

On the Impact of Social Network Profiling on Anonymity

Claudia Diaz¹, Carmela Troncoso¹, and Andrei Serjantov²

¹ K.U. Leuven ESAT-COSIC
Kasteelpark Arenberg 10, Leuven-Heverlee, Belgium
{claudia.diaz,carmela.troncoso}@esat.kuleuven.be

² The Free Haven Project
schnur@gmail.com

Abstract. This paper studies anonymity in a setting where individuals who communicate with each other over an anonymous channel are also members of a social network. In this setting the social network graph is known to the attacker. We propose a Bayesian method to combine multiple available sources of information and obtain an overall measure of anonymity. We study the effects of network size and find that in this case anonymity degrades when the network grows. We also consider adversaries with incomplete or erroneous information; characterize their knowledge of the social network by its quantity, quality and depth; and discuss the implications of these properties for anonymity.

1 Introduction

In the last few years defining and quantifying anonymity in the context of communication networks has been a hot research topic. A substantial set of papers focus on the definition of anonymity, others present designs and analysis of new anonymous communication systems or attacks of existing ones. Yet more focus on the theory of mix systems in order to improve our fundamental understanding of anonymity properties which are possible or practically achievable. This paper takes the fine line between theory and practice and attempts to evaluate the anonymity properties of an abstract anonymous communication system within the practical context of a social network.

We consider the anonymity of users belonging to a social network who communicate with each other via anonymous messages. The attacker is the global passive adversary (she observes the inputs and outputs of the anonymous communication network) and also has knowledge of the users' profiles. First we consider the two sources of information available to the adversary separately, then we combine them and examine what happens as the network grows. Interestingly, it turns out that the details of the mixing algorithm employed by the anonymous communication system play a significant role. Next, we briefly show how additional sources of information can be used by the attacker to further reduce anonymity. Finally, we look at how the quantity, quality and depth of knowledge about the users' relationships affects our results.

Our main contribution is evaluating how the uncertainty in the attacker’s knowledge of user profiles affects anonymity. Indeed, we show that arbitrarily small errors in the profiles can lead to arbitrarily large errors in the anonymity probability distribution and hence point to the wrong subjects in the anonymity set. We develop the intuition behind this result and evaluate the errors in the anonymity probability distributions in the context of the social network. We conduct our experiments by simulation which helps us examine realistic scenarios.

2 Related Work

This paper belongs to a growing body of work focusing on the anonymity analysis of anonymous communication systems. A substantial part of this literature consists of papers evaluating the effectiveness of mix-based anonymity systems in a theoretical setting; e.g., [6, 11, 18]. Such work often involves assumptions such as “users pick their communication partners uniformly at random” which help with the mathematics of calculating anonymity, and hence aid our understanding and intuition, but do not necessarily hold in practice. Furthermore, the authors often examine properties of the anonymous communication systems and shy away from incorporating models of users. This paper takes a more practical approach by assuming a social network, deriving the attacker’s knowledge about users based on the fact that they belong to such a network and then evaluating the performance of the anonymous communication system in the context of this knowledge. Furthermore, we evaluate how errors in the information gained from the social network influence the correctness of the anonymity (and thus, the attacker’s confidence in her result).

In order to evaluate anonymity in a practical setting, it is necessary to incorporate a priori information the attacker might have about communication patterns of users. We briefly mention a number of papers that explore related research problems. Diaz et al. [9] assume that some information on user properties is known, such that the user base can be partitioned in different groups that share a similar profile. Clauß et al. [3, 4] propose a framework and metrics for systems where the adversary has some information on user attributes. In these papers the focus is on user properties or profiles, and little effort is made to combine the knowledge gained through traffic analysis with the profile information available to the attacker. In [3, 4], it is mentioned that the communication layer information gained through traffic analysis can be modeled by means of attributes, but no concrete example is given of how this could be realized. Finally, Diaz et al. [13] showed a toy example where the combination of user sending profiles and data gathered through traffic analysis resulted in higher anonymity, contradicting what had been claimed in [4]. However, no general methodology was given in [13] for computing anonymity metrics when several sources of information are available. The most closely related paper which attempts to combine knowledge about profiles with traffic analysis information is [8] where a lot of the Bayesian theory we use is presented, but only a brief demonstration of the technique is given. Here we give a number of practical examples and evaluate the impact of errors in the profiles on anonymity.

Perhaps the most related piece of related work in terms of the spirit of the analysis and in the style of the results obtained is one of Dingledine and Matthewson [16]. They employ simulations in order to evaluate the effectiveness of statistical disclosure attacks on a model of an anonymity system; i.e., they attempt to recover profiles from the communications data while we build assumptions about profiles from the social network and then add the communications data on top.

3 Preliminaries

3.1 System and Attacker Model

We consider a system where a set U of N users send messages to each other through an anonymous communication channel modeled as a mix¹. Since Chaum [2] first proposed mixes for achieving anonymous communication in 1981, multiple designs have been proposed in the literature both for low-latency communication, e.g. [14] and for high-latency, message-based communication [5, 7, 15].

The adversary we consider can observe all input messages arriving to the mix (and their respective senders), as well as all output messages leaving the mix (and their recipients), but not the internal operations of the mix. Naturally, the messages are encrypted so the content is hidden. Although the attacker does not know the correspondence between inputs and outputs, she is able to compute the probability distributions linking every input with all possible outputs and vice versa.

In addition to observing the mix inputs and outputs, the adversary has a priori knowledge of the users' sending behavior. We assume users to be linked via a social network, and that users send messages to those who are in their *profile*; i.e., their set of "friends." We have used various methods to generate the user sending profiles, which are described in detail in Appendix A.

