Privacy-Triggered Communications in Pervasive Social Networks

Murtuza Jadliwala[†], Julien Freudiger[†], Imad Aad[‡], Jean-Pierre Hubaux[†] and Valtteri Niemi[‡]
[†]*School of Computer and Communication Sciences, EPFL, Switzerland*
[‡]*Nokia Research Center, Lausanne, Switzerland*
Email:firstname.lastname@{[†]epfl.ch,[‡]nokia.com}

***Abstract*—Pervasive social networks extend traditional social networking by enabling users to share information in a peer-to-peer fashion using their wireless mobile devices. Contrary to traditional online social networks, privacy protection in such networks depends heavily on users' context (time, location, activity, etc.) and their sensitivity to the shared data and context. Existing privacy-preserving mechanisms do not adapt well to different data, context and user sensitivities. In this work, we follow a fresh approach for privacy preservation, called *privacy-triggered communications*; it allows users in such pervasive networks to dynamically regulate their communications based on their context and on the evolution of their privacy in that context. Our initial results show that this is a feasible strategy for privacy management in pervasive social networking scenarios.**

I. Introduction

Wireless peer-to-peer communications using WiFi or Bluetooth has been a popular way for mobile phone users to exchange information with other devices in the neighborhood. Recently, Nokia has also introduced Nokia Instant Community (NIC) [1], a platform for power-efficient and always-on device-to-device messaging which uses wireless, ad-hoc, IEEE 802.11 connections to seamlessly carry small-sized local contextual information between devices. Such wireless peer-to-peer networks, referred by us as *pervasive social networks*, enable a wide array of *context-aware* applications such as local-area social networking [2], dating [3], personal safety [4] and micro-blogging [5].

Location privacy concerns in social networking applications have been growing and users are increasingly becoming privacy-aware, as shown by a recent survey [6]. The success of these upcoming pervasive social networking platforms will depend on the efficiency and simplicity of the privacy-preservation mechanisms. Sharing data locally in a peer-to-peer fashion avoids the privacy issues resulting from information revelation to untrusted third-parties, but leaks personal information to wireless eavesdroppers. In particular, by overhearing the content and pattern of communications between mobile devices, it may be possible to identify and track their owners. Various approaches have been proposed to address the location privacy issues in infrastructure-based and wireless peer-to-peer networks. Some popular solutions propose to remove identifiers [7] from exchanged messages or to change them over time [8], [9] in order to prevent an adversary from linking data to users. Others rely on cryptographically secure authentication techniques, such as group signatures, anonymous credentials and secret handshakes, to authenticate and/or encrypt communications.

One limitation of these approaches is that, in addition to being heavy, they do not consider factors such as the device context and users' sensitivity to privacy and shared content, which is crucial in pervasive systems. For example, a user may not require the same privacy at his home as compared to his office. Similarly, stricter anonymization is needed while sharing sensitive data as compared to public information. In addition, source and destination identifiers are generally not encrypted for ease of message routing. Other techniques such as identifier removal and/or changes are ineffective against privacy leakage due to message context.

In this paper, we propose a novel approach that allows users to dynamically fine tune their privacy requirements in order to *control their participation* in pervasive social networks. Our approach, called *privacy-triggered communications*, enables users to make communication decisions based on their current privacy level, the content of communications and the current context. In contrast with existing approaches, users can automatically regulate what they share with others by making dynamic privacy-based decisions as they move in the network. More specifically, users' mobile devices identify appropriate moments to share data with nearby nodes: for example, share intimate data only when their privacy level is sufficiently high. An analogy can be drawn here to a password strength indicator tool, which provides a visual indication of the strength of the password chosen by the user. Our proposed approach is similar; users are provided with a view of how good their current privacy is, which they then use to regulate their communications. We propose mechanisms for effective privacy-dependent message triggering in pervasive communication systems.

