

Utilizing Online Communities to Facilitate Physical World Interactions

Mike Perkowitz

Department of Computer Science & Engineering
University of Washington
Seattle, WA 98195-2350
map@cs.washington.edu

Matthai Philipose, Joseph F. McCarthy

Intel Research
1100 NE 45th Street
Seattle, WA 98105
{matthai,mccarthy}@intel-research.net

ABSTRACT

Online interactions provide a rich source of content from which reputations and communities can be created. With the proliferation of various networked devices through which people can access the digital world from an increasing variety of physical contexts, there is a new potential to utilize these reputations and communities to create interaction opportunities in the physical world. We are creating *blogger bridges* to enable members of a community of online journalers to safely notify others in the community of their proximity—using a method of *progressive revelation* via networked devices—in order to facilitate serendipitous encounters among community members who are gathered together in a particular location. We describe the kinds of communities that can support such bridges, define a mechanism for enabling progressive revelation, discuss possible technical solutions, and discuss plans for deployment and evaluation in a real community of use.

Keywords

Computer-supported cooperative work, human-computer interaction, ubiquitous computing, situated computing, community software, privacy, trust.

1. INTRODUCTION

The Internet has spawned a variety of mechanisms for people to interact using computer-mediated communication (CMC) technologies. Examples of such mechanisms include electronic mail, text chat, electronic bulletin boards and newsgroups, and electronic commerce. These new ways of interacting “online” have created opportunities for communities to quickly form around shared interests on a global scale [15, 12, 17].

One form of online interaction that is gaining increasing popularity is the weblog (“blog”), where individuals maintain an online journal over time, filling it with personal writings, commentary, or hyperlinks. Communities of bloggers link to and comment on each others’ blog entries. Anyone can start a blog, using either freely downloadable software such as Greymatter or Movable Type or using a web service like LiveJournal or Blog*Spot.

Until recently, most online community interactions—indeed, most online interactions, in general—have taken place via desktop computers. With the proliferation of

mobile telephony and wireless connectivity, the range of physical contexts in which these interactions can occur is vastly increased [16]. As more and more devices of various kinds become connected to the Internet, we will see an abundance of new opportunities for bridging the gap between online communities and the physical world [7].

We are designing *blogger bridges* to establish links between online communities and their members as they navigate the physical world, creating serendipitous interaction opportunities for people who previously may have only had opportunities for interaction in the digital world. We introduce mechanisms for *progressive revelation*, so that people can gradually—and safely—reveal something about themselves and their presence to other members of their online community who happen to be in their physical proximity, as well as verify that the people they meet are who they claim to be.

The rest of this paper will describe the kinds of communities that can support such bridges, define a mechanism for enabling progressive revelation, discuss possible technical solutions, and discuss plans for the deployment and evaluation in a real community of use.

2. SCENARIO

Imagine the following scenario. Sven maintains a blog online. He’s created a profile for himself that lists his interests, including baseball and stamp collecting. He’s also provided links to his “friends” – other bloggers that he knows (offline or online) whose blogs he reads and whose opinions he trusts.

One afternoon, Sven is relaxing at a coffee shop when his mobile phone vibrates to indicate he has a message. His phone indicates that it has detected another blogger in the vicinity that he might wish to meet. Though the blogger’s identity is not revealed, Sven is informed that they have several friends in common and are both interested in stamp collecting.

Since he is not very busy, Sven replies that he is “willing to meet”. The other blogger apparently does the same, and Sven receives another message with the blogger’s name – Arlene – and picture. They bond over a coffee, chat about the vagaries of politics and rare stamps, and keep in touch thenceforth.

3. DESIGN

What is required to make this scenario possible? The above scenario carries a number of implicit assumptions about the interaction between Sven and Arlene. We consider three essential ingredients in a successful blogger bridge: the community, the interaction protocol, and the device.

3.1. The Community

Sven is not willing to risk wasting his time – or worse – chatting with just any random person. He elects to meet Arlene both because they share some common interests and because their common friends implicitly “vouch for” Arlene. Similarly, Arlene may be hesitant to reveal herself to a stranger in public, but is willing to take a chance on a stamp-collecting friend of a friend. We discuss the features of online communities that support such trust and describe various types of communities.

3.1.1. Community features

This essential trust and knowledge of another person is grounded in their shared online community. To provide a basis for this trust, the community should ideally support three things: profiles, reputations, and social networks.

