

<https://sprite.utsa.edu/>

Copyright may be reserved by the publisher.



Big Brother Knows Your Friends: on Privacy of Social Communities in Pervasive Networks

Igor Bilogrevic¹, Murtuza Jadliwala^{2*}, István Lám⁵, Imad Aad³, Philip Ginzboorg⁴, Valtteri Niemi⁴, Laurent Bindschaedler¹, and Jean-Pierre Hubaux¹

¹ LCA1, EPFL, Lausanne, Switzerland

² EECS Department, Wichita State University, USA

³ Nokia Research Center, Lausanne, Switzerland

⁴ Nokia Research Center, Helsinki, Finland

⁵ Faculty of Electrical Engineering and Informatics, BME, Hungary
firstname.lastname@{epfl.ch, wichita.edu, nokia.com}, lam@crsys.hu

Abstract. Wireless network operators increasingly deploy WiFi hotspots and low-power, low-range base stations in order to satisfy users' growing demands for context-aware services and performance. In addition to providing better service, such capillary infrastructure deployment threatens users' privacy with respect to their social ties and communities, as it allows infrastructure owners to infer users' daily social encounters with increasing accuracy, much to the detriment of their privacy. Yet, to date, there are no evaluations of the privacy of communities in pervasive wireless networks. In this paper, we address the important issue of privacy in pervasive communities by experimentally evaluating the accuracy of an adversary-owned set of wireless sniffing stations in reconstructing the communities of mobile users. During a four-month trial, 80 participants carried mobile devices and were eavesdropped on by an adversarial wireless mesh network on a university campus. To the best of our knowledge, this is the first study that focuses on the privacy of communities in a deployed pervasive network and provides important empirical evidence on the accuracy and feasibility of community tracking in such networks.

1 Introduction

Every day, mobile operators collect large amounts of users' data that is mined for commercial and performance goals, such as billing, throughput, coverage and usage statistics. In addition to the explicit information (such as cost, duration, location) that can be derived from the communications, operators and infrastructure owners are able to gain additional knowledge based on the communication and contextual patterns, without any action from the user for this regard [20,23].

* The co-author was with EPFL when this work was accomplished.

Users' home/work locations [20,23], activities [31], interests [33] and social networks [10,30] can be inferred from their location and social interactions, much to the detriment of not only their own privacy, but also to that of their peers.

More recently, telecom manufacturers have also added support for seamless, low-cost, wireless device-to-device communications, such as Nokia Instant Community [37], AirDrop by Apple [2] and FlashlinQ by Qualcomm [9], thus complementing existing infrastructure-based communications. The possibility of real-time data sharing among devices, without the need for infrastructure, enables people to form localized and short-lived groups or *communities* of users, which can emerge in scenarios where the infrastructure is inadequate, expensive, untrusted or hostile [36,16]. Although still an emerging research subject in the wireless domain [39], pervasive communities and their structured networks of interactions are able to significantly improve the performance of opportunistic networks [26,8], by leveraging on the structural properties and patterns of the evolving user interactions. In the literature, there are several routing and packet-forwarding algorithms [25,27,8] that exploit the underlying evolving social interactions to improve the network performance, mostly based on the frequency of recorded Bluetooth encounters. Similarly, social communities have been studied from the behavioral perspective [10,14,19], in order to analyze people's preferences and group formation characteristics. The undoubted value of friendship networks and social ties to service providers such as Facebook and Twitter has also dramatically increased their monetary value [7], as more and more targeted advertisements and tailored services are being proposed to groups of users with similar attitudes and interests.

In spite of the soaring interest for the analysis and exploitation of pervasive communities in the wireless domain, in regard to privacy very little has been achieved. Privacy of communities and their members is a major concern in regions where the ability to keep such information from being inferred by unscrupulous third-party providers or suppressive governments is critical [36,16]. Furthermore, the increased availability of public WiFi hotspots and the rapid deployment of low-power and low-range cellular base stations (femtocells) [15] makes such inference even more accurate, as more precise user proximity data can be collected, regardless of the kind of upper-layer protocols and applications. The risks of unsolicited user profiling, data censorship, racial discrimination and political repression, based on users' physical proximity derived from short-range communications, are a major concern. Because most of the existing literature on communities in wireless networks has been primarily focused on performance or human behavior, to the best of our knowledge there is no single empirical work that has addressed the issue of the privacy of communities in deployed wireless networks.

In this paper, we address the problem of community privacy by taking a comparative analysis of the exposure of social relationships and encounters in a deployed wireless peer-to-peer (P2P) network. Over a four-month trial (March-June 2011) with 80 participants, we studied and quantified the extent of leakage of private community information by users, by providing empirical evidence

about the network or infrastructure owner’s accuracy of reconstruction of the social communities of people. Our work is unique in three respects:

- We provide the first privacy analysis of the extent of exposure of community information in a deployed wireless network.
- We experimentally evaluate and compare the wireless sniffing stations owner’s accuracy of reconstruction of the social communities of people, based on the observed traffic patterns, with the local proximity and encounter data that is collected by the mobile devices.
- We characterize the evolution of the social interactions among the participants and evaluate the strength of their interactions by implementing three different social interaction measures that take into account the number, the proximity, the recency and aging effects of social relationships in the underlying wireless network.

The remainder of the paper is organized as follows. In Section 2 we introduce and detail the trial framework, its system and network models, whereas in Section 3 we outline the community and privacy analysis. In Section 4 we present the results of the analysis of communities and their privacy *vis-à-vis* the external adversary. We discuss the related work in Section 5. We conclude the paper and suggest ideas for further work in Section 6.

