A Mobile Peer-to-Peer Search and Retrieval Service for Social Networks

Isaí Michel Lombera, Louise E. Moser, Peter Michael Melliar-Smith, Yung-Ting Chuang
Department of Electrical and Computer Engineering
University of California, Santa Barbara
Santa Barbara, CA 93106 USA
imichel@ece.ucsb.edu,moser@ece.ucsb.edu,pmms@ece.ucsb.edu,ytchuang@ece.ucsb.edu

***Abstract*—This paper describes iTrust over SMS, a peer-to-peer search and retrieval service for social networks of mobile devices. iTrust over SMS enables mobile devices to collaborate with other mobile devices to distribute, search for, and retrieve information using SMS. With iTrust over SMS, there is no centralized search engine or centralized control and, thus, iTrust over SMS is less vulnerable to filtering and censorship by governments, corporations, or other organizations. In this paper, we describe the iTrust over SMS search and retrieval service, as well as participation in the iTrust over SMS social network. We present the user interface of iTrust over SMS, which enables a user to share information with other users, together with use cases of the iTrust over SMS service. We also show how frequent usage of the iTrust over SMS search and retrieval service encourages further participation in the social network.**

***Keywords*-mobile service; mobile device; social network; peer-to-peer network; mobile search and retrieval**

I. Introduction

Social networks and mobile services are transforming the daily lives of ordinary people. Social networks allow individuals to share information and opinions, and mobile devices enable near universal access to information-sharing services. Existing social networks such as Facebook, Twitter, Myspace, *etc.* require (for the most part) that *they* be the intermediary between the individuals in the social network. Such social networks are not completely trustworthy because they are controlled, managed, and metered by centralized services that store and grant access to information and, thus, are subject to filtering and censorship. Furthermore, information sharing in such social networks is somewhat impersonal in that the user discloses information to the centralized authority, rather than sharing the information directly with other users.

Peer-to-peer social networks are more trustworthy than such centralized social networks; individuals can share information directly among themselves without reliance on any intermediary. An individual might find that his/her opinions are not shared by other members of the social network, and that he/she cannot impose those opinions on the other members. In peer-to-peer social networks, we trust the communities of users, rather any one individual member of the social network.

The iTrust over SMS network is a decentralized peer-to-peer network that provides robust and effective search and retrieval among mobile nodes via SMS. iTrust over SMS aims to avoid the censorship and filtering inherent in centralized search and retrieval services. iTrust over SMS aims to ensure the spread of information, which runs counter to the idea of keeping secrets (*i.e.*, privacy). Nonetheless, we are investigating techniques to mask the message content, as well as the source of a message if it has been relayed by intermediary nodes.

Because a user and his/her mobile device are intimately intertwined, the collections of photographs, music albums, contact lists, e-mail messages, and other information stored on the mobile device can be used to present the user to other users. This *personalization*, made possible by the mobile device and the peer-to-peer network, transforms the experience of information sharing from a routine interaction between the client and the central server into an adaptive information sharing service between peers.

In the remainder of this paper, first, we give a general overview of the iTrust over SMS mobile peer-to-peer (P2P) network and the iTrust over SMS search and retrieval service. Second, we present the Android user interface for iTrust over SMS, which we use to illustrate how a typical user can share information with another peer. Third, we describe use cases for iTrust over SMS, and discuss how iTrust over SMS provides incentives and motivations for sharing information that benefit both the user and the network. Finally, we present related work and, then, conclusions and future work.

II. iTrust over SMS Mobile P2P Network

In the iTrust over SMS mobile peer-to-peer (P2P) network, there is a *global membership* of nodes that have installed the iTrust over SMS library, which are termed the *participating nodes*. Each node maintains a list of such participating nodes, which constitute the node's *local membership*.

In the iTrust over SMS network, some nodes, the *source nodes*, produce information, and make that information available to other participating nodes. The source nodes produce metadata that describes their information, and distribute

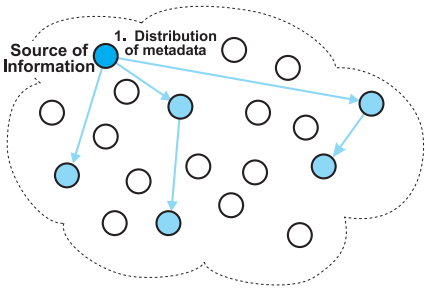

Figure 1. A source node distributes metadata, describing its information, to randomly selected nodes in the network.

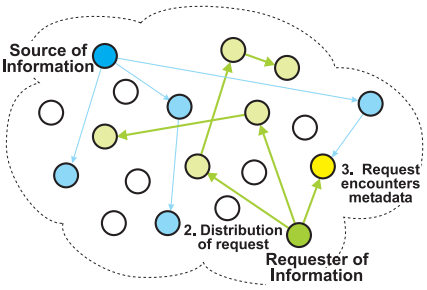

Figure 2. A requesting node distributes its request to randomly selected nodes in the network. One of the nodes has both the metadata and the request and, thus, an encounter occurs.

