

Using Ambient Social Reminders To Stay In Touch With Friends

Ross Shannon, Eugene Kenny and Aaron Quigley

1. Abstract

Social interactions among a group of friends will typically have a certain recurring rhythm. Most people interact with their own circle of friends at a range of different rates, and through a range of different modalities (by email, phone, instant messaging, face-to-face meetings and so on). When these naturally recurring interactions are maintained effectively, people feel at ease with the quality and stability of their social network. Conversely, when a person has not interacted with one of their friends for a longer time interval than they usually do, a situation can be identified in that relationship which may require action to resolve. Here we discuss the opportunities we see in using ambient information technology to effectively support a user's social connectedness. We present a social network visualisation which provides a user with occasional recommendations of which of their friends they should contact soon to keep their social network in a healthy state.

2. Introduction

When modelling the social interactions among a group of friends, certain recurring rhythms are identified between participants. Within this group, a single person may have a range of different rhythms with each of their friends, due to the similarity of their schedules, the differing strengths of those friendships, and a range of other social factors (Viegas et al., 2006). When these rhythms are maintained well—that is, the person interacts with each friend at the regularity that they normally do—the health of that friendship will feel natural. If on the other hand the friendship falls out of rhythm, through neglect or unfortunate circumstance, and the two people do not see each other or otherwise interact, this gap will be felt, though perhaps not always understood.

We refer to this as a person's *social rhythm*, and it describes the rate and regularity with which they interact with the various people they know. It is an intuitive, fuzzy metric; if asked how often you interact with a certain friend of yours, you may reply with “about twice a week” or “most days”, not something more specific like “once every 37 hours.” These frequencies will differ among subgroups of a social network: interactions with family members may have a different regularity than with work colleagues, and some friends may have special significance and be seen much more often.

Still others may have a very low level of engagement—only being seen at annual events like birthdays or academic conferences.

A person's ability to effectively regulate their own social rhythm relies on their perception of time running like clockwork, but the human mind's perception of the passage of time is capricious at best (Harrison et al., 2007). Numerous studies have pointed to the fallibility of this ability, due to stress, anxiety, caffeine intake and a range of other factors (Chavez, 2003). Without external prompts, keeping up with friends—especially peripheral friends, who are not part of one's close social circle—can become a matter of chance and circumstance. Because social interactions are inherently vague and intuitive, there is no single point in time at which one is motivated to rekindle a relationship in decline. We believe that explicit cues based on historically observed rhythms will help alleviate this problem, just as they have been shown to support a user's health in other studies (Consolvo et al., 2008). We will discuss these issues in depth in the next section.

Intuitively, you may have experienced a digital or physical artefact that you come across arbitrarily which spurs you into thinking of a friend and then contacting them. For example, seeing a photograph of you and a friend may prompt you to send them a message to talk about a shared experience. Similarly, hearing a friend's name or reviewing past correspondences with them may remind you to contact them (Viegas et al., 2004). It is along these lines that we seek to provide subtle reminders of a friend at the right time, to induce a user into re-establishing contact. We have developed a visualisation for this purpose, which we present in section 4.

3. Attention and awareness in social networks

One aspect of human memory is the remembrance of past experiences, known as “retrospective memory.” A second form of memory, “prospective memory”, works in the opposite direction and can be thought of as remembering to remember something (a task or object) at a certain time in the future (Winograd, 1988). For example, remembering to call a friend after work at 6 o'clock, or remembering to bring a book you have borrowed with you when going to visit a colleague.

Though the workings of prospective memory are not yet fully understood, the cognitive process is thought not to require external artefacts to trigger a memory (Meier et al., 2006), but can certainly be aided by such objects, like shopping lists. Setting an alarm on a phone or other device that is triggered at a certain time of the day is also effective, as it takes the burden of remembering when to do a task off the person's mind.

Facebook, a prominent social networking website¹, offer a feature they call a “news feed”, which is a way to keep track of activity within your network of friends. The news feed presents a reverse-chronological list of events, such as photos being uploaded, public messages being exchanged, or status messages being updated. This gives the user a constant stream of activity and information about the members of their social network, though as the number of friends in your network increases, this list offers a view of a decreasing subset of your friends’ activity on the website.