2 Trial Setup

During four months (March-June 2011), we conducted a large-scale trial with 80 participants on the EPFL university campus, in order to collect encounter and proximity data. Similarly to previous data collection campaigns [18,22,13], we programmed and distributed 80 Nokia N900 smartphones to the volunteering participants, sampling a coherent population of master’s students and instructors of two classes taught during the spring semester. The participants were asked to carry their device with them as frequently as possible, and they were allowed to use it as their primary phone. The complete description of the goals and methods of the questionnaires and interviews is described in [1]. At the end of the trial, we obtained useful information from 66 devices, amounting to almost ten GBs of collected log data and over 8 million packets captured by the adversarial network. The remaining 14 devices were either not used regularly or did not collect the data properly, hence they were excluded from the analysis.

2.1 Device Configuration

The Nokia devices were configured with both standard infrastructure-based communications, such as cellular and WiFi, as well as with a novel WiFi-based P2P technology, called *Nokia Instant Community* or NIC [37]. Users could connect to both standard Internet services using the WLAN or cellular interface of the device, as well as to an experimental context-aware wireless P2P messaging platform – in order to exchange information with their physical neighbors in a P2P




Fig. 1. Trial setup and deployed eavesdropping network controlled by the adversary.

fashion (Fig. 1(a)). Moreover, several campus and course-related applications were developed in order to stimulate and encourage the usage of the devices throughout the duration of the trial. In order to enhance the context-awareness of the pre-installed applications, the devices were running background services that collected and stored, at regular intervals of [1-30] seconds, information such as the list of neighbors, the associated Received Signal Strength Indicator (RSSI) and the time stamp in the local memory. Whenever a participant connected to the Internet with the device, the new encounter logs were uploaded on a centralized database storing all device logs. To preserve users' anonymity, we removed all personal identifier information (such as the mapping between MAC address - IMEI - participant ID) from the database.

2.2 Adversarial Model and Infrastructure

We emulate a practical adversary who monitors a fixed area using a limited number of wireless sniffing stations. Specifically, the adversary is the owner of a deployed wireless mesh network of 37 APs (Asus WL-500gP APs running OpenWRT Linux) in a specific region of the campus [3], covering one level of six interconnected buildings which have a very high user (student) density (Fig. 1(b)). The coverage area includes the classrooms in which the two classes that the students attended took place. We assume that the adversary passively eavesdrops on the participants' communications, and that he⁶ periodically uploads the eavesdropped data to a centralized server, populating a unified log database for each AP.

In order to perform the pervasive community reconstruction attack discussed in the following section, we assume that the adversary collects the 3-tuple (Time stamp, Source MAC, RSSI) from the messages sent by the participants' smartphones. As encryption is sometimes used to protect the confidentiality of network and application-layer data in real networks, we assume that the adversary does not have access to such data. This reinforces the practicality and better embodies real-world limitations that an external adversary might have, being much

⁶ For conciseness and without loss of generality, we refer to the adversary in the masculine form, although both masculine and feminine forms apply.

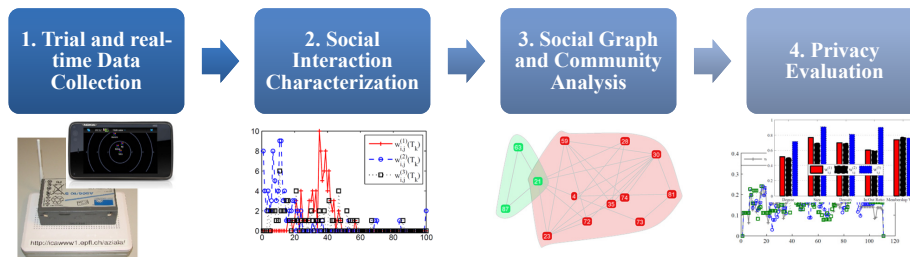


Fig. 2. Flowchart of the pervasive community privacy evaluation process.

weaker than the omniscient Dolev-Yao adversary [12]. Moreover, the information collected by the adversary is present in almost all kinds of wireless networks and technologies (such as Bluetooth, WiFi and cellular), which enlarges the applicability and scope of the results. In this work, we assume that the adversary does not have direct access to any information stored on the mobile devices, and that all devices are honest (i.e., not colluding with the adversary). As part of our future work, we will consider a stronger adversary that can collude and gain access to some of the mobile devices as well.

3 Community Analysis

In order to evaluate the extent of community information leakage in our setting, we first need to define the analytical framework that captures the pervasive community information from the collected data. In this section, we introduce some background on communities in wireless networks and describe how we evaluate communities and their privacy in our trial. A flowchart of the entire process is depicted in Fig. 2.

3.1 Background

In society, people tend to organize themselves in social groups or communities, such as family, work colleagues and hobby groups, where members usually have stronger similarity traits with other members than with non-members [17]. From a graph-theoretic perspective, people and their relationships can be represented by an undirected graph $G = (V, E, W)$, where the vertex set V corresponds to people, the edge set E expresses the existence of a relationship between people, and the weight function W quantifies the intensity of such relationship. In their simplest form, communities can then be represented as subgraphs $\{C_i = (V_i, E_i, W_i,)\}_{i=1}^M$, where $C_i \subseteq G$ and M is the number of communities C_i .

Several community detection (or clustering) algorithms are present in the literature, and they work on either unweighted/weighted and undirected/directed graphs. Although hierarchical clustering [21] and modularity-based algorithms [32] – surveyed in [17] – have been applied to community detection, most of

them lack a fundamental characteristic that is intrinsic to social communities. People are often members of several communities at the same time, such as friends, family members and work colleagues, and most of the aforementioned algorithms assign a single vertex to only one community. In order to allow a vertex to be assigned to multiple (possibly overlapping) communities, Palla *et al.* [35] developed a technique, the *Clique Percolation Method (CPM)*, which allows different communities to share vertices. The idea is that communities are formed by the union of adjacent k -cliques (complete graphs with k vertices), where two k -cliques are adjacent if they share $k - 1$ vertices. Due to the social nature of our trial and the experimental setting, we use the CPM algorithm to detect pervasive communities based on physical proximity and encounter data.

