Social Serendipity: Proximity Sensing and Cueing

Nathan Eagle and Alex Pentland
MIT Media Laboratory

Abstract. We describe a system that detects wireless devices, queries a centralized server that links the devices with their owners' online profiles, and instigates serendipitous interactions. The system is currently running on Symbian mobile phones that support Bluetooth and is comprised of two applications. The initial application, BlueAware, senses unique Bluetooth identifiers (BTIDs) and logs them to a text file. The Serendipity application queries a server with a discovered BTID, and the server associates the BTID with an individual's online profile. A similarity metric is generated between the two people and depending on both users' settings, the server alerts them of their proximity and common interests. The system enables a suite of additional applications including the facilitation of networking in conference and business environments, the augmentation of current models of airborne pathogen dissemination, and the provision of a novel dating service.

1 Introduction

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. Unfortunately, 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.

Mobile phones have reached a point of mass adoption. Over a half billion mobile phone were sold in 2003 [17]. Such an infrastructure of handheld communication devices is ripe for novel applications, especially considering their continual increase in processing power. 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.

The recent universality of mobile communication devices, combined with the growth of online introduction systems, facilitates an opportunity to generate entirely new types of datasets and 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.

We have developed an architecture that leverages this phenomena to facilitate dyadic interactions of two physically proximate agents through a centralized server. Although the agents do not necessarily need to be human, we have initially developed our system to run on Symbian Bluetooth phones. However, the applications are lightweight and can be ported to many additional devices. The client essentially logs discovered devices, sends the device ID to a profile server and, when merited, receives information that cues an interaction between two agents.

2 Social Proximity Sensing

Human proximity sensing systems are traditionally associated with a machine-human interface incorporating technologies such as IR motion sensors or machine vision. However such sensing systems can only function in a fixed or limited area. In contrast, social proximity sensing has almost always involved wearable devices that can detect other proximate people. Over the last decade there have been many instantiations of social proximity sensing, from badges to keychain electronics. Below is by no means a comprehensive review, but rather a sample of the diverse projects in this burgeoning field.

2.1 Badges

Although primarily used for location-based applications, electronic badges[1] can also sense social proximity. The exposed manner in which they are worn allows line-of-sight sensors, such as infrared (IR), to detect face-to-face interactions. They have been primarily thought of for use in conferences and within the office environment.

Office Environment. Office buildings within research labs are popularly used to test these badge systems and conduct social science experiments due to the cooperative nature of the people in the labs.

The ActiveBadge / ParcTab / Bat. Initially developed over fifteen years ago as a technology to enable telephone systems to route calls to an individual's current location, there have now been many experiments tracking people at the office place using electronic badges. Recent developments in ultrasound tracking have greatly

[1] a broad term that, in this instance, implies a device worn overtly over the clothes in plain view.

improved the ability to localize the badge, enabling a wide range of just-in-time information applications [16, 15, 1].

Sociometer. The sociometer is a wearable computer that can accurately infer a person's interactions with others in face-to-face conversations, allowing (for instance) inference of social influence and status [5]. Despite the significant sensor resources and sophisticated machine learning techniques required, it will only be a matter of a few years before this functionality is realized on standard mobile phones.

Conferences. Conferences are an obvious place to use this technology because badges are already prevalent and accepted. Large corporations spend hundreds of thousands of dollars sending employees to industry conferences, yet there is little to show in terms of return on this investment. Quantifying the projects and people that individuals were exposed to not only makes it easier to write trip reports, but also captures more of a conference's value.

nTag. nTag is a one of the pioneers in the commercial electronic badges market and designed a badge to improve networking of event participants. Profiles of the participants are transmitted from a PC over IR to the badge. When two badges are aligned with one another, text on the badges can provide introductions and display items the participants have in common. For additional functionality, the badges can also be enabled with radio frequency identification (RFID) [20].

IntelliBadge. IntelliBadge uses RFID to capture the location of participants. Because the devices have no visible output, public displays are used to support a variety of applications including traffic monitoring between conference halls, and determining how far a participant has walked during the conference [7].

SpotMe. SpotMe is not a traditional badge, but rather a small Linux-based device that uses short-range RF to communicate with similar devices, in order to provide services such as introductions, information about other conference participants, and searches for specific individuals [21].

