

Decentralization: The Future of Online Social Networking

Ching-man Au Yeung¹, Ilaria Liccardi¹, Kanghao Lu², Oshani Seneviratne², Tim Berners-Lee²

¹ School of Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, UK

² Decentralized Information Group, Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

Introduction

Social networking forms an important part of online activities of Web users. Web sites such as Facebook, MySpace and Orkut have millions of users using them everyday. However, these sites present two problems. Firstly, these sites form information silos. Information on one site is not usable in the others. Secondly such sites do not allow users much control over how their personal information is disseminated, which results in potential privacy problems.

This paper presents how these problems can be solved by adopting a decentralized approach to online social networking. With this approach, users do not have to be bounded by a particular social networking service. This can provide the same or even higher level of user interaction as with many of the popular social networking sites we have today. In addition, it also allows users to have more control over their own data. A decentralized social networking framework described is based on open, technologies such as Linked Data [Berners-Lee 2006], Semantic Web ontologies, open single-signon identity systems, and access control. The use of URIs as identifiers throughout allows the decentralized framework to be distributed and extensible, as users, applications and data to be linked to by referring to their URIs.

Motivation




Figure 1: The economist laments "The problem with today's social networks is that they are often closed to the outside web."

Reproduced from www.economist.com

(illustration by David Simonds) [Economist 2008]

activities in external web sites [Malik 2007]. However with decentralized control, no one service has the sole access to data and the capability to enforce arbitrary decisions like that.

On popular social networking sites such as Facebook, MySpace and Orkut, users are given the impression that they are in control of their own data. But this is not always the case. Facebook provides users with an option of deactivating an account. However, it is not possible to completely erase all personal information from the site [Aspan 2008]. While users have privacy settings in which they can specify who may access their personal details, there were instance where this was not the case. For example, it was reported that more than half a million images were leaked from MySpace without any consent from the users [Poulsen 2008]. In addition, MySpace still draws concerns about child stalking [Owyang 2008].

Most of the social networking sites we have today (as of 2008), also hinder creativity because they impose restrictions on how new applications using the social graph can be created. Developers are limited to a set of functions offered by the underlying API, and privacy rules which prohibit much of the data use by third party applications. Although privacy rules in general are considered good, it should also be noted that the companies operating the social networking site have the exclusive control of the users' personal data to be used for their financial advantage. We also cannot expect much interoperability across different social networking sites. In other words, applications cannot be reused. Figure 1 aptly illustrates this aspect. It depicts how people are tied up within a particular social networking site and would like to "jump out of the walled gardens" to enjoy the benefits of the "other" social networking sites, such as to share their data with their friends who may be members of other social networking sites.

These limitations of existing social networking services, motivated us to consider a decentralized framework for online social networking. As Figure 2 illustrates, there is a varying degree of "Closed Data/Centralized Control" to "Open Data/Decentralized Control" with respect to social networking. Data silos offered by proprietary social networking sites such as Facebook operate on the basic web protocols such as HTTP, HTML, CSS and Javascript, but they do not expose their data to the Data Web in a structured format such as RDF. However, sites such as [LiveJournal](http://www.livejournal.com), [Advogato](http://www.advogato.org) and many others output the data in RDF and even allow links from outside of the boundaries of their sites (see <http://esw.w3.org/topic/FoafSites> for a comparison matrix of many of the popular social networking sites with respect to the number of people subscribed to the service, the ability to 'link out' and OpenID integration). We are aiming to provide a framework which goes even beyond, and will let users output their FOAF information and edit it through open protocols such as [WebDAV](http://www.webdav.org). We have demonstrated how to implement this by reconfiguring the Apache Server configurations or by using [SPARUL \(SPARQL Update\)](http://www.sparul.org) with systems such as [Algae](http://www.algae.org) and [ARC2](http://www.arc2.org) (see <http://esw.w3.org/topic/EditingData> for more examples).

Existing social networking services are centralized and the companies providing the services have the sole authority to control all the data of the users. It is not a trivial task for a user to reuse his own data, including his social network, messages with friends and photos among other applications, as there are not many robust mechanisms to port all the data from one platform to another. Essentially, "People are getting sick of registering and re-declaring their friends on every site." [Fitzpatrick and Recordon 2007]. Moreover, users usually have little control over how and what information about them are presented to their friends online. Presentation of their information largely depends on the design of the social networking service the users are using.