When participating in an online community, a person need not reveal his real-world *identity*; instead, one establishes a *persona* in that community. The persona may be nearly identical to one’s real-world identity, or it may be partly hidden or even invented. This disconnect between the persona and real identity allows people the freedom to reveal unusual interests or espouse controversial views online without fear of repercussions. A *profile* is the presentation of this persona; it may include a list of interests as well as details like age and location. Profiles are, of course, voluntary (and may be invented), but they allow people to advertise their interests and detect compatibilities with others at a glance.

A *reputation* is the accumulated evidence over time as to one’s character and personality. It could consist of a person’s own writings, others’ opinions of him, or any accumulated evidence of participation in the community. Whereas the profile may be a deliberately crafted advertisement, the reputation is less intentional and perhaps more revealing. The existence of a reputation over a long period of time establishes a person’s consistency and staying power, and a long history of participation is difficult to fake. A reputation allows others to verify that a person is who he says he is. Is he an established member of the community? Can he be trusted? Does his behavior match the claims of his profile?

A *social network* is the web of connections between people involved in a community, chiefly representing friendship or trust between people. A person with no friends is suspicious (as might be a clique of people with no outside friends), while a friend of a friend might enjoy the benefit of the doubt.

Finally, in order to provide the basis for serendipitous real-world encounters, a community must have a *critical mass* of members who co-habit a common region (such as a metropolitan area) in the physical world.

3.1.2. Types of online communities

Various types of online communities meet some or all of these requirements. Loosely, the common types include weblogs, message boards, chat rooms, and ecommerce sites.

A person’s *weblog* is a collection of writings and links updated regularly over time, like a public journal. The weblog serves as an effective personal profile over time, and some bloggers have been writing for years and are quite well known both in the blogger community and in the wider media. Links to and from other weblogs establish an implicit reputation and social network. Furthermore, some centrally-hosted weblog systems (e.g., LiveJournal) support explicit profiles and social networks (in the form of “friend” lists). Weblogs are widespread and growing rapidly.

A typical *message board* is concerned with some general topic of interest and consists of many discussions on various related subjects. This includes Usenet newsgroups, fan sites for bands, sports teams, etc., and feedback forums for online magazines and the like. A typical community will have many regular members who are well-known to each other. A person’s history and reputation is apparent in his posts (and others’ responses) over time, though it may be difficult for a newcomer to discern such reputations (cf. NetScan [17] for one approach to solving this problem). Such communities often support user profiles as well, though explicit social networking is less common. Slashdot is a community of this type that also features an intricate reputation system.

Chat rooms are sites where people converse in real time. As with message boards, many chat rooms have regular members who are well known to each other. Some research systems have implemented mechanisms for preserving and visualizing post histories [3, 18]. However, most widely available systems do not provide such mechanisms and provide only minimal user profiles, rendering reputation almost entirely implicit, (although, as with message boards, members are often aware of each others’ reputations).

Two different kinds of communities arise around *ecommerce sites*. First, buyers and sellers form a community focused primarily on doing business with each other. These communities (such as Ebay) use reputation as a way of minimizing the risks of making transactions with strangers [6]. Therefore, reputations are explicitly in the form of a record of transactions and their outcomes; a sort of social network may be inferred from the set of people with whom one has done business. Second, many ecommerce sites inspire a community of people discussing and reviewing the products on offer. Members establish reputations based on their reviews, and some of these

communities (Amazon.com, for example) support member profiles and friend lists [6].

3.2. The Interaction Protocol

The interaction between Sven and Arlene is a multi-step process. Sven must find out enough about Arlene to make his decision before he gives away much of his own information; Arlene, similarly, wants to know more about Sven before she reveals herself. Although this may appear to be an impasse, we propose an interaction protocol that allows Sven and Arlene to gradually reveal themselves while minimizing risks. Each person has a real-world identity and at least one online persona, but they may be hesitant to reveal the link between identity and persona, for the reasons discussed in Section 3.1.1. The task of the interaction protocol is to allow Sven and Arlene to learn enough about each other to decide what to reveal and to verify that they are who they say they are.

Our approach to this problem is called *progressive revelation*. Progressive revelation is a protocol (actually, a class of protocols) for exchanging information with an initially untrusted agent while maintaining control over how much is revealed. A small piece of shared information may establish enough trust to exchange something larger, which may eventually lead to sufficient trust to reveal identities and meet. Naturally, the participants must be able to trust that revealed information is true; cryptographic signatures can be used to verify the link between an unknown person and their claimed persona. Returning to the scenario in Section 2, Sven may be unwilling to reveal his identity to just any stranger. However, he is willing to reveal his interest in stamp collecting to another stamp collector, as when the system informs him that another stamp collector is in his vicinity. He might still be hesitant to reveal his full identity to someone just on the basis of a shared hobby, but once he learns that he and Arlene have some mutual friends, he decides he can trust her enough to drop the mask.