2.2 Private Devices

Despite the benefits of wearing an exposed electronic badge that can be easily spotted by other individuals, there are situations where this may not be preferable. Other methods of social proximity sensing have recently been introduced to allow users to keep the device in their pocket, backpack or purse.

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 [22].

Cell Tower Locators. Several wireless service providers now offer location-based services to mobile phone subscribers. Subscribers can agree to expose their location to other friends, and the resolution which these services can pinpoint individuals is currently improving.

Reality Mining. The Realty Mining project consists of sixty 802.11b PDAs which stream audio from close-talking microphones. Software was written for conversation detection, analysis, and modeling. Inference of face-to-face interactions resulted in the successful detection of influence, topic, context and group meeting subtleties [8].

3 Social Software Introduction Systems

We are continually aided by desktops, laptops, handheld computers and mobile phones, yet these innovations were primarily designed to empower the individual. Social software has been defined as software that augments and mediates a user's social and collaborative abilities [6]. This broad definition incorporates everything from email and instant messaging to the "Track Changes" feature in Microsoft Word.

Although the roots of social software predate the personal computer, it has recently received increasing attention. Some of the more popular examples of social software act as a `friends of friends' introduction service while some use matchmaking algorithms developed to find singles with similarities in their profiles. Corporate knowledge management (KM) applications, which attempt to identify experts and quantify the tacit knowledge in an organization, have also begun to incorporate social network information into their services.

Today KM has turned into a 5 billion dollar industry [10], while online dating is the most lucrative form of legal, paid online content. Over 40 million Americans browsed online personal ads during the month of August 2003, and number of user profiles in social software introducer systems exceeds 10 million [9]. However, the majority of these profiles are not typically accessed in social environments, but rather in front of a personal computer. Table 1 shows a sample of the numerous companies that allow users to create their own profiles and publicize their social circle.

Table 1. Types of Social Software that are used as Introduction Systems

Application	Example
Dating: Online Personals	Match, Yahoo Personals, Udate, Spring Street
Social Interests	MeetUp, Friendster, mySpace, Tribe, Orkut
Business Networking	LinkedIn, Ryze
Enterprise	Knowledge Management: Tacit
	Customer Relationship Management

We see that there is convergence of communication devices and introduction systems: our system enables new mobile phones to identify a proximate stranger and retrieve information from his existing online profile. This can enable real-time interventions into social networks and could have significant implications to organizations and even to broader society.

4 BlueAware : Mobile Phone Proximity Logs

BlueAware was designed to passively run in the background on many Bluetooth phones currently on the market. 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 [18]. Bluetooth was primarily designed to enable wireless headsets or laptops to connect to phones, but as a by-product, devices are becoming aware of other Bluetooth 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. Additionally, the automatic logging of social networks creates a unique dataset for both the social science and medical communities.

4.1 Technical Description

The key technology element is 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. The software application we have developed, BlueAware, records and timestamps the BTIDs encountered in a proximity log, and makes them available to other applications. The BlueAware application is automatically run in the background when the phone is turned on, making it essentially invisible to the user.

4.2 Implementation

Two additional Symbian classes were developed for the BlueAware application, one class detects Bluetooth addresses and the other writes information to a proximity log. The program was designed to automatically begin running in the background when the phone is turned on, but there is also an icon provided which allows the user to view the content of the proximity log. The proximity log is split into two files, BTdevices and BTlogtime. BTdevices is an index of every device encountered by the phone. Each individual encounter is timestamped and linked to the indexed number of the device in BTlogtime. Assigning a single number to each device eliminates the need to log the BTID and device name after every timestamp. This method was shown to dramatically decrease the size of the proximity log.[2]

Table 2. Proximity logs generated over one minute

BTlogtime.log	**BTdevices.log**
0 <2004/2/29 12:59:54>	0[name]KIT-KAT [addr]00e09885e388
1 <2004/2/29 12:59:54>	1[name]BARAHONA-IBM [addr]0050f2e17654
2 <2004/2/29 12:59:54>	2[name]Nokia3650 [addr]0060573c1c4e
3 <2004/2/29 12:59:54>	3[name]NORTHOLT [addr]000d888ea558
4 <2004/2/29 12:59:54>	4[name]Rabbit [addr]000a953109af
4 <2004/2/29 13:0:36>	
3 <2004/2/29 13:0:36>	
1 <2004/2/29 13:0:36>	
2 <2004/2/29 13:0:36>	

Refresh Rate vs. Battery-Life. Continually scanning and logging BTIDs can expend an older mobile phone battery[3] in about 18 hours. 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 twice each minute, providing at least 36 hours of standby time.