Likewise, the users have to agree to the policies of the social networking sites when using their services, even though they may involve usage of their data for targeted advertising. In addition, very often users need to explicitly opt-out of certain applications if they are more conscious about the privacy of their data. For example, Beacon, which is a part of Facebook's advertisement system, upset many users because it published news stories on their friends' news feeds about their

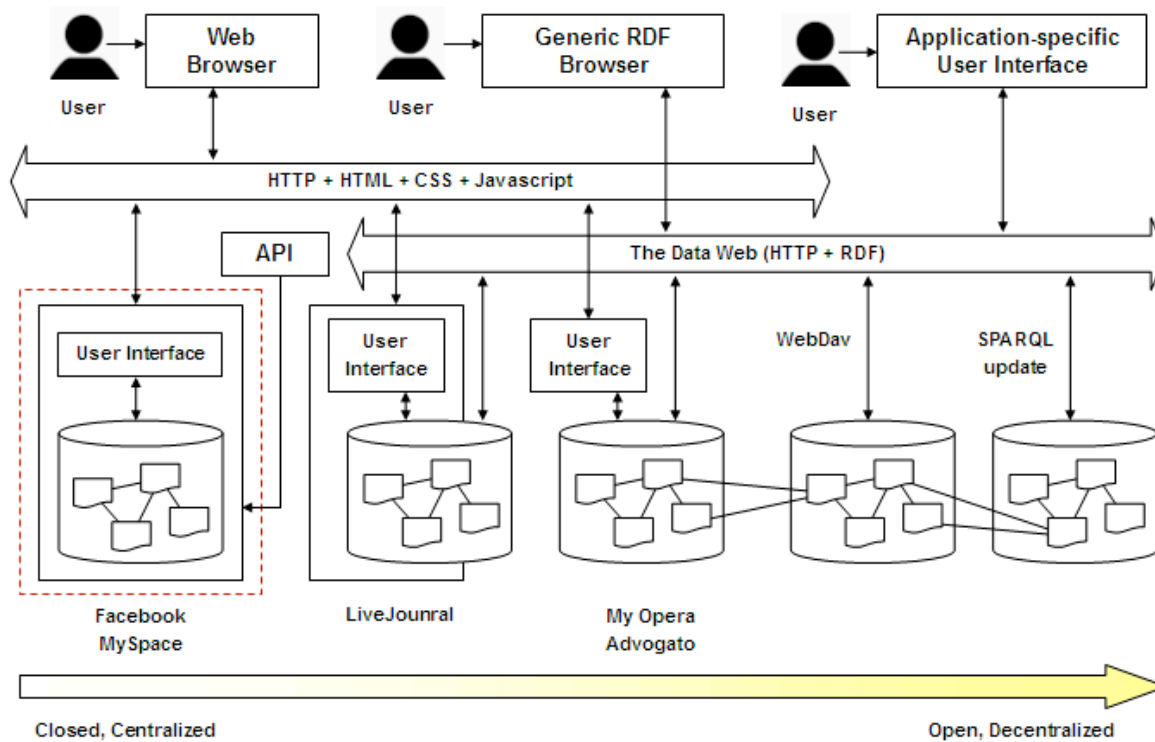


Figure 2: The spectrum of online social networking

We believe such a decentralized setting gives back the users the control of their own data. In particular, in the following three respects:

- **Privacy:** Users in decentralized social networks decide who to show the information to and what restriction there is on the data.
- **Ownership:** As the information is stored on a trusted server or on the local computer, users have complete ownership of the data. They would not lose their data suddenly as the proprietary service hosting their data decides to shutdown without much notice. [Wilson 2008]
- **Dissemination:** Information is disseminated according to users' preferences and friendship relations.

