MIT Media Laboratory Technical Note 580, May 2004
To Appear: IEEE Pervasive, Special Issue on Smart Phones

Mobile Matchmaking: Proximity Sensing and Cueing

Nathan Eagle and Alex (Sandy) Pentland

MIT Media Laboratory
20 Ames St.
Cambridge, MA 02139
{nathan, sandy}@media.mit.du

Abstract. We introduce a system of sensing a social environment and cueing informal interactions between nearby users who don't know each other, but probably should. Proximate people are detected and identified using Bluetooth hardware addresses and matched from a database of user profiles. We show how inferred information from the mobile phone can augment existing profiles and we present a novel architecture for instigating face-to-face interaction designed to meet varying levels of privacy requirements. Finally, additional features are discussed as in response to results from an on-going user study involving one hundred subjects.

1 Introduction

Mobile phones have been adopted faster than any technology in human history and now are available to the majority of people on Earth who earn more than $5 a day. More than one billion mobile phones were sold during 2003, six times as many as the number of personal computers sold that year [9, 14]. Such an infrastructure of handheld communication devices is ripe for novel applications, especially considering their continual increase in processing power. And while digital communications has enabled everything from telecommuting to long-distance relationships across different continents, it has done little to encourage interactions of collocated people. In this paper we describe an architecture that leverages technology designed for communication at a distance to connect people across the room, rather than across the country.

The recent universality of mobile communication devices, combined with the growth of online introduction systems, facilitates an opportunity to generate entirely new types of applications. Although never intended as such, many devices that incorporate low-power wireless connectivity protocols such as Bluetooth that can be used as beacons to identify a user to others nearby. Our application leverages this phenomenon to facilitate dyadic interactions of two physically proximate people through a centralized server. A survey of fifty mobile phone users showed that if it becomes possible to instigate introductions to nearby strangers with similar interests using their phone, 90% of the respondents would use the service regularly. We present such a system, and have named it Serendipity.

2 Bluetooth Proximity Detection

Although hyped for sometime, the RF protocol Bluetooth is finally seeing mass-market adoption in mobile electronics; currently over one million Bluetooth devices are sold each week [10]. Although its primary use is to enable wireless headsets or laptops to connect to phones, as a by-product, Bluetooth devices are becoming aware of other devices carried by people nearby. This "accidental" functionality provides mobile communication devices with the capabilities of online introduction systems, except the introduction is situated in an immediate social context, rather than asynchronously in front of a desktop computer.

BlueAware is a MIDP2 application designed to passively run in the background on many Bluetooth phones currently on the market. The key technological element behind this social scanning application resides in the fact that mobile phones with personal area network capabilities, such as Bluetooth, continuously transmit a unique identification code (BTID) that can be received by other devices. BlueAware records and timestamps the BTIDs encountered in a proximity log, similar to the Jabberwocky project developed by Paulos et al. [6]. If a device is detected that has not been recently recorded in the proximity log, the application automatically sends the discovered BTID over the GPRS network to the Serendipity server.

Privacy Driven Features. BlueAware was designed to automatically begin running in the background when the phone is turned on, alerting the user to its presence with a dialogue box at startup. These types of alerts were incorporated into the system to adequately remind users the application is indeed logging Bluetooth devices. Additionally, the application was designed with a user interface that allows the users to read and delete the specific data being collected, as well as to stop the logging completely.

Refresh Rate vs. Battery-Life. Continually scanning and logging BTIDs can expend an older mobile phone battery in about 18 hours. [1] While continuous scans provide a rich depiction of a user's dynamic environment, most individuals are used to having phones with standby times exceeding 48 hours. Therefore BlueAware was modified to only scan the environment once every five minutes, providing at least 36 hours of standby time.