The first problem this brings up is that at any time, a user watching the activity on their news feed are mostly kept updated on the latest and loudest of their friends. Those friends who post status updates about every detail of their day will be featured much more frequently than those friends who broadcast information about themselves less frequently or not at all, even though it is quite likely that it is these quiet friends that are most likely to fall off a user’s mental radar.

Second, none of the popular online social networks implement any concept of a friendship’s inherent strength. All friends are presented equally, despite some presumably being more important than others to the person at the centre of the network. With many people having identified hundreds of friends on the website (Ellis et al., 2006), many of whom may be peripheral to them in everyday life (Fogg et al., 2008), it is easy for some more important friends to be lost among the throng. Indeed, a user may end up receiving many updates from friends whom they would be happy to hear from much less frequently.

Together, these factors have the effect of selectively emphasising a person’s friends in proportion to their engagement with the social networking website, rather than in proportion to how frequently a user personally interacts with that friend. The phrase “out of sight, out of mind” describes the deleterious effects of this vicious cycle: as a friend is remembered less frequently, they become less likely to be contacted in future.

Previous studies have analysed social rhythms in socio-technical systems, although the focus of these studies was on the general trends of social rhythms apparent on a very large scale. Golder et al. studied interactions between college students on Facebook, and found that students’ social calendars were heavily influenced by their school schedule (Golder et al., 2007). Leskovec et al. analysed all conversations conducted over Microsoft’s Messenger instant messaging service in the

¹ www.facebook.com

month of June 2006, and concluded that users of similar age, language and location were most likely to communicate frequently (Leskovec et al., 2007).

Online social networking sites are used in part to maintain social connections which were originally forged offline (Ellison, 2007). “Dunbar’s number” is a proposed upper bound to the number of people an individual can maintain stable social relationships with. Though we may suppose that this number would be a function of our circumstances and available free time, it is in fact related to the size of the neocortex. Among humans, this bound stands at approximately 150, and is due to the cognitive overheads involved in remembering and being able to meaningfully interact with others (Dunbar, 1992). Although social networking applications have long allowed users to have many more than this number of “friends” identified within the system (Boyd, 2007), it is unlikely that a user would report that they are friends with all of these people in the traditional sense (Boyd, 2006).

Despite the large number of friends identified by the average Facebook user, a person’s social network cannot be described by data from any one source. Though the majority of a user’s friends may indeed be present in an online social networking website, they will also have friends that they interact with purely offline, or mostly by phone or email. These ongoing social interactions are equally valid in characterising a user’s circle of friends.

Because of the range of communication options available to us, reactivating links between people is relatively easy, if we are prompted at the right time. These social networking websites in particular present a low-cost way for people to evolve, maintain and reinforce a wide network of friends and acquaintances. The issue becomes one of identifying which friends are core to the network, and capturing information about the historical regularity of contact with all friends, so that we can accurately deliver helpful recommendations to the user.

Our application provides the user with suggestions of actions they can take to maintain the stability of their social network through a visual interface. This encourages users to contact their friends regularly, but also helps them to identify problems with certain friends early, so that they can take steps to correct a deviation before it becomes more pronounced. Thus, if a user tends their network well, they will have stronger ties with a wider and more diverse set of friends.

4. Visualising social interactions

There have been many visualisations generated of social networks, particularly since the rise of social networking websites and the rich data sets they present. Many visualisations use a familiar

node-link diagram of a graph (Heer et al., 2005). Visualisations in this style will often present the graph from an “ego-centric” perspective, where the user being analysed is shown at the centre of the view, with their friends arrayed around them. In this project, because we are not interested in the network links that exist among friends, we can dispense with this network view, and focus on the strength of the connections between a user and their immediate network of friends, and therefore allow them to answer questions about the health of their network at a glance.