We suggest that this provides a better infrastructure for social resource discovery as users can go beyond boundaries of social networking sites to look for other users with similar interests, similar to decentralized P2P sharing systems [Wang and Sun 2008]. Such a framework also allows to develop applications for use in online social networking activities. In fact, NoseRub [Noserub] have already worked in a similar direction by proposing a protocol for decentralized social networking which uses existing standards like OpenID, RSS and FOAF. We believe such notion can be extended by incorporating more Semantic Web technologies and access control policy management mechanisms to provide a better social networking platform on the Web.

Decentralized Online Social Networking

In a decentralized social networking framework, a user does not need to join any particular social networking service such as Facebook or MySpace. Instead, the user chooses a server which he trusts to host his own data such as his FOAF (Friend-Of-A-Friend) [Brickley & Miller 2007] file, his activity log and his photo albums. Given that we refer to these files with their URIs, they can actually be stored on different servers.

A user's FOAF file plays a central role in the framework we propose. The FOAF specification provides a format for specifying "friend" relationships among people. It is becoming increasingly popular, and has an expanding community [Golbeck and Rothstein 2008]. The FOAF vocabulary is used to provide the basis for data interoperability among disparate social networks. Of course, the FOAF ontology can be extended or complemented by other ontologies such that richer information can be included in one's FOAF file.

As Figure 3 illustrates, an individual can obtain an identity on the web in the form of a URI, we refer to as a "Web ID". This Web ID could point to a reference in the user's FOAF file stored on a server that the user trusts. For example, Alice's Web ID could be something like "http://alice-trusted-server.com/card#alice". In addition to giving a person an identity with its URI, a FOAF file also contains personal details and a list of acquaintances linked by the "foaf:knows" property. Unlike in the traditional data siloed social networking sites we have today, each of these foaf:knows links serve as a pointer to Web IDs of the people that the user knows. In this way, users can have their data stored in different places and still be linked to each other, which avoids centralized data storage and control by proprietary social networking services.



Figure 3: Social network graph constructed using FOAF data

By using FOAF in a decentralized social networking framework, the Web ID of a user can be used as an access point to his data. Other users who want to access this user's social network (friends list), his status, his photos, or to write on his personal message board, will go to his FOAF file and obtain the corresponding URIs. By storing the data in a trusted server chosen by the user, the users are given more control over the data. There's also an option to create much fine-grained access control policies using policy languages such as AIR [Kagal et al. 2008] to help restrict access to his data or applications. Unlike in a centralized social networking site, users will have to authenticate themselves against different servers when they want to access the restricted data of their friends. This can be done for example by using the OpenID protocol [Recordon and Fitzpatrick 2006] by letting users create online identities making use of the existing protocols such as URI, HTTP, SSL and Diffie-Hellma etc or using the users' FOAF+SSL certificates. Such a decentralized framework also allows higher customization of the applications and interfaces. For example, users can create their own homepage which shows their social network, online activities and photos. Notice that once a FOAF compatible server software is developed, it can be installed on different servers and they can link to one another as illustrated in Figure 4.