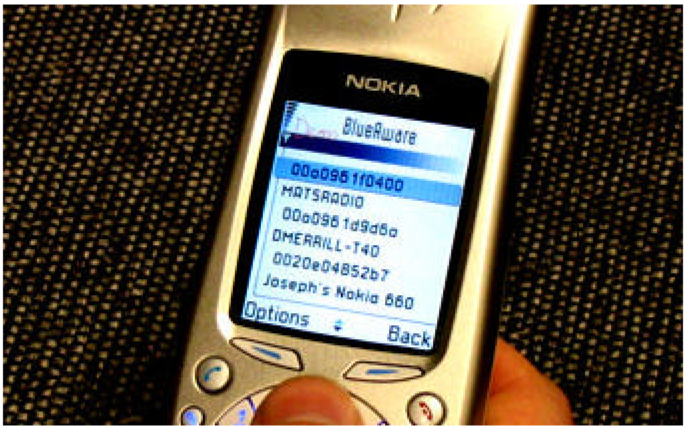

[2] It is important to keep the file size low because the phone stores it in memory during write operations.

[3] 2-year old battery of a Nokia 3650

Fig. 1. BlueAware running in the foreground capturing data on the visible Bluetooth devices in the user's proximity

5 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. Using telephone hardware technology that is already widely deployed, we have developed a system that distinguishes and characterizes individuals, and thus can instigate social interactions.

5.1 Technical Description

Serendipity, a Symbian application developed to enable serendipitous interactions between people, uses the cellular data network to access a central server containing information about individuals in a user's proximity. If a device is detected that has not been recently recorded in the proximity log, Serendipity automatically connects over the GPRS network via a PHP interface to an Apache web server. The phones transmit the discovered BTID with the current threshold for the level of similarity that should generate an alert. The server accesses a database of individual profiles that link a device's BTID to its owner.

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.

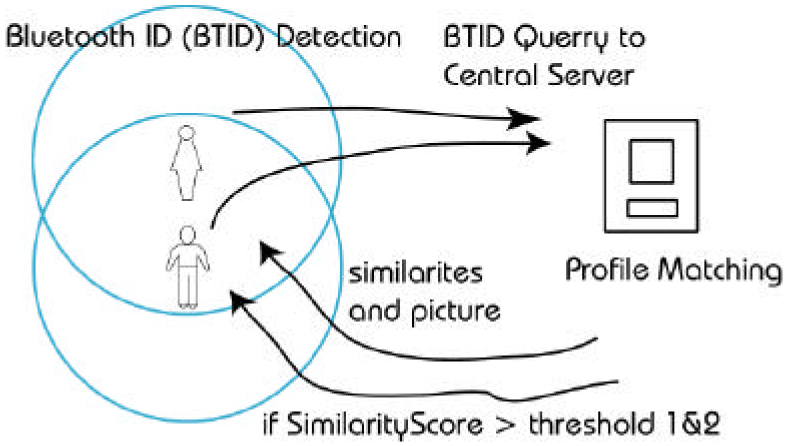

Fig. 2. Serendipity: Bringing social profiles to the social scene

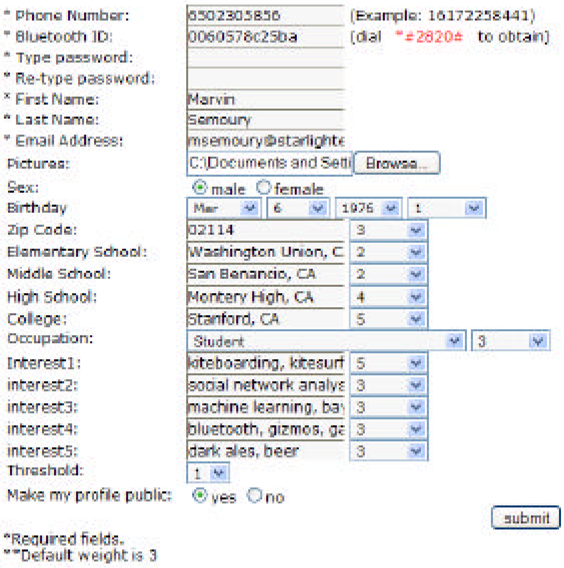

Fig. 3. Profile Creation and Editing

5.2 Implementation

Our initial implementation of Serendipity requires an additional piece of functionality beyond BlueAware. This functionality enables the phones connect to a php document on the server by opening:

http://18.85.44.254/serendipity/php/btid_processed.php?myBTID=BTID&myThres=Thres&rBTID=rBTID

A php script then takes the BTID and threshold variables from the phones and connects to a MySQL database. The current implementation queries a 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.