3.2 Anonymity with One Source of Information

We draw on the literature, more specifically [12] and [17] for our definition of anonymity. The basic idea of these metrics is to use the Shannon entropy [19] of the probability distribution linking subjects to a message or action (normalized entropy in the case of [12]). The entropy of this probability distribution gives a measure of the uncertainty concerning the identity of the subject who originated/received a message. Entropy-based anonymity metrics take into account both the number of users in the system and their probabilities of being linked to a particular action, and anonymity increases both with the number of users and the uniformity of the probability distribution linking them to messages.

The goal of our adversary is to identify the recipient of messages arriving to the mix (*recipient anonymity*) or the sender of messages leaving it (*sender*

¹ Our analysis and experiments apply to any abstract anonymous communication channel for which probabilistic relationships between inputs and outputs can be derived.

anonymity). Therefore, the adversary makes hypotheses of the type “hypothesis h_j is true if u_j is the sender (recipient) of this outgoing (incoming) message,” and computes the probability $\Pr(h_j)$ that h_j is true. Given that every message has one sender and one recipient, the probabilities $\Pr(h_j)$ sum to one (i.e., $\sum_{j=1}^N \Pr(h_j) = 1$).

In this paper we use the *effective anonymity set size* [17] as the metric for sender and recipient anonymity. For a given message entering (leaving) the mix, the recipient (sender) anonymity A is given by the Shannon entropy of the probability distribution of each of the hypotheses h_j being true; i.e., $A = -\sum_{j=1}^N \Pr(h_j) \log_2(\Pr(h_j))$.

Let us first illustrate how anonymity is computed when only one source of information is available to the attacker. If the attacker knows the sending profiles of users, but cannot observe the inputs and outputs of the mix, the recipient anonymity of a message sent by user u belonging to the user population U is given by the entropy of her sending profile. That is, if u chooses user u_j as her recipient with probability $\Pr(u \rightarrow u_j)$, then the recipient anonymity provided by u 's profile is $A_p = -\sum_{j=1}^N \Pr(u \rightarrow u_j) \log_2(\Pr(u \rightarrow u_j))$. Conversely, when u receives a message, the anonymity of the sender is given by $A_p = -\sum_{j=1}^N \Pr(u \leftarrow u_j) \log_2(\Pr(u \leftarrow u_j))$, where $\Pr(u \leftarrow u_j) = \frac{\Pr(u_j \rightarrow u)}{\sum_{k=1}^N \Pr(u_k \rightarrow u)}$ is the probability of u_j being the sender of a message received by u . In the remainder, we denote the sending profile of a user u as $P(u \rightarrow U) = \{\Pr(u \rightarrow u_j), \forall u_j \in U\}$ and its recipient profile as $P(u \leftarrow U) = \{\Pr(u \leftarrow u_j), \forall u_j \in U\}$.

Alternatively, we can consider an adversary who can see the inputs/outputs of the mix but does not have a priori knowledge of user profiles. The probability of an input (output) message matching each of the outputs (inputs) depends on the type of mix, overall traffic load and the timing of messages. Let us consider a timed pool mix. Pool mixes work in cycles called *rounds* that comprise three steps (1) **collect**: it collects messages from senders for a period of time T ; (2) **store**: upon being received, messages are decrypted with the mix's private key (which allows it to retrieve the destination address), and stored in an internal memory called *pool*; and (3) **flush**: once the timeout T has expired, a fraction of the messages are randomly selected and sent to their destinations, while the rest is kept in the pool for the next round.

The probabilities of matching the mix inputs and outputs are computed as follows [10]. Let m_r be the number of messages contained in the mix in round r (prior to the mix flushing), and s_r be the number of messages sent by the mix in round r . If a message M arrived to the mix in round r , its probability $\Pr(M = O_{r',i})$ of matching each of the $s_{r'}$ outputs $O_{r',i}$ that left the mix in round r' is:

$$\begin{aligned} \Pr(M = O_{r',i}) &= 0 \text{ if } r' < r \\ \Pr(M = O_{r',i}) &= \frac{1}{m_{r'}} \text{ if } r' = r \\ \Pr(M = O_{r',i}) &= \frac{1}{m_{r'}} \prod_{k=r}^{r'-1} \left(1 - \frac{s_k}{m_k}\right) \text{ if } r' > r \end{aligned}$$

The recipient anonymity A_m provided by the mix to message M is given by the entropy of the probabilities $\Pr(M = O_{r',i})$. The computation of the probabilities $\Pr(I_{r',i} = M)$ linking an output M to all possible inputs $I_{r',i}$ is analogous, and their detailed derivation can be found in [10]. Note that probabilistic relationships between inputs and outputs can also be derived for other types of mixes such as Stop-and-Go [15].

3.3 Anonymity with Several Sources of Information

Bayesian inference is an approach to statistics in which all forms of uncertainty are expressed in terms of probability. It starts with an initial set of beliefs represented by an a priori probability distribution, which is updated as new evidence is collected. The distribution indicates how likely it is for a hypothesis to be true.

