Self-organized Anonymous Authentication in Mobile Ad Hoc Networks

Julien Freudiger, Maxim Raya, and Jean-Pierre Hubaux

LCA1, EPFL, Switzerland
firstname.lastname@epfl.ch

Abstract. Pervasive communications bring along new privacy challenges, fueled by the capability of mobile devices to communicate with, and thus "sniff on", each other directly. We design a new mechanism that aims at achieving location privacy in these forthcoming mobile networks, whereby mobile nodes collect the pseudonyms of the nodes they encounter to generate their own privacy cloaks. Thus, privacy emerges from the mobile network and users gain control over the disclosure of their locations. We call this new paradigm *self-organized location privacy*. In this work, we focus on the problem of self-organized anonymous authentication that is a necessary prerequisite for location privacy. We investigate, using graph theory, the optimality of different cloak constructions and evaluate with simulations the achievable anonymity in various network topologies. We show that peer-to-peer wireless communications and mobility help in the establishment of self-organized anonymous authentication in mobile networks.

1 Introduction

The current model of wireless communication relies heavily on infrastructure: Two mobile phones have to go through cellular base stations to exchange calls and data for which users pay, even if they are only a few meters apart. But as more mobile devices become equipped with ad hoc (peer-to-peer) communication technologies, such as WiFi and Bluetooth, the coexistence of both peer-to-peer and infrastructure-based communications is inevitable. Moreover, the recent surge in mobile social networks [1,2] reinforces the need for mobile devices, such as phones, to be able to talk to each other without going through the infrastructure. These peer-to-peer communications enable *context-based* applications, such as dating [3], gaming [4], as well as distributed location-based services [43]. But these communications also make possible the continuous tracking of the location of these devices. Thus, whereas the standard privacy threat model focuses on protecting users with respect to the infrastructure (be it the cellular network or the Internet), pervasive communications will expand it to the whole set of mobile devices.

The promised ad hoc sharing of information might turn into a pervasive nightmare if undesired communications cannot be filtered out: For example, if mobile nodes cannot verify the source of information, they are susceptible to mobile

Y. Chen et al. (Eds.): SecureComm 2009, LNICST 19, pp. 350–372, 2009.

spam. To thwart rogue devices from polluting the network, nodes should authenticate each other: The existence of an authentication feature (and the implied procedure to obtain the appropriate credentials) makes it more difficult for attackers to join the network in the first place and thus increases the cost of misbehavior. Hence, by verifying the authenticity of their interlocutor before exchanging information, mobile nodes reduce the amount of undesired data. For example, users of context-based applications would obtain authentication credentials by subscribing to the service. They could subsequently verify that received messages were sent by other subscribers to the service. But if this is done without appropriate precautions, the authentication mechanism would then reveal the identity of the nodes, thus rendering the privacy problem particularly challenging.

The location privacy of mobile devices is guaranteed if and only if devices are anonymous and untraceable. Hence, our quest for location privacy in the upcoming generation of mobile computing becomes an attempt to devise mechanisms for *untraceable anonymous authentication*. More specifically, mobile devices must authenticate themselves directly to other devices without revealing privacy-sensitive information. In this paper, we assume that the authentication mechanism should not rely on the constant presence of a central authority because of the scalability and accessibility problems that this would cause. We also leave the privacy threat of authentication towards the infrastructure out of the scope of the paper. We show that the seeming disadvantage, privacy-wise, of peer-to-peer communications can actually be turned into an advantage, thus allowing each node to create its own privacy cloak without the need for a central privacy coordination service. We coin this new paradigm *self-organized location privacy*.

In this work, we focus on the analysis of anonymity as it is a prerequisite for untraceability: If nodes cannot be anonymous, they cannot be untraceable. The key enablers of our solution are the groups of users themselves and a cryptographic construction called *ring signatures* [34] that allows a node to authenticate itself to other nodes by using a ring of pseudonyms, instead of its pseudonym alone. This ring constitutes the *anonymity set* of the node and can be constructed out of the pseudonyms of the node's past and present encounters without any interactive protocols. Hence, our mechanism provides *self-organized anonymous authentication*. The advantage of this approach is that each user only owns a *single authenticated pseudonym*. But we show that rings alone are insufficient to protect user privacy: By analyzing the different pseudonyms used in rings, an eavesdropper can link - with a sufficiently high probability - some rings to users. As described in this paper, the problem gets worse if the network of nodes grows. Hence, it is crucial to construct rings using mechanisms that maximize user anonymity. We develop a graph-theoretic model to evaluate different *ring construction strategies* and derive the optimal (in terms of achieved anonymity) ring constructions. Leveraging on each node's local knowledge and history of encounters, we devise self-organizing methods to achieve, in practice, near-optimal anonymity. We show with simulations that mobility and peer-to-peer communications are beneficial for the emergence of self-organized location privacy.

The paper is organized as follows: In Section 2, we review the state of the art. In Section 3, we present the system and threat models assumed throughout the paper. After introducing our proposed solution in Section 4, we analyze in Section 5 the achievable anonymity using a graph-theoretic model and evaluate the solution in Section 6. Finally, in Section 7 we discuss the cost of our approach and present remaining challenges. In that section, we also provide preliminary results addressing the untraceability requirement before concluding in Section 8.

2 Related Work

There are several techniques available to achieve anonymous authentication.

A large body of work focuses on the use of multiple pseudonyms [15] and, in particular, in mobile scenarios [7,24]. Instead of using a single pseudonym, mobile devices are preloaded with a set of pseudonyms and change over time the pseudonym used for sending messages. To impede an adversary from linking old and new pseudonyms, the change of pseudonyms must be spatially and temporally coordinated among mobile nodes in regions called *mix zones* [8]. The analysis in [7,20,21] shows that the achieved location privacy depends on the node density and on the unpredictability of node movements in mix zones. The main drawbacks of mix zones is that they are inefficient when the node density in the mix zone is low and can be costly in terms of pseudonym management. A related technique uses frequently changing pseudonyms, silent periods, and power control to hide privacy-sensitive information [26]. As we will see, our approach alleviates the problem of low densities in mix zones by relying on the history of encounters of mobile nodes, instead of strictly using their current neighbors. In addition, we alleviate the problem of pseudonym management by allowing a single pseudonym per device.

Another solution relies on *group signatures* [16] that allow a group member to sign on behalf of a group without revealing the identity of the signer. Nowadays, highly efficient group signatures schemes exist with constant size signatures and efficient signing and verification even when nodes are revoked [10,13,31]. But group signatures require a group manager to add and revoke group members, thus making the flexibility of groups dependent on the availability and computational capacity of the group manager. In contrast, with ring signatures, nodes can change the members of their rings without central coordination.

Anonymous credential systems (e.g., Idemix [11]) allow mobile nodes to anonymously authenticate to third parties with the help of an online credential issuer. The online availability of a credential issuer is often not possible in wireless networks. To circumvent the issue, techniques based on unclonable identifiers, such as e-tokens [12], allow nodes to anonymously authenticate themselves a given number of times per period. However, such techniques lack flexibility, in particular in the case of a prolonged unavailability of the credential issuer.

To the best of our knowledge, we are the first to investigate the potential of ring signatures to achieve anonymity and untraceability in mobile networks. Until now, most of the work focused on proving properties of ring signatures [41] or on the anonymity and unlinkability of the signature generation process [28].

Recently, ring signatures were proposed in [29] as a building block for anonymous routing in MANET but without investigation of the ring creation process.

3 Preliminaries

3.1 System Model

We assume a mobile network with n mobile nodes and a single offline Certification Authority (CA) run by an independent trusted third party. We focus on scenarios where the mobile nodes are autonomous entities equipped with WiFi or Bluetooth-enabled devices that communicate with each other upon coming in range. In other words, we describe a pervasive communication system in which mobile nodes automatically exchange information upon meeting.