As with the social networking websites, a weakness we have identified with existing social network visualisations is that they frequently treat all edges in the network as being weighted identically. That is, an edge is either present or not present; there is no gradation to the strength of each link, and all links are drawn with equal length. In real life, we know that friendships do not behave like this. The social links between people become weaker over time and grow stronger through positive interactions. Representing this dynamism requires additional sources of data.

4.1. *Data sources*

Ambient systems can leverage the vast amounts of data available from the physical and virtual worlds. We now leave digital traces of most of our social interactions: all of our email is archived on a server somewhere, our instant messages are logged locally and remotely, posts to social network profiles are publicly visible, and so on. Even co-location data can be recorded if the users have a capable mobile device, allowing the identification of events like two people meeting in a bar or at a sports event.

Though all of this data is attractive, for this initial version of our application, we decided to focus on records of mobile phone interactions, which we are able to download as a spreadsheet from our telecommunication provider’s website. These records gave us access to traces of a user’s incoming and outgoing phone calls, as well as SMS text messages for the preceding month. We manually relate each phone number in the records to the corresponding friend from the user’s Facebook network, which allows us to refer to each friend by name in the display.

The software has been built to be agnostic to the nature of the interactions, so adding support for emails in future, for example, is a matter of writing a client to connect to the user’s inbox and find mails that they have sent or received from their friends. These, along with other discrete interactions like instant messaging conversations or comments left on each other’s Facebook pages, can then be entered into the system.

4.2. Visualisation of social rhythms

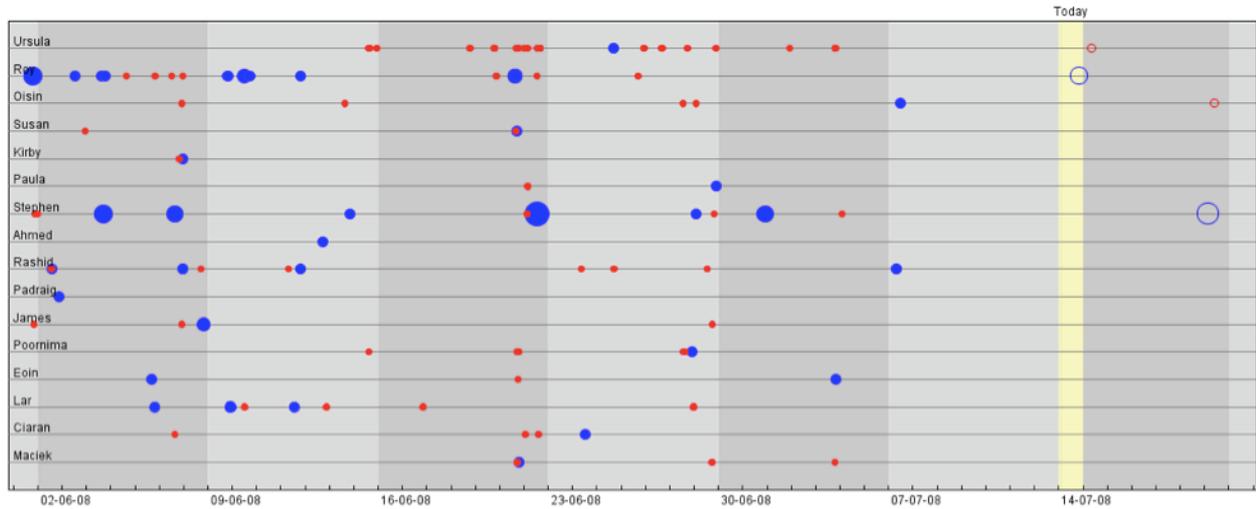


Figure 1: Our visualisation of a single user’s social network, showing a record of their interactions with a subset of their friends over the course of six weeks. Blue circles are phone calls; the size of the circle reflects the length of the call. Red circles are SMS text messages. Suggested future social interactions are indicated by hollow circles on the right, giving the user time to act on those suggestions when it is convenient.