After the pervasive communities have been discovered, several privacy-sensitive statistics can be obtained from the community structure, their overlap and their members. We describe the relevant statistics in Section 3.3.

3.2 Trial Framework

In order to model the collected encounter data using a graph, hereafter we describe the type of information that is used in order to define the existence and intensity of relationships between users.

Trial Data In our trial, we have two sources of proximity information: (i) the local device logs collected by the mobile devices and containing encounter (list of neighbors, the time stamps and the RSSI values of received packets), and (ii) the adversarial (sniffing) logs containing the headers of the packets sent by the mobile devices, which include the time stamps and RSSI values of received packets at the sniffing stations, as well as the device ID of the sender.

We use these two data sources in order to formulate the “strength” or intensity of the social relationships between users and to define the weights of the edges connecting the respective vertices in the social graph $G = (V, E, W)$. There are two types of proximity information in our network: device-to-device RSSI data (collected on the devices) and device-to-AP RSSI data (collected by the adversary). From the local device logs, we can directly obtain the device-to-device proximity information because the recorded RSSI values on the receiving device depend on the real distance to the sending device. However, this is not exactly the case for the RSSI values recorded by the adversarial network, as they depend on the distance between the sending device and the receiving sniffing station, and not the receiving mobile device. Therefore, the adversary needs to derive the device-to-device proximity information from the device-to-AP RSSI values. Hence, we first need to estimate the position of a device, and then compute the device-to-device proximity information in order to determine the weights between vertices of the social graph.

To this end, we developed a robust localization algorithm based on RSSI trilateration [5], which determines the estimated position of a received packet based on the RSSI at all sniffing stations that received that packet. Using the position

estimate, we then compute the distance and RSSI between mobile devices, as described later in this section.

Social Interaction Intensity We define three distinct weight functions $\{w_{i,j}^{(d)}\}_{d=1}^3$ between the vertices $i, j \in V$, taking progressively into account the proximity, the intensity and the aging and recency of the relationships between users. We divide the timeline of the trial into discrete time intervals $\{T_k\}_{k=1}^N$, where $N = 120$ days, and for each day T_k we define the weights $w_{i,j}(T_k)^{(d)}$ between users i, j .

The first and simplest weight function is the (shifted, non-negative) average of the RSSI value between a pair of users i, j for each day T_k , defined as

$$w_{i,j}(T_k)^{(1)} = \left(\frac{1}{c_{i,j}(T_k)} \cdot \sum_{q=1}^{c_{i,j}(T_k)} RSSI_{i,j}(T_k, q) \right) - r_{min}$$

where $c_{i,j}(T_k)$ is the sum of the number of packets received by i (and sent by j) and received by j (and sent by i) during the day T_k , $RSSI_{i,j}(T_k, q)$ is the RSSI value of a packet q received by a user i (and sent by j) or received by j (and sent by i) during the day T_k , and r_{min} is the minimum RSSI value that was recorded during the trial. For instance, we fix $r_{min} = -100$ dBm as no RSSI values lower than -100 dBm have been recorded by any device. Apart from the intensity, this weight function does not consider the duration of the encounters (as it normalizes the intensity by the number of packets) between users or any aging or recency effect.

The second weight function takes into account the duration of the encounters through the sum of the (shifted, non-negative) RSSI values between users i, j , for each day T_k . It is defined as

$$w_{i,j}(T_k)^{(2)} = c_{i,j}(T_k) \cdot w_{i,j}(T_k)^{(1)} = \sum_{q=1}^{c_{i,j}(T_k)} (RSSI_{i,j}(T_k, q) - r_{min})$$

As the devices who are in continuous radio contact automatically exchange more context messages than the non-connected devices, this weight function takes into account the duration of the contacts, in addition to their intensity.

As communities of mobile devices are dynamic and evolve over time, the third weight function captures the natural evolution of social relationships between individuals, where past experience, recency and current state determine the intensity of interactions among people [34]. In this way, two users that have spent much time together in the past, but have not met on a given day, would still keep a relationship during that day (which is not the case for $w_{i,j}(T_k)^{(1)}$ and $w_{i,j}(T_k)^{(2)}$), even if its intensity is lower due to the aging effect – thus avoiding strong temporal fluctuations. Inspired by the formulations in [34,38], we define the third weight function as

$$w_{i,j}(T_k)^{(3)} = \mathbf{1}_{c_{i,j}(T_k) > 0} \left(\tau \cdot w_{i,j}^{(3)}(T_{k-1}) + (1 - \tau) \cdot \gamma_{i,j}(T_k) \right) \quad (1) \\ + (1 - \mathbf{1}_{c_{i,j}(T_k) > 0}) \cdot \left(w_{i,j}^{(3)}(T_{k_e}) \cdot \theta_{i,j}(T_k, T_e) \right)$$

where

$$\begin{aligned}\gamma_{i,j}(T_k) &= \frac{1}{\alpha} \cdot w_{i,j}(T_k)^{(2)} \\ \theta_{i,j}(T_k, T_e) &= \exp\left(-\frac{\lambda(T_k - T_e)}{1 + \sum_{r=0}^{\min(T_k - T_e, T_e)} m_{i,j}(T_e - r)}\right) \\ m_{i,j}(T_k) &= \begin{cases} 1 & \text{if } \gamma_{i,j}(T_k) > \beta \\ 0 & \text{otherwise} \end{cases}\end{aligned}$$

and $\mathbf{1}_{c_{i,j}>0}$ is the indicator function, $0 \leq \tau \leq 1$ is the aging coefficient, $\alpha > 0$ is the normalization factor, $0 < \lambda \leq 1$ is the temporal decay value, $0 \leq T_e \leq T_{k-1}$ is the last day before T_k when users i, j exchanged messages, $m_{i,j}(T_k) \in \{0, 1\}$ is the recency factor that indicates whether a meeting took place during T_k or not, and $\beta \geq 0$ is the meeting threshold value. The idea behind the formulation is the following: If users i, j exchanged at least one message on a day T_k , then the weight of their edge is an exponential moving average of the aged weight – accumulated up to the day before (T_{k-1}) – and the recent day’s weight; on the contrary, if i, j did not exchange any message on day T_k , the current day’s weight is a function of the previously accumulated weight, the frequency of their encounters just before the last encounter and the amount of time between the last time i, j had exchanged messages (T_e) and the current day T_k .