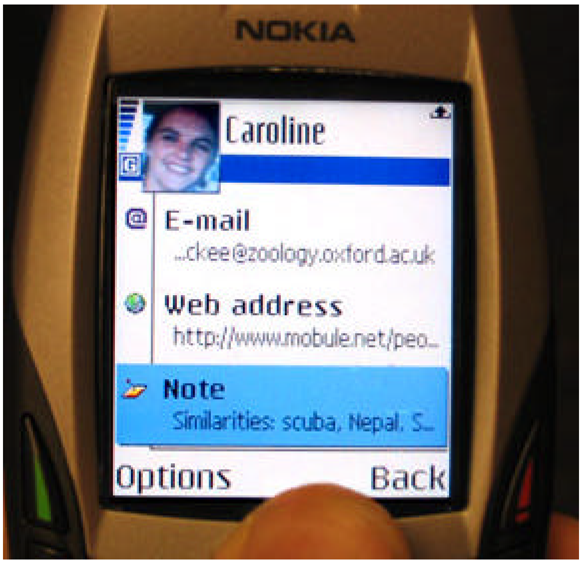

Fig. 4. Serendipity User Screen - Server sends back information in vcf business card format

GPRS, Bluetooth Access Points and SMS. GPRS data networks enable additional functionality for Serendipity, however tend dissipate within large buildings. We have enabled Serendipity to also connect to Bluetooth access points for applications within the enterprise and for testing within our lab. For phones with limited network access we are also developing an SMS interface to the server.

Dissemination. Unlike other social software introduction systems, Serendipity does not require users to login to a website to register for the service. A current user just needs to send the Symbian .sis installation file to a friend's compatible phone over Bluetooth or IR. When the application is opened and installed, the phone automatically connects to the network and the server creates a new profile with the friend's phone number, BTID, and a link to the original user as a 'friend'.

Scan Mode. Users also have the option to create public profiles that can display on all other users' phones regardless of similarity score. An additional variable added to the php script allows users to be in scan mode and receive all public profiles of other people nearby. While most interactions instigated by Serendipity require information to be sent to both users, scan mode allows users to simply see public profiles of nearby people without disclosing information about themselves.

6 End User Applications

Bridging social software introduction systems with current mobile phone technology enables a diverse suite of applications. Assuming the Serendipity application can be ported to the majority of phones today, 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 building 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.

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 strangers that they have only seen on a computer screen. As discussed in section 8, Serendipity allows for a large variation in the privacy protection. Phones can be set in invisible mode so that they will never be contacted, but they can still see other people's 'public profiles'. They can also be set to only establish potential links to others within a trust network. Users become part of a trust network when they are within one degree from an individual's social circle, ie: a friend-of-a-friend. Trust networks can also be expanded to incorporate several degrees of separation. Finally, a user can have the option of creating a public profile. This will allow anyone running Serendipity in scan mode to view this person's information. The public profile can be different from the main profile and acts as a type of proximity webpage.

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. 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.

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 [19]. 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.

Social Relationship Inference. Being able to infer relationships between people based on proximity and interaction behaviors will augment a variety of existing services that currently require users to manually quantify existing relationships. Research is being pursued on developing devices that are not only aware of each other, but also infused with a sense of social curiosity. By continually monitor proximity data, the phones can begin to learn patterns in an individual's behavior. This enables inferences to be made regarding whom the user knows and the type of relationship between everyone in their social network. Applications for these services include:

Automated Relationship Expander. Relationship inference can be used to automatically update the profiles for social software such as LinkedIn or Friendster. Time regularly spent with an individual typically implies some type of relationship with them. And inference as to the relationship type should take time and location into account. Friday night encounters should imply a different relationship than Tuesday afternoons at the office.