Prior to entering the network, each mobile node registers with the CA that preloads a single *public/private key* pair (K_i, K_i^{-1}) and a digital certificate in the nodes' device. The CA verifies the identity of each user upon registration. The public key K_i serves as the identifier of node i and is referred to as its *pseudonym* P_i. The private key K_i^{-1} permits mobile node i to digitally sign messages, while the digital certificate validates the authenticity of the signature.

We assume that mobile nodes automatically exchange information as soon as they are in communication range. To do so, mobile nodes advertise their presence by periodically broadcasting proximity beacons containing the node's authenticating information (i.e., the sender attaches its pseudonym to signed messages). When a node receives a beacon, it verifies the authenticity of the sender before reading the message.

3.2 Threat Model

We assume that a *passive* adversary $\mathcal{A}$ aims to track the location of mobile nodes. In practice, the adversary can be a rogue individual, a set of malicious mobile nodes or may even deploy its own infrastructure (e.g., by placing eavesdropping devices in the network). In the worst case, $\mathcal{A}$ obtains a complete coverage and tracks nodes throughout the entire network. We characterize this type of adversary as *global*.

$\mathcal{A}$ collects identifying information (i.e., pseudonyms) from the entire network and attempts to break the anonymity provided by ring signatures in order to track the location of mobile nodes. If the adversary is successful, it can implicitly obtain the true identity of the owner of a mobile node from the analysis of its mobility [27]. Hence, the *location privacy* of mobile nodes cannot be taken for granted.

Finally, we assume that the key-pair generation process cannot be altered or controlled by the adversary.

3.3 Problem Statement

The location of mobile nodes can be tracked based on the information leaked from authentication messages. To thwart this threat, we define the following design goals:

- **Anonymous authentication:** The nodes should be able to authenticate to each other without being identifiable. Anonymous authentication permits mobile nodes to verify the origin of received messages without revealing their identity (neither to the receiver, nor to an eavesdropper).
- **Self-Organization:** The anonymity of nodes should not depend on the constant presence of a central authority because of the scalability and accessibility problems this would cause (the CA distributes pseudonyms to nodes prior to their entrance in the network but is not always accessible). With self-organization, the cost of anonymity management is distributed among all the nodes.

It should be noted that we do not consider *accountability* as a design goal. Indeed, like many Internet applications, the peer-to-peer wireless scenarios we study do not require it.

4 Self-organized Anonymous Authentication

In this section, we describe the techniques that permit the emergence of self-organized anonymous authentication.

4.1 Overview

With standard asymmetric cryptography, nodes authenticate themselves to others by signing their messages with their private key and providing the public key for signature verification, thereby revealing their identity. Instead, self-organized anonymous authentication, explained in detail in the next section, allows a node to select a set of pseudonyms called a *ring* and then sign its messages with a *Ring Signature* (RS) [34]. A RS preserves the *cryptographic anonymity* of the signer because it cannot be distinguished among the members of the ring.[1] Besides, rings are setup-free: The knowledge of the pseudonyms of the other nodes is sufficient to create a ring without any interaction. Hence, unlike group signatures, RSs have no group managers and do not require any coordination among ring members. Finally, two signatures generated by the same signer with the same ring are *cryptographically unlinkable*. Of course, to be able to generate a RS, each node must always use its own pseudonym in its ring, thus guaranteeing the authentication requirement.

The pseudonyms used for constructing rings can be collected by downloading sets of rings from online databases, much like PGP keyrings, or, in the case of the mobile network considered here, by recording the pseudonyms of neighboring nodes in a *history* S_i. Each node constructs a ring of pseudonyms by selecting a subset of pseudonyms from its history of encounters. This allows nodes to have an anonymity set without any central coordination: Rings are *dynamically* and

[1] In this paper, anonymity and untraceability are evaluated with respect to the pseudonyms used in rings and not with respect to the signature generation process, thus the distinction "cryptographic".

independently created by mobile nodes. A node i can thus authenticate itself to other nodes at time t by sending a message m with a ring signature $RS_{i,t}(m)$ created with ring $R_{i,t}$. It is worth making a clear distinction between the notions of "node" i, "pseudonym" P_i and "ring" $R_{i,t}$. Mobile nodes are indexed by a counter i (that does not refer to any ordering of the nodes). A mobile node i is represented in the network by its pseudonym P_i. In order to avoid being tracked by its pseudonym, i actually uses a set of pseudonyms of other nodes together with its own pseudonym P_i, to create its ring $R_{i,t}$ at time t.

Authenticating the source of information is a crucial primitive in pervasive communication systems to limit the spread of undesired data. By signing a message with a ring of pseudonyms, a signer proves its membership to a club of nodes (e.g., a mobile social network). The verifier can then be sure that a message originates from a member of the club. Of course, all members of a ring should have the appropriate credentials - the pseudonyms have been certified by the CA - and belong to the club; otherwise their presence in the ring would invalidate its authenticity. For simplicity of presentation, we consider in the rest of the paper that there is a single club of members encompassing all the legitimate nodes of the network.

4.2 Anonymous Authentication with Ring Signatures

Ring signatures were formalized by Rivest, Shamir and Tauman in [34] as an anonymous signature scheme. A ring signature allows a member of an ad hoc collection of users, i.e., the *ring*, to prove to any verifier that a message was sent by a member of the ring. An authenticated message does not leak the identity of its signer. Every node i has a ring $R_{i,t}$ at time t that is composed of a finite subset of the collection $\mathcal{P}$ of all pseudonyms in the network: $R_{i,t} = \{..., P_i, P_j, ...\}$. Let $\mathcal{R}_t = \{R_{1,t}, R_{2,t}, ...\}$ be the set of rings in the network at time t. Based on the pseudonyms in their local histories, mobile nodes decide which pseudonyms to use in their rings. We call this the *ring construction strategy*.

Ring signatures can be constructed upon any type of public key cryptographic primitive [6]. What is common to these schemes is that ring signatures are based on combining functions:

$$\mathcal{C}_{\mathcal{H},v}(T_0, T_1, ..., T_{r_i-1}) = v \tag{1}$$

where $\mathcal{H}$ is a secure cryptographic hash function, v is a random glue value, r_i is the size of the ring (constant over time) and T_k, $k = 0, ..., r_i - 1$, are randomly generated values except for one that requires the knowledge of a secret key to solve (1).

For efficiency reasons, we consider the ring signature scheme presented in [41] in which the combining function $\mathcal{C}$ relies on bilinear pairings and the public key cryptosystem is identity-based (i.e., ID-based cryptography [36]). In ID-based cryptography, the knowledge of the identifier (i.e., pseudonym) of a node is sufficient to validate the authenticity of its signature. This reduces the communication overhead because it avoids the use of certificates accompanying signatures

generated by traditional cryptosystems such as RSA and ECC. The Achilles' heel of ID-based cryptosystems has always been their slower speed compared to other cryptosystems. But the recent introduction of efficient algorithms for computing pairings starts showing its feasibility on mobile devices [40]. We will elaborate more on the corresponding costs in Section 7. In ID-based cryptography the CA must be replaced by a Private Key Generator (PKG). A common critique of ID-based cryptography is that the PKG must be trusted to generate/protect private keys, and can forge signatures on behalf of the nodes (i.e., the key escrow problem). But for the applications considered here, we assume that the PKG (i.e., the CA) is trusted.

