Steven Lehrburger
When Strangers Meet – Kio Stark
NYU Interactive Telecommunications Program
17 April 2009

# ITP Thesis Project Proposal: Constellate

*a walled garden, but one with many gates*

*a tool for consolidating social relationships on the Internet*

The problem with the plethora of social websites is not that people have to keep track of and maintain multiple profiles/identities - we do that in real life too (when we act differently in different situations, talk about different things with different people, etc) and we're pretty good at it. The problem is instead that we have to keep track of multiple profiles/identities for every person with whom we interact in more than one context. It is this social/cognitive task which we find so disorienting.

The thesis project I propose to design, develop, launch, and market is an online service that will solve the problem of the distributed/duplicated social graph. This problem - that we have to re-add all of our friends for each new social web/mobile service we want to try - is faced by both application users and developers. For users, Constellate will act as a single location for maintaining information about relationships, including functionally for folksonomic grouping of contacts and both symmetric and asymmetric relationships. For developers, Constellate will provide a robust API that they can use to enable them to easily import and synchronize their data with their users' entire existing social graph.

In the real world, management of our social graph – the network of all of the

people we know combined with how well and in what contexts we know them – and the ways in which content flows across it is intuitive. We know different people in different contexts, and these contexts often have defining physical boundaries. Examples of these contexts include the office at work, the house where one's parents live, the coffee shop where you meet your friends, and the grocery store. Previously, most of the content the sharing of content we did (where content includes spoken words, written words, four by six photo prints, home videos, mix tapes, etc.) happened distinctly in one of those contexts. If one had taken several rolls of film during a vacation, for example, then one would show all the photos to those friends that were also on the trip, but might carefully curate the selection of photos to be shown to one's coworkers. Previously, the sharing of our personal content was transparent and intuitive to as well as controllable by the individuals doing the sharing.

This is no longer the case: the rise of the popularity of the Internet as a medium for sharing those aforementioned types of content has stripped much of our control over who sees what. Some controls do exist, but they are insufficient. One can post an album of vacation photos on Facebook, but there are only rudimentary ways to share all the photos with a select group (such as the friends you were with) and to share a selection of the photos with another group (such as coworkers). Facebook does offer tools for organizing friends in Lists, and these Lists, once made, can be used to control the sharing and filter the receiving of content. Friendship on Facebook is binary – either you are friends with someone or you are not – and while this is not a sufficiently nuanced representation of real-world relationships, Friend Lists can be used to organize contacts into user-defined groups. These Lists cannot be created in a hierarchy, however: the list

of people in one's freshman dormitory at college cannot be automatically a subset of everyone who attended the same college in that year.

Furthermore, the tools that Facebook provides for the creation and management of these Lists are paltry at best. Information about which of a user's Friend Lists a particular Friend is a member of is not displayed on that Friends primary profile page or elsewhere. Although it's possible to quickly select multiple Friends to add to a single List, it is not possible to quickly add a single Friend to multiple Lists. When the List features were first introduced many users already had several hundred Friends. Thus, in order to effectively organize these Friends those users had to go through every single one of their Friends once for each Friend List that the user wished to create, and then decide if that Friend should be on that List; this cognitive chore was overwhelming in the absence of well-designed tools.

Furthermore, these Friend Lists are limited in scope to the Facebook website/environment. Although Facebook does offer functionality for connecting a Facebook account to an account on other services, Facebook Connect[1] is primarily used by these other services for pushing updates about a user's use of that service to his/her Facebook News Feed[2], and *not* for importing/syncing information about their contacts. Vimeo, for example, does supports Facebook Connect for these feed updates[3] but does not support the importing of Facebook Friends[4]. In addition, even if this latter type of functionality did exist on Vimeo via Facebook connect, users might not want to use it due

---

[1] Facebook. *Help Center > Facebook Connect*. http://www.facebook.com/help.php?page=730

[2] boyd, danah. 2006. "Facebook's 'Privacy Trainwreck': Exposure, Invasion, and Drama." Apophenia Blog. September 8. http://www.danah.org/papers/FacebookAndPrivacy.html

[3] Vimeo. *Settings / Extending Vimeo*. http://www.vimeo.com/settings/extend

[4] Vimeo. *Find your Gmail and Yahoo contacts*. http://www.vimeo.com/contacts/import

to the former automatic publishing of feed updates. If a person enjoys trolling on reddit, he or she might want to find which of his Friends on Facebook also enjoyed trolling on reddit, but he or she certainly would not want to broadcast updates of that trolling to all of his or her Friends on Facebook. Facebook has created a walled garden through its attempts to integrate all of a user's activity on the Internet under a single umbrella of personhood. This garden does not service as an environment conducive to all types of content sharing, however, and thus encourages the existence of additional online environments.[5]

Thus, users are faced with the nearly-insurmountable hassle of recreating their entire social graph on each new service they wish to try. This problem is especially pronounced on less popular services such as Flickr, Goodreads, and Vimeo: users simply do not have the time or energy to establish contacts with several hundred people, nor do they have the time or energy to keep their contacts on each of these services up to date as they make new contacts and lose old ones. People have friends who are using the same services that they are using all across the Internet, but they don't know about these friends because they have no way to find their accounts.