Role-Based Access Control. Role-Based Access Control (RBAC) is a technique used to assign user permissions that correspond to functional roles in an organization [14]. 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. The final test of Serendipity will be its public release on www.mobule.net. We hope that not only will the software prove to be robust, but also quite popular within the realms described above, as well as those unanticipated.

7 Future Social Science Experiments

While BlueAware enables a wide suite of applications for the end user, the system can also be of equal importance to the social science community. It is common practice in studies involving social networks and human behavior to use self-report surveys as a primary means of data collection. However, the inherent sparsity and bias of the data have been limiting factors in the research. It is this absence of rich data that also hinders the machine learning and agent-based modeling communities from constructing more comprehensive predictive models.

Although by no means a replacement for survey data, BlueAware has significant advantages as a complementary method of capturing data on social networks and their dynamics. By continually logging and timestamping information about individuals in a user's proximity, the dynamics of large-scale human behavior can be documented. Additionally, only when behavior of a large group is established can interventions such as Serendipity be accurately evaluated. By building a model of the people with whom an individual typically interacts, the effects of an introduction service can be properly verified. If deployed within a large group of people working closely together, correlations between the phone log and proximity log could also be used to provide insight behind the factors driving mobile phone use. Furthermore, a dataset providing the proximity patterns of large groups of people has implications within the computational epidemiology communities, and may help build more accurate models of airborne pathogen dissemination, such as SARS.

7.1 Initial Experiments

One central question of this research is to verify the accuracy of the BlueAware application for quantifying social networks. The main technical issues are:

- the ability of Bluetooth to penetrate through some types of walls, so that people not physically proximate may incorrectly be logged as such
- the spread pattern of Bluetooth, meaning that people behind the subject are logged in exactly the same manner as those in front of the subject
- the range of Bluetooth is sometimes greater than the size of the interacting group (or dyad), so that nearby people may sometimes be logged incorrectly

In laboratory testing we have found that such errors are relatively infrequent and are generally spread evenly across all subjects. These problems therefore produce a small `background noise' against which the true proximity relationships can be reasonably measured. However, social interactions within a laboratory are not necessarily typical of a broader cross-section of society.

If testing in a more general population shows that the level of background noise is unacceptable, there are various technical remedies available. For instance, the temporal pattern of BTID logs allows us to identify various anomalous situations. If someone is not involved in a specific group conversation but just standing nearby, then they will often enter and leave the log at a different time than the members of the group. Similar geometric and temporal constraints can be used to identify other anomalous logs. However to really develop and evaluate these technical remedies we require a large, long-term dataset collected in more natural conditions. To provide such a dataset we are currently preparing extensive longitudinal studies involving several hundred Symbian phones running BlueAware.

First-year Business School Students. We propose to provide all 350 first-year Sloan School of Management MBA students with Nokia 6600 mobile phones running BlueAware. In return for free telephones students will be required to fill out web-based surveys regarding their social activities and the people they are proximate to throughout the day. Comparison of BlueAware logs with survey data will allow us to verify and quantify BlueAware's ability to accurately map social network dynamics. Our method will be to examine correlations between web-based surveys regarding subjects' social networks and the people they are proximate to throughout the day. Through the surveys we can validate whether reported frequency of interaction is correlated to the number of logged BTIDs, as well as perform relationship inference based on the proximity information. We are also interested in quantifying how the social dynamics of the first week of school influences the social network by the end of the semester. Informal surveys of current Sloan students indicate that nearly 100% of the students would choose to be part of this research program.

Workplace / Research Lab. People in research labs are often subjects of initial experiments due to the nature of being at the origin of the innovation. We are initially planning on outfitting our eight-person group with BlueAware-enabled Nokia phones. We will run several experiments over the course of one month and look for correlations in social closeness and proximity, as well as attempt to infer relationship and roles.

Boarding School. Boarding schools present a unique opportunity to capture data on a fairly closed network. This can be especially important when simulating pathogen dissemination across a host population, as will be discussed in the next section.