II. System Model

In pervasive social networks, users exchange information, based on relationship, interests, affiliations and context, with other co-located users using their mobile devices in a peer-to-peer wireless fashion (e.g., WiFi or Bluetooth). Such communications are analogous to chatting on a public Internet forum, but happen independently of any infrastructure or chat server. As a result, they complement infrastructure-based communications using cellular or WLAN. We assume

978-1-4577-0351-5/11/$26.00 ©2011 IEEE

that communications are either *non-interactive* or *interactive*. In non-interactive communications, a user broadcasts information without any specific request from other users, for example, airing a point-of-view at a political rally. Whereas, in interactive communications, users multicast relevant information only in response to specific queries from other users, for example, sending local restaurant recommendations to nearby requesting users. The wireless ad-hoc interface of the devices possess a smart and efficient beaconing mechanism [10] for *neighborhood discovery*; there are provisions in the standard such that a group of ad-hoc devices in the vicinity of each other, can efficiently distribute the beaconing task uniformly within the group. We assume that any communication consists of: 1) an identifier or pseudonym (e.g., MAC address) belonging to the message originator and 2) the message itself in encrypted or unencrypted form, depending on whether the communication is private or public, respectively. If the message is not public, identifying information of the target user(s) is also included. To prevent trivial tracking and message linking to users, we assume that device identifiers or pseudonyms are regularly changed at pre-defined intervals. Similarly, user security credentials are also regularly updated.

III. Privacy Threats

In this work, we focus mainly on location privacy threats due to an eavesdropping adversary (can be a legitimate node) whose goal is to identify and track users based on the observed communications. By collecting messages together with their location, the adversary will try to obtain the true identity of the device owner from an analysis of pseudonymous location traces [8]. This can also be used to link past communications to specific users. Users require *source anonymity*, i.e., they do not want eavesdroppers to learn that a particular message comes from a specific user. For example, participants at a political rally may not want the government to figure that a particular, possibly anti-government, message came from a specific user. By observing the exchanged messages, the adversary also attempts to learn sensitive information about users and their social networks, for example, association with particular religious or political groups. The strength of the adversary depends on the means available to him, i.e., he can be a local eavesdropper with a single radio or a more sophisticated one with multiple sensing stations. We assume that factors such as user mobility, multi-path fading and interference will provide enough radio signal uncertainty to prevent physical layer identification attacks. We assume that security and integrity of data, especially in private communities, can be addressed using appropriate cryptographic measures.

IV. Privacy-Triggered Communications

We implement our privacy-triggered networking concept as a *privacy-wrapper* (middle-ware) composed of a set of cross-layer libraries (Fig. 1 (a)). The privacy-wrapper can be utilized by any wireless peer-to-peer application through an API. The privacy-wrapper continuously monitors communications from the device and its neighborhood, and it dynamically determines whether the context provides enough privacy for users to share their information. In other words, the context and its associated privacy level act as a trigger for the *networking activity* of nodes.

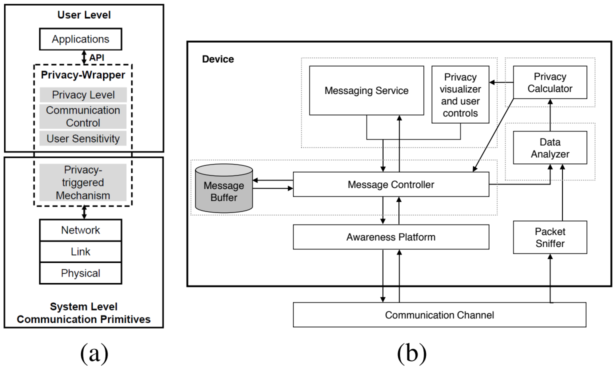

Figure 1. Privacy-wrapper (a) Conceptual view, (b) Architecture

Privacy policies and controls can also be implemented as user-friendly policy management tools [11] or as operating system-level libraries [12]. The approaches proposed in [11] are too application-specific whereas the approach in [12], contrary to ours, focuses on enforcing a system-wide privacy policy by placing privacy-related controls in the operating system. Instead, our approach focuses on providing dynamic, application-independent privacy controls to the users, which would enable them to regulate their participation in pervasive social networks. Our approach consists of three parts: 1) privacy measurement, 2) capturing user sensitivity, and 3) privacy-based communication triggering.