Let h_j be the hypothesis that user u_j is the sender (or recipient) of a given message received (or sent) by user u , and $\Pr(h_j)$ the prior probability of this hypothesis being true. Let E be some evidence or observation that gives us additional information on the truthfulness of h_j , and $\Pr(E|h_j)$ be the probability of observing evidence E conditioned to h_j being true. Bayesian inference can be used to compute the posterior probability $\Pr(h_j|E)$ of h_j , given that we have obtained evidence E . We denote this probability distribution by $P(H|E) = \{\Pr(h_j|E), 1 \leq j \leq N\}$:

$$\Pr(h_j|E) = \frac{\Pr(h_j) \Pr(E|h_j)}{\sum_{k=1}^N \Pr(h_k) \Pr(E|h_k)}$$

In our setting, we consider that both sender profiles and mix input/output observations are available to the adversary. The prior probability $\Pr(h_j)$ is given by the sending profiles of users, and corresponds to $\Pr(a \rightarrow u_j)$ in the case of recipient anonymity, and to $\Pr(a \leftarrow u_j)$ for sender anonymity (as explained in the previous section). The conditional probability $\Pr(E|h_j)$ is computed as follows. For recipient anonymity (analogous for sender anonymity), let R_j be the set of messages received by user u_j . Given that u sent message M to u_j (i.e., h_j is true), the probability $\Pr(E|h_j)$ of observing the evidence E corresponds to the probability of the mix matching M to one of the messages received by u_j :

$$\Pr(E|h_j) = \sum_{O_{r',i} \in R_j} \Pr(M = O_{r',i})$$

Bayesian inference can be applied recursively if new independent evidence E' becomes available to the adversary. We show results that introduce an additional source of information in Sect. 5.2.

4 Analysis

4.1 Intuition

The attackers' knowledge about the communication partners of users inside the social network comes from two sources—observing the mix and her a priori

knowledge of the user profiles. Naturally, if we have a perfectly anonymous communication layer, the anonymity of the system comes only from the attacker's (lack of) information about the profiles. Conversely if the attacker has no information on the profiles of the users, she is restricted to observing the communications layer; i.e., the mix. The more complex setting when the attacker has knowledge of both is examined below.

Consider the case of users belonging to a vast social network and hence knowing a tiny fraction of the overall user population. In our model the attacker can see the inputs and the outputs of the mix and knows the profiles of all the users, so the only mixing that will take place is that between senders who share potential recipients or between recipients who share potential senders. Hence if the network grows and users' connectivity remains constant, anonymity falls. On the other hand, higher traffic load and number of users increase the anonymity provided by the mix. In Sect. 5.1 we show the tradeoff between these two effects.

The increasing popularity of blogs and, more generally, the availability of user-generated content makes it easy to gather a corpus of text linkable to an individual. Different people have different writing styles and patterns (such as word frequency or preferred grammatical constructions), and statistical tests that detect these patterns can be used to help identifying the authors of anonymous text. We study in Sect. 5.2 how the results of such a test can be combined with profiles and traffic analysis information, and its impact on sender anonymity.

The attacker's knowledge of the social network can vary in its quantity, quality and depth. She may know only of existence of links between individuals, the extent of those links, lack knowledge of links in some part of the network and hence have to make do with approximations or, worst of all, assume wrong information. We assess the impact of each of these on anonymity in Sections 5.3, 5.4 and 5.5. Before proceeding to the results of the analysis, we give details of our experimental setup.

4.2 Experimental Setup

We performed the analysis in the setting of a social network with a population of users arranged in a small-world network constructed following the Watts-Strogatz algorithm [21]. We also performed experiments on a scale free network [1] created with preferential attachment and the same number of average users, and the only noticeable difference was a larger variance in the results, which is due to the more uneven distribution of links per node in these networks. Unless indicated otherwise, we consider in our experiments 1000 users with an average of 20 friends each, arranged in a small world network with parameter $p = 0.1$ (i.e., highly clustered).

Users send messages only to their *friends* (i.e., users linked to them in the social network) with the probability specified in their profile. For the purposes of our experiments, we have developed several sets of user profiles with slightly different probability distributions. A detailed summary of the profiles used and the algorithms used to generate them can be found in Appendix A.

We chose a Mixmaster [5, 20] mix, as it is the most widely deployed high-latency network for anonymous email. The time intervals between users sending messages follow an exponential distribution with parameter λ , common to all users. We have chosen $1/\lambda$ to be 25 times greater than the timeout of the mix, so if users send messages on average once a day, the expected delay is between 30 mins and 1 hour. In every experiment we simulate 130 rounds of mixing. We then extract the information which could have been observed by the attacker and compute the sender and receiver anonymity of each message.

5 Results

5.1 Growing the Network

In this section we consider the anonymity of users as the social network is scaled up. To help develop the intuition we show the anonymity calculated from traffic analysis (mix input/output observations) and knowledge of the profiles separately. As the network grows, the anonymity provided by the mix increases as shown in Fig. 1(a)(Mix) simply because more traffic goes through it. As for the anonymity provided by the profiles (corresponding to Uniform profiles in Appendix A), we can see in Fig. 1(a)(Profile) that it remains constant, because we assume that the connectivity does not increase with the network (though in a real network it might increase slightly), which becomes more sparse. Interestingly, Fig. 1(a)(Combined) shows that the combined anonymity decreases with the network size. As we shall see, variations in parameters that have a positive (mix) or no (profiles) effect on anonymity when sources of information are considered separately, can have a negative impact when all information is put together.

In this particular case the decrease in anonymity with network size is due to an interaction between profiles and mix function. Consider a random user Alice. The attacker is aware of her sender profile, so only users who share friends with Alice contribute to her anonymity. Alice and her friends send and receive on average the same number of messages whether the network is large or small. At the same time, the Mixmaster function [5] that determines the fraction f of messages sent per round increases with the traffic load until it reaches its limit²—note that in Fig. 1(a)(Combined) anonymity stabilizes beyond that point. Therefore, the larger network induces the mix to flush a higher fraction of messages, which consequently in the mix for fewer rounds. This effect, in fact, decreases the amount of mixing, because friends of Alice who sent or received messages in the rounds before or after her contribute less to her anonymity³.