7.2 Computational Epidemiology

Computational epidemiology is the study of modeling disease propagation. In order to understand and control the spread of pathogens, it is essential to establish some of the key parameters associated with disease transmission. Determining the basic reproductive ratio (R_0) of a disease, for example, is the primary objective for many epidemiological studies. R_0 defines the number of secondary cases produced by an infected individual in an entirely susceptible population. Ideally, health policies can attempt to change the parameters involved in its formulation in order to control a pathogen's spread. Unfortunately, R_0 is notoriously difficult to measure, and must be derived indirectly. Mathematical models have played an important role in assessing R_0 and understanding the dynamics of disease transmission in human populations. For example, the proportion of individuals requiring vaccination for the eradication of a disease may be formulated using R_0.

The majority of epidemiological models are based on a compartmental, SIR framework; the host population is partitioned into those that are susceptible, infected, or immune to a particular pathogen [2]. These deterministic models assume that the rate at which new infections are acquired is proportional to the number of encounters between susceptible and infected individuals, and leads to an effective reproductive ratio that is dependent on a threshold density of susceptibles [11]. Thus, R_0 is dependent not only on parameters intrinsic to the disease such as latent and infectious periods, but also on contacts between infectious and susceptible hosts. *Compartmental models of this kind implicitly assume that the host population is well mixed, such that the probability of infection is equal for all.*

Social network structures are clearly not always well mixed, however, and the complexities of host interactions may have profound implications for the interpretation of epidemiological models and clinical data. Standard mean-field models do not account for heterogeneities of risk between individuals due to the finite number, variability, and clustering of social contacts. Studies have shown that network structure can significantly affect the processes occurring on social networks, including the dynamics and evolution of infectious diseases. Some have investigated the effect of network structure on the evolution of disease traits such as infectious period and transmission rates, as well as invasion thresholds for epidemics, for example [13]. Others have explored the role of spatial contact structure in the evolution of virulence [3, 4, 12].

These models have used hypothetical, extreme network structures as caricatures of real host contact networks. During the recent outbreak of SARS in Singapore, contact tracing of infectious individuals showed a 'superspreader' pattern of disease transmission, although unlike for STDs, the mechanisms behind this are not understood [4]. Furthermore, many directly transmitted diseases are distributed globally, in different types of society, with different characteristic interactions between people.

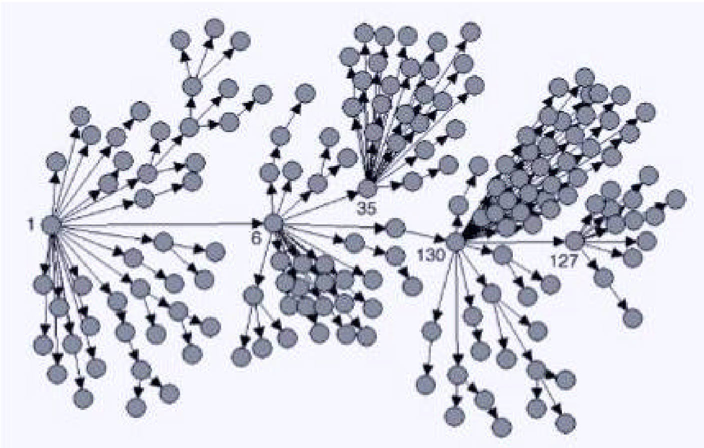

Fig. 5. Probable cases of SARS, from reported sources of infection [4]

The accurate quantification of the host contacts, and therefore the associated variability in the probability of infection, is clearly of great importance. Hypothetical models are valuable for understanding the kind of effect different social network structures would have on disease spread, however we suggest that the data captured by applications such as BlueAware would give a much more realistic interpretation of human social network dynamics. With detailed data on mixing parameters within a social network, epidemiologists will be armed with more information to make predictions about our vulnerability to the next SARS, as well as greater insight into preventing future epidemics.

8 Privacy Implications

BlueAware 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. However, we have made an effort to only capture information relevant to the application. For example, all location information has been purposefully omitted from this system because of the concerns over privacy.

BlueAware. Because BlueAware runs in the background on mobile phones and logs BTIDs, all phones used in our study will have stickers explicitly stating that BlueAware is running and logging other Bluetooth devices in range. While all subjects in our experiment will have given their explicit consent to participate, they will be also collecting data from 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. Additionally, the BTIDs not involved in the experiment will be hashed to provide further user privacy.