A. Measuring Privacy

Before enabling privacy-based communication decisions, it is essential to accurately quantify user location privacy. In the context of pervasive social networks, users should be able to measure how identifiable and traceable they are based on their communications. One commonly used metric for measuring anonymity in pseudonymous communications is k*-anonymity* [13]. A user is k-anonymous if a passive adversary is unable to distinguish communications of the user in question from at least $k-1$ other users (called user's anonymity set). Other metrics for measuring privacy can also be used. For example, Diaz et al. [14] and Serjantov et al. [15] propose a privacy metric based on the *entropy* of the anonymity set of the message originator. Similarly, there are other metrics that measure the adversary's error in correctly identifying users' events [16] or that measure the extent to which a user can be tracked with high certainty [17].

Our approach of privacy-triggered communications is *not* restricted to any specific privacy metric(s); system designers should employ a metric that is simple, reflects users' intuition of privacy and effectively protects against the assumed adversary. Currently, we have implemented the standard k-anonymity metric that measures privacy using the neighborhood density of the device. We capture the strength of the adversary by measuring the maximum distance between the device and its neighbors, called *confusion distance*; more packed a neighborhood, more difficult it is to differentiate users from each other. In this way, we are able to capture both the user's anonymity set as well as the difficulty of the adversary in distinguishing the user from others in its anonymity set. One drawback of such a simple metric is that an adversary can easily have k devices in the neighborhood of a user in order to track him. In order to overcome such problems, a more sophisticated metric, which takes into account other neighborhood parameters and complex adversarial strategies can also be used. Trivial attacks, such as the one above, can be thwarted by a dynamic k value (unknown to the adversary) chosen based on the user sensitivity as discussed next.

B. User Sensitivity

Privacy metrics, by themselves, do not capture the sensitivity of a user to a given context and to the network observables in that context. For example, in a known environment (e.g., office), a given neighborhood density context may be considered as high privacy, whereas a similar density in an unknown environment may be considered as low privacy by the same user. Dey et al. [18] defined context as "any information that can be used to characterize the situation of an entity, where an entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and the application themselves. Context is typically the location, identity and state of people, groups and computational and physical objects". Determining and managing user sensitivities to the various external contexts in pervasive environments is non-trivial. There are two approaches that can be followed for this, namely *manual* and *automatic*. In manual approaches, users can create and customize profiles, one for each significant context, that captures their privacy sensitivity in that context. A profile can contain, in addition to context information such as location, time, etc., privacy metric parameters such as the required minimum and maximum anonymity set sizes and minimum and maximum confusion distances. Due to the difficulty in accurately estimating the privacy parameters for a context, sensitivity profiles can also be alternatively expressed as preferred locations or points-of-interest [19]. Network parameters can then be extracted from these locations and used as a baseline during privacy computation. After creation of profile(s), users manually choose the sensitivity profile that best matches his current context (or is most appropriate for current communications). The system dynamically adjusts or scales the users' privacy level based on the chosen profile. One drawback of the manual approach is that profile management can become cumbersome leading to usability issues. Alternatively, the process of user sensitivity determination can be partially automated. Users are provided with a fixed set of customizable sensitivity profiles. Based on the current device context (e.g., location, neighborhood), an appropriate profile is automatically chosen by the system. Automation improves usability, but it is risky as erroneous observations can result in the selection of incorrect profiles.

C. Privacy-Triggering Mechanisms

We propose two mechanisms for triggering communications based on privacy.

1) Threshold-based Technique: In the threshold-based technique (Algorithm 1), users assign privacy and time validity thresholds to their communications. The privacy threshold indicates the minimum level of privacy required by the user before that communication can be broadcast on the network whereas the time validity threshold gives the time period after which the communication is no longer valid (or significant). Unless the privacy threshold is met, the communication is not sent out on the network; it is instead queued in a local message buffer (B). Buffered communications that are still valid are immediately scheduled for delivery when the corresponding privacy threshold is met by the user.