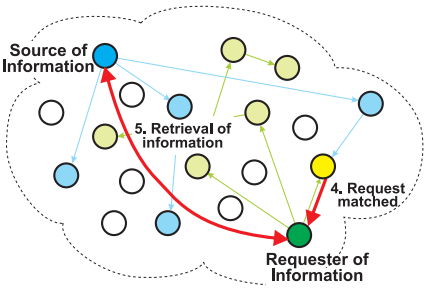

Figure 3. A node matches the metadata and the request and reports the match to the requesting node, which then retrieves the information from the source node.

that metadata to a subset of the participating nodes chosen at random, as shown in Figure 1.

Other nodes, the *requesting nodes (searchers)*, request and retrieve information. Such nodes generate requests (queries) that refer to metadata (keywords) for the desired information, and distribute their requests to a subset of the participating nodes chosen at random, as shown in Figure 2.

The participating nodes compare the metadata in the requests they receive with the metadata they hold. If such a node finds a match (which we call an *encounter*), the matching node returns the node address (mobile phone number) of the node holding the document and the document identifier of the associated information to the requesting node. The requesting node then uses the node address and document identifier to retrieve the information from the source node, as shown in Figure 3.

The metadata include a list of keywords for the information, as well as the node address (mobile phone number) of the source of the information and the document identifier. Metadata generation is dependent on the application, and may be manually provided by the end user, or automatically generated by appropriate packages such as Apache Tika, Apache Lucene, *etc*. Metadata matching may be an exact match or a partial match, and may involve synonyms using dictionaries such as WordNet.

Each node to which a search request is distributed may relay the query to yet another node. Network flooding is avoided by a combination of techniques. The relaying probability is chosen so that the metadata and the requests are distributed to about $2\sqrt{n}$ nodes, in a global membership of n nodes. Unique relaying ensures that a node never relays metadata or a query that it has previously received.

A requesting node builds up its local membership by adding the node to which its search request was relayed and from which it receives a response, and also by adding the source node given in the response. Similarly, a node to which a search request is relayed adds to its local membership the node that relayed the request and also the requesting node. Likewise, a source node adds to its membership a node that retrieves a document from it.

III. iTrust over SMS Android User Interface

To make it easier to understand the common use cases of a typical iTrust over SMS user, first we describe the Android user interface for iTrust over SMS. For this purpose, we present example screen shots for the Android application. Implementers of other Instant Messaging applications may choose to provide similar functionality, in addition to their already existing features, by studying these examples as they fully illustrate the features of iTrust over SMS.

The iTrust over SMS user interface for Android comprises five distinct Java classes, each of which consists of both a layout file written in XML and an activity file containing event handlers written in Java. The layout file specifies the location and style (color, font, *etc.*) of widgets placed on the mobile device screen, as well as attribute identifiers for Android. The Android identifiers can be used for various purposes, such as binding Java resources to program subroutines during run-time, or even simple string value replacement (*e.g.*, internationalization). An event handler is triggered when a user interacts with a widget in the layout (*e.g.*, a button tap triggers an event handler for the *onClick* method). From the user's perspective, an activity is simply the layout of widgets on the screen that allows interaction with iTrust over SMS; for this reason, we use the terms activity and screen interchangeably in the rest of this section.

Below, we briefly describe and illustrate each of the screens that a user can use to distribute, search for, and retrieve information in the iTrust over SMS network.

A. Existing Search

Figure 4 shows the default screen when a user starts the iTrust over SMS Android application. This screen lists the searches that the user has made from the mobile device. If the number of searches exceeds the space on the screen, the screen automatically allows vertical scrolling to accommodate the display of more searches.

The screen in Figure 4 lists all searches that were explicitly *initiated* on the mobile device; searches from other nodes that passed through this node by way of query relaying are *not* shown here. This design choice fits the expectation of

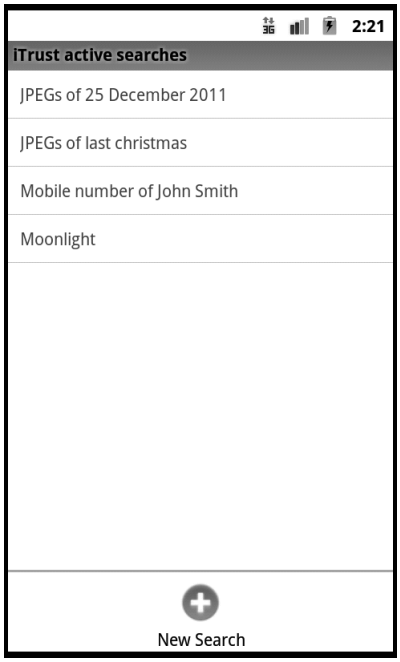

Figure 4. Screen that lists all searches sent from the local iTrust node.