² The maximum fraction of messages sent by Mixmaster is $f = 0.65$. In our setting, this is reached when there are around 2500 users.

³ Friends who sent messages during the same round as Alice contribute the same amount as in the case of the smaller network.

5.2 Adding Extra Information

In this section we briefly show how Bayesian inference can be used to incorporate additional sources of information. Consider, for instance, a writing pattern recognition test. Let us assume that the attacker can run a test on the messages at the output of the mix and compare the writing to available text from the potential senders. This test outputs a true positive result with probability p_t and a false positive with probability p_f , and therefore produces as result a set of positives U_p and a set of negatives U_n .

Based on the evidence E' produced by the test, the adversary can derive for each user u_j the probability $\Pr(h_j|E')$ that she was the true author of the text. Users testing negative (i.e., $u_j \in U_n$) have probability $\Pr(h_j|E' = 0)$ of being the writer, while those testing positive (i.e., $u_j \in U_p$) are the originator of the message with probability $\Pr(h_j|E' = 1)$.

The posterior probability distribution $P(H|E')$ is computed applying Bayesian inference as explained in Sect. 3.3. The evidence E' is a vector with zeros for users who tested negative and ones for those who tested positive. The prior $P(H)$ corresponds to the (already existing) probability distribution that combines the profile and traffic analysis information. Assuming that E' contains k positives for a population of N users, $P(E'|H)$ is computed as follows:

$$\Pr(E' = 0|h_j) = (1 - p_t) \binom{N-1}{k} p_f^k (1 - p_f)^{N-k-1}$$

$$\Pr(E' = 1|h_j) = p_t \binom{N}{k-1} p_f^{k-1} (1 - p_f)^{N-k}$$

We made experiments where we considered two tests that give correct answers with different degrees of accuracy. The high accuracy test had a true positive rate $p_t = 0.8$ and a false positive rate $p_f = 0.01$, while in the low accuracy one these values were $p_t = 0.5$ and $p_f = 0.1$. The results are shown in Figure 1(b), where we can see how the new information provided by the test reduces (on average) sender anonymity. Note however the outliers: in some instances, the additional information provided by text recognition test does not help reducing anonymity. We further investigate this effect in Sect. 5.6.

5.3 Quantity of Profile Knowledge

In the previous section we compared anonymity in these cases: (i) the adversary knows the profiles of all users, but cannot perform traffic analysis; (ii) the adversary does not know any profiles, but can observe the mix; and (iii) the adversary has access to all profiles and communication data. Here, we look at sender and recipient anonymity towards adversaries who can observe all traffic through the mix but only know a fraction of the user profiles (generated following the Uniform description in Appendix A). We assume that the attacker has perfect knowledge of some profiles, and knows nothing about the rest. Whenever the attacker does not know a profile, she will consider it as uniform.

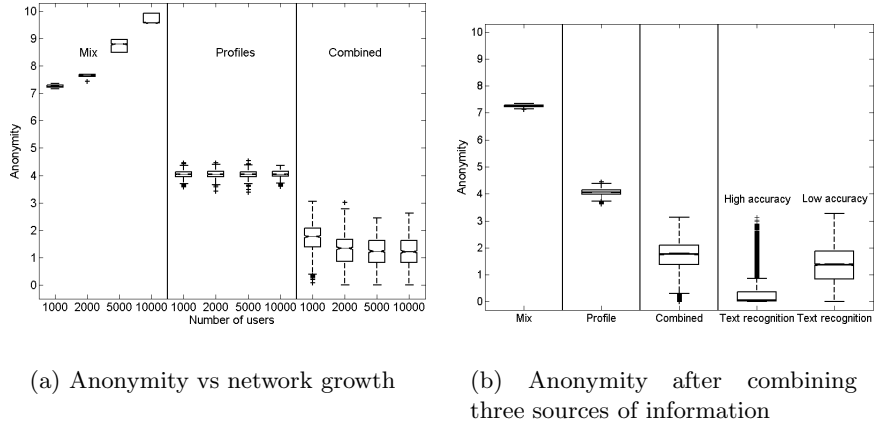


Fig. 1. Sender anonymity when various sources of information are available

In Figure 2(a) we show the results for recipient anonymity with respect to the percentage of known profiles. On the left hand side of the figure (Mix) we show the anonymity A_m provided by the mix, which is independent from the quantity of profile knowledge and thus invariant for all experiments. The center and right hand side show, respectively, the anonymity A_p of the profiles and the combined A_c . Recipients of users with unknown profiles are unaffected by the percentage of known profiles, and enjoy the maximum anonymity that the mix can offer. For them, the profile anonymity is $A_p = \log_2(N)$ and the combined recipient anonymity is $A_c = A_m$. Conversely, recipients of profiled users do not benefit from unknown profiles, and their recipient anonymity is the same regardless of how many other profiles are known. The aggregation of these two sets of recipient anonymity results can be clearly seen in the box plots of Fig. 2(a). Note the sudden jump in the median when half the profiles are known, and the values of the quartiles and outliers.