A process of progressive revelation may be assisted by a *trusted third party* who knows the identities of both participants and may therefore be able to tell them what they have in common (e.g., mutual friends) without revealing their identities. This third party need not be an actual person but could be, for example, an internet server holding information about members of the community. When two people encounter each other in the real world, their mobile phones could contact the server for information about each other that will help them decide whether to meet. If the third party is not available at the time of encounters (for example, if the device has no network connection), a similar purpose may be served by pre-filtering. A person might be willing to specify in advance that members of the community who meet certain criteria can be given certain information about him. This pre-approval could take the form of encrypted tokens that can be exchanged during an encounter. For example, Sven might be willing to reveal himself to all friends of friends;

these people would be provided encrypted tokens, and Sven's mobile phone would automatically reveal information to anyone with a token.

A protocol in which each person may choose whether or not to participate is of little use if people feel pressured to go along. For example, when in public, people are free to reject conversation from others, but will often put up with a fair amount of unwanted conversation to avoid appearing rude. This effect is heightened when the conversationalists are acquainted and concerned about acquiring a reputation for rudeness. Our interaction protocol must therefore take into account such social considerations. For example, we should support *plausible ignorability* – the ability to appear to have not noticed someone rather than explicitly rejected them.

3.3. The Device

The interaction between Sven and Arlene is mediated by Sven's (and presumably Arlene's) mobile phone. In what follows, we examine the properties that such a device must satisfy to make the interaction successful. For each property, we discuss how well existing devices satisfy these properties. We consider four broad classes of devices for this purpose:

- laptops with wireless connectivity,
- mobile phones with wireless connectivity,
- personal digital assistants (PDAs) with only local peer-to-peer wireless connectivity
- simple custom devices (such as Motes [8]) with local connectivity

Since we are targeting interactions in the physical world, it is important that the device involved be *portable*. The degree of portability we need depends upon the ubiquity of the interactions we want to support. At one extreme, we may want to support an “any time, any place” model. For this purpose, the device would need to be one that a person can carry in a pocket, is always on, and has a long battery life. Mobile phones, PDAs and tiny custom devices such as Motes clearly have an edge over laptops here. On the other hand, we may wish only to support people who are at particular gathering spots. In this case, it suffices if the device is one they can carry around and turn on only occasionally; a wirelessly connected laptop at a WiFi hotspot may be quite acceptable.

Since the goal of our system is to facilitate serendipitous interactions within a large population, and since our framework cannot introduce two people without an intermediary device, it is essential that any such device (or class of devices) we choose be *broadly available* to many members of this population. The greater the number of people that have the device, the higher is the chance that some two or more of them can meet up. In the case that we select a *set* of devices, we will need to provide a layer of abstraction above these devices so that every device can

connect with the other via to execute higher-level protocols. Common devices such as mobile phones and laptops are appealing from this point of view, whereas custom devices such as Motes and Lovegetys less so.

Because we must be able to detect others nearby, it is important that the devices are *capable of determining proximity* of other devices. There are two common ways in which this can be done. First, an established infrastructure can keep track of the location of each device and notify them if another device is nearby. Existing technologies such as those based on mapping WiFi hotspots [1], GPS [5], and E-911 [2] data for mobile phones can provide location at a resolution of ten to a hundred feet. These systems require users to trust the system with their location. Second, if the devices have the capability of transmitting and receiving signals within a fixed radius, they can use a pre-determined fixed channel to detect the presence of other devices. Technologies such as Bluetooth and ad-hoc radios such as those in Motes support such proximity detection to within a few feet. Although these latter technologies preserve more privacy, they are much less widespread than the former.

Devices must provide user interfaces that are *expressive but discreet*. Given two co-located people with the appropriate devices, each device needs a way to *notify* its owner of the presence of the other. The key requirement of the notification system is that it preserve plausible ignorability i.e. accepting or rejecting the notification should not expose the identity of the notified person. Technologies that are compatible with ignorability include vibration-based notifiers for pocket sized devices such as mobile phones and flashing icons for laptops. Technologies such as flashing LEDs and audible signals are less plausibly ignored.