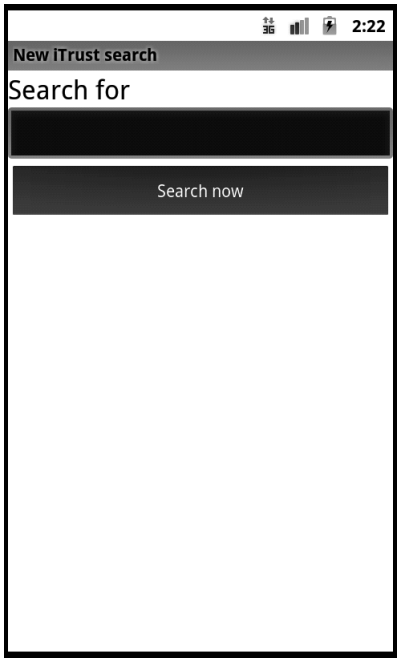

Figure 5. Screen to initiate a new search query from the local iTrust node.

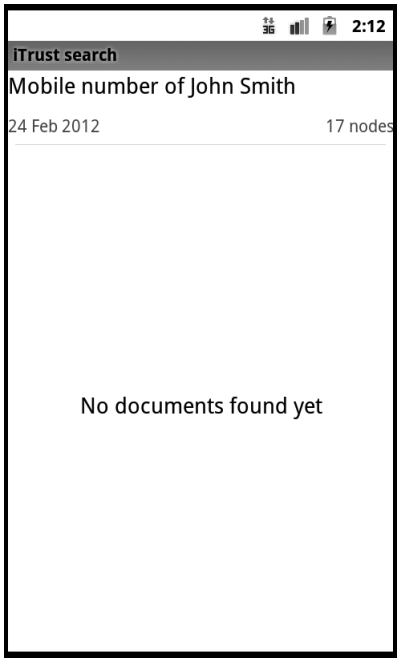

Figure 6. Screen showing the detailed information for a particular iTrust search.

the typical user, who is concerned with the searches that he/she started and not the searches that other users started.

Tapping on any search list item immediately displays detailed information about that search, as shown in Figure 6. For example, tapping on the search item *Mobile number of John Smith* immediately switches to detailed information about that particular search query.

At the bottom of Figure 4 is a pop-up menu that is enabled by pressing the *menu* hard/soft button present on the Android mobile device. By default, this menu and the *New Search* menu item are not visible on the screen. However, a user can press the *menu* button, which causes the menu to pop-up (pressing the *menu* button again causes the menu to disappear). When the *New Search* button is pressed, the user is taken to Figure 5.

B. New Search

Figure 5 shows the screen used to initiate a search across the iTrust over SMS network. A user is brought to this screen by tapping the *New Search* menu item found on the screen in Figure 4. The design is purposely simple and feature limited. Just as a typical Web user prefers a single text box to enter a query, the typical mobile phone user prefers a simple interface to enter search queries.

Tapping the *Search now* button takes the user back to the list of active searches sent from the device, as shown in Figure 4. Meanwhile, the search is automatically serviced by the iTrust over SMS library and relayed by the iTrust over SMS search service.

C. Search and Retrieval Details

When the user taps on a search item on the screen in Figure 4, that particular search request (query) is displayed in detail on the screen in Figure 6.

Near the top of the screen in Figure 6 is the text of the search request followed by two important fields: the date/timestamp when the search was initiated and the number of nodes to which the search request was relayed. The date/timestamp enables the user to recall how old the search is; the date is *not* used for priority ranking. The number of nodes displayed is the number of nodes to which the request was directly sent by this node (although, because SMS is used, only a best-effort service is provided). This number is the *minimum* number of nodes to which the search request is distributed; each such node may relay the query to yet another node.

Below the data/timestamp and the number of nodes, and separated by a thin line, is the space reserved for showing a list of matches reported back to the node. When another node has an encounter or match and reports back to the node originating the search, this space displays a tappable list of items along with the node address of the node where the match occurred. Simply tapping on the list item triggers an automatic fetch of the document by the iTrust over SMS retrieval service; the resource is then displayed on the screen or optionally saved for later viewing. When no information is found, the screen displays the *No documents found yet* notice to the user.