Figure 1 shows a visualisation that we generated from the phone data combined with the user’s Facebook network information. Each row represents the social interactions a person has engaged in with one of their friends via their phone: blue dots indicate phone calls, with the size of the dot reflecting the length of the call, while red dots indicate text messages and are uniformly-sized. Weeks are delineated by differing background colours to provide users with an indication of their longer-term habits at a glance. Our visualisation is built using Processing (Reas et al., 2003), a Java-based visualisation framework.

The current day is highlighted, and the next week is visible on the right of the display. Cues for future interactions are displayed in this area as hollow circles. Their colour and size indicate the type of interaction suggested, based on a prediction algorithm that we have written for this purpose.

Predicted social interactions are drawn on the day that our algorithm has calculated to be most likely for them to occur, but the user can see them a week in advance. This gives the user several opportunities to act on the information being presented to them at an appropriate time. The intent is not to interrupt the user, but simply to plant the seed of memory so that they can act upon it when it is convenient.

If the user does not interact with their friend in any way before the suggested interaction, an “X” is marked at this position and this is counted as a “miss”. The prominent marking of these events (or non-events, if you will) serve to draw the user’s attention to these more critical cues. If the user then contacts the friend in question within a week of this event, the marker is removed and a regular blue or red marker is placed at this point.

5. Properties of ambient information

Neely et al. have previously explored their hypothesis that some context sources are more applicable to being presented in an ambient manner than others (Neely et al., 2007). The reasons they described are precision, criticality, periodicity, interpretability and self-descriptiveness.

In addition, we propose three aspects of the reminders in our visualisation which make them appropriate for delivery by ambient information systems: they are passive, dynamic and simple. *Passive* means that changes in the information do not always require immediate attention; users can take note of reminders but choose not to act on them until later. *Dynamic* means that the data changes over time; if the display remains at the periphery of a person’s attention, they can monitor for changes while concentrating on other activities. *Simple* means that the information can be digested easily; at a basic level, a reminder simply consists of the name of a friend whom they should contact soon. Other information may be present, such as a suggested contact time or medium, but this only serves to augment the primary information.

These three properties correspond well to the interaction, reaction and comprehension model proposed by McCrickard et al. (2003). Not all notification systems are as well-suited to an ambient implementation. Consider as a counterpoint the visualisation an air-traffic controller uses to direct planes at an airport, which satisfies none of the above criteria: the information requires immediate response, as planes must be given clearance to land or take off as quickly as possible; must be monitored constantly; and there are typically a huge number of variables to take into account for each notification, such as the plane’s location, scheduled departure/arrival time, current velocity, etc. It would of course be possible to create an ambient display which delivers information about planes arriving and leaving an airport; while passengers might find this interesting and informative, air-traffic controllers would have no use for it.

6. Ambient applications

The implementation described above is used as both an interactive display, where a user filters the information processed by the system manually to achieve insights into their social trends, and as a passive information display, which allows a user to get a feel for the general health of their social environment in an instant. An implementation which more closely follows the traditional definition of an ambient display could adopt a similar presentation to the Whereabouts Clock developed at Microsoft Research (Sellen et al., 2006). This is a glanceable ambient display placed on a wall in a home, which displays the current location of all members of the family. A similar display which displays a collection of avatars representing some of the user's friends which harnesses the information traces we discussed previously would have an ideal marriage of these properties.

Since the critical information for the user—reminders indicating when a friendship is stagnating—is atomic and relatively simple, it could be used in conjunction with a number of lo-fi data delivery methods. The user could subscribe to receive suggestions as text messages on their mobile phone, or through email or twitter tweets, informing them of the person they need to catch up with.

One can imagine a future scenario where all devices in the home are connected, and digital photo frames could be updated on a frequency predicated by the requirement of the user to update that friendship. Facebook-enabled photo frames have already been released—it is simply the randomisation algorithm that needs to be more intelligent.

7. Conclusions

We have presented a visualisation of a user's interactions with members of their social network, and describe how this kind of information can help a user to keep their social network in a healthy state. Given sufficiently careful treatment, infrequent notifications can become a useful addition to an ambient information display. We have postulated that certain traits are desirable in an ambient reminder system; these are a long possible response time, variance in the timing and meaning of reminders, and simple, easy to interpret reminder information.

