HAN LAYOUT 7/3/08 12:56 PM Page 50

WhozThat? Evolving an Ecosystem for Context-Aware Mobile Social Networks

Aaron Beach, Mike Gartrell, Sirisha Akkala, Jack Elston, John Kelley, Keisuke Nishimoto, Baishakhi Ray, Sergei Razgulin, Karthik Sundaresan, Bonnie Surendar, Michael Terada, and Richard Han, University of Colorado at Boulder

Abstract

One of the most compelling social questions, which until now was left unanswered by current technology, is "Who's that?" This question is usually asked about a new unfamiliar person in a social setting who piques your curiosity. This article presents WhozThat, a system that ties together online social networks with mobile smartphones to answer this common and essential social question. WhozThat builds an infrastructure that shares social networking IDs locally, using wireless technology, while also leveraging a wireless connection to the Internet's social networks to bind identity with location. In addition, WhozThat offers an entire ecosystem on which increasingly complex context-aware applications can be built; that is, once the environment knows who one is, the environment can adapt its content based on the individual's identity or even the collective tastes of a local group. We describe a prototype of the basic WhozThat system and also describe a more advanced service we built, a context-aware music player. We also discuss security and privacy issues introduced by mobile social networks.

The mobile social networking revolution is upon us and could have as profound an effect in enriching local social interaction as the Internet has had in enriching online information access and discourse. The key observation in this article is that the explosive phenomenon of online social networks can be harnessed using mobile devices to answer the compelling question that frequently appears in local social contexts: "Who's that?" It is often the case that people want to find out more about those who are around them; for example, who is that speaking to a group of people in a corner of the room, or who is that who just walked into the room? Standard solutions include asking those around you, looking at name tags, introducing yourself, and so on, none of which leverage the power of technology to help answer these compelling questions and thereby enrich the social interaction.

Online social networks have exploded in popularity [1–3]. As of December 2007, Facebook had over 59 million users [4]. It is estimated that over 85 percent of four-year college students have a Facebook profile, presenting a very usable penetration rate and providing an incredible resource for applications that might leverage this data. These online social networks provide a wealth of personal contextual information, including name, picture, contact information, gender, relationship status/interests, activities/hobbies, musical preferences, literature interests, group membership, and, of course, friendship information concerning user interconnection. Social networks provide a variety of mechanisms for users to share these rich sets of contextual data with other users, including searching for other users with similar interests, as well as a means to establish and maintain communication with other users. Social networks can be seen as a natural evolution of the Internet, where the first big wave facilitated a person's access to information; for example, Web servers and peer-to-peer networks providing news and information content, as well as ways to buy products, whereas this next big wave is focused on facilitating person-to-person communication.

WhozThat is motivated by the idea that bringing this rich contextual information from online social networks into the real world of local human interactions substantially enriches local social interaction. Imagine if you knew more about the people around you in a social gathering, such that you could more easily strike up a conversation with someone with whom you were interested in talking. By being informed via mobile technology of the identity of the person with whom you are seeking to interact and consulting information obtained from that person's public social networking profile, you could more easily initiate a conversation, perhaps introducing yourself and saying, "I noticed we have a shared interest in this hobby or that cause." *The ability of mobile social networking (MoSoNet) technology to substantially lower the barriers to social discourse by minimizing unfamiliarity could revolutionize human social interaction.*

Consider the following vision of what could be achievable by MoSoNet technology. A man walks into a bar and is informed through his cell phone of the identities of the people in the establishment, perhaps discreetly through his Bluetooth earpiece. Knowing his personal relationship information from his social network profile, his cell phone also notifies him that one of the people located elsewhere in the bar is a woman named Sara, a friend of a friend whom he has met in the past. He remembers her saying that she enjoyed listening to folk music and sends her a message asking if he could buy her a drink, to which she responds in the affirmative, texting "Meet me at the corner table." The man spots her and

 0890-8044/08/$25.00 © 2008 IEEE

approaches the table, at which time his favorite song, "Simple Twist of Fate" by Bob Dylan, begins playing over the bar sound system. Not coincidentally, she, too, loves the song, and while the two of them talk, another of their favorite Bob Dylan songs begins to play.