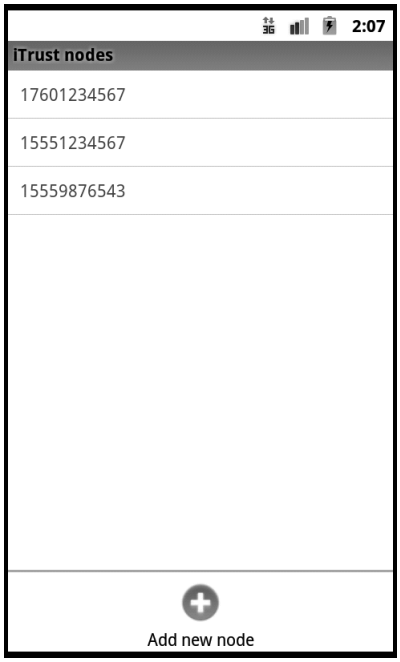

Figure 7. Screen to add new nodes to the local iTrust membership.

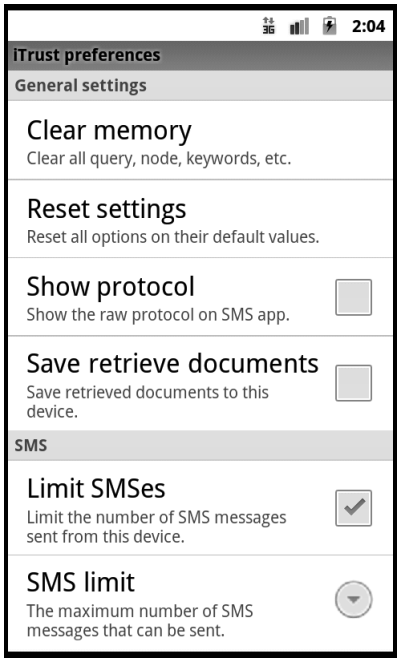

Figure 8. Screen that configures local iTrust preferences (top half).

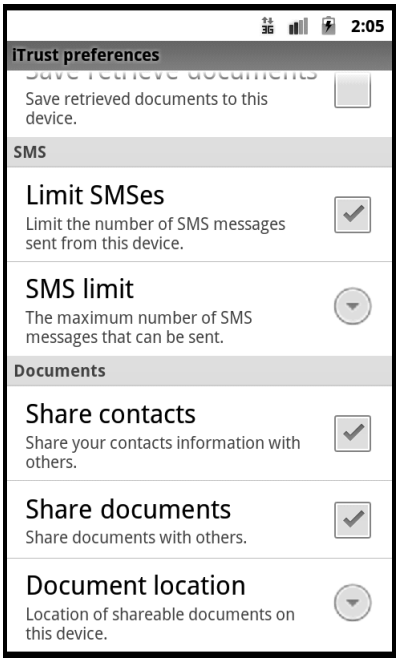

Figure 9. Screen that configures local iTrust preferences (bottom half).

D. Nodes

Figure 7 is similar to Figure 4 except that, instead of displaying the list of searches, it displays the list of nodes (or membership) of the local node. The pop-up menu near the bottom of the screen allows the user to enter a new node address in a pop-up dialog text box (not shown here). Explicit addition of node addresses by the user is not a common occurrence, but addition of node addresses is often performed automatically by the iTrust over SMS library; therefore, a new screen is not required for this task as it was for adding new searches in Figure 5. A simple entry dialog box suffices.

As searches are relayed through the membership of the iTrust over SMS network, the originating query node address is saved by the iTrust library on each node that receives the relayed query. Thus, node addresses may be added to a node's membership without the node's making direct contact with those nodes.

E. Preferences (top)

The preferences screen shows the configurable settings that a user may modify to change the behavior of the iTrust over SMS service running on his/her mobile device. Because the preferences activity is longer than some of the other activities, it must be vertically scrolled on the mobile device, as shown in Figure 8 (top of the activity) and Figure 9 (bottom of the activity). The first two preferences categories, General Settings and SMS Settings, are shown in Figure 8, and are discussed below.

1) General Settings: This preference category enables a user to configure various options that directly affect the iTrust over SMS search and retrieval service on the mobile device or local node.

The *Clear memory* preference deletes all information on the local node including node addresses, saved documents, saved searches, metadata and any other information generated by iTrust over SMS and stored on the local device. Note that this action is applicable only to documents stored on the local node; if another node already fetched a document from the local node, the fetched copy is not deleted.

Tapping the *Reset settings* preference restores all preferences to their default state to what they were when the application was first installed. No searches, fetched documents, node lists, *etc.* are deleted or altered in any way. The *Reset settings* preference is a subset of the *Clear memory* preference that does *not* alter any shareable or iTrust over SMS information.

The *Show protocol* check box toggles the ability to show the underlying iTrust over SMS messaging protocol inside the Android Instant Messaging application. By default, the protocol is not shown to the Instant Messaging application, but the user may enable this option (*e.g.*, for debugging).

Also, by default, the *Save retrieve documents* check box disables the option of saving, to local storage, each document retrieved by the iTrust over SMS retrieval service. Toggling the preference saves each retrieved document into a predefined location; the user may then review the document offline or if the source node is no longer available.