Let $\mathcal{G}$ be a Gap Diffie Hellman (GDH) group of prime order q. When a mobile node i wants to send a message m at time t, it first constructs a ring $R_{i,t}$ by selecting r_i pseudonyms (including its own pseudonym) out of its history. The ring signature is an $r_i + 1$ tuple of random values $T_k \in \mathcal{G}$ for $k = 0, 1, ..., r_i - 1$ and of $c_0 \in \mathcal{G}$:

$$(c_0, T_0, T_1, ..., T_{r_i-1}) \tag{2}$$

where c_0 is an initialization value for the ring creation; it contains the hash of the message m. T_k are randomly generated values except for one (only known to user i) that solves (1) with $v = 0$ and requires the knowledge of the secret key K_i^{-1}. We denote $RS_{i,t}(m) = (c_0, T_0, T_1, ..., T_{r_i-1})$ the ring signature on a message m sent by node i. To avoid replay attacks, the message m also contains a timestamp. The entire packet sent over the air looks as follows:

$$m, R_{i,t}, RS_{i,t}(m) \tag{3}$$

4.3 Anonymous Communications

Upon receiving a message, a node validates its signature before reading it. The receiver can reply to the message to initiate a communication session. To do so, two nodes establish a security association through an authenticated key exchange, e.g., ring signatures can be used in conjunction with the Diffie-Hellman protocol [29].

However, in order to allow for bidirectional communications, mobile nodes must be identifiable in the short term. Much to the detriment of privacy, mobile nodes already make use of long term identifiers, such as MAC (Medium Access Control) addresses, to communicate on the data link. For example, in IEEE 802.11, the MAC addresses are 48-bit values included in frames to identify the source or destination of a frame. Hence, whereas rings can provide an appropriate layer of anonymity at the application layer, the MAC addresses have to be anonymized to serve uniquely for short term communications. One approach consists in changing the MAC address [24] every time the ring changes, to preserve the anonymity created by the ring while still being able to identify nodes in the short term. The MAC address can be generated randomly, taking into account that collisions must be avoided. In [23], the authors suggest another approach based on an identifier-free link layer protocol. Basically, their solution

increases the difficulty of profiling users from the link layer by obscuring long term explicit identifiers.

Finally, it must be noted that, at the physical layer, the wireless transceiver has a wireless *fingerprint* that can identify mobile devices in the long term [33]. However, this requires a costly installation for the adversary and stringent conditions on the wireless medium. A more generic approach consists in the analysis of the signal power of mobile devices to track their locations. It is still an open problem to determine how much identifying information a sophisticated adversary can extract from the physical layer.

5 Anonymity Analysis

In this section, we evaluate the anonymity provided by rings, considering a passive adversary. We show how to optimally construct rings to maximize the achievable anonymity.

5.1 Attack Description

A global and passive adversary observes the rings used by the nodes to authenticate each other over time (Fig. 1). Based on this information, it attempts to de-anonymize rings signatures.

Given a ring alone, an adversary is unable to determine the identity of the ring owner because of the cryptographic anonymity of ring signatures. However, if an adversary obtains all the rings used at time t in the network, it can infer the most probable owner of each ring by analyzing the ring members. For example, node i constructs a ring $R_{i,t}$ of size r_i. It uses its pseudonym P_i and selects $r_i - 1$ pseudonyms out of its history. If no other ring in the network uses pseudonym P_i, the adversary can conclude that ring R_i corresponds to pseudonym P_i (e.g., node u_4 in Fig. 2 (a)). A methodic analysis of ring members can thus reverse the *anonymity* provided by rings. Repeating this attack for each t, the adversary can track the locations of mobile nodes. In this section, we focus on the analysis of anonymity, which, as explained above, is a prerequisite to untraceability. The adversary will thus analyze snapshots of rings (columns in Fig. 1). Without loss

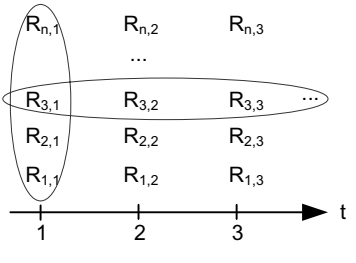

Fig. 1. Rings over time. An adversary will observe sets of rings changing over time and try to track the locations of mobile nodes. $R_{i,t}$ is the ring of node i at time t.

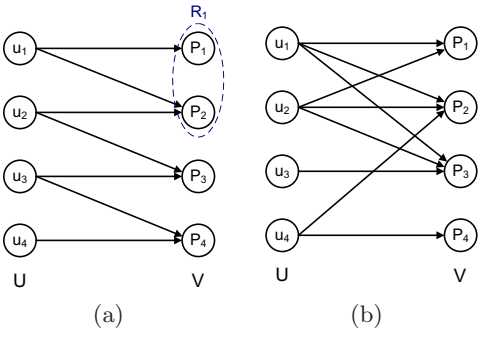

Fig. 2. Two examples of bipartite graphs G. (a) Rings are $R_1 = \{P_1, P_2\}$, $R_2 = \{P_2, P_3\}$, $R_3 = \{P_3, P_4\}$, and $R_4 = \{P_4\}$. (b) Rings are $R_1 = \{P_1, P_2, P_3\}$, $R_2 = \{P_1, P_2, P_3\}$, $R_3 = \{P_3\}$, and $R_4 = \{P_2, P_4\}$.

of generality, we write in the following $R_{i,t} = R_i$ and $RS_{i,t}(m) = RS_i(m)$. The adversary can also try to defeat the untraceability of rings by linking resembling rings over time to the same ring owner (rows of Fig. 1). Section 7 will give preliminary results on the untraceability analysis.

5.2 Graph-Theoretic Model

A set of rings can be modeled with a bipartite graph $G = (U \cup V, E)$, where $U = \{u_i\}_{i=1}^{n}$ is the set of nodes, $V = \{P_j\}_{j=1}^{n}$ is the set of pseudonyms and $E \subseteq U \times V$ is the set of edges. A graph is *bipartite* if its vertices can be partitioned into two sets such that no edge connects vertices in the same set. If a pseudonym P_j is in R_i, then we say that the node u_i using R_i is connected to P_j and we create an edge $(u_i, P_j) \in E$. We consider a balanced graph, that is, there are $|U| = |V| = n$ nodes in the system. There are $|E| = e$ edges directed from U to V. We denote d^{in} the in-degree of a node in V, i.e., the number of edges directed towards the node. Similarly, d^{out} denotes the out-degree of a node in U, i.e., the number of edges directed away from the node (the size of the ring). Two possible bipartite graphs are illustrated in Fig. 2. Graphs are *simple* if there are no multiple edges between two nodes.

After modeling rings with a graph G, $\mathcal{A}$ aims to discover which among the pseudonym $P_j \in R_i$ corresponds to the node u_i. To do so, $\mathcal{A}$ must find the most likely mapping of pseudonyms in V onto nodes in U. In graph-theoretic terms, $\mathcal{A}$ is looking for an *assignment* of nodes in V to nodes in U in the bipartite graph G. An assignment is a *matching* if no two edges share a common vertex. A *perfect matching* is a matching that covers all vertices of the graph. $\mathcal{A}$ must thus find the most probable perfect matching.

To do so, $\mathcal{A}$ assigns probabilities to all edges of the graph: $p_{j|i}$ is the probability that pseudonym P_j in V corresponds to node u_i in U. Hence, the graph G is weighted with probabilities computed by the adversary. Finally, $\mathcal{A}$ can find

the most probable perfect matching by computing the maximum-weight perfect matching over the weighted bipartite graph G.

Measuring Anonymity. The individual anonymity of a node u_i (i.e., the uncertainty of an adversary about the identity of node u_i [35]) can be measured by:

$$H_i = -\sum_{j=1}^{r_i} p_{j|i} \log_2(p_{j|i}) \tag{4}$$

which is the entropy of the random variable $p_{j|i}$ and where $r_i = |R_i|$ is the size of the ring of node u_i.

A priori, the adversary will choose probability $p_{j|i}$ equal to $1/d_i^{out}$ as each outgoing edge is equally likely to be chosen. In Fig. 6(a), the entropy yields with this approach a non-zero anonymity for all but the last node. However, by doing so, the adversary focuses on the anonymity of individual nodes and overlooks some important properties of the system as a whole [19]. A clever adversary would eliminate many possible assignments by working backwards from vertices with a degree of one and the entropy would then yield a zero anonymity for all nodes.