Unlike in the case of receiver anonymity, the percentage of known profiles affects the sender anonymity of all users, profiled or not, in the same way. This is because recipient profiles $P(u_i \leftarrow U)$ are computed using all sender profiles (see Sect. 3.2), and unpredictability of some users' sending patterns introduces uncertainty for all messages. The results of our experiments are shown in Fig 2(b)—as more profiles become available to the attacker, the sender profile and combined anonymity decrease.

Note that although the behaviour of sender and recipient anonymity is different when the adversary has partial knowledge, the values are the same for the extremes—i.e., sender and recipient anonymity are symmetric (in their distribution of values) both when all profiles are known and when all profiles are unknown, but not when some profiles are and some are not.

Finally, note that in our experiment all users have non-uniform sending profiles (they only send messages to their friends), so the adversary's assumption of uniform behaviour for unknown users introduces errors in her results. We further

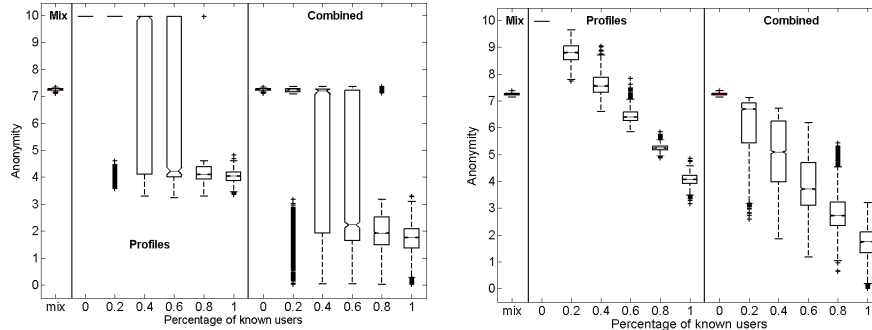


Fig. 2. Receiver (left) and sender (right) anonymity depending on the quantity of profile knowledge

elaborate on the implications of having (or assuming) wrong information in the next section.

5.4 Quality of Profile Knowledge

Human behaviour is hard to model and predict, and even the most sophisticated adversary with access to vast amounts of information can only at best approximate user behavioural profiles. Therefore, we can reasonably assume that in a real world scenario there is going to be some difference between the profiles guessed or predicted by the adversary and the actual user sending patterns. Furthermore, due to the lack of available real-world data, little is known about how user sending profiles might actually look like, or how they evolve in time. For this reason, it is worth looking at the implications for the anonymity adversary of making wrong behavioural assumptions, such as assuming uniform sending profiles. In this section we study how noise in the profiles propagates and find that small errors in the profiles may lead to big errors in the end results.

There are many ways for the adversary to construct her guessed profiles. They can be obtained, to mention some examples, by studying the links between users in online social networks such as Facebook or LiveJournal, by analyzing user sending patterns when messages are sent over a non-anonymous channel (assuming that the user does not always use the mix for sending her messages), or by applying statistical disclosure attacks [8] to previous mix communications of the user. The profile construction method and the quality of data available to the adversary determine not only the accuracy of the profile, but also the nature of the “error” with respect to the real profile. For example, users may be linked in Facebook to acquaintances to whom they rarely or never send messages; they may have friends to whom they only communicate through an anonymous channel (and therefore do not appear in their non-anonymous communications); and the profiles obtained through disclosure attacks are noisy versions of the real sending patterns. Such a wide range of possibilities makes it hard to predict

the type of profile errors we can expect in a real world scenario, and has led us to consider various kinds of erroneous profiles.

One important thing to note is the independence between error magnitude and actual anonymity value. Small errors in the final result indicate that the probability distribution obtained by the adversary is roughly similar to the one she would obtain had she used the true profiles; while large errors indicate that the adversary’s view on who are the likely senders or receivers of a message is very different from the actual distribution computed with the real profiles—regardless of the entropy of the actual (guessed and true) distributions. Anonymity gives a measure of the adversary’s uncertainty on who are the likely senders or recipients messages given that all available information is correct; while errors model the uncertainty of the adversary concerning the accuracy of her anonymity results, assuming that some information may not be correct.

In order to measure and compare the magnitude of the errors in the profile and final result making abstraction of the nature of the error, we use as metric the Euclidean distance $dist(x, y) = \sqrt{\sum_i (x(i) - y(i))^2}$ between true and guessed probability distributions. We have chosen Euclidean distance for its simplicity and well understood meaning, and because it provides clear bounds for the final error—the maximum distance between two probability distributions occurs when they are orthogonal; its value is $\sqrt{2}$ and the minimum distance is 0.

Let us illustrate with a toy example our method for quantifying the impact of errors and the meaning of our results. Consider a simple scenario as the one depicted in Fig. 3, with a population $U = \{A, B, C, \dots, Z\}$ and a unique (threshold) mixing round. User A sends with uniform probability $\Pr(A \rightarrow u_j) = 1/4$ to each of her four friends $\{B, C, D, F\}$, and with $\Pr(A \rightarrow u_j) = 0$ to the other users. The attacker, however, has a noisy version of A ’s profile, and believes that she chooses uniformly from the set $\{B, C, D, Z\}$. The attacker sees a single round of a threshold mix where A sends a message which comes out to either F or Z . Naturally, it was F as Z is not in A ’s true set of friends. The attacker, however believes it is Z , because he thinks that Z rather than F is in A ’s set of friends. Hence he wrong profile has led the attacker that Z is the recipient with probability one. We note that in this example, the receiver anonymity computed by the attacker when considering the wrong profile is zero ($A_{attacker} = 0$), as is the one she would obtain if she had precise knowledge of A ’s sending behavior ($A_{true} = 0$). However, the probability distribution obtained by the attacker is very different from the true result, and consequently her error is large. As the distance between the true and wrong results is much larger than the distance between the true and wrong profiles, this example provides the intuition that small errors in the profile may lead the attacker to completely wrong results.