The weight functions can be directly applied to the local-device proximity information, as the available proximity information (time stamps, RSSI values from neighboring devices and their IDs) are sufficient for their computations. However, an intermediate step is required in order to compute the weights by using the external (adversarial) proximity information (time stamps, RSSI values from devices to sniffing APs and device IDs). In the following we show how to use the external proximity information in order to compute the edge weights.

User-Distance Estimation by the Adversary As the adversary does not have access to device-to-device proximity data, he can decide to only use the estimated positions of a user i in a day T_k , defined as $P_i(T_k) = \{p_i(T_k, 1), \dots, p_i(T_k, b)\}$, where b is the number of subintervals of a day T_k and $p_i(T_k, z) = (x_i(T_k, z), y_i(T_k, z)) \in \mathbb{R}^2$ is the estimated position of user i in the subinterval z of day T_k . Moreover, because there is a possibility that a user’s packet may not be detected in each subinterval z , due to mobility or radio interference, we assume that the last position estimate $p_i(T_k, z_{last})$ of a user i is valid in f subsequent subintervals, if no $\{p_i(T_k, z_{last} + 1), \dots, p_i(T_k, z_{last} + f)\}$ are available (Fig. 3).

With such information, the adversary computes the edge weights as follows:

- (1) $\forall z \in \{1, \dots, b\}$, compute $p_i(T_k, z)$ for all users i observed on day T_k .
- (2) $\forall z \in \{1, \dots, b\}$, compute the estimated Euclidian distance $d_{i,j}(T_k, z) = \|p_i(T_k, z) - p_j(T_k, z)\|$ between any two users i, j observed on day T_k .




Fig. 3. Users' positions estimates by the adversary. In this example, the adversary has the position estimate of user u at $z = 1$ but not at subsequent subintervals. In this case, u 's last position estimate (at $z = 1$) is assumed to be valid in f subsequent subintervals. Here $f = 1$.

- (3) $\forall z \in \{1, \dots, b\}$, compute the estimated RSSI value according to the adapted Haka-Okumura model for indoor radio propagation [6]

$$\widehat{RSSI}_{i,j}(T_k, z)[dBm] = P_t + 20 \log \left(\frac{\lambda}{4\pi} \right) + 10n \log \left(\frac{1}{d_{i,j}(T_k, z)} \right)$$

where $P_t = 20$ [dBm] is the transmission power of the mobile device, $\lambda = 0.125$ [m] is the wavelength, $n = 4.8$ is the path-loss exponent suited for office environments such as the university buildings under observation. The $\widehat{RSSI}_{i,j}(T_k, z)$ value replaces $RSSI_{i,j}(T_k, q)$ in the weight functions $w_{i,j}^{(d)}(T_k)$, where $z \in \{1, \dots, b\}$.

Weight Distributions Due to the different features of a social relationship that each weight function models, their numeric values fall in different domains. For example, if $\alpha = 100$, $\beta = 1$, $\lambda = \tau = 0.5$ we have $0 \leq w_{i,j}^{(1)} < r_{min}$, $0 \leq w_{i,j}^{(2)} < 2.5 \cdot 10^5$ and $0 \leq w_{i,j}^{(3)} < 600$. It is therefore necessary to put them on the same scale for the identification of communities, as simply comparing the absolute values of the three weight functions is pointless. Hence, rather than comparing absolute values, we compare the weight distributions relative to the maximum of each weight function for each day T_k . To this end, we select an equal number of bins $I^{(d)}$ for each weight function $w_{i,j}^{(d)}(T_k)$. We then count the number of weight values that fall inside each such bin for all weight types, and we compare the distributions.

Fig. 4(a) and Fig. 4(b) show the relative edge weight distribution for a day T_k , by using the internal (local device) and external (adversarial estimate) input data, respectively. We see that, compared to the adversarial data, the local device data yields more pronounced characteristics for all three weight types and provides a more discriminating information set for the subsequent community detection phase, whereas the external data is less feature-rich due to the presence of uncertainty in the estimates of the proximity between users. This means

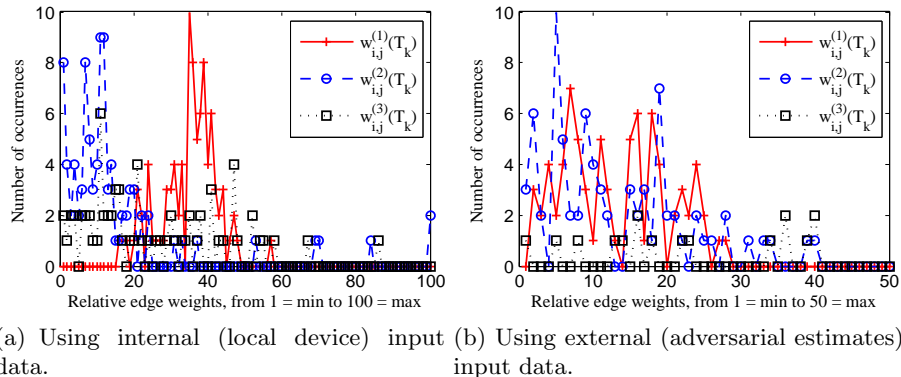


Fig. 4. Relative edge weight distribution for different input data sets.

that the adversary will likely struggle to infer with high accuracy the community characteristics for that day. We quantify such inaccuracies in Section 4.