[1] 2-year old battery of a Nokia 3650 in a sparsely populated Bluetooth environment

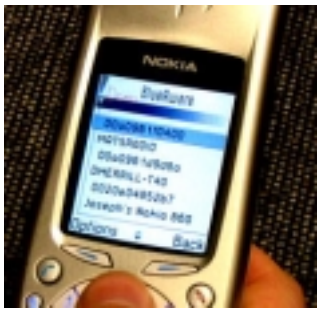

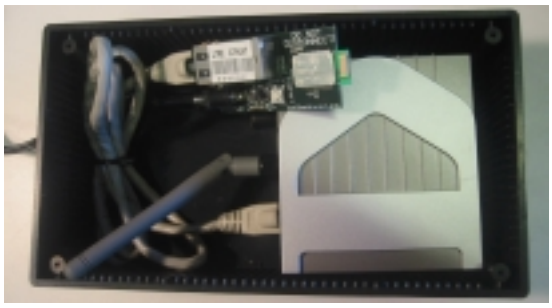

Fig. 1 Methods of detecting Bluetooth devices. BlueAware running in the foreground on a Nokia Series 60 phone (left). BlueDar, a Bluetooth beacon coupled with a WiFi bridge (right).

BlueDar. A variation on BlueAware is BlueDar. BlueDar was developed to be placed in a social setting and continuously scan for visible devices, wirelessly transmitting detected BTIDs to the Serendipity server over an 802.11b network. The heart of the device is a Bluetooth beacon designed by Mat Laibowitz incorporating a class 1 Bluetooth chipset that can be controlled by an XPort web server [4]. We integrated his beacon with an 802.11b wireless bridge and packaged them in an unobtrusive box. An application was written to continuously telnet into multiple BlueDar systems, repeatedly scan for Bluetooth devices, and transmit the discovered proximate BTIDs to our server. Because the Bluetooth chipset is a class 1 device it is able to detect any visible Bluetooth device within a working range of up to twenty-five meters.

3 Serendipity : Situated Introductions

Today's social software is not very social. From standard CRM systems to Friendster.com, these services require users to be in front of a computer in order to make new acquaintances. Serendipity embeds these applications directly into everyday social settings: on the bus, around the water cooler, in a bar, at a conference.

3.1 Previous Work

We are continually aided by desktops, laptops, handheld computers and mobile phones, yet these innovations were primarily designed to empower the individual. However, over the last decade there have been many instantiations of social proximity sensing using pocket-sized devices. Below is by no means a comprehensive review, but rather a sample of the diverse projects in this burgeoning field.

Lovegety. The Lovegety's introduction in Japan in early 1998 was the first commercial attempt to take introduction systems away from the desktop and into reality. Users input his and her responses to a couple questions into the Lovegety; the device then alerts both users when a mutual match has been found. Gaydar, a similar product

specifically targeted for the gay community, was launched soon afterwards in the United States [12].

Cell Tower / SMS Locators. Several wireless service providers now offer location-based services to mobile phone subscribers using celltower IDs. Users of services such as Dodgeball.com can expose their location to other friends by explicitly naming their location using SMS [13].

Experience Ubicomp Project. Using inexpensive RFIDs with traditional conference badges, the Experience Ubicomp Project was able to link profiles describing many of the conference participants with their actual locations. When users would approach a tag reader and display, relevant 'talking points' would appear on the screen [5].

Social Net. Social Net is a project using RF-based devices (the Cybiko) to learn proximity patterns between people. When coupled with explicit information about a social network, the device is able to inform a mutual friend of two proximate people that an introduction may be appropriate [8].

Hummingbird. The Hummingbird is a custom mobile RF device developed to alert people where they were in the same location in order to support collaboration and augment forms of traditional office communication mediums such as instant messaging and email [3].

Jabberwocky. Jabberwocky is a mobile phone application that performs repeated Bluetooth scans to develop a sense of an urban landscape. It was designed not as an introduction system, but rather to promote a sense of urban community [6].