Given that it is hard to predict the type of error the adversary is most likely to make, we have tested multiple instances of erroneous profiles. These include: (i) adding a *tail* to the profile distribution so that the probability of sending to non-friends appears greater than zero—yet significantly smaller than the one assigned to friends; (ii) introducing *Gaussian* noise; (iii) *eliminating* or (iv) *swapping* friends; and (v) assuming *uniform* behaviour. Appendix B provides a




Fig. 3. Example of how small errors in the profile can induce large errors in the attacker’s results

detailed overview of the types of errors we have considered and the algorithms used to generated them.

The results of our experiments are shown in Figures 4(a) and 4(b). In both figures, the X axis represents the distance between the true user profiles (with which the messages were generated) and the erroneous profiles considered by the attacker; i.e., $D_p = \text{dist}(P_{\text{true}}(A \rightarrow U), P_{\text{attacker}}(A \rightarrow U))$. The Y axis expresses the distance between the probability distributions the attacker would obtain with the correct and wrong profiles; i.e., $D_c = \text{dist}(P_{\text{true}}(H|E), P_{\text{attacker}}(H|E))$. The grey dots include results of experiments generated with the five error methods previously mentioned, and we have highlighted in black the results for two types of errors: adding a *tail* to the profile distribution (Fig. 4(a)) and assuming *uniform* profiles (Fig. 4(b)). We can see that the errors induced by adding a *tail* to the profile are relatively benign compared to other types in the background, as they take mostly low values in Y (note that this is the type of error obtained when learning users’ profiles with a statistical disclosure attack). On the other hand, whenever the adversary (due to lack of information) assumes users send uniformly, she obtains a distribution that substantially deviates from the correct result—to the extent that she cannot have any confidence on whether or not she is getting a good approximation to the correct anonymity set. This is aggravated when we consider errors coming from swapping or eliminating friends, which cover most of grey area.

5.5 Depth of Profile Knowledge

In some practical scenarios (e.g., Facebook) the adversary may guess the friendship graph but lack enough data to estimate the strength of links between friends. We say that the adversary’s guessed profiles lack *depth* when she cannot estimate the frequency with which friends are chosen as recipients, in spite of accurately distinguishing friends from non-friends (to whom users never send messages). In these circumstances, the best the adversary can do is to consider that recipients are picked uniformly at random from the set of friends. This is a special case of

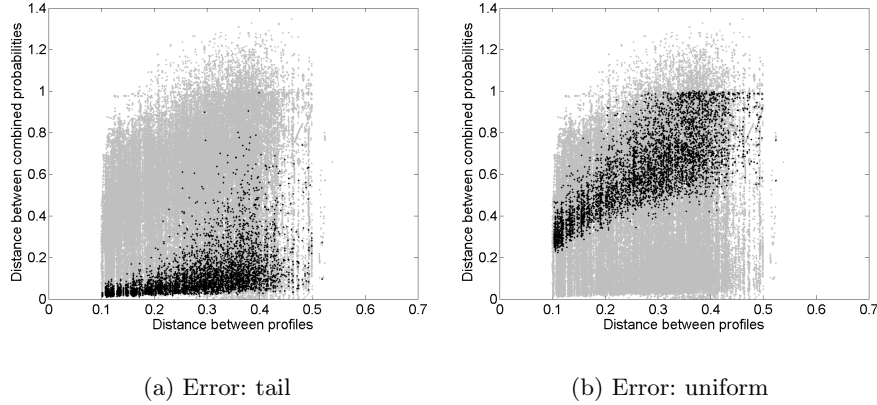


Fig. 4. Euclidean distance between true and guessed probability distributions vs distance between true and guessed profiles (quality of profile knowledge)

erroneous profiles like those analyzed in the previous section, but we have chosen to present it separately for two reasons: first, because of its practical relevance (such profiles would be reasonably easy to construct); and second, although the profiles are noisy, correctly identifying friends (and non-friends) already provides very valuable information to the attacker.

To better illustrate the impact of the attacker’s assumption, we consider that users choose their partners of communication having strong preferences for some of them (Skewed in Appendix A). In Fig 5 we show how the error in the combined probability increases proportionally to the error in the profile. When the true profile of a user is close to uniform⁴, the assumption of the attacker is not far from the truth—the distance D_p between both profiles is small, and so the distance D_c between the combined distributions. As D_p increases, so does D_c , but as a rule of thumb we could say that the error D_c is most likely to be smaller than the original error D_p . The contrast with the previous section’s results (considering profiles uniform in the whole population) indicates that an adversary who correctly identifies friendship links obtains two advantages: she eliminates non-friends from the anonymity sets, effectively decreasing anonymity; and she has higher confidence in her result, because the true and guessed distributions are comparatively closer to each other.

5.6 How Often Does Additional Information Reduce Uncertainty?

It was pointed out in [13] that in some cases additional information may result in higher anonymity, even if on average anonymity decreases as more information

⁴ Because of the algorithm used to generate the profiles (see App. A), recipient profiles are on average more uniform than sender profiles, this explains why the values in Fig 5(b) are smaller than in Fig 5(a).