Algorithm 1: Threshold-based algorithm

```
Data: $d(t)$: user privacy at $t$, $d_b$: min. reqd. privacy for msg
      $b$, $t_b$: validity period for msg $b$.
for every time instant $t_{cur}$ do
    while $B$ is not empty do
        for each $b \in B$ do
            if $d(t_{cur}) \geq d_b$ then
                if $t_{cur} \leq t_b$ then
                    Schedule $b$ for delivery;
                else
                    Delete $b$ from $B$;
                end
            end
        end
    end
end
```

2) Probabilistic Technique: Although the threshold-based approach is simple, it is restrictive as it would be difficult for users to decide a value for the privacy threshold; the threshold may change depending on users future states in terms of privacy and communications. The threshold-based approach does not take into account the past states of the user and also the possible future states, before making communication decisions. For example, users can be identified based on the messages that are transmitted often due to lower privacy requirements, which is not good for subsequent messages that require higher privacy.

In order to overcome these problems, we also propose a probabilistic technique for privacy-triggering. Similar to the threshold-based technique, the device consists of a message buffer holding a fixed number of messages that need to be communicated based on privacy and the probabilistic technique is implemented as a decision mechanism in the middle-ware. At each time instant, the middle-ware is consulted to determine the message(s) (from the buffer) that can be forwarded in that time instant. This message(s) is chosen based on the optimal policy computation by the decision mechanism of the middle-ware. A policy decides which messages in the buffer are forwarded at each time instant. An optimal policy maximizes the overall utility or reward of the device (user) for forwarding messages.

The device (user) location privacy in a pervasive social network is stochastic in nature as it depends on non-deterministic factors such as user context, user mobility and the mobility, activity and context of the device neighborhood. Thus, the system's privacy evolution can be modeled as a stochastic process, such as a discrete time process. Let $B = \{b_1, b_2, \ldots, b_n\}$ be the n messages in the device buffer. Let $s \in S$ represent the state of a device at any time instant. The state of a device is represented in terms of its privacy and other factors (e.g., remaining energy). The objective is to determine an *optimal policy* $\mu^*(s)$ that, as a function of the current device state, determines a message (or a set of messages) $b \in \mathcal{P}(B)$[1] that can be forwarded in that state. Assuming that time is divided into discrete time intervals, an optimal policy can be evaluated at the beginning of fixed time slots to determine the message(s) to be forwarded in that time slot. At each such decision epoch, a decision is consulted by the device from the privacy-triggering middle-ware system. The middle-ware determines message(s) that, if communicated in the current time epoch, would maximize an overall discounted *reward* for the user.

We now characterize a device's state in terms of its privacy in the environment. As discussed in Section IV-A, there are various metrics for quantifying location privacy in a pseudonymous wireless peer-to-peer network. Let us assume that the device privacy $d \in [d_{min}, d_{max}]$ can be measured through active measurements on the device's wireless peer-to-peer interface using one of the metrics mentioned in Section IV-A. We can map the range of the privacy values of the device as a set of quantized states $S_d = \{s_{d_1}, s_{d_2} \ldots s_{d_J}\}$. Similarly, let $S_o = \{s_{o_1}, s_{o_2} \ldots s_{o_L}\}$ represent the states based on the quantized values of some other significant device parameters, for example remaining energy of the device. Then, the set of all possible device states S can be represented as $S = S_d \times S_o$.

The variations in the measured device privacy depends on various factors such as the privacy metric, device mobility, device neighborhood dynamism, users' context and sensitivity to specific contexts and current device message traffic. Similarly, other variations such as device energy would depend on device usage. We capture these variations in device states using a *finite-state Markov chain model* as follows. Let $p_{i,j}$ denote the probability to transition from state $s_i \in S$ to state $s_j \in S$ at time instant t:

$$p_{i,j} = P(s(t+1) = s_j | s(t) = s_i)$$

where, $s(t)$ denotes the state of the device at time instant t. Thus, we can define the state transition matrix as $M = [(p_{i,j})]$ and the corresponding transition equation as

$$\tilde{p}(t+1) = \tilde{p}(t) \cdot M$$

where, $\tilde{p}(t)$ represents the probability vector over the entire state space S at the time instant t.