2) SMS Settings: This preference category restricts the iTrust over SMS service on the SMS telephony service. Because many mobile service providers charge a fee for each SMS message sent or received by a mobile device, this category allows a user to control his/her data usage fees more effectively.

The *Limit SMSes* check box allows a user to enable or disable the SMS messages transmitted from the mobile device (there is no realistic user control for restricting incoming SMS messages that does not rely on the mobile service provider to some degree). If this preference is enabled, the preference *SMS limit* can be tapped and a dialog box pops up requesting the maximum number of SMS messages that may be sent by the iTrust over SMS service on this node. If the *Limit SMSes* preference is disabled, the *SMS limit* preference is likewise disabled and the maximum number of SMS messages sent is ignored.

F. Preferences (bottom)

Figure 9 displays the *Documents* category in the preferences activity. This category deals with the metadata distribution service of iTrust over SMS; configuration of metadata on the local device is managed in this category.

1) Document Settings: The *Share contacts* check box disables the creation of metadata (to be shared with other nodes) from information stored in the local device's contact list. For example, if a user in the iTrust over SMS network had *John Smith* in his/her contact list and if this preference is enabled, then another user searching for information about John Smith (as in Figure 6) would have an encounter or match. By default, this preference is enabled.

Likewise, the *Share documents* check box allows a user to share metadata and documents with any node that sends a query to the local node. This preference has dual functionality. Enabling the option shares both metadata about a document (during distribution) and the document itself (during retrieval). Similarly, disabling this option disables both the sharing of metadata and the related document.

The *Document location* preference, when tapped, pops up a dialog box asking for the location where the shareable documents are kept in local storage. Disabling the *Share documents* preference also disables the *Document location* preference.

IV. Use Cases

The use cases for iTrust over SMS extend those for iTrust over HTTP [17], but are adapted for typical mobile phone users. Although mobile phones are increasing in computational power and the ability to display more information on the screen, they are far smaller than laptops or desktops and, as such, necessitate a smaller simpler interface. For example, even though mobile users can perform Google searches on their mobile phones, they rarely venture past the first match, whereas desktop users commonly view second, third, or more matches. In the use cases below, we explain the use case context, and analyze how the iTrust over SMS service responds or adapts to requests.

A. Sporadic Searcher

We define a sporadic searcher to be a user who only occasionally uses the iTrust over SMS search and retrieval service; searches are relatively infrequent and retrieved documents are typically small. Such searchers are not likely to have many documents to distribute to other nodes, and the documents are not likely to be large. An example of a sporadic searcher might be someone who has no data service plan or who primarily makes only telephone calls on his/her mobile device.

The sporadic searcher mainly interacts with the screen in Figure 4 to view active searches, and occasionally interacts with the screen in Figure 6 to retrieve documents. Searches (Figure 5) are infrequent, and other activities (Figures 7, 8, 9) are rarely used. The iTrust over SMS service accommodates the sporadic searcher, and defaults to Figure 4 when the application starts but, otherwise, does not adapt to the user. Specifically, it does not attempt to increase the node's membership by sending messages with node addresses. Because of the distributed nature of iTrust over SMS, it is difficult to decide, from a single node's perspective, whether its membership is sufficiently large.

The sporadic searcher is differentiated mostly by the need to address the *bootstrapping* problem when there are relatively few nodes in the local membership, but also by the relative lack of information or documents held by the sporadic searcher. Early social networking services also suffered from the bootstrapping problem. Social networks have limited value if only a few of the user's friends participate in the network; most centralized social networks require manual addition of friends, or suggest friends based on personal information.

In iTrust over SMS, the user can manually add nodes to the local membership via the user interface; however, more likely, the iTrust over SMS library automatically adds nodes to the local membership (it does not merely suggest that they be added), if those nodes are not already in the local membership. A common way of building a node's membership is that the iTrust library adds a matching node and a source node to the membership of a requesting node (searcher), it adds a requesting node to the membership of a node that receives the request, and it adds a retrieving node to the membership of a source node. This design choice increases a node's membership, by adding nodes that hold documents that match the user's search criteria and that provide interesting information from the user's perspective. Moreover, it allows sporadic searchers to *auto* promote themselves to casual searchers by simply searching more often and, thus, increasing their memberships.

B. Casual Searcher

A casual searcher is a user who uses the iTrust over SMS search and retrieval service to share information at a moderate frequency, size, and variety of shared documents. The casual searcher has a moderate number of documents stored on his/her mobile device, such as e-mail messages, contact information, personal photographs or videos, music and other documents. The amount of personal information stored correlates well with the usage of the device by typical smart phone users. For example, most smart phones have a basic built-in digital camera, which the smart phone user uses to take personal photographs when convenient; in contrast, a photographic enthusiast takes many more pictures but with a better-quality, stand-alone digital camera. Likewise, the typical smart phone user might store text documents or e-books but not literature manuscripts, home or amateur videos but not professional videos, e-mail messages but not work documents, *etc.*

The casual searcher mainly interacts with the screens in Figures 4, 5 and 6 to search for and retrieve documents. Sharing documents is handled automatically by iTrust over SMS, but the casual user may configure node settings using the screens in Figures 8 and 9. Like the sporadic searcher, the casual searcher is accommodated by iTrust over SMS by first showing the default activity in Figure 4 when the application starts.