Fig. 5. Receiver (left) and sender (right) anonymity error depending on depth of knowledge

becomes available. In this section we present some results showing under which conditions we can expect these cases to appear. In all experiments we used Mixmaster (i.e., the anonymity A_m provided by the mix is invariant), and a small world network with 1000 users that send to friends with probability Pr_f and to non-friends with Pr_{nf} , such that $0 \leq \text{Pr}_{nf} \leq \text{Pr}_f$. The details of the generation of profiles is available in Appendix A, under the name Step. We study the results according to two variables: the number F of friends per user, and a parameter $0 \leq D \leq \infty$ that tunes the difference between Pr_f and Pr_{nf} , such that $D = 0$ implies $\text{Pr}_{nf} = 0$, and $D = \infty$ implies $\text{Pr}_{nf} = \text{Pr}_f$.

To better understand how sending behaviour affects anonymity, we have studied separately the frequency of cases where the combined anonymity A_c is higher than the anonymity of the mix alone A_m or the profile A_p , and its variation with the parameters F and D . The results in Figs. 6(a) and 6(b) show, respectively, the percentages of messages for which $A_c > A_m$ and $A_c > A_p$, which we denote $f_{c>m}$ and $f_{c>p}$.

To interpret the results, note that increasing F and/or D leaves A_m constant; increases A_p (because it makes the profile more uniform); and A_c increases as well as a result of more uniform profiles. When $D = 0$ users *only* send to friends—i.e., the recipient anonymity set is reduced drastically—and A_c is always lower than A_m and A_p . For $0 < D < 1$ and small F , A_p has increased only slightly, while A_c benefits mostly from messages sent to non-friends—these are “rare⁵ events” in which the hints coming from the mix and the profile are “contradictory.” Given the profile always points to the highest probability friends, when the mix points to (less probable) non-friends as most likely recipients, the mix and profile distributions compensate instead of reinforcing each other, making the combined distribution more uniform than one or both originals—i.e., $A_c > A_p$ and/or $A_c > A_m$. This also explains the high $f_{c>m}$ for larger values of D . Once F and/or D grow to make $A_p > A_m$, it becomes harder for A_c to catch up with

⁵ Note that for $D = 1$ half the messages are sent to non-friends, even if the probability of picking a concrete non-friend is small.

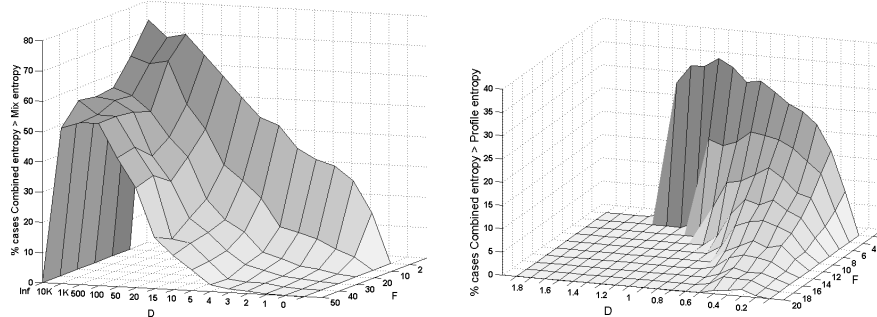


Fig. 6. Percentage of cases where the combined anonymity is higher than the anonymity of the mix only $f_{c>m}$ (a) and profile only $f_{c>p}$ (b)

it (we can see in the Fig. 6(b) that $f_{c>p} = 0$ for $F > 15$ and/or $D > 1.25$). When A_p hits its maximum with *perfectly* uniform profiles at $D = \infty$, the profiles stop bringing any additional information and $A_c = A_m$. Thus, $f_{c>m} = 0$ at $D = \infty$ in Fig. 6(a).

6 Conclusions and Future Work

In this paper we examined the anonymity of users in the practical context of a social network. We showed the overall anonymity is low and likely does not increase with the size of the social network—if anything, it decreases as the network becomes more sparse.

The positive result of this paper is that it is necessary to trust the social network entirely to provide high quality information about the sender profiles of the users, otherwise big mistakes can be made in the sender and receiver anonymity of messages. Indeed, unless the profile is perfect, the results may be meaningless as we demonstrated occurrences of huge errors in the anonymity probability distribution even when the profile error is small. We have found however that certain types of errors induce more bounded deviations than others in the overall anonymity.

Many issues remain to be addressed, particularly in the practical setting. Particularly interesting to us is the problem of assessing the anonymity of a real social network such as Facebook and its approximation as mapped by the attacker. Although we believe that we modeled the “friendship” between users to a fair degree of accuracy by using a Watts-Strogatz graph, the extent of the linkage and the resulting sender profiles remain a more difficult issue. Only empirical modeling can gauge how much the real social dynamics differ from the theoretical models employed here.

One extremely promising line of research is to set up and evaluate an attack where the adversary continuously updates the social network graph with new

information gained from observing the communication patterns and simultaneously tries to deanonymize the messages. Interestingly, the result of our paper holds in this setting too – whatever the methodology of deriving the social network graph, small errors in the graph may cause large errors in the anonymity of the message. Although complex statistical disclosure attacks may prove efficient at minimizing the errors in the graph, they can never eliminate such inaccuracies which may arise as a result of external factors, for instance changes of user behaviour over time.