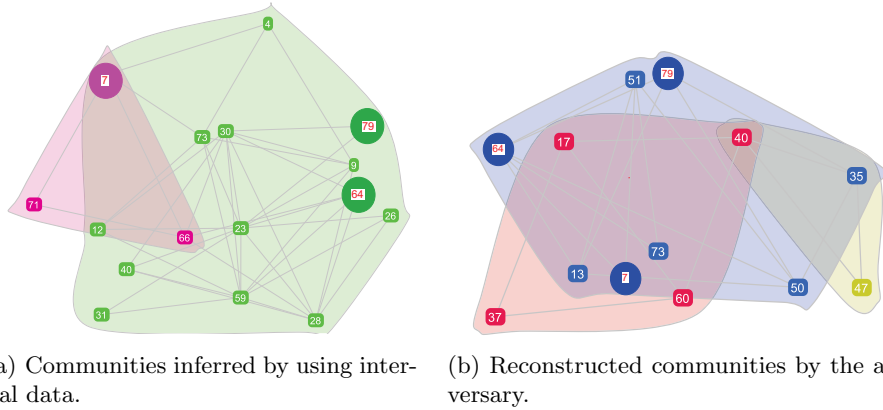
Next, we describe the method we adopted to evaluate the extent of community information leakage and the related privacy measures.

3.3 Communities and Privacy

Having quantified the social interaction intensity as edge weights between any two trial participants, we now outline the community detection process, the suitable community statistics and privacy measures used to evaluate community privacy in our work.

Community Detection In its simplest form, the CPM community detection algorithm is defined for undirected and unweighted graphs [34], thus requiring only connectivity between vertices in order to discover communities. However, in order to consider the “strength” of the interactions between vertices, it was extended to work on weighted graphs by the use of a threshold weight w^* . In its weighted version, the CPM algorithm considers the existence of an edge $e_{i,j}$ between two vertices i, j if and only if the weight $w_{i,j}^{(d)} > w^*$. In order to determine the threshold weight w^* , Palla *et al.* propose to choose a value such that “the largest community becomes twice as big as the second largest one” [35], which is below the critical value w_{crit}^* for which a giant connected component arises [11].

In our experiment, we calibrated the $\{w_q^*\}_{q=1}^T$ threshold values on a per-day basis, instead of keeping the same w^* throughout the trial. Because most of the participants followed one specific class that took place on Wednesdays, and the remaining days they might or might not have followed any common classes, we registered high RSSI proximity values on course days and more sparse values on non-course days. Hence, the per-day threshold $\{w_q^*\}_{q=1}^T$ was better suited for such bi-modal proximity patterns.



(a) Communities inferred by using internal data. (b) Reconstructed communities by the adversary.

Fig. 5. Detected communities on a day T_k based on internal (local device) and external (adversarial estimates) data, respectively. The larger vertices are present in both community sets.

To illustrate the output of CPM, Fig. 5(a)⁷ and 5(b) show an example of the detected communities on a given day, based on the internal data and the observations of the adversary respectively. As it can be seen, some communities detected by the adversary are not present in the internal case; there is however an overlap between the members (the larger vertices) of the two sets of communities. We discuss and quantify this difference in Section 4.

Community Statistics In addition to detecting communities and their members, we compute five privacy-relevant and common community statistics $\{S_{(i)}(T_k)\}_{i=1}^5$ that will be compared in the accuracy evaluation process. In particular, for each day T_k of the trial we compute and compare the following statistics: S_1 is the community degree (the number of edges shared between two communities), S_2 is the distribution of the community size (the number of members of each community), S_3 is the community density (proportion of edges out of all possible edges relative to the sparsest set with $|C_i| - 1$ vertices), S_4 is the ratio of total out- and in-degree of communities and S_5 is the community membership value (the number of communities a vertex belongs to). The difference between the results obtained using the internal and external input data is defined by Eq. (2) as the ratio between the absolute difference of the observed statistics over the maximum value

$$\Delta S_i(T_k) = \frac{|S_i^{ext}(T_k) - S_i^{int}(T_k)|}{\max_{\forall T_k} (S_i^{ext}(T_k), S_i^{int}(T_k))} \quad (2)$$

We have $\Delta S_i(T_k) = 0$ when the adversary's statistics is exactly the same as the statistics obtained using the internal proximity data, and $\Delta S_i(T_k) = 1$ when the

⁷ The figure is obtained by using the CFinder application developed by the authors of the CPM algorithm, freely available on www.cfinder.org.

two statistics have the largest discrepancy (or lowest similarity). We define the adversary’s accuracy in inferring the community statistics as $1 - \Delta S_i(T_k)$.

Community Privacy In addition to the differences in statistics $\Delta S_i(T_k)$, it is crucial to assess the similarity of the community composition in order to ascertain in a comprehensive way the privacy leakage of community information. To this end, we compute the well-established Jaccard index measure [28] for community similarity on each day T_k , which is a statistic that computes the similarity between two sample sets (or communities) C_i, C_j , where values close to zero mean that the adversary did not accurately infer the communities and their members, whereas values close to one indicate a very good adversarial accuracy in inferring the same communities. The Jaccard index is defined as

$$J(C_i, C_j, T_k) = \frac{|C_i(T_k) \cap C_j(T_k)|}{|C_i(T_k) \cup C_j(T_k)|} \quad (3)$$

In order to evaluate the adversary’s accuracy of reconstruction of the communities in our pervasive network, we compute the Jaccard index on each day T_k between the communities $C_i(T_k)$, detected using internal device data, versus the reconstructed communities $C_j(T_k)$, detected using the adversarial estimated proximity information. Given $J(C_i, C_j, T_k)$ for each i, j on a day T_k , we define the Jaccard index matrix $JMat(T_k)$, where each element of the matrix is defined as $JMat(T_k)_{i,j} = J(C_i, C_j, T_k)$, i.e., the Jaccard index for all pairs of communities C_i and C_j . Without having access to the internal data, the adversary has no prior knowledge about which community C_i corresponds to which reconstructed community C_j . Therefore, in order to consider the best possible match for any pair of internal/reconstructed communities for each day T_k , we choose the match $(C_i(T_k), C_j(T_k))$ that maximizes $JMat(T_k)_{i,j}$. We then compute the aggregated Jaccard index over all such best matches as

$$JI(T_k) = \text{avg}_{\forall i} \left(\max_{\forall j} (JMat(T_k)_{i,j}) \right) \quad (4)$$

for each day T_k of the trial where there is at least one community detected by using both the internal and adversarial proximity information.