WhozThat achieves this vision of seamless social interaction through MoSoNet technology by implementing a basic two-step protocol that first shares identities between any two nearby cellular smartphones (e.g., via Bluetooth or WiFi) and then consults an online social network with the identity to import the relevant social context into the local context to enrich local human interaction. Indeed, on this basic identity-sharing protocol, more complex context-aware applications can be built that can adapt their behavior to the identities of those around them. The example we implement in this article is a context-aware music player in a bar that could adapt its song playlist based on the tastes of the people in the bar, that is, on the social profiles obtained from the identities advertised by the smartphones of people in that bar. We envision that in the future context-aware video applications will also emerge, as well as a host of context-aware local services for selling merchandise and advertising events that all could be piggybacked on the basic identity-sharing protocol of WhozThat.

Most current approaches toward integrating social networks with mobile devices have missed the opportunity to bind the rich social context tightly with the local context of interacting people. Most approaches simply extend the Web interface of the social network to the mobile device; that is, you can view the social network through your mobile phone [5]. Recent work, such as CenceMe, has sought to enrich the amount of information flowing in the direction from the mobile device to the social network, for example, the location of the user and perhaps context cues such as whether the person is talking [6]. Although these cues provide more information to the online social network user tracking his/her friends, they do not help the local user answer the fundamental question of who is that unfamiliar person standing in the doorway or by the table conversing with other folks? To answer this question, we feel it is critical to exploit mobile computing technology to import contextual information from social networking sites into the local proximity context. Serendipity [7], a system similar to WhozThat, imports social context into the local context using mobile devices. However, it populates its own database of social context rather than connecting with the broader online social networking phenomenon, and it does not consider the evolution of an ecosystem of more complex context-aware applications as does WhozThat.

The WhozThat Protocol for Mobile Social Networking

The basic identity sharing protocol underlying the WhozThat system is shown in Fig. 1. Our assumption is that smartphones will soon become ubiquitous, and will possess both a local wireless capability e.g., Bluetooth or WiFi) and a wide-area wireless connection to the Internet, through either a cellular data plan such as enhanced data rates for GSM evolution (EDGE), a universal mobile telecommunications system (UMTS), evolution-data optimized/only (EV-DO), general packet radio service (GPRS), high-speed downlink packet access (HSDPA), third generation (3G), or WiFi/WiMAX. We expect this to be a reasonable assumption given current market trends and the explosive popularity of smartphones. As shown, both users' smartphones are equipped with the WhozThat identity sharing protocol, in which each mobile phone periodically advertises or broadcasts an identifier or handle such as the social networking ID of the owner of that mobile device. This handle represents a pointer to the social networking profile stored on a remote site. During the *meet-and-greet phase* in Fig. 1a, for example, when Anne and Dave first enter the bar, their smartphones begin wirelessly sharing their social handles with each other. Through this mechanism, each individual in a room or meeting hall learns the handles of other individuals in the same proximity.

Then the users query the online social networking sites using these handles and retrieve the social networking profiles of the associated identities. This second *context creation phase* of the protocol is required to link to stored social context on remote online social networking sites such as MySpace, LinkedIn, and Facebook, and import that information into the local context. As shown in Figs. 1a and 1b, Anne's Internet-enabled smartphone would use the social networking handle provided by Dave's phone to retrieve information about Dave from an online social networking site, and similarly Dave's phone automatically would retrieve information about Anne. The result is that both Anne and Dave now have access to each other's interests, namely, the information that they have chosen to reveal. In general, the result of the context-creation phase is that each individual in a given area now has the social context information to decide whether to initiate a conversation or interaction with another individual or group in the same vicinity.

The WhozThat protocol is advantageous for its simplicity, energy/bandwidth efficiency, agnosticism, and extensibility. The two-phase protocol is powerful yet simple to implement, which should substantially ease its deployment on mobile devices. The protocol is also lightweight for mobile devices by expending only a small amount of energy and bandwidth. Modest time/energy is spent periodically transmitting the social ID and small queries to the online social networking site, whereas the rest is consumed in the relatively low-energy task of receiving the fetched social context data. Our design is agnostic in multiple dimensions. Because we employ application-level handles or IDs, our design is agnostic to the local wireless link; unlike, for example, Serendipity, which is bound to Bluetooth addresses [7]. Also, our design is agnostic to any particular social network and in fact, spans all social networks, that is, Anne's and Dave's personal profiles could reside on different social networks, yet the protocol can handle diversity and share, retrieve, and display content from different sites. Finally, a great advantage of our design is that it is fully extensible, so that an array of more advanced applications and services can be integrated into the basic information-sharing mechanism as described next.