3.2 Mobile Match-Making

Serendipity consists of a central server containing information about individuals in a user's proximity and several methods of matchmaking. These profiles are similar to those stored in other social software programs such as Friendster and Match.com. However, Serendipity users also provide weights that determine each piece of information's importance when calculating a similarity score. The similarity score is calculated by extracting the commonalities between two users' profiles and summed using user-defined weights. If the score is above the threshold set by both users, the server alerts the users that there is someone in their proximity whom might be of interest. The thresholds and the weighting scheme that defines the similarity metric can be set on the phones and correspond to the existing profile types such as meeting, outdoors, silent mode, etc. When it has been determined that the two individuals should have an interaction, an alert is sent to the phones with each user's picture and a list of talking points.

3.3 Implementation

Serendipity receives the BTID and threshold variables from the phones and queries a MySQL database for the user's profile associated with the discovered BTID address. If the profile exists, another script is called to calculate a similarity score between the two users. When this score is above both users' thresholds, the script returns the commonalities as well as additional contact information (at each user's discretion) back to the phones.

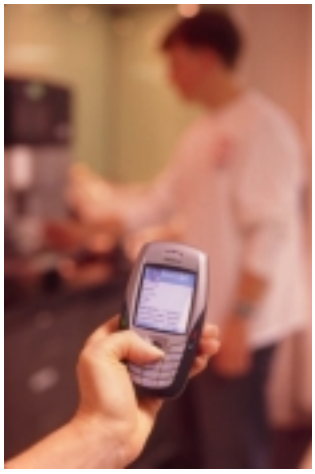

Fig. 2. Serendipity User Screen - Server sends back information in either vcf business card format or a MMS picture message

Feedback. By replying to the introduction message with a number value from one to ten, users can give feedback about the value of the introduction. Although this information is currently only being used as guidance for the system designers, it lays the foundation for a future personalized matchmaking architecture based on reinforcement learning for each individual user.

3.4 Relationship Inference

For select users who have opted in, we are also collecting a variety of additional data using the Context application [15] including celltower IDs and communication logs. Figure 3 depicts the patterns in cell tower transitions and Bluetooth devices over the course of a day. It is clear the user is at home during the hours of 1-9am, and then typically heads into lab. We can see the majority of devices detected during these regular office hours, with the exception of one outlier, a device the user comes into contact with outside the workplace in the evening. We use information about proximity to other mobile devices to do inference about a user's social network.

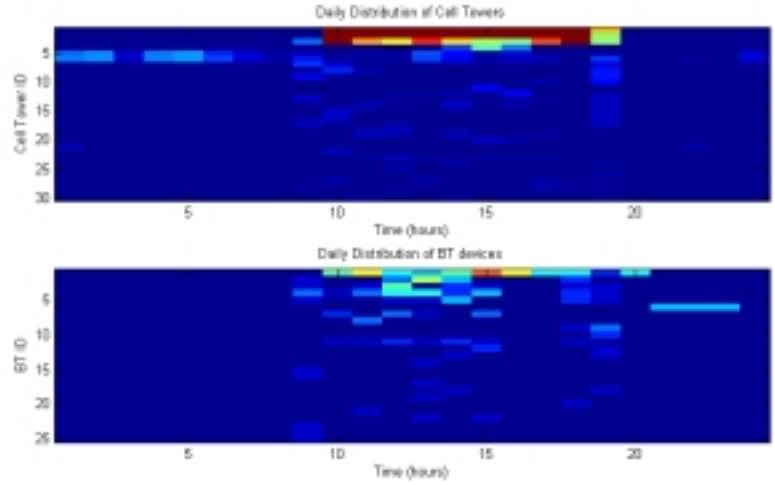

Fig 3. A sample user's daily distribution of observed cell towers transitions and Bluetooth devices. The 'hot spot' in mid-day is when the user is in the office.

By looking at the context of the proximity between two people, much can be gleaned about the nature of the relationship. For example, being nearby someone at 3pm by the coffee machines confers much different meaning than at 11pm at a local bar. We have trained a Gaussian mixture model to detect patterns in proximity between users and then correlate these patterns with the type of relationship.