In the next section we quantify the community privacy leakage by computing the accuracy measure $1 - \Delta S_i(T_k)$, and similarity $JI(T_k)$ for each day T_k and weight function $\{w_{i,j}^{(d)}\}_{d=1}^3$, comparing the results obtained using the internal (local device) and external (adversarial) input data respectively.

4 Privacy Evaluation

In this section we provide the experimental evaluation of the privacy of pervasive communities through a comparative analysis of the adversary’s accuracy of reconstruction of both community statistics and memberships. First, we evaluate the privacy across the three weight functions $\{w_{i,j}^{(d)}\}_{d=1}^3$ (inter-weight accuracy), by comparing the similarity between communities and the accuracy of




Fig. 6. Adversary’s accuracy of reconstruction of the pervasive communities for the three weight functions.

their statistics obtained by using the internal (local device) proximity information with the external (adversarial estimates) data collected by the set of wireless sniffing stations. This will allow us to observe the evolution of the accuracy while increasing the sophistication of the weight functions, taking progressively into account several features of human and social behavior such as proximity, intensity, aging and recency of social relationships. Second, we perform an intra-weight comparison for the more realistic weight function $w_{i,j}^{(3)}$, in order to characterize the effect of the aging factor τ on the similarity and accuracy of community reconstruction attained by the adversary.

Fig. 6 and 7 show the adversarial reconstruction similarity and accuracy results with respect to the communities detected using internal data, for the inter-weight and intra-weight scenarios respectively. For Fig. 6(a) and 7(a), a value of $JI(T_k) = 0$ means that on day T_k there were no communities detected either using the internal proximity data or the external one. The complete list of the experimental parameters – selected in order to provide as much information as possible – can be found in the Appendix, which is provided as a supporting file to this document.

4.1 Inter-Weight Accuracy

By observing Fig. 6(a), we first notice that the adversary is able to correctly reconstruct communities and identify their members in 20% – 40% of the cases, compared to the communities detected by using internal proximity data. In general, we observe that there is a significant difference in terms of similarity results between the first two weight functions $w_{i,j}^{(1)}, w_{i,j}^{(2)}$ and the third function $w_{i,j}^{(3)}$. The former two functions are solely based on the observations made on each particular day and independently of what happened in the previous days. Therefore one noticeable characteristic is the increased fluctuations in the similarity from one day to the other, which is a much less visible aspect for the latter weight function. As $w_{i,j}^{(1)}, w_{i,j}^{(2)}$ are very exposed to the periodicity of the course schedule of the participants, the adversary’s similarity of reconstruction of the

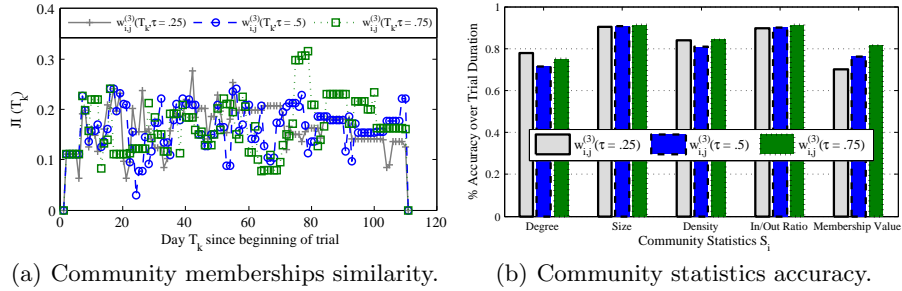


Fig. 7. Adversary’s accuracy of reconstruction of the pervasive communities for three different values of the aging factor τ .

actual communities and their members greatly depends on the amount of data collected by his wireless mesh network. We notice that for the days when most students attended a particular class, the reconstruction similarity is higher (up to 40%) than for days in which students do not attend classes together. Hence even the two basic weight functions are able to provide a sensible similarity to the adversary when the users’ movements are tracked by several sniffing stations simultaneously.

Contrary to $w_{i,j}^{(1)}$ and $w_{i,j}^{(2)}$, $w_{i,j}^{(3)}$ is able to capture more proximity information and allow the CPM algorithm to detect communities on the days in which the other two weight functions were unable to provide a sufficient amount of data. At the same time, however, the peaks of similarity tend to be much lower (25%) compared to the other functions. This suggests that $w_{i,j}^{(3)}$, while being able to produce more community information with scarce data, performs worse in the identification of the members in each community.

Regarding the difference in community statistics, depicted in Fig. 6(b), we observe a better accuracy for $w_{i,j}^{(3)}$ compared to $w_{i,j}^{(1)}$ and $w_{i,j}^{(2)}$. In four out of five community statistics, $w_{i,j}^{(3)}$ has an almost 40% better accuracy compared to the other functions, which indicates that the former function provides better results on a higher structural community level rather than on an lower, individual community member level.

In general, we observe that all three weight functions are better able to produce accurate community statistics (Fig. 6(b)) than to identify the correct community members (Fig. 6(a)). In particular, $w_{i,j}^{(3)}$ shows that it is possible to achieve very accurate community statistics only by relying on externally collected data, thus shrinking the discrepancy between the community statistics based on internal data and adversarial’s estimates down to 9%. This result indicates that, by collecting and analyzing radio information passively and without access to the devices themselves, an adversary is able to breach the privacy of community information very successfully, although the more fine-grained identification of members of any given community remains a more challenging task.