$\mathcal{A}$ can thus first consider all assignments m_k of the elements of V onto U before computing $p_{j|i}$ a posteriori [7,38]. The probability of an assignment m_k is given by:

$$p(m_k) = \prod_{l \in m_k} w_l$$

where w_l is the weight of edge l in G. The weight is the a priori probability $p_{j|i}$ and we write $w_l = 1/d_i^{out}$ where node u_i is the origin of edge l. Because all the weights of the edges leaving a node are equal, all perfect matchings are equally likely and we have: $p(m_k) = p(m)$.

Hence, the probability of a perfect matching, i.e., the probability that an assignment is perfect knowing the set of all perfect matchings M, is:

$$p(m_k|M) = \frac{p(m_k, M)}{p(M)} = \frac{p(m_k)}{p(M)} = \frac{p(m_k)}{\sum_{k=1}^{|M|} p(m_k)} = \frac{1}{|M|}$$

where $m_k \in M$ for $k \in [0, |M|]$, and $p(M)$ is the sum of probabilities of all perfect matchings. The a posteriori probability $p_{j|i}$ is finally computed by considering all perfect matchings containing the pair (u_i, P_j).

$$p_{j|i} = \sum_{m_k \in M | (u_i, P_j) \in m_k} p(m_k|M) \tag{5}$$

In words, the number of perfect matchings going over an edge determines the weight of an edge. Hence, the anonymity of a node not only depends on its out-degree but also on the distribution of perfect matchings, i.e., the structure of the bipartite graph. Considering again the example in Fig. 2 (a), there is a single perfect matching in the graph, and consequently the anonymity of each node

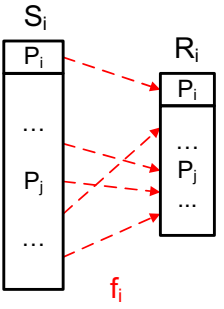

Fig. 3. Ring construction. To construct its ring R_i, node u_i uses its pseudonym P_i and selects, according to its strategy f_i, $r_i - 1$ pseudonyms out of its history S_i.

is null. To further illustrate the result, consider the example in Fig. 2 (b) and observe node 4. To compute $p_{4|4}$, we consider all the perfect matchings with the pair $(4, 4)$. In fact, every perfect matching in the graph contains that pair because $d_4^{in} = 1$, hence $p_{4|4} = 1$. The same analysis is true for node 3 as $d_3^{out} = 1$. Still, nodes 1 and 2 have a non-zero anonymity. In other words, both the in-degrees of nodes in V and out-degrees of nodes in U affect the distribution of perfect matchings and determine the probability $p_{j|i}$.

Complexity. The analysis presented above is difficult to carry out in practice because of its complexity: All perfect matchings must be found. Itai *et al.* introduce in [25] a polynomial time algorithm to find all perfect matchings in a bipartite graph. The algorithm starts from a perfect matching to iteratively produce them all in $\mathcal{O}(e \cdot (\sqrt{n} + |M|))$ time. The algorithm remains hard to use, as the number of nodes n can be extremely large and the number of matchings $|M|$ increases exponentially with the number of nodes. The adversary could thus focus on small sets of rings using, for example, the divide and conquer approach presented in [22].

5.3 Ring Construction Problem

Fig. 3 illustrates the ring construction process: Each node u_i obtains a ring R_i of size r_i by using its pseudonym P_i and choosing $r_i - 1$ pseudonyms from its history S_i. Rings must be carefully created to obtain high anonymity. The *ring construction strategy* of node u_i gives the criteria to include a pseudonym from u_i's history in its ring. We define it as a function $f_i : 2^{\mathcal{P}} \rightarrow 2^{\mathcal{P}}$. The selected pseudonyms must belong to u_i's history: $f_i(X) \subseteq X$ where $X \subseteq \mathcal{P}$ is a set of pseudonyms. The number of selected pseudonyms must not exceed the ring size: $|f_i(X)| \leq r_i$. In this paper, we consider that all nodes use the same ring construction strategy: $f_i = f$.

The *ring construction problem* consists in finding the ring construction that maximizes anonymity. To do so, we must obtain the optimal graph that maxi-

mizes the achievable anonymity for all nodes:

$$\max_{G}(H_i)\ \forall u_i \in U \tag{6}$$

subject to:

$$1 \leq d_i^{out} \leq d^{\max} \tag{7}$$

$$(u_i, P_j) \in E \Leftrightarrow P_j \in f(S_i) \tag{8}$$

Equation (7) confines the out-degrees d_i^{out} to a maximum $d^{\max}$. The graph construction is constrained by (8): The resulting graph G depends on the information collected by the nodes (their knowledge of pseudonyms). In other words, we seek to obtain a graph that maximizes the level of anonymity of every node constrained by the maximum out-degree and using a distributed construction function f (i.e., self-organization).

Optimal Graph *G*. As illustrated in Fig. 2, the distribution of the node in-/out-degrees affects the anonymity of each node. Let us introduce the following notation: $G_{\setminus(u_i,P_j)}$ corresponds to graph G without the edge (u_i, P_j).

The following Theorem identifies the graphs that provide the maximum anonymity. The proof is provided in the Appendix.

Theorem 1. *Anonymity is maximal if and only if every vertex in the bipartite graph G has the same degree $d^{\max}$ and for each i, the subgraphs $G_{\setminus(u_i,P_j)} \subseteq G$ for all $P_j \in R_i$ are isomorphic to each other. The anonymity of each node is then* $\log_2(d^{\max})$.

Theorem 1 characterizes the optimal graph G that maximizes the achievable anonymity. Basically, all subgraphs of G obtained by removing an edge starting at one node must be isomorphic. The isomorphism property captures the notion of *similarity* between subgraphs: If subgraphs are similar (i.e., have the same structure), it is more difficult to distinguish nodes in G. A large body of work has studied the existence of graph isomorphism and shown that the problem is NP: It belongs to its own complexity class, neither known to be solvable in polynomial time nor NP-complete [9,17]. In other words, theory says that it might be hard to determine whether two graphs are isomorphic. However, in practice, the graph isomorphism problem is easy to solve in polynomial time with heuristics [5]. It is thus possible in principle to determine whether subgraphs are isomorphic and, as a consequence, whether a graph G provides maximum anonymity.

In the following, we will compare the anonymity provided by different graph constructions and will see that regular graphs perform best. In fact, the local structure of the graph (i.e., the way each node is connected) determines whether subgraph isomorphisms can exist. In particular, the regularity of graphs is a necessary condition in our scenario for subgraphs to be isomorphic to each other (see Appendix).

6 Ring Construction Strategies

In the previous section, we examined the achievable anonymity with rings and derived necessary conditions (i.e., regular and isomorphic) to maximize it. In

this section, we evaluate the performance of different graph constructions by means of simulations. The simulations are carried out in C++ using the LEDA library [30] to manipulate graphs. First, we assume that the nodes know the entire network and show the superiority of regular constructions over random graphs. The results are averaged over 20 runs with a running example of 10 nodes, which is sufficient to evaluate the effect of ring construction strategies with a reasonable simulation complexity. Then, we approximate the achievable anonymity on geometric graphs with 100 mobile nodes that only know a portion of the entire network.

6.1 Random Graphs

Let us assume that the nodes are aware of all the pseudonyms in the network (i.e., $S_i = \mathcal{P}$, $\forall u_i$). With a random graph construction f^{rand}, mobile nodes choose pseudonyms randomly: We consider a bipartite Erdos-Renyi random graph $G(n, p)$ where n is the number of nodes and each edge is included in the graph with probability p independently of others. With such graphs, the in-/out-degree distribution is binomial $Pr(d_i = k) = \binom{n}{k} p^k (1-p)^{(n-k)}$ with average $E[d_i] = np$ and variance $var[d_i] = np(1-p)$. Fig. 4 (a) shows the average distribution of the achieved entropy. We observe that the average anonymity increases with the edge density p, whereas the average variance decreases when p approaches 0 or 1. In other words, with a low or high density of edges, the achievable anonymity has a narrow distribution. As $p \to 1$, the graph becomes complete (i.e., all nodes are connected) and thus optimal in terms of anonymity.

