Informatics*, December 2012, pp. 205–214.

[10] TechChrunch, Facebook announces monthly active users were at 1.01 billion as of September 30th, an increase of 26% year-over-year, September 2012.

[11] C. Wilson, B. Boe, A. Sala, K.P. Puttaswamy and B.Y. Zhao, User interactions in social networks and their implications, in: *Proceedings of the 4th ACM European Conference on Computer Systems, EuroSys'09*, ACM, New York, NY, USA, 2009, pp. 205–218.