4.2 Intra-Weight Accuracy

Fig. 7(a) shows the adversary’s performance in correctly identifying the communities and their individual members when using $w_{i,j}^{(3)}$ with three different values of the aging factor $\tau = \{.25, .5, .75\}$. According to its definition in Eq. (1), we assign an increasing coefficient to the past accumulated weight information $w_{i,j}^{(3)}(T_{k-1})$ in the computation of the current day’s weight function $w_{i,j}^{(3)}(T_k)$. The goal is to study the effect of the “retention” of the intensity from the past on the privacy (or lack thereof) of community information.

One recurring characteristic, present also in the inter-weight comparison, is that the CPM algorithm detects communities in all days of the trial, independently of the amount of information available to the adversary on each particular day, even for a small value of τ . When $\tau = .25$, as expected the similarity fluctuates more when compared to $\tau = .5$, especially at the beginning of the trial.

However, Fig. 7(a) shows that the stabilization of the similarity is not achieved by simply increasing the value of τ from .25 to .75; in fact, for the intermediate value of $\tau = .5$, we notice that the fluctuations are less pronounced than for a smaller or larger value. This suggests that, for relatively small or large values of the aging factor, the similarity achieved by the adversary tends to diverge more frequently from steady values, indicating that a stable value for the aging factor is more likely to be in the middle of the possible values [0.25,0.75], rather than at any of the extremes. When $\tau = .75$, the adversarial similarity increases sharply as the time passes, especially towards the end of the trial. This is somewhat surprising, as we would expect that by increasing the emphasis on the past – rather than on the current weight information – the similarity would be more stable when going through the trial. This is an interesting aspect to consider in further studies on our community data.

When observing the results on the accuracy of the community statistics, as shown in Fig. 7(b), we notice that, among the three considered values of τ , $\tau = .5$ is the least accurate, compared to smaller or larger values of τ . Moreover, in four out of five statistics, the largest value of $\tau = .75$ produces the best accuracy on average over the trial duration. This suggests that, although not converging towards a stable interval for the accuracy in identifying the communities and their members, putting more emphasis on the past accumulated information does increase (on average) the adversary’s accuracy in computing correct community statistics using only passively collected data from fixed WiFi access points.

Overall, the results indicate that although less stable and more accurate at inferring community structures, emphasizing the past yields better accuracy for both community detection, identification of their members and for generic community statistics. This finding in particular is concerning in regard to privacy, as the amount of individual and community data that is collected by external parties might provide very accurate statistics, especially for group and community-targeted services. These results are significant, as they show how the message source ID, contained in almost any kind of radio message, not only is enough to

provide accurate social community statistics, but it is also sufficient to successfully infer almost half of the members of such communities.

5 Related Work

The structural properties of short-lived communities in pervasive networks have been recently investigated from the performance [26,8] and routing [25,27,8] perspectives; the authors of [10,14,19] investigated similar issues on the socio-behavioral level while studying people’s preferences and group formation characteristics. For instance, it is shown that performance of packet-forwarding algorithms could greatly benefit from the human mobility and sporadic nature of inter-contacts [26], as the different connection frequencies between members of the same community with respect to members of other communities could significantly improve intra-community packet-forwarding while not disrupting inter-community communications. Similarly, [27] shows how forwarding performances similar to state-of-the-art algorithms could be achieved at a sensibly lower resource utilization if structural properties of communities are considered.

With respect to privacy, several works on location privacy address the risk and propose protection mechanisms for users’ locations [4,24,29]. These contributions focus mostly on individual mobile users and their current neighbors. However, to the best of our knowledge, there is no prior study on the increasingly important issue of pervasive community privacy and its evaluation on a deployed network. This work constitutes the first building block for analyzing community privacy issues in pervasive networks.

6 Conclusion and Future Work

In this paper, we have addressed the important aspect of community privacy in pervasive networks. We have conducted an experimental analysis of the adversary’s accuracy of reconstruction, on one hand, of the communities and their individual members and, on the other hand, of the generic community statistics that are less dependant on the correct identification of individual users inside such communities.

Through a fine-grained characterization of the intensity of social contacts among people, we quantified the accuracy in both community reconstruction and community statistics for the whole duration of the trial, showing that even basic social intensity functions capture very accurately the generic statistics, such as the degree of a community, its size and density of links. However, reconstructing more specific information about the composition of each community and their individual members remains more challenging, even when using a more comprehensive model for characterizing the intensity of social relationships, which considers recency, aging, and contact frequency in addition to proximity and duration. As a result, there is a substantial risk that accurate community information may be easily collected, inferred and misused by external third-parties, much to the detriment of users’ community privacy.

Our results provide empirical evidence about the two distinct levels of community information leakage to external observers, who may be able to infer with high accuracy the different social groups and generic communities of people in pervasive networks, while being much less accurate in determining the affiliation of any particular individual to a community. As part of our future work, we intend to pursue the analysis of this dual flow of community information leakage and derive mitigation mechanisms in order to reduce information leakage and the gap between the accuracy of both generic statistics and specific people’s affiliations to communities. We also intend to study the adversary’s accuracy of classification of the communities and their members based on the type of their relationship, such as friends, classmates, study group and strangers.