We compare the performance of random and r-regular graphs in Fig. 4 (b) by computing their minimum and mean anonymity. We observe that regular graphs have a near-optimal behavior as they approach the maximum achievable anony-

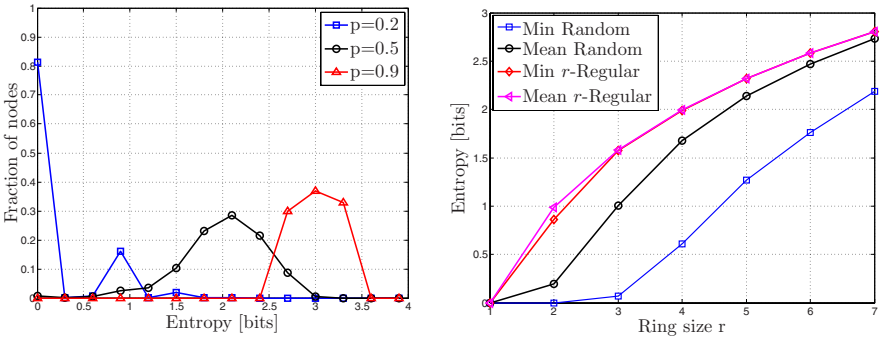

Fig. 4. Comparison of random and regular graphs. (a) Entropy distribution of random graphs with an increasing edge density p. The x-axis is divided into bins of size 0.3 and the y-axis represents the fraction of nodes in each bin. (b) Minimum and mean entropy levels of random and r-regular graph constructions.

mity $\log_2(r)$. Random graphs perform poorly, illustrating the importance of regular degree distributions: Nodes with a low in-/out-degree have lower anonymity, which is hardly compensated by nodes with a higher degree as the anonymity is logarithmic. The node degree variance induces a larger anonymity variance. Thus, to guarantee a minimal level of anonymity, the mean degree must be even larger.

As $n \to \infty$, the node degree distribution is approximated by a Poisson distribution with parameter $\lambda = np$. As the variance of the degree distribution equals the mean, it will be large and reduce the average anonymity. Bollobas in [9] notably investigates the asymptotic distribution of the degree sequence of graphs and proves that random graph constructions do not permit to guarantee predictable minimal and maximal degrees: the minimum and maximum degrees are *essentially determined* (by a function whose exact value is unknown). Bollobas further demonstrates that for some values of p (Theorem 3.5, [9]), there is a minimal degree $d^{min} \geq 2$. However, in this case, the maximum degree is not finite. In other words, as the graph grows larger, the degree sequence of random graphs is unpredictable and the performance gap with regular graphs increases.

6.2 K^{out} Graphs

We evaluate whether introducing a structure in the graph construction increases the achievable anonymity. We impose the same fixed out-degree $d^{out} = K$ to every node to obtain a K^{out} graph [9]. We consider various ring construction strategies with and without the help of a central entity.

Centralized Algorithm. The central entity is a network coordinator that knows the in-/out-degrees of each node and generates regular graphs from K^{out} graphs. Bollobas in [9] suggests a pairing model (i.e., a ring construction f^{reg}) to construct regular graphs with a centralized algorithm: Every vertex of the graph is connected to K nodes uniformly at random forming Kn pairs. If there are no multiple edges between two same nodes, the resulting graph is a random regular graph.

Distributed Algorithm. In the absence of a central entity, the nodes must decide individually with whom to connect. Each vertex $u_i \in U$ uses its pseudonym P_i and randomly selects $K-1$ vertices from V. As all $\binom{n-1}{K-1}$ choices are equiprobable, the probability that a pseudonym P_j is chosen by another node is the ratio of assignments containing P_j over all possible assignments:

$$\begin{aligned} p &= Pr(\text{“Node } u_i \text{ picks } P_j \text{ after } K-1 \text{ tries”}) \\ &= \frac{\binom{n-2}{K-2}}{\binom{n-1}{K-1}} = \frac{K-1}{n-1} \end{aligned} \tag{9}$$

The node in-degree distribution is then $Pr(d_i^{in} = k) = \binom{n}{k} p^k (1-p)^{n-k}$. Thus, the average in-degree distribution is $E[d_i^{in}] = \frac{n}{n-1}(K-1)$ and the variance is

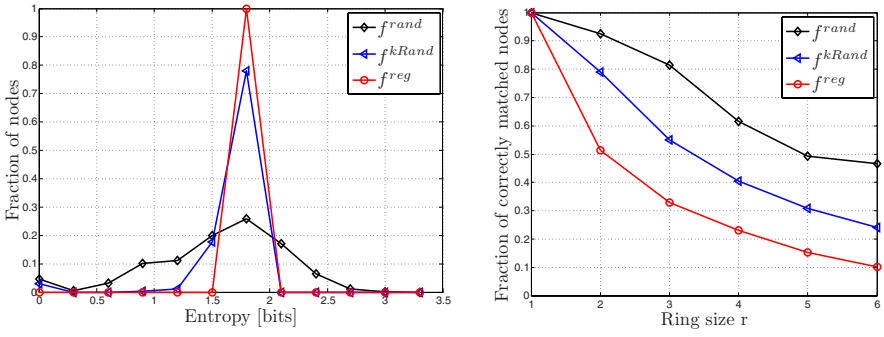

(a) (b)

Fig. 5. Comparison of various ring constructions f. (a) Entropy distributions. The x-axis is divided into bins of size 0.3 and the y-axis represents the fraction of nodes in each bin. (b) Fraction of correctly matched pseudonyms to rings.

$var[d_i^{in}] = \frac{n}{n-1}(K-1)(1-\frac{K-1}{n-1})$. Consequently, with the distributed algorithm, the ring construction strategy f^{kRand} heavily depends on the degree K.

Fig. 5 (a) compares the entropy distribution for various ring constructions. We consider equivalent graphs constructions with $n = 10$ nodes, $p = 4/10$, $K = 4$ and $d = 4$ for d-regular graphs. We observe that f^{kRand} obtains a narrower distribution of entropies than f^{rand}, thus illustrating the importance of a regular out-degree. If random graphs obtain the maximum entropy among all constructions, they also have a smaller minimum entropy and lower average entropy. f^{reg} obtains very good results, close to the maximum achievable entropy $\log_2(d)$. This is due to the regularity of the in-degree distribution. With f^{reg} and $d = 4$, the majority of the nodes are indistinguishable as their entropy is $2^2 = 4$ equal to d. Like with random graphs, as $n \to \infty$, the in-degree distribution of K^{out} graphs is Poissonian: The mean and variance approach $K - 1$. In other words, the difference between regular and K^{out} graphs will increase as K becomes large.

In Fig. 5 (b), we observe that the proportion of successful matchings of pseudonyms onto nodes (i.e., the adversary success ratio) varies significantly among graphs. In the worst case, the adversary cannot infer information statistically and thus makes random attempts. The probability of success of the adversary is then equal to $1/r$. For regular graphs, $\mathcal{A}$'s success is limited and approaches its worst case strategy. In other words, $\mathcal{A}$ would do better by randomly matching rings to pseudonyms. With random constructions however, the adversary can infer significant information: Even with rings composed of 6 nodes, 5 out of 10 nodes in the example are correctly matched.

6.3 Geometric Graphs

As discussed above, the introduction of a structure in the ring construction dramatically increases the achievable anonymity. Still, the nodes were aware of

all the pseudonyms in the network. In practice, mobile nodes will only have access to information gathered from the network, i.e., the rings of their encounters. In this section, we evaluate how the topology affects the achievable anonymity. In particular, as regular graph provide high anonymity, we study several ring construction strategies to obtain a regular graph G. To take the network topology into account, we consider a geometric graph G_g in which each vertex is associated with a physical device. Two vertices are connected (i.e., learn each other's rings) if and only if they are within distance $D(u, v) \leq \Gamma$ of each other, where Γ is a fixed radius (i.e., the unit disk graph model). The geometric graph G_g models the connection between nodes. We assume that the nodes are homogeneous (i.e., identical devices) and equipped with omnidirectional antennas. We consider both static and mobile scenarios.