These same problems exist for the developers of new social web services. These services often need a critical mass of users to thrive. This is especially pronounced for location-based social applications such as Loopt and Foursquare that are only fun for users when these users have enough other friends using them at the same time. Note only does each web service have to somehow convince enough users to add their friends as contacts, they must also each individually store and maintain a separate representation of the social graph.

---

[5] Jorge Ortiz. unpublished conversation. 24 April 2009.

Furthermore, the current distributed and duplicated nature of the social graph on the Internet makes it more difficult to keep track of less-well known people and to meet new people. We are more likely to establish relationships with people with whom we have repeated and sustained interactions, and these relationships are further strengthened by interactions across a variety of contexts. Just like there is no way to re-use representations of existing relationships across multiple sites, there is also no way to establish new relationships based off of interactions across multiple sites. A person might be commenting on the same photos on Flicker and conversing with the same people on Twitter, but they have no way to be aware of their mutual friend or their mutual presence because they are both hiding their identity behind a screen name. Although this anonymity is crucial for users to feel comfortable and is a simple way to protect one's identity, it can also inhibit formation of new and desirable relationships.

It will beneficial to examine additional precedents before detailing the proposal for Constellate. The catalyst for the conversation about the social graph began on the blog Tropophilia[6], which reported on the controversy between blogger Robert Scoble and Facebook. Scoble used an early version of a program called Plaxo Pulse to rapidly scrape information about his Facebook Friends from their profiles for importation to an address book on Plaxo, and Facebook promptly noticed and disabled his account. Jarred wrote that "[Facebook's] competitive advantage is not the information they store, but the services they provide around and through that information. And like Battelle, I believe that Facebook will recognize this and give users an easy way to export the social graph." Facebook has yet to do this over a year later, and has not indicated that they plan to do so.

---

[6] Jarred. "Who Owns the Social Graph?" *Tropophilia*. 5 January 2008. http://tropophilia.com/2008/01/05/who-owns-the-social-graph/

Other attempts have been made to provide centralized and extensible representations of a user's identity. OpenID[7] advertises itself as an "easy way to use a single digital identity across the Internet" and is interesting as an example of an open web standard that has gained at least some success. It is only a tool for management of login credentials, however, and doesn't have a social aspect. Chi.mp[8] is a web service that aims to provide robust identity management and content aggregation tools. It provides users with a free .mp domain name, allows users to associate their account with third-party services, aggregates content from those services, and provides a centralized and organized repository of to a user's profiles elsewhere on the Internet. Although an interesting competitor to Facebook's identity management functionality, Chi.mp does not provide tools for social graph management that are more useful than those offered by Facebook. It does, however, provide a notable export feature so that users can potentially take their data with them to other sites.

I propose Constellate as a web service to solve the shortcomings described above with current representations of and tools related to the social graph. Constellate will provide a single interface for users to manage and organize their contacts online. Constellate users will represent their social graphs on the service but will not need to concern themselves with a representation of their identity or other sharing of content. Instead, they will be identified by a username, their (unique) email address, and an optional full name. Users will then be able to enter, import from other services (such as Gmail), or upload from desktop contact management software (such as Address Book or Outlook) information about their relationships. Users can maintain a representation for

---

[7] OpenID. http://openid.net/

contacts *without the requirement of them signing up for the service*, and as long as these contacts are associated with an email address then they will automatically be linked to a real account if/when that user signs up for Constellate and claims the address. Constellate will also provide users with emailing invitations to their friends to sign up, but this should neither be required nor pushed (note that all users will need to do to sign up is click the link, since Constellate will automatically have their address, and can include a random password in the invitation email). Constellate will also allow users to associate an image with each photo, based on the idea that people can more quickly recognize a face than a written name. Users will always be able to use their uploaded image when referring to another person, regardless of what image that person chooses for him or herself. This consistency in perception of other people will make Constellate more user-friendly.

After a user has added contacts in the service they can then begin to assign their contacts to Circles to organize their relationships. Each Circle has a short name that the user can define. These Circles of people can be created hierarchically – just as a circle can be contained within a larger circle – and can each person can be in as many Circles as the user likes – just as circles can overlap (as in a Venn-diagram). These Circles will be used later when Constellate is paired with services so that the user can specify which Circles can be accessed by a specific service, and Constellate will allow Circles to be used in complex ways for the careful specification of privacy preferences. If Flickr were paired with Constellate, for example, a user could say that only those in the 'photography' Circle could see his or her photos, or could say everyone who was not in the 'family' Circle could see the photos, *or* could say that everyone who was not in the 'family' Circle could see the photos *unless* they were also in the 'photography' Circle.