If the two parties agree to explore the possibility of meeting, they then need to be able to *exchange information* about each other. They may want to exchange text messages to get comfortable. Crucially, however, at some point they may want to reveal deeper information about each other. They will then need to select and send revelatory information about themselves. Given that revelatory information may be as simple as the name of a mutual friend, the minimum requirement is a small text display, which is all a Mote or even many mobile phones can provide. On the other hand, the information to exchange may include segments of blogs, images and entire web pages. The larger screens and keyboards provided by laptop-style devices allow for more expressivity. The increased expressiveness, however, comes at the cost of discreetness: it may be possible to determine that a person is responding to one's messages just by observing their pattern of interaction with the keyboard. For maximal discreetness, a thumb-based input device (such as those on most phones) combined with a earphone may be the best option.

Executing progressive revelation protocols require participating devices to provide *appropriate computational capabilities*. If connectivity to the internet is intermittent, the device must provide sufficient storage for names of potential friends and encrypted tokens as well as a summary of vetted revelatory information about the owner of the device. Devices such as Motes and PDAs that only have local peer-to-peer connectivity but no internet connectivity may prefer this option. An alternative is for the device to provide low-latency, high-throughput connection to the internet so that all necessary computation can be performed offline, and information shipped on to the device as needed. Laptops in WiFi hotspots and net-connected phones and PDAs can use this model. In either case, providing end-to-end data privacy and authentication will likely require both devices to perform some form of encryption. Resource-limited devices such as Motes may not be able to provide the required resources at all, power-limited devices such as mobile phones and PDAs may be limited in the amount of support they provide, whereas a laptop may provide complete support.

In summary, many devices satisfy the requirements of our application to varying extents. No one device excels on all fronts. The choice of device (if one needs to be made) depends on the importance placed on each property.

4. RESEARCH AGENDA

We plan to proceed in three general steps: gathering more information, deployment, and further research.

4.1. Gathering More Information

Before we can deploy and evaluate our proposed system, we need to both verify some of our assumptions and learn more about our target user community. We have identified LiveJournal as a promising community, for several reasons: (1) LiveJournal has user profiles with explicit interests, represented as keywords; (2) it supports social networks through explicit friend lists; and (3) it has reached critical mass, with hundreds of thousands of users. The questions we wish to answer include:

- Are LiveJournal users interested in meeting each other?
- Are shared interests a compelling reason to meet?
- Are shared friends a compelling reason to meet?
- What is the density of LiveJournal users in the Seattle area? At the University of Washington?
- What kinds of devices do LiveJournal users already own?

We intend to address these questions with surveys distributed to as many LiveJournal users as possible, as well as with individual interviews and with analysis of user profiles and friend lists.

4.2. Deployment

We intend to implement and deploy prototypes to enough users to both test the technology and observe the social effects. Accordingly, we need a community of sufficient

density in a particular area to allow for serendipitous encounters. The exact community as well as the nature of the technology used will depend greatly on what we learn while gathering information.

4.2.1. Community

We expect to deploy to some subset of the LiveJournal community. The most likely group is LiveJournal users who are students, faculty, or staff at the University of Washington. LiveJournal use appears to be widespread, and we intend to verify this through surveys and analysis of user profiles.

4.2.2. Device

Our two most likely device options are Bluetooth-equipped mobile phones and laptop computers with WiFi access. Bluetooth mobile phones are becoming increasingly popular; Bluetooth can be used for proximity detection, and the screen and buttons of a phone should be sufficient for our user interface needs. Wireless hotspots are also becoming increasingly common, and we may find that enough UW students use wireless-equipped laptops to make a hotspot-based solution promising.

4.2.3. Evaluation

Our deployment is intended to address several different kinds of questions: (1) do the devices work as intended? (2) does our interaction protocol work as intended? (3) what is the user experience? (4) what are the social effects in the larger community? Answering these questions will require interviews and some qualitative evaluation as well as some quantitative measurements.

4.3. Further Research

We expect our deployment to reveal a number of areas for improvement. Furthermore, once we establish the plausibility of the basic idea, we intend to pursue further research in several directions.

Two possible additional tests of our approach are a citywide deployment and deployment at some kind of large, densely-packed event. Both will present technical challenges and may suggest changes in implementation. A larger-scale deployment will provide more insight into how the system scales, both in its technical and social aspects. Exploring different venues such as a large festival will also suggest new directions, such as the opportunity to tie interaction into ongoing events in the environment.

We would also like to explore various kinds of devices and how to integrate them into a single system. Making connections between, for example, mobile phones and laptops at hotspots presents technical and social questions.

We would also like to analyze a person's profile and online content (such as all of his blog entries). One use would be to infer interests and predict compatibility between people, as well as verify one's stated interests. For example, although Arlene lists stamp collecting among her interests, analysis of her blog may reveal that she never talks about it, possibly making it a poor basis for a friendship. More

generally, such analysis could be used to verify any claim a person makes about himself, by checking whether it is backed up by his actual writings. We may also be able to discover connections between people (from mutual links, for example) even in communities without explicit friend lists.