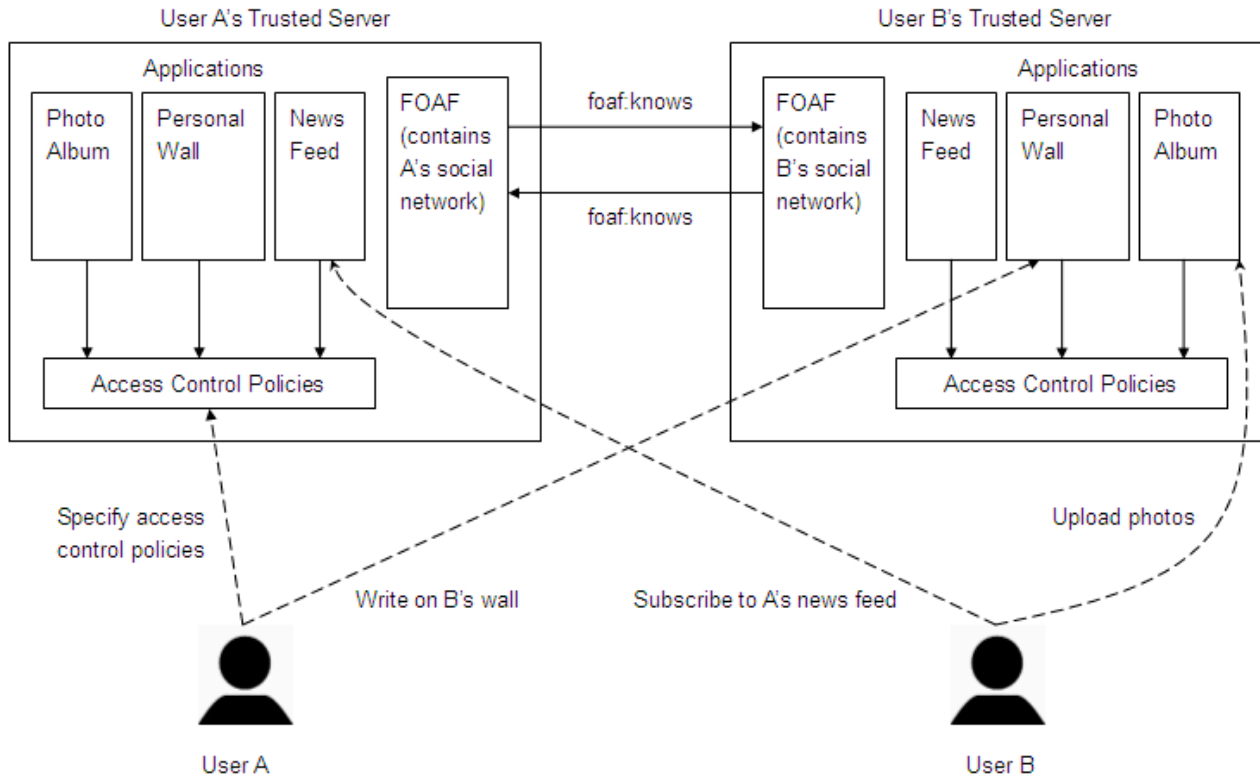


Figure 4: A framework of decentralized online social networking

Now we describe how three popular applications in current social networking sites look in a decentralized version of online social networking and mention their advantages.

Personal Wall

A wall in Facebook is a space where you accept information from others. It is semi-private in the sense that you should be able to determine who can write and view the information on this space. Within the decentralized social networking infrastructure, a user can choose a server he trusts to host the wall application and specify access control policies such that only certain users will be allowed to perform certain actions.

For example, we can have the following options on a wall:

- **Friends' Update**: only friends who are listed in the FOAF file can write on the wall. This is similar to the default policy in Facebook. But users who want to post something on a Facebook user profile must first obtain a Facebook account. However, with a decentralized social network setting, the only dependency on a third party site is to provide authentication. As long as the user has access to a web-editable space to record the FOAF file and keep other related files to maintain state of his friends' network, there should not be any dependency on an external centralized data silo.
- **Universal Update & Private Access**: this allows anyone to write on a user's wall but only the owner of the wall to be able to see the changes. This is analogous to a normal Webmail, where only the owner of the mail can see the messages and has the ability to control what he sees by applying filters.
- **Universal Update/Read (except robots)**: this policy gives the read and write rights for the wall, so that anybody on the Web can write on the wall, except automated programs. This is similar to the policy in Wikipedia.

In addition, a user can be allowed to create more sophisticated access control policies, such as restricting certain words or phrases to be posted to the wall.

Photos

Tagging users in photos is another interesting application on existing social networking sites. The activities of a user (called news feed in Facebook) may be automatically updated when someone has tagged himself in a photo. In a decentralized setting, a user can be given more control of the news to be included in his own feed.

Suppose User A wants to tag User B in a photo which is uploaded to User A's trusted server. Since User A owns the photo, he can maintain any tag on his trusted server. The framework makes sure that User B will be notified when User A tags User B by sending a SPARQL update to Server B. User B can configure the settings on the server software which will accept or ignore the notification updates. It is up to User B to maintain knowledge of the existence of the photo uploaded by User A on his trusted server. User B can choose to generate corresponding news feed for this tagging event.

News feed

The "News Feed" is one of the controversial features in Facebook which draws concerns of privacy in social networks [Arrington 2006]. Many users complain that the feature reveals too much of their activities on the social networking site. In a decentralized setting, users will be allowed more control over their own data, including the log of their activities. A news feed similar to the ones in Facebook can be provided by publishing an RSS feed and including its URI in the user's FOAF file.