Static Scenario. In static scenarios, nodes learn the pseudonyms in the rings of their direct neighbors. Indirectly, they also learn the pseudonyms of the neighbors of their neighbors as they are passed along. Given a history S_i and a ring size $r_i = r$, $\forall u_i$, the probability that node u_i chooses pseudonym P_j from its history in its ring is: $Pr(P_j \in R_i) = \min(1, r/|S_i|)$. For first-hop neighbors of u_i, the probability of learning P_j corresponds to the probability that u_i uses P_j, i.e., $Pr(P_j \in R_i)$. In other words, the ring size determines the propagation rate of pseudonyms in the network. For a x-hop neighbor u_l of u_i where x is larger than 1, the probability that pseudonym P_j is used by all rings on a path $\Delta_{i,l}$ from node u_i to node u_l is: $Pr(P_j \in R_k, \forall k \in \Delta_{i,l}) = \prod_{\forall k \in \Delta_{i,l}} \min(1, r/|S_k|)$. However, nodes belonging to disconnected sets of the graph G_g are isolated from each other and have zero probability of learning each other's pseudonyms. Hence, the propagation of pseudonyms is limited by the graph connectivity as well, reducing the potential size of anonymity sets. The topology of the network thus critically affects the achievable privacy. With the ring construction f^{static}, we consider that mobile nodes randomly choose pseudonyms from their local history to construct their rings.

Mobile Scenario. We examine how mobility can lessen the negative impact of topology: As nodes move in the network, they discover a larger portion of the set of pseudonyms $\mathcal{P}$. We consider the *restricted random waypoint model* introduced in [14]. In the *random waypoint model*, a mobile node moves on a continuous plane from its current position to a new position by randomly choosing its destination coordinates, its speed and the amount of time it will pause when it reaches the destination. After its pause, a node chooses a new destination and speed. This is repeated for each node until the end of the simulation time. In the restricted model, the choice of destination points is restricted with some probability ϕ to a set Ψ of fixed points on a plane. With probability ϕ a node randomly chooses a point from Ψ, and with probability $1 - \phi$, a node will choose a random point on the plane. This model is close to reality as users do not choose their destinations randomly, but instead meet at cafés, bus stops, etc.

In this mobile environment, we evaluate various ring construction strategies that aim at obtaining the most regular graph. These strategies capitalize on the frequency and freshness of the appearance of pseudonyms.

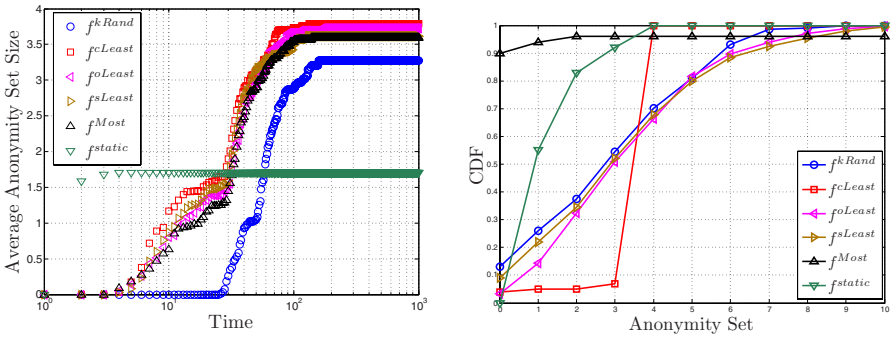

(a) (b)

Fig. 6. Average anonymity set size for several ring construction strategies f in a mobile scenario with $\phi = 0.5$ and $r = 4$. (a) Average anonymity set size over time. (b) Cumulative distribution function of the anonymity set.

Least Popular Strategy. With the *least popular ring construction strategy*, each node maintains a counter for each pseudonym and selects the pseudonyms with the lowest counter value (i.e., the least popular). We consider three variations of the strategy: In f^{cLeast}, a central server informs the nodes of the in-degree of members of their histories. In f^{sLeast}, nodes choose in their histories the pseudonyms that were used the least often. In f^{oLeast}, nodes choose in their histories the pseudonyms that were used the least often in the rings of others.

Most Popular Strategy. We consider a *most popular ring construction strategy* f^{Most} in which the most popular nodes are chosen with the help of a central server.

Random Strategy. With the *random construction strategy* f^{kRand}, nodes choose their ring members randomly from their local history.

We ran 20 simulations on a 500×500 m^2 torus with $n = 100$ nodes, transmission range = 25m, pause = 20s, history $|S| = 10$, ring size $r = 4$, $|\Psi| = 5$ and $\phi = 0.5$. For simplicity and clarity, instead of computing the entropy, we compute the anonymity sets of mobile nodes, which corresponds to the in-degree distribution of graph G_g.

Fig. 6 (a) shows the evolution over time of the average anonymity set size of mobile nodes. We observe that the achieved anonymity set in mobile scenarios surpasses by far the static scenario but takes longer to converge. In general, we observe a percolation region ($[10, 10^2]$ seconds) where the anonymity set of the nodes increases quickly, and then a region of convergence ($[10^2, 10^3]$ seconds). f^{static} reaches a small anonymity set and is topped by all mobile strategies. Comparing mobility scenarios, we observe that f^{cLeast}, f^{oLeast} and f^{sLeast} improve the average size of the anonymity set with respect to the f^{kRand} (10% to 20% improvement). We notice that f^{oLeast} performs slightly better than f^{sLeast} as it takes better advantage of mobility (i.e., nodes have a better global knowledge

of rings) and approaches the performance of the centralized algorithm f^{cLeast}. The f^{Most} approach seems to perform as well as the least popular approaches.

Fig. 6 (b) gives the cumulative distribution function (CDF) of the anonymity set showing the fraction of nodes per anonymity set size. The spread of the curve represents the variance across mobile nodes in the anonymity set sizes. We observe that f^{cLeast} performs well: It has a small variance as the majority of the nodes has an anonymity set size equal to 4 (i.e., the ring size). f^{oLeast} and f^{sLeast} have a smaller variance than f^{kRand}. Notably, with the f^{oLeast} approach, fewer nodes (20% less) have a small anonymity set. Finally, we observe that the f^{Most} approach actually performs worse than all other strategies: As a small number of nodes become extremely popular, the majority of nodes (90%) has a small anonymity set. Hence, although the average anonymity set size is large (Fig. 6 (a)), only a few nodes are actually anonymous, while others are easily identifiable. As social networks (usually modeled with scale-free graphs) tend to have this form, social networks based ring constructions would perform poorly.

In conclusion, the knowledge of the *popularity* of a pseudonym helps to achieve high anonymity (i.e., the least popular strategy). The nodes can thus independently aggregate information about their encounters and achieve anonymity in a self-organized way (without harming the anonymity of other nodes). Hence, peer-to-peer communications between mobile nodes enable privacy to emerge in ad hoc wireless networks.

7 Discussion

In this section, we present preliminary results on the untraceability of rings, explain their resilience to Sybil attacks, detail how revocation works and finally discuss the cost of ring signatures.

7.1 Untraceability

Untraceability of rings is also required in order to achieve self-organized location privacy. Similar to mix zones [7], mobile nodes can change their rings simultaneously upon meeting in the network. An external adversary will have to infer the most probable matching of old and new rings. Mobile nodes are untraceable if the adversary is unlikely to successfully match rings. Unlike the multiple pseudonym approach, in self-organized location privacy rings are correlated over time. Hence, by analyzing the similarity of ring members over time, an adversary could statistically estimate the matching of rings and track mobile nodes in the network. For example, if ring members remain constant, an adversary trivially tracks the whereabouts of mobile nodes. Ring members must thus vary: Except for the pseudonym P_i of the ring creator u_i, a ring $R_{i,t+1}$ can be entirely different from the previous ring $R_{i,t}$. Still, if all but one pseudonym are systematically updated, an adversary tracks mobile nodes by identifying persistent ring members.