We implemented the basic information-sharing protocol for WhozThat on both software emulators and actual mobile phones. Figure 2a shows two emulated phones programmed in Java Micro Edition (JME) with the WhozThat protocol. The two phones are able to exchange social networking IDs over an emulated Bluetooth link, retrieve the associated social context or profile from an online social networking site, and display that information locally on the phone. Figure 2b shows the same WhozThat protocol implemented on an actual Nokia N80 phone, where again the phone can receive a social ID via Bluetooth, and retrieve and display the corresponding user name from the social networking site using a wide-area wireless Internet connection.

Evolving the Ecosystem for WhozThat Context-Aware Applications

Our observation is that there are major opportunities to evolve the WhozThat ecosystem beyond the basic one described previously. Figure 3 illustrates at least three directions to evolve WhozThat:

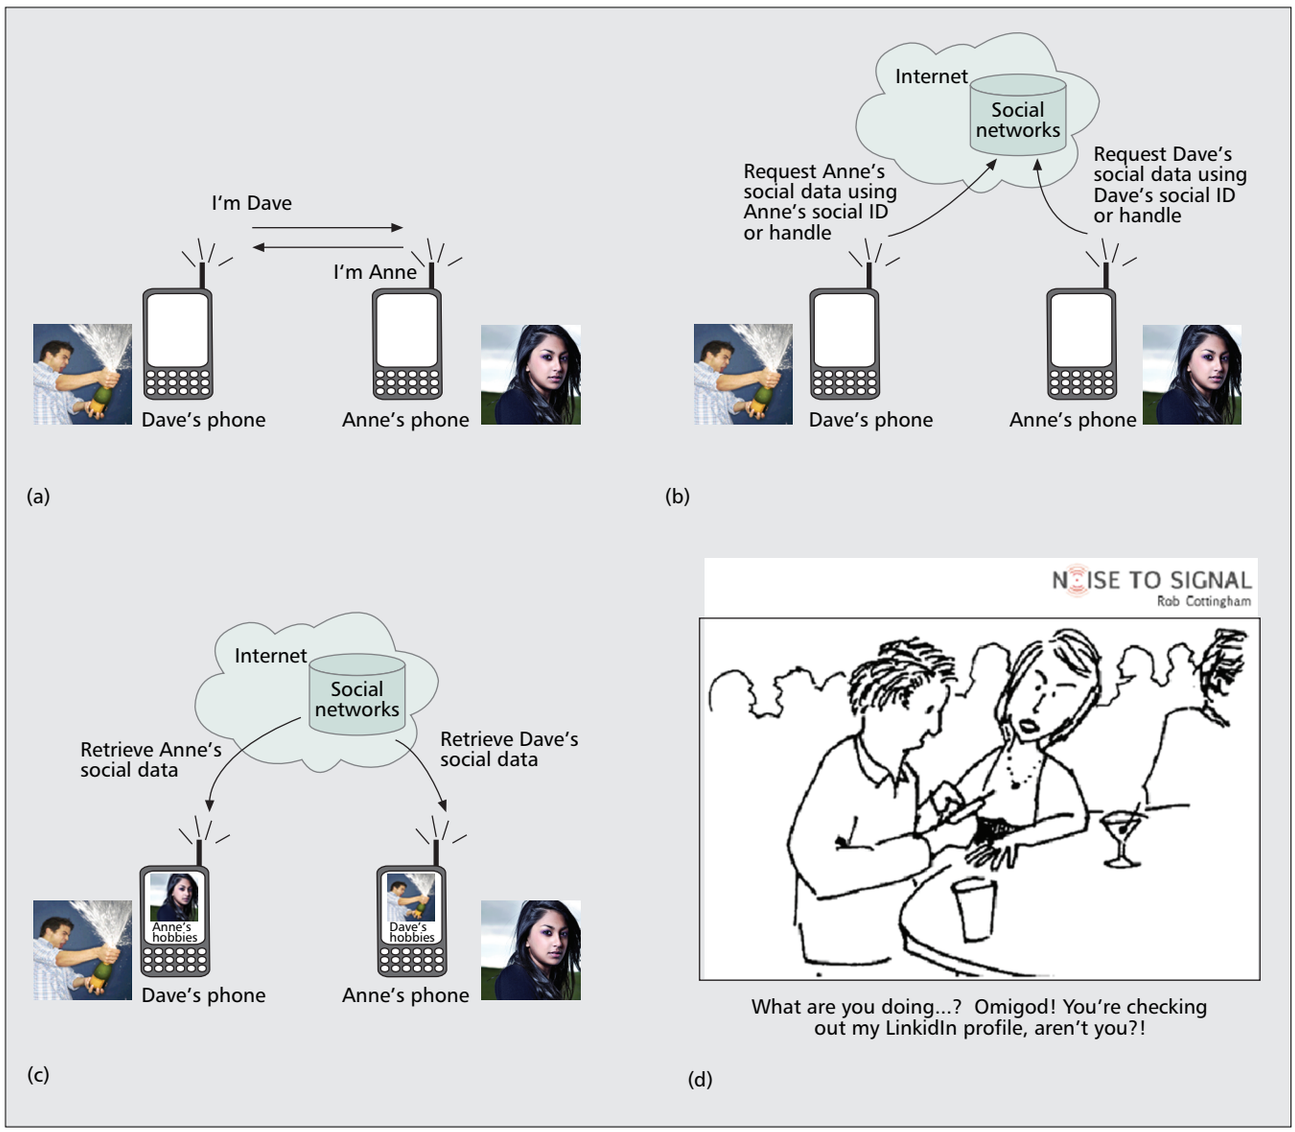

■ Figure 1. *WhozThat's basic identity-sharing protocol for Mobile Social Networks (MoSoNets): two nearby smartphones exchange social networking IDs and then look up the related profiles at online social networking sites, enriching local context by importing social context. a) Anne and Dave are in the same vicinity, and their phones exchange social IDs/handles; b) Anne requests Dave's social data from online social networking sites, and Dave requests Anne's data as well; c) Anne and Dave receive social information about one another, and can then more easily meet/communicate if desired; d) Anne and Dave meet up.*

- Context-aware local services
- Multihop mesh networking
- Gateway services

In addition, though not shown, we describe a further direction that leverages the ad hoc mesh to enable peer-to-peer applications, for example, spontaneous distributed voting/polling and chat. Each of these four directions represents a significant new class of applications that substantially extends the power of the WhozThat paradigm.