The labels for this model came from a survey taken by all of the experimental subjects at the end of two months of data collection (some users came late to the study, but were included anyway). The survey asked who they spent time with both in the office and out of the office, and who they would consider to be in their circle of friends. We compared these labels with estimated location (using cell tower distribution and static Bluetooth device distribution), proximity (measured from Bluetooth logs), and time of day.

What we found was that we could identify office acquaintances, outside friends, and people within their circle of friends with good accuracy. The cross-validated Gaussian mixture model was able to correctly predict 90% of the "close friends" labels with a 15% false acceptance rate, calculated over the 1034 potential dyads. Initial examination of the errors indicates that the inclusion of communication logs combined with a more powerful modeling technique, such as Support Vector Machine, will have considerably greater accuracy.

Some of the information that permits inference of friendship is illustrated in Figure 4. This figure shows that our sensing technique is picking up the common-sense phenomenon that office acquaintances are frequently seen in the office, but rarely outside the office. Conversely friends are often seen outside of the office, even if they are co-workers.

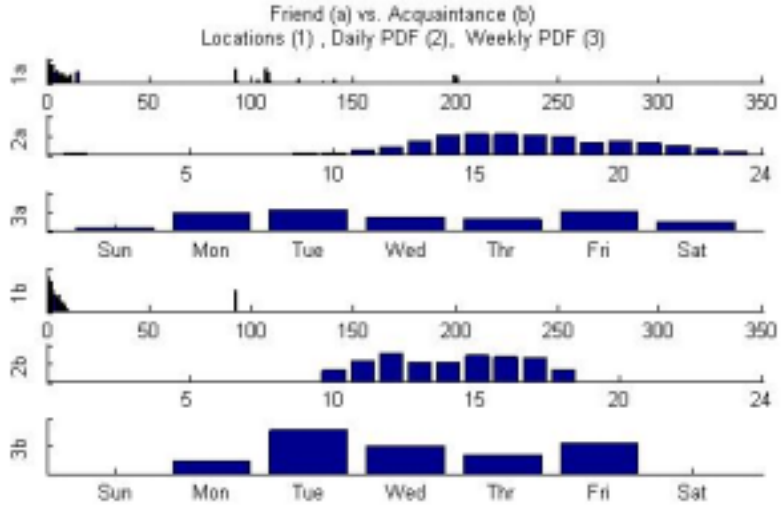

Fig 4. Plotted is data from one subject who listed (a) as a friend and (b) as an office acquaintance. The top three graphs show the frequency of proximity for the friend (a) while the bottom three graphs show the frequency of proximity for the acquaintance (b).

3.5 Privacy Driven Features

After the initial deployment of Serendipity in May, we have now incorporated several features that help better protect a user's privacy. As will be discussed in the privacy section, mobile social software application must be designed to work with the varying privacy concerns of a diverse user community.

Proximity Webpages. The application provides the user the option to view any information a proximate person has deemed public, regardless of their similarity score. While most interactions instigated by Serendipity require information to be sent to both users, proximity webpages allow users to simply see public profiles of nearby people without disclosing information about themselves.

Alternate Introduction Mediation Techniques. Although the current matching algorithm simply looks at similarity thresholds and scores described above, there are many other methods of matchmaking. One such approach described by Terry et al, relies on a mutual friend to make the introduction [8]. Such a method can be incorporated into Serendipity by alerting the mutual friend rather than the two individuals. Alternatively, to preserve a user's privacy and to minimize disruption we also have enabled a feature of sending only an anonymous text message alert that there is a person nearby who shares similar interests; both users must respond "yes" to actuate the dissemination of any personal information.

4 User Studies

The Serendipity system has been tested and iterated upon for almost one year. There are currently one hundred Serendipity users split between two university departments. Although the study has officially only been running for several weeks, Serendipity has undergone a preliminary deployment of forty users during a day-long conference earlier this year. The feedback from the users during this initial trial was incorporated into the current version of the system – now deployed on one hundred phones on the MIT campus.

Fig 5. Executives introduced at a conference with Serendipity.