In order to defeat an attack on untraceability by an external adversary $\mathcal{A}$, ring members must evolve with time depending on both past ring members and

new encounters. Therefore, on top of the self-organization involved to achieve anonymity, mobile nodes must coordinate the *evolution* of their ring members to obtain untraceability. One possible way to coordinate the evolution of rings is to *cluster* ring members. The clustering coefficient of a vertex is used to quantify how close the vertex and its neighbors are to being a complete graph [39]. In our case, the clustering coefficient of a node measures the number of common ring members it shares with nearby nodes. The clustering coefficient of ring members results in an overlap of rings, which hardens the attack by $\mathcal{A}$. Mobile nodes can cluster their rings in a self-organized way by favoring pseudonyms recently observed: Newly acquired pseudonyms have a higher probability of being chosen in a ring. Preliminary results have demonstrated the success of this approach. We leave the formal investigation of this method for future work.

7.2 Sybil Attacks and Revocation

If a single node can present multiple identities, it can control a substantial fraction of the system and thereby undermine its security. These Sybil attacks [18] are not possible if there is a central entity to vouch for a one-to-one correspondence between entity and identity. In our model, the offline CA attributes a single pseudonym to every node after proper identification and rings are only used for authentication purposes. Hence, rings are unaffected by Sybil attacks. Actually, as privacy is generated by the nodes, RSs can be viewed as a *Sybil defense* that exploits the redundancy of mobile networks to generate a self-cloak.

Typical misbehavior remains possible in peer-to-peer wireless networks: For example, a mobile node can engage in denial of service attacks. However, the CA can exclude misbehaving nodes by revoking their keying material (as a signer must own a private key to generate a ring signature). Thus, keys can be blacklisted using certificate revocation lists (CRLs) like traditional revocation algorithms [42].

7.3 Cost

As ring sizes affect the anonymity level, users will tend to create the largest possible rings. But as ring signatures incur a communication and computation overhead, ring sizes will be bounded by the acceptable performance overhead.

Computation Overhead. RS *computational requirements* depend on the underlying trapdoor permutation, i.e., with ID-based ring signatures, one bilinear pairing computation is required for each member of the ring. In other words, for a node u_i, the *signature cost* C_{sign} is:

$$C_{sign} \approx r_i \cdot C_{BP} \tag{10}$$

where C_{BP} captures the cost of a Bilinear Pairing. The verification of a message has the same complexity. Using FPGA hardware accelerators for bilinear pairings [37], one bilinear pairing takes 61μs. In total, for a ring of size $r_i = 10$, $C_{sign} = 610\mu$s. Without hardware accelerators, the efficiency of ring signatures

in mobile phones depends on software optimizations: Currently, one bilinear pairing takes 478ms on a 225MHz ARM9 processor [40]. If this is not usable, the computation cost will eventually decrease as mobiles' hardware improves.

Transmission Overhead. To sign messages with ring signatures, only the first authenticated message between two nodes must contain the ring. Subsequent messages will thus have a smaller overhead. A RS is an $r_i + 1$ tuple: $(c_0, T_1, ..., T_{r_i-1})$. Each of those tuples is taken out of the group $\mathcal{G}$ of prime order q. Hence, the total size of the signature is $(r_i + 1) \cdot \mathcal{M}$ bits where $\mathcal{M} = \log_2(q)$. On top of the signature, a ring R_i is composed of r_i pseudonyms of $\mathcal{M}$ bits. Hence, the size of the signature grows linearly with the size of the ring. The *transmission cost* C_{trans} is:

$$C_{trans} \approx (r_i + 1) \cdot \mathcal{M} + r_i \cdot \mathcal{M} = (2r_i + 1)\mathcal{M} \tag{11}$$

For example, assume that a node u_i with pseudonym P_i creates a ring of size $r_i = 10$. For 128-bit security, NIST [32] recommends $\mathcal{M} = 283$ with elliptic curves defined over a *binary* underlying finite field of characteristic two ($\mathcal{F}_{2^{\mathcal{M}}}$). The communication overhead of the first message of each node is then $21 \cdot \mathcal{M}$ bits and $11 \cdot \mathcal{M}$ for subsequent messages.

8 Conclusion

We introduced the *self-organized location privacy* paradigm to solve the problem of location privacy in wireless mobile networks. With this approach, the network protects the location privacy of its nodes in a self-organized manner relying on Ring Signatures. Using graph theory, we theoretically measured the efficiency of the approach to provide anonymous authentication and derived its optimum. We examined numerically different ring construction strategies at the mobile nodes and showed that regular constructions achieve near-optimal anonymity. Then, we demonstrated that enabling nodes to communicate with each other increases their respective privacy levels by means of simulations. Despite their lack of knowledge of the entire network, mobile nodes achieve a high anonymity level by relying, for example, on the popularity of pseudonyms. In particular, choosing to connect to the *least popular pseudonyms* tops the achievable anonymity. Another particularly interesting result is that mobility helps in establishing self-organized anonymous authentication by improving the network awareness of every node without compromising their anonymity.

Future Work. We will investigate the effect of stronger adversary models on the achievable anonymity, such as an adversary that compromises members of the network. We will also extend the study of the effect of social networks on the construction of rings. In particular, social networks could provide information (e.g., the social graph) to improve the efficiency of ring constructions. Finally, we intend to complete our preliminary study on the untraceability of rings.

Acknowledgments

We would like to thank Levente Buttyan, Rafik Chaabouni, Mario Cagalj, Marcin Poturalski, and Serge Vaudenay for their insights and suggestions on earlier versions of this work, and the anonymous reviewers for their helpful feedback.

References

1. http://www.techcrunch.com/2007/09/11/the-holy-grail-for-mobile-socialnetworks
2. http://www.aka-aki.com/
3. http://en.wikipedia.org/wiki/Bluedating/
4. http://www.gamemobile.co.uk/bluetoothmobilegames/
5. http://cs.anu.edu.au/~bdm/nauty/
6. Abe, M., Ohkubo, M., Suzuki, K.: 1-out-of-n signatures from a variety of keys. In: Zheng, Y. (ed.) ASIACRYPT 2002. LNCS, vol. 2501, pp. 415–432. Springer, Heidelberg (2002)
7. Beresford, A.R.: Location privacy in ubiquitous computing. Ph.D. thesis, University of Cambridge (2005)
8. Beresford, A.R., Stajano, F.: Mix zones: User privacy in location-aware services. In: PerSec (2004)
9. Bollobas, B.: Random Graphs. Cambridge University Press, Cambridge (2004)
10. Boneh, D., Boyen, X., Shacham, H.: Short group signatures. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 41–55. Springer, Heidelberg (2004)
11. Camenisch, J., Van Herreweghen, E.: Design and implementation of the Idemix anonymous credential system. In: CCS (2002)
12. Camenisch, J., Hohenberger, S., Kohlweiss, M., Lysyanskaya, A., Meyerovich, M.: How to win the clone wars: efficient periodic n-times anonymous authentication. In: CCS (2006)
13. Camenisch, J.L., Lysyanskaya, A.: Dynamic accumulators and application to efficient revocation of anonymous credentials. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, p. 61. Springer, Heidelberg (2002)
14. Capkun, S., Hubaux, J.-P., Buttyan, L.: Mobility helps peer-to-peer security. IEEE Transactions on Mobile Computing (2006)
15. Chaum, D.: Untraceable electronic mail, return addresses, and digital pseudonyms. Communications of the ACM 24(2) (1981)
16. Chaum, D., van Heyst, E.: Group signatures. In: Davies, D.W. (ed.) EUROCRYPT 1991. LNCS, vol. 547, pp. 257–265. Springer, Heidelberg (1991)
17. Corneil, D.G., Gotlieb, C.C.: An efficient algorithm for graph isomorphism. J. ACM 17(1), 51–64 (1970)
18. Douceur, J.R., Donath, J.S.: The sybil attack. In: Druschel, P., Kaashoek, M.F., Rowstron, A. (eds.) IPTPS 2002. LNCS, vol. 2429, p. 251. Springer, Heidelberg (2002)
19. Edman, M., Sivrikaya, F., Yener, B.: A combinatorial approach to measuring anonymity. Intelligence and Security Informatics (2007)
20. Freudiger, J., Raya, M., Felegyhazi, M., Papadimitratos, P., Hubaux, J.-P.: Mix zones for location privacy in vehicular networks. In: WiN-ITS (2007)
21. Freudiger, J., Shokri, R., Hubaux, J.-P.: On the optimal placement of mix zones. In: PETS (2009)