Context-Aware Local Services: A Context-Aware Music Player

The exchange of social network IDs in WhozThat allows us to build local context-aware applications and services. These services listen to the announced IDs and *adapt their behavior* to the people that are located near the service. For example, Fig. 3 illustrates a context-aware music jukebox player that may be located in the same bar that Sara, Dave, and Anne currently are patronizing. If we know the social network IDs of WhozThat users around us, we are able to automatically build media playlists that reflect the preferences of those users, in this case a playlist that reflects the joint tastes of Sara, Dave, and Anne. We built and demonstrated this context-aware playlist application in our lab, and the following explains our process of building this application within the WhozThat framework.

The social network ID of a user provides a handle that potentially can be used to access all of the information that a user publishes on his/her social network profile, subject to the profile access privileges specified by the user. For example, a Facebook user like Dave may choose to publish a list of his favorite music on his Facebook profile. The Facebook Web services API enables us to access a wide array of user-specific information available on the user's profile [8]. After we obtain the user's favorite music information, which usually consists of a list of the user's favorite artists, we are able to build a music playlist that contains tracks from each favorite artist.

Several online music databases and services allow us to

(a) Two emulated JME phones exchange social IDs via Bluetooth and retrieve the appropriate social networking profiles.

(b) A Nokia N80 phone receives a social ID over Bluetooth and retrieves the appropriate social networking profile.