4.1 Initial Deployment

Serendipity was initially deployed in early May 2004 at an elite conference consisting of senior corporate executives and professors. Personal profiles were created for forty of the conference participants who picked up their assigned phone upon arrival in the morning. Over one hundred introductions were made over the course of the day, primarily during the inter-session coffee breaks. As it was the first time the system was deployed, a significant amount was learned about these types of situated introductions that helped refine the system in subsequent versions.

The conference setting necessitated several modifications from our original design of Serendipity. Because all the subjects were proximate to each other during the talks, it was necessary to develop a method for preventing introductions to be made while the talks were progressing. Simply hard-coding the conference break schedule into phones was not advisable due to the uncertainty in the talk lengths as well as the fact that it would then also prevent introductions between people who both happened to be outside during a particular talk. Instead, we were able to use several personal Bluetooth devices of our research group to prevent these unwanted introductions. We had volunteers disperse themselves throughout the auditorium each carrying a visible Bluetooth device whose name was changed to "BLOCK". Any of the forty phones

inside the auditorium during the talks were able to detect at least one of these "BLOCK" devices. When this name was detected, the Serendipity application was paused and no information was recorded about devices in proximity or sent to our server.

While we succeeded in preventing introductions during the talks when we knew they were not appropriate, we had not taken into account the density of people mingling during the breaks. Several users complained of receiving multiple introductions to people within only a few minutes of each other. This led to a social disruption as one conversation was just getting underway, another conversation was initiated. One user solved the problem by simply turning his phone off while in conversation and then turning it back on when he was ready to meet someone else. In our subsequent version of the software we formalized this feature as "Hidden Mode" as well as imposing a maximum of receiving one introduction every ten minutes.

Some other surprising results included many users who were working for large corporations appreciating being introduced to other coworkers in the same company. For a couple of the participants, the introduction component of the application was not clear; they did not know what the picture messages about people nearby were meant to accomplish. However, besides the comments about the disruption of multiple introductions, the initial user feedback was primarily positive. Most of the initial subjects did not voice any privacy concerns, however this turned out to be not the case for a longer longitudinal study that is scheduled to last for the duration of the 2004-2005 academic year.

4.2 Campus Deployment

Currently Serendipity is running on the phones of one hundred users on an academic campus. Seventy of the users are either students or faculty in the same technical lab, while the remaining thirty are incoming students at the business school adjacent to the laboratory. We are currently receiving information from the devices regarding the other subjects typically observed over the course of the day. The profiles of users from the technical lab are currently bootstrapped from information available within their public project directory. Users also have the opportunity to input personal information and change any aspect of their profile. Although only started recently, the reactions of the initial users have been overwhelmingly positive. The most enthusiastic response has come from the introduction between specific engineers and business school students interested in the commercial potential of their research projects. There has also been positive response when introducing members of the technical lab to each other. On average the lab members are acquainted with only three to five other subjects in the study. Five percent of the subjects have elected not to participate in the matchmaking process due to primarily to time issues (not wanting to be interrupted) as well as privacy concerns.

Fig 6. A small portion of the profiles on stored on mobule.net.

4.3 BlueDar Deployment

BlueDar networked Bluetooth scanning units have been dispersed throughout several social settings on campus including the student lounge, coffee machine, and local bar. Above each device is a piece of paper explaining BlueDar's functionality and the type of data being captured. However, in this university setting we have found that only one of approximately 150 people (excluding participants in the study) have a visible Bluetooth device, far fewer than the number of people actually carrying Bluetooth-enabled gadgets. This implies that potential users will have to decide to make their device visible in order to participate - something that many have been reluctant to do due to power consumption and security concerns. Although we anticipated significant privacy issues being vocalized from the installation of BlueDar in public places, it appears that the explanatory flyer with a description of the data being captured has mitigated much of the unease associated with the device.

5 Privacy Implications

BlueAware, BlueDar, and Serendipity introduce a significant number of privacy concerns if deployed outside of a carefully controlled experiment with human subjects approval. It is clear these privacy implications need to be reviewed in extensive detail before releasing this service to the general public.