Acknowledgements: This work is supported by Science Foundation Ireland through an Undergraduate Research Experience and Knowledge grant (UREKA), and under grant number 03/CE2/I303-1, "LERO: the Irish Software Engineering Research Centre."

8. References

- boyd, d., & Ellison, N. B. (2007). Social network sites: Definition, history, and scholarship. *Journal of Computer-Mediated Communication*, 13(1), article 11.
- boyd, d. (2006). Friends, Friendsters, and MySpace Top 8: Writing community into being on social network sites. *First Monday*, 11(12).
- Chavez, B. R. (2003). *Effects of Stress and Relaxation on Time Perception*. Masters thesis, Uniformed Services University of the Health Sciences, Bethesda, Maryland.
- Consolvo, S., Klasnja, P., McDonald, D. W., Avrahami, D., Froehlich, J., LeGrand, L., Libby, R., Mosher, K., & Landay, J. A. (2008). Flowers or a robot army?: encouraging awareness & activity with personal, mobile displays. In *UbiComp '08: Proceedings of the 10th international conference on Ubiquitous computing*, NY, USA, 2008 (pp. 54–63). ACM.
- Ellison, Nicole, C. S., & Lampe, C. (2006). Spatially Bounded Online Social Networks and Social Capital: The Role of Facebook. In *Annual Conference of the International Communication Association*, Dresden, Germany, 2006.
- Fogg, B., & Iizawa, D. (2008). Online Persuasion in Facebook and Mixi: A Cross-Cultural Comparison. *Persuasive Technology*, 2008 (pp. 35–46).
- Golder, S. A., Wilkinson, D., & Huberman, B. A. (2007). Rhythms of Social Interaction: Messaging within a Massive Online Network. In *3rd International Conference on Communities and Technologies (C&T 2007)*.
- Harrison, C., Amento, B., Kuznetsov, S., & Bell, R. (2007). Rethinking the progress bar. In *UIST '07: Proceedings of the 20th annual ACM symposium on User interface software and technology*, New York, NY, USA, 2007 (pp. 115–118). ACM.

- Heer, J., & boyd, d. (2005). Vizster: Visualizing Online Social Networks. In *IEEE Symposium on Information Visualization (InfoVis 2005)*. Minneapolis, Minnesota, October 23-25.
- McCrickard, D. S., Chewar, C., Somervell, J. P., & Ndiwalana, A. (2003). A model for notification systems evaluation- assessing user goals for multitasking activity. *ACM Transactions on Computer-Human Interaction, 10*(4), 312–338.
- Meier, B., Zimmermann, T., & Perrig, W. (2006). Retrieval experience in prospective memory: Strategic monitoring and spontaneous retrieval. *Memory, 14*(7), 872–889.
- Neely, S., Stevenson, G., & Nixon, P. (2007). Assessing the Suitability of Context Information for Ambient Display. In *Workshop on Designing and Evaluating Ambient Information Systems at Pervasive 2007*.
- Reas, C., & Fry, B. (2003). Processing: a learning environment for creating interactive Web graphics. In *SIGGRAPH '03: 2003 Sketches & Applications*, NY, USA, 2003 (p. 1). ACM.
- Sellen, A., Eardley, R., Izadi, S., & Harper, R. (2006). The whereabouts clock: early testing of a situated awareness device. In *CHI '06: extended abstracts on Human factors in computing systems*, New York, NY, USA, 2006 (pp. 1307–1312). ACM.
- Viegas, F. B., boyd, d., Nguyen, D. H., Potter, J., & Donath, J. (2004). Digital artifacts for remembering and storytelling: posthistory and social network fragments. *Proceedings of the 37th Annual Hawaii International Conference on System Sciences*, IEEE Computer Society.
- Viegas, F. B., Golder, S., & Donath, J. (2006). Visualizing email content: portraying relationships from conversational histories. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, New York, NY, USA, 2006 (pp. 979–988). ACM.
- Winograd, E. (1988). Some observations on prospective remembering. *Practical aspects of memory: Current research and issues, 1*, 348–353.