Because the casual searcher sends queries frequently, the membership can be moderately large due to adding the matching node and the source node to the searcher's membership, adding the searcher and the relaying node to the matching node's membership, and adding the searcher to the source node's membership. Frequent searches make the casual searcher relatively well-known among other nodes in the iTrust over SMS network.

Increasing a node's membership requires an increase in the number of nodes to which the metadata and the requests are distributed in order to maintain the same number of responses to a search query. An adaptive method [6] that we have developed for iTrust over HTTP can also be used for iTrust over SMS. It increases dynamically, and strategically, the proportion of *queried* nodes in the node's membership (rather than the total number of nodes in the node's membership). It uses an algorithm that detects whether the number of matches corresponds to an analytically expected number of matches.

Increasing one's own membership and increasing one's presence in other nodes' memberships can improve access to information as well as the speed with which matches are made. Doing both provides a kind of "instant gratification," which is desirable for the mobile user demographic. Thus, by making more searches and increasing their memberships, casual searchers can *auto* promote themselves to become avid searchers.

C. Avid Searcher

An avid searcher has a plethora or abundance of shareable (and likely very desirable) information, and has or seeks hours of music or video and entire collections of shareable documents. At present, this behavior transcends the typical smart phone user; therefore, the avid searcher population is smaller than the casual searcher population.

However, a crucial difference between the avid searcher and the casual searcher is that the avid searcher typically retrieves not only the document for the first match but also the documents for the second, third or more matches. Because an avid searcher is likely to retrieve all documents for which the metadata match, the order of the match responses is less important than that for the casual searcher for which the first match response is the most important.

As for the previous types of searchers, the screen in Figure 4 serves as the default activity when the application starts. However, for the Avid searcher, the *Search details* activity shown in Figure 6 is typically used more often than the *New search* activity shown in Figure 5. The remaining activities shown in Figures 8, 9 and 7 are still seldom used.

Importantly, the avid searcher becomes more and more *instantly gratified* as the match responses return ever faster; however, there is a physical limit to the speed of SMS (which is determined by the specific mobile service provider). Repeatedly reaching this limit might have the effect of pushing the avid searcher behavior back down to that of the casual searcher and, indeed, might create a churn of casual searchers entering and leaving the avid searcher status.

D. Pure Searcher

The pure searcher is any searcher who searches and retrieves documents but, unlike the other searchers previously discussed, does not contribute (distribute) documents to other nodes in the iTrust over SMS network.

The pure searcher does not distribute documents by not storing documents locally, ignoring search queries, or ignoring retrieval requests. Not storing documents or ignoring search queries is made possible using the preferences shown in the screens in Figures 8 and 9. Such preferences are optional, because there might be legitimate reasons not to share local documents with others (political oppression, copyright laws, *etc.*).

The pure searcher can still distribute metadata on shareable information and, thus, send its node address to other nodes for inclusion in their memberships. Consequently, iTrust over SMS works as intended, until the final step when a searcher attempts to retrieve the document, at which point the source node simply ignores the retrieval request. There are no preferences to enable this behavior and, indeed, iTrust over SMS does *not* support this option. To achieve this behavior, one would have to modify the iTrust over SMS source code and create a mimic iTrust over SMS service.

By *not* sharing documents, the pure searcher is not sending its node address to other nodes during query relaying, match reporting, or document sharing and, thus, it pays the price of having a smaller chance of being included in other nodes' memberships. This membership penalty might encourage the pure searcher to distribute shareable documents and become a sporadic or casual searcher.

The pure searcher in iTrust over SMS is similar to leechers in other peer-to-peer networks, such as Gnutella, that provide little or no benefit to the community. Leechers are discouraged but are sometimes unavoidable, particularly when there are new nodes with small memberships or nodes that hold only a few documents locally (such as sporadic searchers).

E. Other Use Cases

Some nodes might freely distribute local documents and never search for documents; such behavior mostly occurs because of an abundance of resources or general good-will. Other nodes might simply relay queries, allowing the built-up membership and search queries to be used for other purposes either benign, nefarious, or somewhere in between. Lastly, malicious nodes might actively or passively attack other nodes, again not necessarily by any direct user action.

V. Related Work

In a study of mobile search behavior, Kamvar *et al.* [13] found that most mobile searchers use the search service for a short period of time, do not engage in exploration, and have a specific topic in mind. In a subsequent study [14], they found that the diversity of mobile search topics is rather limited. Evans and Chi [9] have provided an analysis of the activities of individuals conducting search over social networks, with a focus on foraging and sense making.