Subscribing to the feed is subject to certain access control policies specified by the user. In the simplest form, an access control policy allows users who are specified as friends in the FOAF file (by using the foaf:knows property) to subscribe to the feed. It is also possible to serve different information depending on the identity of the user who submits the request. In current social networking sites, such functionalities, although possible, depend on the design of the sites which the users have little control.

A Possible User Interface to Decentralized Social Networking

▼ Dan Brickley

This is you

Basic Information

danbri homepage
my.opera homepage
danbri.livejournal weblog
my.opera weblog

Openid <http://danbri.livejournal.com/>

Nickname danbri
danbri

Acquaintances

Aaron Swartz
Amy van der Hiel
Art Barstow
Benjamin Joffe
Charles McCathieNeville
Damian Steer
Dan Connolly
Dave Beckett
Dean Jackson
Edd Dumbill
Eddie Lopez
Eric Miller
Eva Méndez
Gregory J. Rosmaita
Jan Grant
Jim Ley
Joe Brickley

Tip: Do you have **web ID**?

[Make or set A Web ID](#)

See Also <http://foaf.qdos.com/reverse/?path=http://danbri.org/foaf.rdf#danbri>

Figure 5: The Tabulator FOAF pane

We have developed several components which can be used in a decentralized online social networking setting. [Tabulator](#) [Berners-Lee 2008] represents our effort in terms of a possible user interface. It is designed to be a generic data browser and editor for linked RDF data on the Web with the motivation of providing a natural and a seamless experience for browsing and editing data. The Tabulator also provides a system of Views and Panes which allows the user to see data in different perspectives. The [FOAF pane](#) is one such application specific view which lets a user to explore his 'Friends Network' or the 'Acquaintances' as shown in the Figure 5. The FOAF pane becomes visible when a user is viewing any FOAF data, and lets you register who you are in the graph by letting Tabulator know your "Web ID". If the user does not have a "Web ID", tabulator provides a mechanism to enter the user data to bootstrap his "online social network". The user can choose to save his personal details on a trusted server as long as it supports the WebDAV or SPARQL Update protocols, and the Tabulator Firefox extension provides editing capabilities on the browser itself to modify the data.

There are a number of technical challenges when emulating most functionalities offered by centralized social networking sites in a decentralized social networking setting. For example, in a centralized social network, users' profile picture thumbnails are kept in a predetermined format and in a specific location. However, in a decentralized social network users have the freedom to specify their "foaf:depiction" in whatever manner, and the User Interface should be intelligent enough to cache thumbnails to

display the profile images in a consistent form. Another challenge is verifying reverse links. This is because when it is decentralized, it becomes very easy to claim that you know a certain user by specifying his "Web ID" use the "foaf:knows" property. However, this claim may not be true and it would be worthwhile to provide robust verification mechanisms. We are doing some research on the aspects of providing access control to resources on the Web, authenticating users based on custom policies and providing accountability on the Web to address many of these technical challenges.

Conclusions

Online Social networking sites are extremely popular. Those sites provide the means for many Web users to maintain contacts, communicate and exchange information with each other. While existing social networking sites offer a lot of interesting functionality, they bring potential problems related to privacy, information accountability and ownership of information.

Decentralized social networks have the potential to provide a better environment within which users can have more control over their privacy, and the ownership and dissemination of their information. Therefore, online social networking will be more immune to censorship, monopoly, regulation, and other exercise of central authority [Agre 2003]. More importantly, a decentralized approach to online social networking breaks the boundaries between social networking sites by providing users more freedom to interact with each other.

One major challenge of realizing decentralized online social networking is its adoption by users. Users who are already participating in existing social networking sites will need to migrate their data to decentralized social networks to break away from the traditional data silos offered by the current social networking sites. It is understandable that users would be resistant to change and may not be keen on changing applications even if this provides all the features they need. In fact it has been noted that "even though individuals express concerns and awareness about Internet privacy, they are still willing to engage in risky online activities" [Campbell 2001]. Good user interfaces, tools for importing and exporting data, and ease of setting up the software will definitely be part of the solution to this problem as suggested by the framework we propose.