---

[8] Ch.imp. http://chi.mp/

The Constellate site itself will provide users only with a variety of intuitive tools for creating new relationships, associating data with existing relationships, and organizing relationships into Circles. A possible interface tool could be a clickable grid with names and/or faces for adding people to Circles, multiple colors to represent Circles, and potentially keyboard shortcuts for switching quickly between the Circle to be assigned with each click. This will allow users to quickly and painlessly organize their contacts. It should be stated explicitly that this is not a proposal for a new social network in the traditional sense. Constellate will not serve or even host any content about the users. Users will not have a 'profile' with their 'favorite movies' – it does not aim to replace Facebook. Users will not be able to upload photos to Constellate, and they'll still have to go to a respective site (such as Flickr) to comment on the photos of others. There are many services competing in each of the different content spaces - photos, video, status updates, music, etc – and Constellate does not enter this competition. Instead, Constellate will provide existing sites with an alternative to designing, implementing, and building their own systems that their users will use to manage their social graphs.

Additionally, a user will be able to see what services each contact is revealing to him or her, but Constellate will not attempt to aggregate that data. Constellate will also provide an innovative privacy feature that will allow users to view their own information *from the perspective of some other person*. A user could specify some general person in some combination of Circles, or some specific person, and see what of his or her own content he or she would be able to see if logged in as that other person. Users will then be able to share content with confidence that only the intended recipients will see it.

Developers will be able to integrate Constellate into their own web services using

a robust API, and will be encouraged to do so *instead* of their own social tools and not in addition to them. The API will be designed so that other web services do not even need to store and synchronize their own copy of social graph information, and can instead query the Constellate API whenever a piece of information is needed. If developers find that the delay required by an additional HTTP call is too significant, Constellate will provide an optional installable black-box application that developers can integrate into their own services. The data will be hosted locally and the application will work just like the API, synchronizing itself automatically with the primary Constellate databases. Constellate will also use OAuth so that it is easy for users to safely and securely add Constellate to third-party services. OAuth describes itself as "an open protocol to allow secure API authorization in a simple and standard method"[9].

Constellate also has the potential to solve the aforementioned challenges related to interactions between strangers on the Internet. Constellate could offer a service similar to that offered by the now-defunct location-based social networking service Dodgeball. Dodgeball allowed users to input their social networks and locations and would automatically broker interactions between individuals who were both nearby and had a mutual friend[10]. Similarly, Constellate will know about the actions that users are performing elsewhere on the Internet, and will also know about the social relationships between users. For example, if two users are looking at the same photo on Flickr, Flickr can query the Constellate API and learn that these two users are mutual friends. It can then ask both of them if they would like an introduction to each other in the contact of that photo and that shared friend – note that this requires users to opt-in on a case-by-case

[9] OAuth. http://oauth.net/
[10] Dodgeball. http://www.dodgeball.com/home

basis, and allows users to easily conceal their web activity – and if both users consent Flickr can then automatically provide an introduction. Thus Constellate can strengthen connections between creators and viewers of content, turning the web from an amorphous space occupied by mutually unaware strangers into a more friendly space in which viewing a website implies a real sense of presence.

The social dynamics of encouraging both developers and users to invest their time and data in Constellate will be complex. Because Constellate would be designed specifically not to be a platform for content, it would not have obvious ways to monetize the service. Targeted advertisements have not been demonstrably profitable even when the social network *does* have incredible amounts of user data - Facebook, despite its robust social graph and thorough information about user preferences, still serves irrelevant advertising and is not known to be profitable. Constellate would not be able to charge users for the service at the entry level, and although a paid 'pro version' is imaginable it would likely be unsuccessful - if people were charged to maintain more than a certain number of contacts or leverage that information across more than a certain number of services, they would just remove old contacts/services when they wanted to add new ones, which defeats the purpose of the software and would discourage developer participation. Finally, charging developers seems unfeasible because their universal participation is of primary importance to its success - if there are some services that are either unwilling or unable to pay to integrate with Constellate, user participation will be discouraged.

I propose instead a Wikipedia-like philosophy instead. Constellate will be a service owned by everyone and to be used by everyone, free of charge. Information about

relationships and people will remain private, but the users will still own their own data, and will be free to export it using open formats and delete their accounts whenever they wish. A not-for-profit philosophy will also encourage the essential participation from the developers of other web services – Constellate will not be built to increase the value of these third-party sites by providing a shared platform for services to collaborate on the development of the social graph while still healthily competing on content.

In conclusion, Constellate will provide both the users and developers of web services with a centralized social graph – a unified database that users can populate once and update easily, and a unified database that developers can integrate to enhance their own services. By isolating shared social graph information from content, Constellate will allow further specialization and improvement of existing web services, making the Internet a less complex and happier place.