BlueAware / BlueDar. While all subjects in our experiment will have given their explicit consent to participate, data is also being collected about devices carried by people who are not directly participating in the experiment. However, we are operating under the assumption that when a device is consciously turned to 'visible'

mode, the user is aware and accepting of the fact that others can detect his or her presence.

Serendipity. The privacy concerns involving Serendipity are numerous. Providing a service that supplies nearby strangers with a user's name and picture is rife with liability and privacy issues. Utmost care must be made to ensure this service never jeopardizes a user's expectation of privacy. As discussed above, several measures have been taken to assuage some of these concerns. Whether it is through proximity webpages, anonymous SMS chat, or simply limiting interactions to users within a friends-of-friends trust-network, it is clear that Serendipity needs to make as many privacy-protecting tools available as possible in order to maintain user diversity, and most importantly, keep everyone safe.

6 Future Applications

Bridging social software introduction systems with current mobile phone technology enables a diverse suite of applications. Conference participants will be able to find the right people during the event; large companies interested in facilitating internal collaboration could use Serendipity to introduce people who are working on similar projects, but not within one another's social circles; single individuals could go to a bar and immediately find people of potential interest.

6.1 Enterprise

Although static employee surveys can be easily analyzed, the output reflects a severely limited view of an organization's social network. We propose that the dynamics of the social network can be inferred from proximity data. Examples of the possible significance of BlueAware include the ability to automatically build a network model of the individuals within an organization, in order to quantify the effects of, for instance, a management intervention. Additionally, incorporating Serendipity into the workplace could instigate synergistic collaborations by connecting people who may be working on similar material, or someone who may have related expertise to another employee's current problem. Finally, forming groups based on their inherent communication behavior rather than a rigid hierarchy may yield significant insights to the field of organizational behavior. We are currently in ongoing discussions with a large technology company to install several BlueDar units within one of their local campuses at the end of the year, integrated with an informal knowledge management system [2].

6.2 Dating

The growth of online dating has soared over recent years as the stigma associated with personal ads diminishes. Serendipity provides users an alternative to encounters with

people that they have only seen on a computer screen. Although we need many more users than our current number in order to test the efficacy of Serendipity as a dating tool, we are dialoguing with several online dating companies about the possibility of integrating a similar system in their own product line involving millions of active participants.

6.3 Conferences

It has been well established that there is a need for introduction systems at events such as large conferences and trade-shows [1]. Salesmen can generate their own proximity webpages similar to the one described above to publicize their products and expertise (rather than interests and photos). Conference participants can customize their profiles to only be connected with individuals who can address their specific area of interest. As we have shown during the initial deployment in May, Serendipity can be an effective tool for networking at conferences.

6.4 Beyond Serendipity

Technology-driven societal change is a hallmark of our era; new infrastructure of intelligent mobile devices are influencing culture in ways that are unplanned and unprecedented. For example SMS text messaging now generates a significant fraction of many service providers' revenue, yet it is a protocol originally developed by cellular network operators as a way for their service technicians to test the network [11]. Similarly, Serendipity's main use may not involve any of the previously mentioned applications but rather something less expected. Perhaps by leveraging trust networks the system could dramatically change the trade-offs of hitchhiking. Additionally, providing notifications of nearby resources (e.g., taxis, restrooms), or coordinating mobile platforms with embedded computers (e.g., cars, buses) could facilitate other ridesharing and car-pooling.

Human-Machine Interactions. By equipping physical infrastructure with embedded computing and a Bluetooth transceiver, a variation on this system can be used to notify human users of nearby resources or facilities. For instance, the system can notify the user of an approaching free taxi, or a nearby public restroom. If instead of human users we consider mobile platforms with embedded computers (e.g., trucks, buses) we can envision other applications. For instance, busses could wait until passengers from other busses had gotten on-board, or delivery vehicles could more efficiently service pickup/drop-off requests.