Acknowledgements

The authors express gratitude towards Dan Brickley, Lalana Kagal and Ralph Swick for their comments and contribution to the ideas presented in this paper.

References

- ^ [Agre 2003] Philip E. Agre. P2P and the Promise of Internet Equality. Communications of the ACM, Vol.46, No.2, pp.39-42, 2003.
- ^ [Arrington 2006] Michael Arrington. Facebook Users Revolt, Facebook Replies. September 6, 2006. <http://www.techcrunch.com/2006/09/06/facebook-users-revolt-facebook-replies/>.
- ^ [Aspan 2008] Maria Aspan. How Sticky Is Membership on Facebook? Just Try Breaking Free. The New York Times. February 11, 2008.
- ^ [Berners-Lee 2006] Tim Berners-Lee. Linked Data. 2006. <http://www.w3.org/DesignIssues/LinkedData.html>
- ^ [Berners-Lee 2008] T Berners-Lee, J. Hollenbach, Kanghao Lu, J. Presbrey, E. Prud'ommeaux, mc schraefel. Tabulator Redux: Browsing and Writing Linked Data. Linked Data on the Web Workshop at WWW08. <http://sunsite.informatik.rwth-aachen.de/Publications/CEUR-WS/Vol-369/paper11.pdf> 2008
- ^ [Brickley & Miller 2007] Dan Brickley, Libby Miller. FOAF vocabulary specification. <http://xmlns.com/foaf/spec/>. 2007.
- ^ [Campbell 2001] J. Campbell and R.C. Sherman and E. Kraan and Z. Birchmeier, Internet Privacy Awareness and Concerns amon College Students, APS Toronto, June 2001 . 2001.
- ^ [Economist 2008] Everywhere and nowhere. May 19, 2008. http://www.economist.com/business/displaystory.cfm?story_id=10880936. 2008.
- ^ [Golbeck and Rothstein 2008] Jennifer Golbeck, Matthew Rothstein. Linking Social Networks on the Web with FOAF: A Semantic Web Case Study. In Proceedings of the Twenty-Third Conference on Artificial Intelligence (AAAI'08), Chicago, Illinois, July 13–17, 2008m, pp.1138-1143. 2008
- ^ [Fitzpatrick and Recordon 2007] Brad Fitzpatrick and David Recordon. Thoughts on the Social Graph. <http://bradfitz.com/social-graph-problem> 2007.
- ^ [Kagal et al. 2008] Lalana Kagal, Chris Hanson, Daniel J. Weitzner. Using Dependency Tracking to Provide Explanations for Policy Management. In Proceedings of the 9th IEEE International Workshop on Policies for Distributed Systems and Networks (POLICY 2008), 2-4 June 2008, Palisades, New York, USA, 54–61. IEEE Computer Society.
- ^ [Malik 2007] Om Malik. Is Facebook Beacon a Privacy Nightmare? November 6, 2007. <http://gigaom.com/2007/11/06/facebook-beacon-privacy-issues/>.
- ^ [Noserub] <http://noserub.com/>
- ^ [Owyang 2008] Jeremiah Owyang. The Many Challenges of Social Network Sites. February 11th, 2008. <http://www.web-strategist.com/blog/2008/02/11/the-many-challenges-of-social-networks/>
- ^ [Poulsen 2008] Kevin Poulsen. Pillaged MySpace Photos Show Up in Massive BitTorrent Download. Wired. January 23, 2008.
- ^ [Recordon and Fitzpatrick 2006] David Recordon and Brad Fitzpatrick. OpenID authentication 1.1. 2006. <http://openid.net/specs/openid-authentication-1.1.html>
- ^ [Wang and Sun 2008] Fang Wang and Yaoru Sun, Self-organizing peer-to-peer Social Networks. Computational Intelligence, Vol. 24, No. 3. (August 2008), pp. 213-233.
- ^ [Wilson 2008] Kelly Wilson. Hometown Has Been Shutdown. <http://www.peopleconnectionblog.com/2008/11/06/hometown-has-been-shutdown>.