The forwarded message(s) $b \subseteq B$ in each time instant affects the privacy (or other factors such as energy) of the device, thus taking it from the current state to some other state. In other words, the action[2] b chosen by the system will modify the dynamic system and these state transitions can be represented by a controlled Markov chain model $M(b), b \subseteq B$ (or $b \in \mathcal{P}(B)$). The transition probability $p_{i,j}(b)$ of $M(b)$ represents the probability of transiting from state s_i to state s_j provided action b has been taken in s_i, i.e., message(s) b has been forwarded in state s_i.

After defining the Markov model on the finite state space S and the finite action space $\mathcal{P}(B)$ of the device, we can now formulate the decision control problem of choosing optimal actions (or messages to be forwarded) as a *Markov Decision Process (MDP)*. In order to do that, we first need to define a *reward function*. We define a real-valued reward function $R : S \times \mathcal{P}(B) \rightarrow \mathbb{R}$, defined over the set space S and action space $\mathcal{P}(B)$, which quantifies the preference or utility (of the user) obtained by taking action b (in this case, forwarding messages b) when the device is in the (privacy) state s.

Let X_t and Δ_t be the random variables denoting the device state and the action chosen at time period t, respectively. Now, for any stationary policy μ, initial state s and the reward function R, the *total reward* $V_\mu(s)$ is given as:

$$V_\mu(s) = \sum_{t=0}^{N-1} \beta^t \cdot E_{\mu,s}\{R(X_t, \Delta_t)\} \qquad (1)$$

where, $E_{\mu,s}\{R(X_t, \Delta_t)\}$ is expected reward corresponding to the stochastic process $\{X_t, \Delta_t, t \geq 0\}$, $\beta \in (0, 1]$ is the discount factor over future decisions and N is the finite horizon. We assume the stationarity of policies over the finite horizon N. N can be a fixed time period, say from 9 a.m. in the morning to 9 p.m. in the evening.

[1] $\mathcal{P}(B)$ is the powerset of B

[2] Readers should note that as the buffer is finite, the action space of the system is also finite

We then define the *optimal value function* for the discounted expected total reward as

$$V(s) = \sup_{\mu \in \Pi} V_\mu(s), s \in S \qquad (2)$$

where, Π is the set of all possible policies. In other words, $V(s)$ is the best reward one can achieve from an initial state s, with a discount factor β when there remain N horizons. Any policy μ^* is *optimal* if $V_{\mu^*}(s) = V(s)$ for all $s \in S$. Now, if $V_\mu(s)$ is well-defined (and finite) for all μ and all $s \in S$, then the optimal value function (2) is a unique solution of the optimality equation (3), also referred as the *Bellman's equation*, under certain conditions [20].

$$V(s_i) = \sup_b \{R(s_j, b) + \beta \sum_{s_j} p_{i,j}(b) \cdot V(s_j)\}, s_i \in S \qquad (3)$$

Thus, in order to determine the optimal policy (or messages to be forwarded in each time instant), we need to solve the finite discounted reward problem, as described by the optimality equation in (3). There are various techniques for solving the optimality equation. We use a dynamic programming-based approach, called the Value Iteration Algorithm, as outlined below.

Algorithm 2: Value Iteration Algorithm

```
$V(s) \leftarrow 0\ \forall s \in S$;
$\delta \leftarrow 0$;
repeat
    foreach $s \in S$ do
        $v \leftarrow V(s)$;
        foreach $b \in \mathcal{P}(B)$ do
            $Q(s,b) \leftarrow R(s,b) + \beta \sum_{s' \in S} V(s') \cdot M(b)[s,s']$;
        end
        $V(s) \leftarrow max_b\ Q(s,b)$;
        $\delta \leftarrow max\{\delta, |v - V(s)|\}$;
    end
until $\delta < \phi$;
for $s \in S$ do
    $\mu^*(s) \leftarrow arg\ max_b\ Q(s,b)$
end
```

The parameter ϕ in the terminating condition of Algorithm 2 is a system-defined constant. The value iteration algorithm is executed and updated periodically by the privacy-triggering middle-ware in order to evaluate the optimal control policy μ^* for forwarding messages. The messages selected by the optimal policy (in the current state) are then forwarded by the device. The MDP invocation on the devices has to be preceded by a *training* or *learning* phase. During the training phase, the state transition matrix $M(b)$ is evaluated by monitoring the system state transitions associated with the forwarded messages for different device contexts such as time, location, day-of-week, etc. We can observe that the proposed approach evaluates an optimal policy for all states $s \in S$, and thus the larger the total number of states and the set of all possible actions in those states, the slower will be the evaluation of the optimal policy. The feasibility of the proposed approach clearly depends on how the system privacy states are quantified and the size of the message buffer.