Church and Smyth [8] have also addressed the information needs of mobile users, and Church *et al.* [7] have developed a Social Search Browser for mobile users. Acero *et al.* [1] have investigated voice search on mobile phones using Web Services. Schusteritsch *et al.* [20] have undertaken work to improve mobile search using SMS text messages.

Adamic and Adar [2], and also Watts *et al.* [22], have investigated the effectiveness of search in social networks, which appears to depend on the structured nature of those networks and a few highly-connected nodes. Many searchers were able to exploit that structure to find information in relatively few steps. In experiments with students where such structure does not exist, such local search strategies were less effective. We are investigating the effects of a local search strategy on iTrust over SMS.

Existing commercial mobile search services include AOL Mobile [3], Google SMS [10], Windows Live Mobile [23], and Yahoo! OneSearch [24]. Those mobile search services use conventional centralized Web search engines, which are subject to filtering and censorship. They provide a limited set of pre-defined topics, and use either special keywords within a search query (*e.g.*, "directions" to obtain directions) or a specialized parser to determine the intended topic (*e.g.*, "INTC" for a stock quote).

The SMSFind system [4], [5] also utilizes conventional centralized Web search engines. It does not use pre-defined topics but, rather, allows the user to enter an explicit contextual hint about the search topic. SMSFind uses information retrieval techniques to extract an appropriate condensed 140-byte snippet as the final SMS search response, which iTrust over SMS currently does not do but which might be a valuable feature for a future version of iTrust over SMS.

The Mobile Agent Peer-To-Peer (MAP2P) system [12] supports mobile devices in a Gnutella file-sharing network using mobile agents. The mobile agent (rather than the mobile device) attaches itself to the peer-to-peer network, and acts as a proxy for the mobile device.

The Distributed Mobile Search Service [16] broadcasts query results locally and forwards them over several hops. It is based on a distributed index that comprises, on each mobile device, a local index cache, containing keywords and corresponding document identifiers, where received query results are cached. iTrust over SMS likewise maintains a local index cache, with metadata keywords as well as node addresses and document identifiers, on the mobile device.

The 7DS system [19] supports information sharing among mobile devices. The 7DS system uses a multi-hop flooding algorithm together with multicasting of queries, which is not trustworthy. In contrast, iTrust over SMS does not use multicasting or flooding, which are too expensive in message cost, but instead relays requests more selectively.

Search in social networks can exploit the trust that members have in each other, and route information and requests based on their relationships. Gummadi *et al.* [11] investigate the integration of social network search with Web search, and conclude that such integration can lead to more timely and efficient search. Tiago *et al.* [21] describe a system for mobile search in social networks based on the Drupal content site management system, using the network of social links formed from the address book on the mobile device, which iTrust over SMS likewise does.

PeopleNet [18] is a social network that exploits physical location to facilitate searching. The authors observed a rapid increase in the number of copies of a query as it propagates in search of data, akin to flooding. Thus, they advocate a swap strategy in which a request migrates but does not replicate itself. iTrust over SMS explicitly manages the replication of queries to achieve a desired probability of finding a match.

Yang *et al.* [25] propose a search mechanism for unstructured peer-to-peer networks, based on special interest groups formed by nodes that share similar interests. iTrust over SMS likewise allows users interested in a particular topic or cause to form a social network, so that they can share information.

VI. Conclusion and Future Work

We have described the iTrust over SMS search and retrieval service, as well as the Android user interface that enables the user to use the service. We have also presented use cases for iTrust over SMS, namely, the sporadic, casual and avid searchers who distribute and search for documents, as well as the pure searchers who search for but do not distribute documents. In the future, we plan to distribute iTrust over SMS to smart phone users, either through the Android Market app store or another such service. We also plan to do extensive performance evaluations of iTrust over SMS using mobile phones in large real-world social networks.

Acknowledgment

This research was supported in part by U.S. National Science Foundation grant number NSF CNS 10-16193.

References

[1] A. Acero, N. Bernstein, R. Chambers, Y. C. Ju, X. Li, J. Odell, P. Nguyen, O. Scholz and G. Zweig, "Live search for mobile: Web services by voice on the cellphone," *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, Las Vegas, NV (March–April 2008), pp. 5256–5259.

[2] L. Adamic and E. Adar, "How to search a social network," *Social Networks*, vol. 27, iss. 3 (July 2005), pp. 187–203.

[3] AOL Mobile, http://www.aolmobile.com

[4] J. Chen, B. Linn and L. Subramanian, "SMS-based contextual Web search," *Proceedings of the ACM SIGCOMM MobiHeld Workshop*, Barcelona, Spain (August 2009), pp. 19–24.