Role-Based Access Control (RBAC) is a technique used to assign user permissions that correspond to functional roles in an organization [7]. By capturing extensive user behavior patterns over time, our system has the potential to infer not only relationships between users, but also their permissions. For example, if two students working in

different labs begin Tuesday collaborations at a coffee shop, they should be granted constrained entrance access to each other's lab.

Public Release of Serendipity. While Symbian Series 60 phones have become a standard for Nokia's high-end handsets, they represent a small fraction of today's Bluetooth devices. We are in the final stages of developing a MIDP (Java) version of the BlueAware application that will run on a wider range of mobile phones. The final test of Serendipity will be its public launch on www.mobule.net. We hope that not only will the application prove to be robust, but also quite popular within the realms described above, as well as those unanticipated.

7 Conclusions

Our society is more connected than ever before due to two parallel paradigm shifts in computing: movement from desktop to mobile computing, and from individual to social software. Mobile phones have become standard attire across the globe. In millions of pockets and purses are wireless transceivers, microphones, and the computational horsepower of a desktop computer of just a few years ago. Today the majority of this processing power goes unused. However, once the emphasis of mobile applications shift towards supporting the desire of individuals to affiliate with others to achieve their personal goals, this will soon change. We are catching glimpses of introduction services with the advent of online dating and knowledge management, yet the real potential of these new applications will be realized by an infrastructure of socially curious mobile devices, allowing us to untether social software from the desktop and imbue it into everyday life. It is our belief that the mobile phone market is at a critical tipping point where the functionality will shift from the traditional telephone paradigm to a much broader social-centric perspective. We hope that this work represents a step further in that direction.

Acknowledgements

The authors would like to express their gratitude to everyone who participated throughout the development of the application including Pedro Yip, Steve Kannan, Doochan Han, Greg Sterndale, Tony Pryor, Jon Gips, Mat Laibowitz, and Igor Sylvester. Additionally, the authors would like to thank members of both Nokia and Nokia Research, particularly Saku Hieta, Peter Wakim, Hari Pennanen, Suvi Hiltunen and Timo Salomaki.

References

1. Borovoy, R. et al. "Meme Tags and Community Mirrors: Moving from Conferences to Collaboration". ACM CSCW 1998.

2. Eagle, N. "Can Serendipity Be Planned?" MIT Sloan Management Review. Vol. 46 No. 1 pp 10-14. Fall 2004.

3. Holmquist, L.E., Falk, J., and Wigström, J., "Supporting Group Collaboration with Inter-Personal Awareness Devices". Journal of Personal Technologies 3(1-2) 1999.

4. Laibowitz, M., "Parasitic Mobility for Sensate Media", MS Thesis in Media Arts and Sciences. Cambridge, MIT (2004).

5. McCarthy, et al. "Proactive Displays & The Experience UbiComp Project", UbiComp 2003.

6. Paulos, E., Goodman, E., "The familiar stranger: anxiety, comfort, and play in public places". CHI 2004 223-230.

7. Sandhu, R., Ferraiolo, D., Kuhn, R., "The NIST Model for Role Based Access Control: Towards a Unified Standard," Proceedings, 5th ACM Workshop on Role Based Access Control, July 26-27, 2000.

8. Terry, M., Mynatt, E., Ryall, K., and Leigh, D., "Social net: Using patterns of physical proximity over time to infer shared interests". In Proceedings of Human Factors in Computing Systems, CHI 2002.

9. Wood, B. et al. 'Mobile Terminal Market Shares: Worldwide, 4Q03 and 2003', Gartner Group, March 2004.

10. http://www.bluetooth.com/news/releases.asp

11. http://www.ezmsg.com/origins_of_free_sms_text_messaging.htm

12. http://www.wired.com/news/culture/0,1284,12899,00.html

13. http://www.dodgeball.com

14. http://www.mobile-mind.com/htm/mmsexcerpt.htm

15. http://www.cs.helsinki.fi/group/context