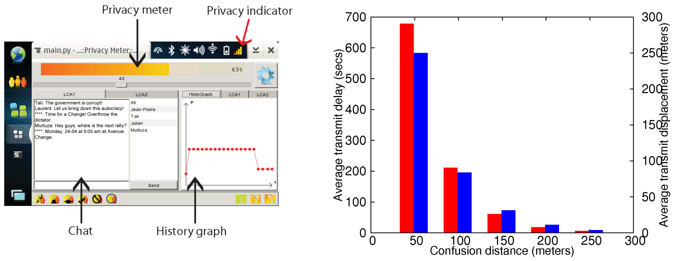

Tali: The government is corrupt!
Laurent: Let us bring down this autocracy!
****: Time for a Change! Overthrow the dictator.
Murtuza: Hey guys, where is the next rally?
****: Monday, 24-04 at 9:00 am at Avenue. Change.

(a)

(b)

Figure 2. (a) Prototype, (b) Simulation results: Average delay (red) and displacement (blue)

V. Implementation

We have implemented the concept of privacy-triggered communications for NIC enabled Nokia N810 mobile devices. The overall system architecture is shown in Fig. 1(b). Applications developed on top of NIC, for example chat, messaging, file sharing, etc., send data to and receive data from neighboring NIC devices on the wireless ad-hoc interface. The NIC platform provides basic user management, community management and communication primitives. Our implementation focuses on development of privacy measurement mechanisms (data analyzer and privacy calculator in Fig. 1(b)), privacy visualization and privacy control tools (privacy visualizer in Fig. 1(b)) and privacy-triggering mechanisms (message controller in Fig. 1(b)). Fig. 2(a) shows a screen-shot of our prototype application on Nokia N810. We have currently implemented privacy-triggered communications as an extension to a simple NIC-based peer-to-peer chat application. This can be easily extended to all applications using NIC. The application consists of a privacy meter, which displays the current level of user privacy in the pervasive environment. A user-context choice control can be used by the user to select and customize a profile that best matches his current context sensitivity. A slider control can be used to select the minimum privacy level required (compared to the existing level) for the upcoming communications. Currently, we have implemented only the threshold-based technique for privacy-triggered communications. In this case, if the required privacy for a message is higher than the current privacy level, it is locally buffered until sufficient privacy is available. Otherwise, messages are sent out on the network immediately.

Our goal in this work is to verify if privacy-triggered communications are feasible and effective for protecting user privacy in upcoming pervasive social networking platforms such as NIC. We plan to do that by deploying around 100 NIC enabled mobile phones on the EPFL campus, which will be carried by students for a period of 3 months. During this period, we will test the usability and effectiveness of privacy-triggered communications in preserving user anonymity for certain communications, as well as the effect of privacy-triggering on the overall Quality-of-Service (QoS).

We prepare for such a large-scale deployment by first carrying out simulation experiments. We simulate the movement and privacy-triggered peer-to-peer communications of 100 mobile devices equipped with 914MHz wireless radios moving along the streets of a $1km \times 1km$ city block at pedestrian speeds ($<$ 3km/h). We implement the threshold-based approach with k-neighborhood as the privacy metric. Simulation results (Fig. 2b) are as expected and show that the average delay and displacement of communications due to privacy-triggering decreases as confusion distance increases, given a privacy threshold of 6 neighboring devices. Interestingly, a confusion distance of $100m$, results in an average communication delay of around 3 min. and an average communication displacement of roughly $75m$, which is reasonable for most social communications. An interesting observation is that places with higher user density (e.g., Points of Interest or POI) had higher communication rates. POI's provide good mixing and are suitable for anonymous peer-to-peer communications. A side-effect of this is that communications become bursty, leading to higher congestion within and near POIs. This also implies that users share lots of data all at once, which is not good. As expected, we can see that privacy comes at the price of a lower QoS. By means of appropriate tools and mechanisms, our goal is to enable users to make their own choice in this regard. Those who prefer high privacy can have it, but at the cost of a lower QoS.