[5] J. Chen, L. Subramanian and E. Brewer, "SMS-based Web search for low-end mobile devices," *Proceedings of the 16th ACM International Conference on Mobile Computing and Networking*, Chicago, IL (September 2010), pp. 125–136.

[6] Y. T. Chuang, I. Michel Lombera, P. M. Melliar-Smith and L. E. Moser, "Detecting and defending against malicious attacks in the iTrust information retrieval network," *Proceedings of the 26th International Conference on Information Networking*, Bali, Indonesia (February 2012), pp. 263–268.

[7] K. Church, J. Neumann, M. Cherubini and N. Oliver, "SocialSearchBrowser: A novel mobile search and information discovery tool," *Proceedings of the 15th International Conference on Intelligent User Interfaces*, Hong Kong, China (February 2010), pp. 101–110.

[8] K. Church and B. Smyth, "Understanding the intent behind mobile information needs," *Proceedings of the 14th International Conference on Intelligent User Interfaces*, Sanibel Island, FL (February 2009), pp. 247–256.

[9] B. M. Evans and E. H. Chi, "Towards a model of understanding social search," *Proceedings of the ACM Conference on Computer Supported Cooperative Work*, San Diego, CA (November 2008), pp. 485–494.

[10] Google SMS, http://www.google.com/sms

[11] K. P. Gummadi, A. Mislove and P. Druschel, "Exploiting social networks for Internet search," *Proceedings of the 5th ACM SIGCOMM Workshop on Hot Topics in Networks*, Irvine, CA (November 2006), pp. 79–84.

[12] H. Hu, B. Thai and A. Seneviratne, "Supporting mobile devices in Gnutella file sharing network with mobile agents," *Proceedings of the 8th IEEE Symposium on Computers and Communications*, Kemer-Antalya, Turkey (July 2003).

[13] M. Kamvar and S. Baluja, "A large scale study of wireless search behavior: Google mobile search," *Proceedings of the 24th ACM SIGCHI Conference on Human Factors in Computing Systems*, Montreal, Quebec, Canada (April 2006), pp. 701–709.

[14] M. Kamvar, M. Kellar, R. Patel and Y. Xu, "Computers and iPhones and mobile phones, oh my!: A log-based comparison of search users on different devices," *Proceedings of the 18th International Conference on the World Wide Web*, Madrid, Spain (April 2009), pp. 801–810.

[15] J. Larsen, K. W. Axhausen, J. Urry, "Geographies of social networks: Meetings, travel and communications," *Mobilities*, vol. 1, no. 2 (July 2006), pp. 261–283.

[16] C. Lindemann and O. P. Waldhorst, "A distributed search service for peer-to-peer file sharing in mobile applications," *Proceedings of the 2nd International Conference on Peer-to-Peer Computing*, Linkoping, Sweden (September 2002), pp. 73–80.

[17] I. Michel Lombera, Y. T. Chuang, P. M. Melliar-Smith and L. E. Moser, "Trustworthy distribution and retrieval of information over HTTP and the Internet," *Proceedings of the 3rd International Conference on the Evolving Internet*, Luxembourg City, Luxembourg (June 2011), pp. 7–13.

[18] M. Motani, V. Srinivasan and P. S. Nuggehalli, "PeopleNet: Engineering a wireless virtual social network," *Proceedings of the 11th ACM International Conference on Mobile Computing and Networking*, Cologne, Germany (August 2005), pp. 243–257.

[19] M. Papadopouli and H. Schulzrinne, "Effects of power conservation, wireless coverage and cooperation on data dissemination among mobile devices, *Proceedings of the ACM Symposium on Mobile Ad Hoc Networking and Computing*, Long Beach, CA (2001), pp. 117–127.

[20] R. Schusteritsch, S. Rao and K. Rodden, "Mobile search with text messages: Designing the user experience for Google SMS," *Proceedings of the 23rd ACM SIGCHI Conference on Human Factors in Computing Systems*, Portland, OR (April 2005), pp. 1777–1780.

[21] P. Tiago, N. Kotiainen, M. Vapa, H. Kokkinen and J. K. Nurminen, "Mobile search – Social network search using mobile devices," *Proceedings of the 5th IEEE Consumer Communications and Networking Conference*, Las Vegas, NV (January 2008), pp. 1201–1205.

[22] D. J. Watts, P. S. Dodds and M. E. J. Newman, "Identity and search in social networks," *Science*, vol. 296, no. 5571 (May 2002), pp. 1302–1305.

[23] Windows Live Mobile, http://home.mobile.live.com

[24] Yahoo! OneSearch, http://mobile.yahoo.com/onesearch

[25] J. Yang, Y. Zhong and S. Zhang, "An efficient interest-group-based search mechanism in unstructured peer-to-peer networks," *Proceedings of the International Conference on Computer Networks and Mobile Computing*, Shanghai, China (October 2003), pp. 247–252.