22. Gierlichs, B., Troncoso, C., Diaz, C., Preneel, B., Verbauwhede, I.: Revisiting a combinatorial approach toward measuring anonymity. In: WPES (2008)
23. Greenstein, B., McCoy, D., Pang, J., Kohno, T., Seshan, S., Wetherall, D.: Improving wireless privacy with an identifier-free link layer protocol. In: MobiSys (2008)
24. Gruteser, M., Grunwald, D.: Enhancing location privacy in wireless LAN through disposable interface identifiers: a quantitative analysis. Mob. Netw. Appl. (2005)
25. Itai, A., Rodeh, M., Tanimoto, S.: Some matching problems for bipartite graphs. Journal of the Association for Computing Machinery (1978)
26. Jiang, T., Wang, H.J., Hu, Y.-C.: Preserving location privacy in wireless LANs. In: MobiSys (2007)
27. Krumm, J.: Inference attacks on location tracks. In: LaMarca, A., Langheinrich, M., Truong, K.N. (eds.) Pervasive 2007. LNCS, vol. 4480, pp. 127–143. Springer, Heidelberg (2007)
28. Lin, H.-C., Yen, S.-M., Chen, H.-S.: Protection of mobile agent data collection by using ring signature. In: International Conference on Networking, Sensing and Control (2004)
29. Lin, X., Lu, R., Zhu, H., Ho, P., Shen, X., Cao, Z.: ASRPAKE: An anonymous secure routing protocol with authenticated key exchange for wireless ad hoc networks. In: ICC (2007)
30. Mehlhorn, K., Naher, St.: The LEDA Platform of Combinatorial and Geometric Computing. Cambridge University Press, Cambridge (1999)
31. Nakanishi, T., Fujii, H., Hira, Y., Funabiki, N.: Revocable group signature schemes with constant costs for signing and verifying. In: PKC (2009)
32. NIST. Recommended elliptic curves for government use. White Paper (1999)
33. Rasmussen, B., Capkun, S.: Implications of radio fingerprinting on the security of sensor networks. In: SecureComm (2007)
34. Rivest, R.L., Shamir, A., Tauman, Y.: How to leak a secret. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, p. 552. Springer, Heidelberg (2001)
35. Serjantov, A., Danezis, G.: Towards an information theoretic metric for anonymity. In: Dingledine, R., Syverson, P.F. (eds.) PET 2002. LNCS, vol. 2482, pp. 41–53. Springer, Heidelberg (2003)
36. Shamir, A.: Identity-based cryptosystems and signature schemes. In: Blakely, G.R., Chaum, D. (eds.) CRYPTO 1984. LNCS, vol. 196, pp. 47–53. Springer, Heidelberg (1985)
37. Shu, C., Kwon, S., Gaj, K.: FPGA accelerated Tate pairing based cryptosystem over binary fields. In: FPT (2006)
38. Tóth, G., Hornák, Z.: Measuring anonymity in a non-adaptive, real-time system. In: Martin, D., Serjantov, A. (eds.) PET 2004. LNCS, vol. 3424, pp. 226–241. Springer, Heidelberg (2005)
39. Watts, D.J., Strogatz, S.: Collective dynamics of small-world networks. Nature (1998)
40. Yoshitomi, M., Takagi, T., Kiyomoto, S., Tanaka, T.: Efficient implementation of the pairing on mobile phones using BREW. IEICE Transactions on Information and Systems (2008)
41. Zhang, F., Kim, K.: ID-based blind signature and ring signature from pairings. In: Zheng, Y. (ed.) ASIACRYPT 2002. LNCS, vol. 2501, pp. 533–547. Springer, Heidelberg (2002)
42. Zheng, P.: Tradeoffs in certificate revocation schemes. SIGCOMM Comput. Commun. Rev. (2003)

43. Zhong, G., Goldberg, I., Hengartner, U.: Louis, lester and pierre: Three protocols for location privacy. In: Borisov, N., Golle, P. (eds.) PET 2007. LNCS, vol. 4776, pp. 62–76. Springer, Heidelberg (2007)

A Proof of Theorem 1

Proof. We first show that each node u_i must have an out-degree $d_i^{out} = d^{\max}$ and then obtain the condition for achieving maximum anonymity. Assume a bipartite graph G' where at least one node $u_i \in U$ has $d_i^{out} < d^{\max}$. We add new edges to G' such that $d_i^{out} = d^{\max}\ \forall u_i$, and obtain the graph G. Because no edges were removed, G will contain at least the same number of perfect matchings as G'. Adding new edges might actually increase the number of existing perfect matchings and consequently increase the anonymity of the nodes. In other words, to maximize their anonymity, each node must choose $d_i^{out} = d^{\max}$.

To maximize the entropy of each node, the random variable $p_{j|i}$ must have a uniform distribution. Given a node u_i, $p_{j|i}$ is uniform if and only if the number of perfect matchings over (u_i, P_j) is the same for all $P_j \in R_i$. A simple way to verify this consists in comparing whether the subgraphs obtained by removing any pair $G_{\setminus(u_i,P_j)} \subseteq G\ \forall P_j \in R_i$ yield the same number of perfect matchings. The number of perfect matchings without (u_i, P_j) will be the same for any pair (u_i, P_j) with $P_j \in R_i$, if and only if all subgraphs $G_{\setminus(u_i,P_j)}$ have the same number of perfect matchings. This will be true if all subgraphs are isomorphic to each other (i.e., belong to the same equivalence class). Consider two subgraphs $G_{\setminus(u_i,P_1)}$ and $G_{\setminus(u_i,P_2)}$. An isomorphism of graphs $G_{\setminus(u_i,P_1)}$ and $G_{\setminus(u_i,P_2)}$ is defined as $\mathcal{I} : \mathcal{V}(G_{\setminus(u_i,P_1)}) \rightleftarrows \mathcal{V}(G_{\setminus(u_i,P_2)})$ where $\mathcal{V}(G_{\setminus(u_i,P_1)})$ is the vertex set of graph $G_{\setminus(u_i,P_1)}$. $\mathcal{I}$ defines an assignment of the nodes of $G_{\setminus(u_i,P_1)}$ onto the nodes of $G_{\setminus(u_i,P_2)}$ such that $\forall(u_i, P_j) \in G_{\setminus(u_i,P_1)}$, there is $(\mathcal{I}(u_i), \mathcal{I}(P_j)) \in G_{\setminus(u_i,P_2)}$. A necessary (but not sufficient) condition for the graph isomorphism to exist in this case is that the graph is *d-regular*: Each vertex has the same degree d. Indeed, if the degrees of vertices of two subgraphs cannot be matched (e.g., a subgraph has a node of degree 5 while the other does not), then it is impossible for the subgraphs to be isomorphic. Hence, we know that the graph will be $d^{\max}$-regular and that the entropy of each node will be $\log_2(d^{\max})$.