References

1. Imad Aad, Murtuza Jadliwala, Igor Bilogrevic, Valtteri Niemi, Jean-Pierre Hubaux, Philip Ginzboorg, and Kari Leppnen. Nokia Instant Community at EPFL: a real-world large-scale wireless peer-to-peer trial. Technical Report EPFL-REPORT-170421, 2011.
2. Apple AirDrop. <http://www.apple.com/macosx/whats-new/>.
3. Aziala-net. <http://icawww1.epfl.ch/aziala/index.html>.
4. A.R. Beresford and F. Stajano. Location privacy in pervasive computing. *IEEE Perv. Comp.*, 2, 2003.
5. Laurent Bindschaedler*, Murtuza Jadliwala*, Igor Bilogrevic, Imad Aad, Philip Ginzboorg, Valtteri Niemi, and Jean-Pierre Hubaux. Track Me If You Can: On the Effectiveness of Context-based Identifier Changes in Deployed Mobile Networks. In *NDSS*, 2012.
6. A. Bose and C.H. Foh. A practical path loss model for indoor wifi positioning enhancement. In *Int. Conf. on Inform., Comm. & Signal Proc.*, 2007.
7. Business Week. Facebooks value tops amazon.com; trails only google on web. <http://www.businessweek.com/news/2011-01-28/facebook-s-value-tops-amazon-com-trails-only-google-on-web.html>.
8. A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott. Impact of human mobility on opportunistic forwarding algorithms. *IEEE TMC*, 2007.
9. M.S. Corson, R. Laroia, J. Li, V. Park, T. Richardson, and G. Tsirtsis. Toward Proximity-aware Internetworking. *Wireless Communications*, 2010.
10. D.J. Crandall, L. Backstrom, D. Cosley, S. Suri, D. Huttenlocher, and J. Kleinberg. Inferring social ties from geographic coincidences. *Proc. Nat. Academy of Sciences*, 107, 2010.
11. I. Derényi, G. Palla, and T. Vicsek. Clique percolation in random networks. *Physical review letters*, 94, 2005.
12. D. Dolev and A. Yao. On the security of public key protocols. *IEEE TIT*, 29, 1983.
13. N. Eagle and A. Pentland. Reality mining: sensing complex social systems. *Pers. and Ubiqu. Computing*, 10, 2006.
14. N. Eagle, A.S. Pentland, and D. Lazer. Inferring friendship network structure by using mobile phone data. *Proc. Nat. Academy of Sciences*, 106, 2009.
15. Femto Forum. <http://femtoforum.org/fem2/pressreleases.php?id=277>.
16. M. Follman. “Bluetoothing” Iran’s revolution. *Markfollman.com*, 2010.

17. S. Fortunato. Community detection in graphs. *Physics Reports*, 486, 2010.
18. N.W. Gong, M. Laibowitz, and J. Paradiso. Dynamic privacy management in pervasive sensor networks. *Ambient Intelligence*, 2010.
19. MC González, HJ Herrmann, J. Kertész, and T. Vicsek. Community structure and ethnic preferences in school friendship networks. *Physica A: Statistical mechanics and its applications*, 379, 2007.
20. M. Gruteser and B. Hoh. On the anonymity of periodic location samples. *Security in Perv. Comp.*, 2005.
21. T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning (2nd Edition)*. 2008.
22. T. Henderson, D. Kotz, and I. Abyzov. The changing usage of a mature campus-wide wireless network. In *Int. Conf. on Mobile comp. and networking*, 2004.
23. B. Hoh, M. Gruteser, H. Xiong, and A. Alrabady. Enhancing security and privacy in traffic-monitoring systems. *IEEE Perv. Comp.*, 5, 2006.
24. J.I. Hong and J.A. Landay. An architecture for privacy-sensitive ubiquitous computing. In *Conf. on Mobile systems, applications, and services*, 2004.
25. T. Hossmann, T. Spyropoulos, and F. Legendre. Know thy neighbor: Towards optimal mapping of contacts to social graphs for dtn routing. In *INFOCOM*, 2010.
26. P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, and C. Diot. Pocket switched networks and human mobility in conference environments. In *ACM SIGCOMM workshop on DTN*, 2005.
27. P. Hui, J. Crowcroft, and E. Yoneki. Bubble rap: social-based forwarding in delay tolerant networks. *IEEE TMC*, 2010.
28. P. Jaccard. Etude comparative de la distribution florale dans une portion des alpes et du jura, 1901.
29. M. Jadhwal, I. Bilogrevic, and J.P. Hubaux. Optimizing mixing in pervasive networks: A graph-theoretic perspective. *ESORICS*, 2011.
30. S. Mardenfeld, D. Boston, S.J. Pan, Q. Jones, A. Iamntichi, and C. Borcea. Gdc: Group discovery using co-location traces. In *Int. Conf. on Social Comp.*, 2010.
31. Y. Matsuo, N. Okazaki, K. Izumi, Y. Nakamura, T. Nishimura, and K. Hasida. Inferring Long-term User Property based on Users. In *IJCAI*, 2007.
32. M.E.J. Newman. Fast algorithm for detecting community structure in networks. *Physical Review E*, 69, 2004.
33. A. Noulas, M. Musolesi, M. Pontil, and C. Mascolo. Inferring interests from mobility and social interactions. In *NIPS Workshop on Analyzing Netw. and Learning w. Graphs*, 2009.
34. G. Palla, AL Barabási, and T. Vicsek. Quantifying social group evolution. *Nature*, 446, 2007.
35. G. Palla, I. Derényi, I. Farkas, and T. Vicsek. Uncovering the overlapping community structure of complex networks in nature and society. *Nature*, 435, 2005.
36. Sam Reeves. Internet is double-edged sword in arab revolts. <http://middle-east-online.com/english/?id=46109>, 2011.
37. Rhiain. Nokia instant community gets you social. <http://conversations.nokia.com/2010/05/25/nokia-instant-community-gets-you-social/>.
38. K. Xu, G.H. Yang, V.O.K. Li, and S.Y. Chan. Detecting dynamic communities in opportunistic networks. In *ICUFN*, 2009.
39. D. Zhang, B. Guo, B. Li, and Z. Yu. Extracting social and community intelligence from digital footprints: An emerging research area. *Ubiq. Intell. and Computing*, 2010.