■ Figure 2. *Implementing WhozThat in software and hardware: a) two emulated JME phones exchange social IDs via Bluetooth and receive the appropriate social networking profiles; b) a Nokia N80 phone receives a social ID over Bluetooth and retrieves the appropriate social networking profile.*

search for information about an artist, including the list of albums released by that artist and the list of tracks for each album [9–11]. We can use this information to construct a music playlist given a list of artists. For example, the Audioscrobbler Web services API [12] enables us to obtain a list of top tracks for a given artist using a music recommendation system based on collaborative filtering. With the information provided by such services, we can programmatically construct a music playlist. A PC running the WhozThat Bluetooth component could construct this playlist after discovering all nearby WhozThat users and begin playing music from this playlist. This context-aware music player PC is represented by the context-aware music jukebox in Fig. 3.

Figure 4 describes the system that we implemented for our context-aware music playlist generation application, called WZPlaylistGen. We have implemented WZPlaylistGen on a PC with the Java Standard Edition (SE) run-time environment installed. WZPlaylistGen:

- Listens over Bluetooth using the WhozThat protocol for advertised social networking IDs
- Retrieves the user's musical preferences from Facebook
- Uses the Audioscrobbler API to generate a playlist
- Plays that playlist in the local environment

In this way, we can demonstrate the feasibility and power of extending the basic WhozThat protocol to accommodate a whole new class of complex context-aware MoSoNet applications that can adapt their behavior to the people in the immediate vicinity.

Our efforts have raised a host of context-awareness research issues. A straightforward implementation of our music playlist generator simply concatenates the top tracks of each favorite artist found for each discovered WhozThat user. There has been much research into understanding what types of music a user would like based on user history and stated music preferences [13, 14]. These methods of automatically generating playlists based on user behavior and seed songs could be used to enhance the playlist generated by our application. We also could use more sophisticated methods to create and manage changes to the playlist over time. For example, a policy could be developed to control how music is removed from the playlist over time. A user policy could be implemented so that tracks for newly discovered users are placed at the top of the playlist, whereas tracks for users discovered some time ago are placed at the bottom of the playlist. Tracks for previously discovered users that are no longer visible to the music player PC could be deleted after some elapsed time. There are also open questions about how to properly account for all of the musical tastes in a particular local group of users. If the musical preferences conflict for two or more users, is there a way to generate a playlist that is appropriate for these users? For example, if one user prefers music from rap artists, whereas another user prefers music from classical artists, how do we create a playlist that best accommodates these preferences if the user who enjoys rap dislikes classical music?

Integrating sensing presence information provided by a system such as CenceMe [6] could enhance context-aware media playback applications like our music player. For example, if a WhozThat user is currently engaged in a conversation with several other nearby users, we may want to reduce the music playback volume or play a quiet instrumental music track to avoid interfering with the conversation.

Although context-aware media playback applications certainly will benefit from MoSoNet technology, there are other classes of context-aware applications for which MoSoNets would be applicable. Eagle et al. suggest that in a physical environment containing embedded Bluetooth-enabled computers, a MoSoNet system could notify human users of nearby resources or facilities, such as informing a user of a nearby public restroom [7]. With access to a user's social network profile, the WhozThat application could automatically inform users of nearby locations or events of interest. The nearest medical personnel to an accident could be informed that their unique medical services are needed. For example, a doctor or emergency medical technician (EMT) walking in the neighborhood could be informed that there is an accident just around the corner, and assistance is required.

MoSoNets also could enable a class of applications similar to Bat Teleporting as described in [15]. When the user approaches a computer running the WhozThat Bluetooth