Finally, we would also like to expand to other communities, weblog-based and otherwise. Even communities of very different types may be able to support our approach, especially if we are able to infer profiles and social networks automatically.

5. CONCLUSION

We have described our motivation and plans for Blogger Bridges: physical devices that can be used to build bridges between people's real-world identities and their online personae in physical contexts in which they may want to mutually reveal aspects of themselves to other, physically co-present, members of their online communities. We introduced the progressive revelation protocol to help ensure that the revealing actions are taken in a gradual and safe way, and highlighted the need to maintain plausible ignorability so as not to embarrass any of the prospective interaction participants.

Our plan is to iterate on the design, implementation, deployment and evaluation of different instantiations of blogger bridges in different communities, seeking to understand how the technology can help enhance users' experience of place (and each other). It is our hope that this use of technology can help people better recognize that we are often surrounded by far more kindred spirits than we are aware.

ACKNOWLEDGMENTS

The authors wish to express their gratitude to David Nguyen and Jenna Burrell for their thoughtful reviews and suggestions for improving this paper.

REFERENCES

1. Byers, Simon, and Dave Kormann. 2003. 802.11b Access Point Mapping. *Communications of the ACM*, 46(5), May 2003, pp. 41-46.
2. Churchill, Elizabeth F., and Sara Bly. 1999. It's All in the Words: Supporting Work Activities with Lightweight Tools. In *Proceedings of the International ACM SIGGROUP Conference on Supporting Group Work (Group '99)*, pp. 40-49.
3. Erickson, Thomas, David N. Smith, Wendy A. Kellogg, Mark Laff, John T. Richards and Erin Bradner. 1999. Socially Translucent Systems: Social Proxies, Persistent Conversation, and the Design of "Babble". In *Proceedings of the 1999 Conference on Human Factors in Computing Systems (CHI '99)*, pp. 72-79.
4. FCC Enhanced 911 Website. <http://www.fcc.gov/911/enhanced>

5. Frauenfelder, Mark. 2000. Revenge of the Know-It-Alls: Inside the Web's Free-Advice Revolution. *Wired* 8.07.
6. Garmin Corporation. About GPS Website. 2001. <http://www.garmin.com/aboutGPS/>.
7. Gershman, Anatole, Joseph F. McCarthy and Andrew E. Fano. Situated Computing: Bridging the Gap between Intention and Action. In *Proceedings of the Third International Symposium on Wearable Computing (ISWC '99)*.
8. Hill, Jason and David Culler. 2002. Mica: A Wireless Platform for Deeply Embedded Networks, *IEEE Micro*, 22(6). Nov/Dec 2002. pp. 12-24.
9. Iwatani, Y. Love: Japanese Style. *Wired News*, June 11, 1998.
10. Milgram, Stanley. 1992. *The Individual in a Social World: Essays and Experiments*. John Sabini and Maury Silver (eds). New York: McGraw Hill.
11. Oldenburg, Ray. 1989. *The Great Good Place: Cafes, Coffee Shops, Bookstores, Bars, Hair Salons, and Other Hangouts at the Heart of a Community*. New York: Paragon House.
12. Preece, Jennifer. 2000. *Online Communities: Designing Usability and Supporting Sociability*. Chichester, UK: John Wiley & Sons.
13. Putnam, Robert. 2000. *Bowling Alone: The Collapse and Revival of American Community*. New York: Simon & Schuster.
14. Resnick, Paul, Richard Zeckhauser, Eric Friedman and Ko Kuwabara. 2000. Reputation Systems. *Communications of the ACM*, 43(12), December 2000, pp. 45-48.
15. Rheingold, Howard. 1993. *The Virtual Community: Homesteading on the Electronic Frontier*. Reading, MA: Addison-Wesley.
16. Rheingold, Howard. 2002. *Smart Mobs: The Next Social Revolution*. New York: Perseus.
17. Smith, Marc A. and Peter Kollock. 2000. *Communities in Cyberspace*.
18. Viegas, Fernanda B., and Judith S. Donath. 1999. Chat Circles. In *Proceedings of the 1999 Conference on Human Factors in Computing Systems (CHI '99)*, pp. 9-16.
19. Wasserman, Stanley, Katherine Faust and Dawn Iacobucci. 1994. *Social Network Analysis: Methods and Applications*. Cambridge University Press.