VI. Conclusion and Future Work

We introduced privacy-triggered communications as a novel technique for privacy-preservation in an upcoming class of social networks called pervasive social networks. Privacy-triggered communications allow users in such networks to dynamically regulate their participation based on their privacy. This work is a first step towards providing tools that consider the wireless context of the users in order to control their privacy. We feel that the successful adoption of pervasive social networks will depend on the ability of the system to accurately capture users' privacy requirements in different environments and contexts. As part of ongoing work, we plan to implement the probabilistic technique for privacy-triggering and other metrics for privacy measurement. These efforts will take advantage of the real-life mobility and human-behavior data available from the planned on-campus trial.

References

[1] Rhiain, "Nokia instant community gets you social, 2010," http://conversations.nokia.com/2010/05/25/nokia-instant-community-gets-you-social/.

[2] A. Ahtiainen, K. Kalliojarvi, M. Kasslin, K. Leppanen, A. Richter, P. Ruuska, and C. Wijting, "Awareness networking in wireless environments: Means of exchanging information," *IEEE Vehicular Tech. Mag.*, 2009.

[3] M. Khiabani, "Metro-sexual," http://bit.ly/theranMetroSexual.

[4] E. Paulos and E. Goodman, "The familiar stranger: anxiety, comfort, and play in public places," in *CHI*, 2004.

[5] S. Gaonkar, J. Li, R. Choudhury, L. Cox, and A. Schmidt, "Micro-blog: sharing and querying content through mobile phones and social participation," in *MobiSys*, 2008.

[6] M. Madden and A. Smith, "Reputation management and social media," http://pewinternet.org/Reports/2010/Reputation-Management.aspx.

[7] M. Gruteser and B. Hoh, "On the anonymity of periodic location samples," in *Security in Pervasive Computing*, 2005.

[8] A. R. Beresford and F. Stajano, "Mix zones: User privacy in location-aware services," in *PerSec*, 2004.

[9] F.-L. Wong and F. Stajano, "Location privacy in Bluetooth," in *ESAS*, 2005, pp. 176–188.

[10] "IEEE 802.11 LAN/MAN Wireless LANS," http://standards.ieee.org/getieee802/802.11.html.

[11] J. Cornwell, I. Fette, G. Hsieh, M. Prabaker, J. Rao, K. Tang, K. Vaniea, L. Bauer, L. Cranor, J. Hong, B. McLaren, M. Reiter, and N. Sadeh, "User-controllable security and privacy for pervasive computing," in *HotMobile*, 2007.

[12] S. Ioannidis, S. Sidiroglou, and A. Keromytis, "Privacy as an operating system service," in *HOTSEC*, 2006.

[13] L. Sweeney, "Achieving k-anonymity privacy protection using generalization and suppression," *Int. Jour. on Uncertainty, Fuzziness and Knowledge-based Sys.*, 2002.

[14] C. Diaz, S. Seys, J. Claessens, and B. Preneel, "Towards measuring anonymity," in *PET*, 2002.

[15] A. Serjantov and G. Danezis, "Towards an information theoretic metric for anonymity," in *PET*, 2002.

[16] L. Fischer, S. Katzenbeisser, and C. Eckert, "Measuring unlinkability revisited," in *WPES*, 2008.

[17] B. Hoh, M. Gruteser, H. Xiong, and A. Alrabady, "Preserving privacy in GPS traces via uncertainty-aware path cloaking," in *CCS*, 2007.

[18] A. Dey, D. Salber, and G. Abowd, "A conceptual framework and toolkit for supporting the rapid prototyping of context-aware applications," in *HCI*, 2001.

[19] T. Xu and Y. Cai, "Feeling-based location privacy protection for location-based services," in *ACM CCS*, 2009.

[20] Q. Hu and W. Yue, *Markov Decision Processes with Their Applications*. Springer, 2008, ch. 2.