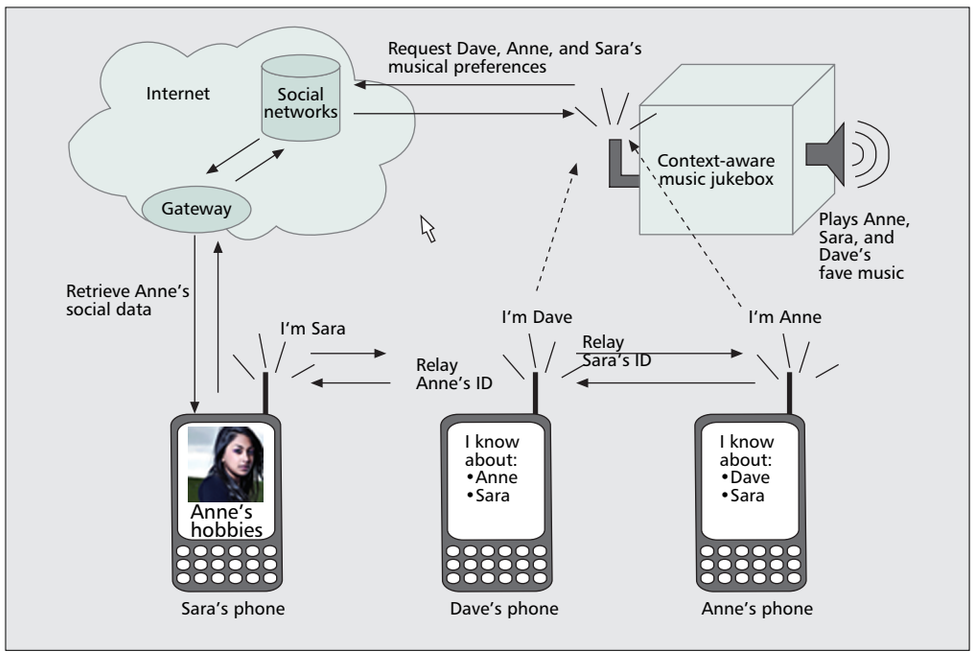

■ Figure 3. *The WhozThat ecosystem can evolve to incorporate the following: 1) local context-aware services, such as. a music player that generates a playlist based on the social IDs advertised by nearby devices; 2) multihop relaying; and 3) gateway services.*

component, the user's social network handle could serve as a unique identifier that allows the user to log in without entering his user name and password. Because the user already was authenticated in the WhozThat client on his or her mobile device, with the appropriate security mechanisms, such as encrypted MAC'ed packets with nonces, we can trust that the social network handle transmitted by his phone has unique and valid credentials. Furthermore, we could customize the desktop computing experience presented to the user after a login based on preferences defined in the user's social network profile.

In general, MoSoNets could be applied to many of the context-aware applications described in [16]. Further research will reveal new classes of context-aware applications enabled by the exchange of social network handles in MoSoNets.

Extending the Ecosystem with Additional Capabilities

There are several other directions that present significant opportunities to expand the ecosystem of WhozThat. Figure 3 illustrates that multihop mesh networking could be integrated into the basic WhozThat protocol to extend its capabilities to enable mobile devices to *relay* or rebroadcast the social IDs of other participants [17]. This capability would allow WhozThat to extend to wireless environments where not all participants are in range of each other, that is, their cell phones cannot directly hear each others' local radio frequency (RF) broadcasts. This same substrate would be capable of relaying not just social IDs, but more generally, application messages of any sort among the participants. Thus, peer-to-peer applications could be supported, such as distributed voting, question-and-answer sessions, data object sharing, and instant messaging. Figure 3 also introduces the concept of a gateway server into the WhozThat ecosystem that could be used to offload complex, compute-intensive, and memory-intensive operations off of mobile devices and into the gateway. Data mining operations that seek to provide a more complete contextual picture, given a social handle, could be offloaded to the gateway, and the final results retrieved by the mobile device.

Privacy and Security

The sharing of social networking IDs or handles raises important security and privacy issues, as does the access to online social networking profiles using these handles. Should the actual social ID be broadcast or should there be some indirection or mechanism for imposing anonymity? Should these handles be authenticated prior to contacting the social networking site to prevent spoofing attacks, and if so, how would they be efficiently authenticated? What kind of infrastructure can MoSoNets employ to mitigate cyber stalking? How do MoSoNets influence personal information disclosure on social networking sites?