Serendipity. The privacy concerns involving Serendipity are numerous. Providing a service that supplies nearby strangers with a user's name and picture is ripe with liability and privacy issues. Utmost care must be made to ensure this service never jeopardizes a user's expectation of privacy. One method of preserving privacy is to provide an additional check box when creating a profile that allows the user to determine whether this information is shared. In addition, it will be possible to limit the potential people who can see a profile to those within the user's friends-of-friends trust-network.

9 Conclusions

As the growth of social software and Bluetooth mobile phones continue in parallel, the two industries will inevitably intersect. Whether used to facilitate networking at a conference, to model airborne pathogen dissemination, or to instigate interactions at a singles bar, infusing mobile devices with a sense of social curiosity has significant potential. By unwittingly carrying devices that can publicize their proximity, Bluetooth users have created an opportunity to capture social interaction data on an unprecedented scale. Additionally, Serendipity untethers social software profiles from computers screens and makes them accessible in situations where they are most useful. We hope to have shown that combining these two broadly deployed technologies will not only enable a windfall of new social network applications for users, but also have meaningful consequences for social and medical science.

Acknowledgements

The authors would like to express their gratitude to Pedro Yip, Steve Kannan and Doochan Han for their work on BlueAware and Serendipity. Both Jon Gips and Caroline Buckee merit additional acknowledgement for their comments and suggestions. Lastly the authors would like to thank Nokia Research, particularly Leo Kretschmer, Saku Hieta and Timo Salomaki.

References

1. Addlesee, M., Curwen, R., Hodges, S., Newman, J., Steggles, P., Ward, A., and Hopper, A. Implementing a Sentient Computing System. IEEE Computer Magazine, Vol. 34, No. 8, August 2001, pp. 50-56.

2. Anderson, R.M., ed. (1982) Population Dynamics of Infectious Diseases, Theory and Applications. London: Chapman & Hall.

3. van Baalen, M. (2002) in The Adaptive Dynamics of Infectious Diseases: In Pursuit of Virulence Management,eds.Dieckman, U.,Metz, J.A.J., Sabelis, M.W. & Sigmund., K. (Cambridge, Cambridge University Press), pp. 85-103.

4. CDC. Update: Severe Acute Respiratory Syndrome --- Singapore, 2003. MMWR 52(18): 405-11

5. Choudhury, T, "Sensing and Modeling Human Networks," Ph. D. Thesis, Program in Media Arts and Sciences, Massachusetts Institute of Technology, Feb. 2004.

6. Coates, T. 'My working definition of social software', May 2003. http://www.plasticbag.org/archives/2003/05/my_working_definition_of_social_software.shtml

7. Cox, D., Kindratenko, V., and Pointer, D. IntelliBadge™: Towards Providing Location-Aware Value-Added Services at Academic Conferences, UbiComp 2003: Ubiquitous Computing, 5th International Conference, Lecture Notes in Computer Science, 2003, vol. 2864, pp. 264-280.

8. Eagle, N and Pentland, A. "Social Network Computing", Ubicomp 2003: Ubiquitous Computing, 5th International Conference, Lecture Notes in Computer Science, 2003, vol 2864, pp. 289-296.

9. Egan, J. 'Love in the Time of No Time', New York Times, Nov 23, 2003. http://query.nytimes.com/gst/abstract.html?res=F20610F63A5C0C708EDDA80994DB404482
http://www1.chinadaily.com.cn/en/doc/2003-11/24/content_284115.htm

10. Gilmour, D. 'How to fix Knowledge Management', Harvard Business Review, 81 (10) Oct 2003: 16-17.

11. Kermack WO, McKendrick AG. (1991) Contributions to the mathematical theory of epidemics--I. Bull Math Biol. 53(1-2):33-55.

12. O'Keefe, K. J., & Antonovics, J. (2002) Am. Nat. 159, 579-605.

13. Read, J.M., & Keeling, M.J. (2003) Disease evolution on networks: the role of contact structure Proc. R. Soc. Lond. B. 270, 699-708.

14. 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.

15. Schilit, B., Adams, N., Gold, R., Tso, M., and Want, R.. "The ParcTab mobile computing system." In Proceedings of the Fourth Workshop on Workstation Operating Systems, pp. 34--39, October 1993

16. Want, R., Hopper, A., Falcao, V., and Gibbons, J., "The active badge location system," ACM Transactions on Information Systems, vol. 10, pp. 91--102, Jan. 1992.

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

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

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

20. http://www.ntag.com

21. http://www.spotme.ch

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