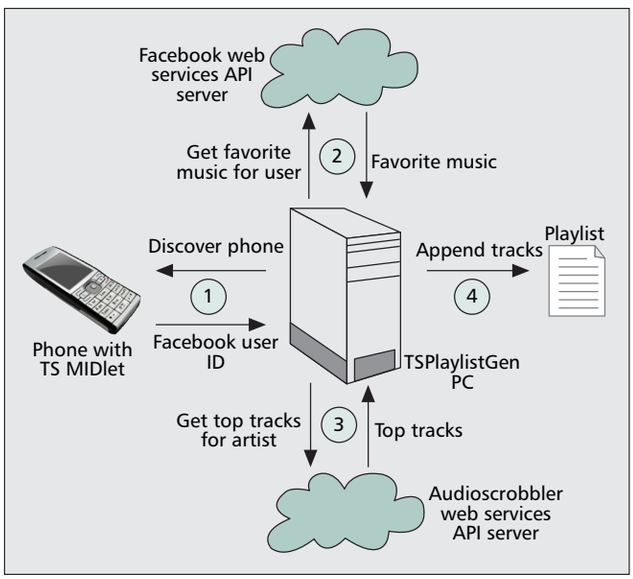

■ Figure 4. *Leveraging WhozThat to support a context-aware music playlist generation system.*

As we embark on this new field of mobile social networks, it is important to be cognizant of many of these questions. For example, privacy is a concern in the modern world, and any technology that opens up a wealth of personal information is bound to be considered risky. Our current work is limited in its consideration of privacy in that it has relied on privacy control mechanisms that already exist in the social network communities. Such privacy mechanisms include allowing only certain communities (say a user's fellow college students on facebook.com) to access your information. Also, Facebook currently enables users to set options for a limited profile that will be seen only by the users they select. However, we have not considered how to limit the information presented to the requestor based on the social context of the requestor-requestee situation. Current approaches base the amount of disclosure on only the friendship relation between the requestor and requestee. However, in certain social contexts, such as at a party whose intent is to foster the meeting of new people, this friendship restriction may be relaxed to permit fuller disclosure of personal information to non-friends. In more unfamiliar environments, it may be desirable for information disclosure to be kept at a minimum. An automated privacy policy for deciding what information to reveal in which social contexts offers intriguing research possibilities. These privacy issues warrant further investigation.

Security is always an issue when information is spread into new contexts. The physical environment makes security issues especially urgent given the intimate nature of technology. It would be very hard to track someone "snooping" information out of the air. The fact that this information would be broadcast continually between nodes means that such security issues would exist at all times and possibly in almost all places. The fact that your cell phone follows you home, and so does that of your neighbor may mean that while a user is having dinner or sleeping, the phones are broadcasting information and taking part in a complex network, sharing this information with the outside world. Opt-out versus opt-in strategies must be considered, and they should be context-dependent. These are just a few examples of how ubiquitous these issues could become.

Conclusion

The age of mobile social networking is dawning. This article shows how the power of online social networks can be harnessed using mobile and wireless technology to help answer the age-old human social question, "Who's that?" We have presented WhozThat, a system that builds a local wireless networking infrastructure using mobile smartphones and cell phones to share each individual's social networking ID, then connects to the Internet via wireless to look up the advertised identities. This simple wireless protocol for electronically sharing IDs could revolutionize the way that humans interact with one another because it lowers the barrier to social interaction caused by unfamiliarity that has limited human discourse for thousands of years. As if this is not remarkable enough, the extraordinary power of this WhozThat paradigm is further demonstrated by its ability to support an entire new class of advanced context-aware applications. Once WhozThat accomplishes the sharing of individual identities, then the environment around an individual or group of users can morph and adapt to individual or collective tastes. We describe a prototype of the basic WhozThat system that we built on mobile phones and the extension of this prototype to support a context-aware music player that adapts its playlist based on the advertised identities of the local individuals. This is just the tip of the iceberg in terms of what is achievable in novel context-aware applications. Imagine being able to be automatically informed, perhaps via a Bluetooth headset, of each person's identity as you walk through a party and mingle. No more awkward moments of forgetting a name. Moreover, chance social meetings where WhozThat finds common interests will be greatly facilitated, connecting people in new ways that were previously not possible or conceivable. One day people may expect that every interaction is preceded by a quick discreet electronic "caption" that introduces one person to another. The revolution in mobile social networks is upon us, and we have sought in this article to convey the excitement that we feel about the prospects offered by the MoSoNet paradigm as a new frontier for research and society. The caveat is that security and privacy must be considered in order to truly realize the full positive potential of MoSoNets.

References

[1] Facebook Developer Resource, http://developers.facebook.com/resources.php
[2] MySpace, http://myspace.com
[3] LinkedIn Professional Network, http://linkedin.com
[4] Facebook Statistics, http://www.facebook.com/press/info.php?statistics
[5] Blackberry Facebook Application, http://www.facebook.com/apps/application.php?id=2254487659
[6] E. Miluzzo *et al.*, "CenceMe: Injecting Sensing Presence into Social Networking Applications," *Proc. 2nd Euro. Conf. Smart Sensing and Context*, 2007.
[7] N. Eagle and A. Pentland, "Social Serendipity: Mobilizing Social Software," *IEEE Pervasive Computing*, vol. 4, no. 2, Apr.–June 2005.
[8] Facebook API — Users.getInfo, http://wiki.developers.facebook.com/index.php/Users.getInfo
[9] MusicBrainz, http://musicbrainz.org/
[10] Last.fm, http://www.last.fm/
[11] freeDB, http://www.freedb.org/
[12] AudioScrobbler, http://www.audioscrobbler.net/data/webservices/
[13] B. Logan, "Content-Based Playlist Generation: Exploratory Experiments," *Proc. 3rd Int'l. Conf. Music Info. Retrieval*, 2002.
[14] E. Pampalk, T. Pohle, and G. Widmer, "Dynamic Playlist Generation Based on Skipping Behavior," *Proc. 6th Int'l. Conf. Music Info. Retrieval*, 2005.
[15] A. Harter *et al.*, "The Anatomy of a Context-Aware Application," *Proc. 5th ACM/IEEE Int'l. Conf. Mobile Computing and Networking*, 1999, pp. 59–68.
[16] B. Schilit, N. Adams, and R. Want, "Context-Aware Computing Applications," *Proc. Wksp. Mobile Computing Sys. and Apps.*, 1994, pp. 85–90.
[17] K. Seada and C. Perkins, "Social Networks: The Killer App for Wireless Ad Hoc Networks?," http://research.nokia.com/files/NRC-TR-2006-010.pdf

Biographies

AARON BEACH (aaron.beach@colorado.edu) is a Ph.D. student in computer science at the University of Colorado (CU) at Boulder. He joined CU-Boulder in fall 2006. He graduated from Northwestern University with a Bachelor's degree in computer science in 2006. His research interests include mobile social networks, wireless security and privacy, wearable computing, and wireless sensor networks.

CHARLES "MIKE" GARTRELL (cgartrel@fusemail.com) will begin his studies for a Ph.D. in computer science at CU Boulder in fall 2008 and is currently an M.S. student in the Computer Science Department at CU Boulder. He graduated from Virginia Tech with a B.S. in computer engineering in 2000. His research interests include wireless sensor networks, cross-layer protocol design, and mobile social networks.

SIRISHA AKKALA, JACK ELSTON, JOHN KELLEY, KEISUKE NISHIMOTO, BAISHAKHI RAY, SERGEI RAZGULIN, KARTHIK SUNDARESAN, BONNIE SURENDAR, and MICHAEL TERADA are all graduate students at UC Boulder in the College of Engineering. They participated in the graduate seminar in mobile computing taught by Professor Han in fall 2007 and contributed valuable material to this article.

RICHARD HAN (rhan@cs.colorado.edu) is an associate professor in the Department of Computer Science at CU Boulder. His research interests span mobile computing systems and applications, wireless sensor networks (WSNs), and wireless security and privacy. He is an NSF CAREER Award winner, an IBM Faculty award winner, and a Best Paper award winner at ACM MobiSys 2006. His research group has published papers on the FireWxNet WSN deployed on wildland forest fires, the open source Mantis sensor OS (http://mantis.cs.colorado.edu), the short preamble X-MAC protocol for WSNs, and secure routing and secure code distribution in WSNs. He is an Associate Editor for *IEEE Transactions on Mobile Computing*, has served on the Technical Program Committees of ACM MobiSys and ACM SenSys, and is general co-chair for ACM MobiSys 2008.