References

1. Barabasi, A.-L., Albert, R.: Emergence of scaling in random networks. *Science* 286, 509–512 (1999)
2. Chaum, D.: Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM* 24(2), 84–88 (1981)
3. Clauß, S.: A framework for quantification of linkability within a privacy-enhancing identity management system. In: Müller, G. (ed.) *ETRICS 2006*. LNCS, vol. 3995, pp. 191–205. Springer, Heidelberg (2006)
4. Clauß, S., Schiffner, S.: Structuring anonymity metrics. In: *Proceedings of the ACM Workshop on Digital Identity Management*, pp. 55–62 (2006)
5. Cottrell, L.: Mixmaster & remailer attacks (unpublished manuscript), <http://www.obscura.com/~loki/remailer/remailer-essay.html>
6. Danezis, G.: The traffic analysis of continuous-time mixes. In: Martin, D., Serjantov, A. (eds.) *PET 2004*. LNCS, vol. 3424, pp. 35–50. Springer, Heidelberg (2005)
7. Danezis, G., Dingleline, R., Mathewson, N.: Mixminion: Design of a type iii anonymous remailer protocol. In: *Proceedings of the 2003 IEEE Symposium on Security and Privacy*, pp. 2–15 (2003)
8. Danezis, G., Serjantov, A.: Statistical disclosure or intersection attacks on anonymity systems. In: Fridrich, J. (ed.) *IH 2004*. LNCS, vol. 3200. Springer, Heidelberg (2004)
9. Diaz, C., Claessens, J., Seys, S., Preneel, B.: Information theory and anonymity. In: Macq, B., Quisquater, J.-J. (eds.) *Werkgemeinschaft voor Informatie en Communicatietheorie*, pp. 179–186 (2002)
10. Diaz, C., Preneel, B.: Reasoning about the anonymity provided by pool mixes that generate dummy traffic. In: Fridrich, J. (ed.) *IH 2004*. LNCS, vol. 3200, pp. 309–325. Springer, Heidelberg (2004)
11. Díaz, C., Serjantov, A.: Generalising mixes. In: Dingleline, R. (ed.) *PET 2003*. LNCS, vol. 2760, pp. 18–31. Springer, Heidelberg (2003)
12. Diaz, C., Seys, S., Claessens, J., Preneel, B.: Towards measuring anonymity. In: Dingleline, R., Syverson, P.F. (eds.) *PET 2002*. LNCS, vol. 2482, pp. 54–68. Springer, Heidelberg (2003)
13. Diaz, C., Troncoso, C., Danezis, G.: Does additional information always reduce anonymity? In: Yu, T. (ed.) *Workshop on Privacy in the Electronic Society 2007*, pp. 72–75. ACM, New York (2007)
14. Dingleline, R., Mathewson, N., Syverson, P.: Tor: The second-generation onion router. In: *Proceedings of the 13th USENIX Security Symposium*, pp. 303–320. USENIX (2004)

15. Kesdogan, D., Egner, J., Büschkes, R.: Stop-and-go MIXes: Providing probabilistic anonymity in an open system. In: Aucsmith, D. (ed.) IH 1998. LNCS, vol. 1525, pp. 83–98. Springer, Heidelberg (1998)
16. Mathewson, N., Dingleline, R.: Practical traffic analysis: Extending and resisting statistical disclosure. In: Martin, D., Serjantov, A. (eds.) PET 2004. LNCS, vol. 3424, pp. 17–34. Springer, Heidelberg (2005)
17. Serjantov, A., Danezis, G.: Towards an information theoretic metric for anonymity. In: Dingleline, R., Syverson, P.F. (eds.) PET 2002. LNCS, vol. 2482, pp. 41–53. Springer, Heidelberg (2003)
18. Serjantov, A., Dingleline, R., Syverson, P.: From a trickle to a flood: Active attacks on several mix types. In: Petitcolas, F. (ed.) IH 2002. LNCS, vol. 2578. Springer, Heidelberg (2003)
19. Shannon, C.: A mathematical theory of communication. The Bell System Technical Journal 27, 379–423, 623–656 (1948)
20. Ulf Moller, P.P., Cottrel, L., Sassaman, L.: Mixmaster protocol - version 2 (2003), <http://www.abditum.com/mixmaster-spec.txt>
21. Watts, D., Strogatz, S.: Collective dynamics of small-world networks. Nature 393, 440–442 (1998)

A User Profiles

In order to create a diverse testbed for our experiments, we defined four different kinds of user profiles. We create a profile $P(u \rightarrow U)$ for each user u in the set U of N users connected through a friendship graph, and we say that u_j is a “friend” of u if they share an edge in the graph and a non-friend otherwise. In the following, we denote the set of friends of u as $f_u := \{u_j \in U \mid Pr(u \rightarrow u_j) \neq 0\}$ (cardinality F), and the set of non-friends as $nf_u := \{u_j \in U \mid Pr(u \rightarrow u_j) = 0\}$ (cardinality $N - F$). The three first types of profiles (*Uniform*, *Random* and *Skewed*) restrict users to sending only to their friends, while the fourth type (*Step*) allows users to send to non-friends with a smaller probability than to friends.

P_U : Uniform. users who send messages according to a P_U profile pick their recipient uniformly at random from their set of friends. $P(u \rightarrow U)$ is defined as:

$$P(u \rightarrow u_j) = \begin{cases} \frac{1}{F} & \text{if } u_j \in f_u \\ 0 & \text{if } u_j \notin f_u \end{cases}$$

P_R : Random. in this setting users send non-uniformly to their friends, but they have no particularly strong preferences for any of them. Hence, this profile can be considered a noisy version of P_U . P_R is created by generating a random number between 0 and 1 for each friend, and normalizing the resulting distribution.

P_K : Skewed. users whose profile is P_K usually have strong preferences for a small subset of their friends who are chosen as recipients very frequently, while the others only appear sporadically. The algorithm to generate P_K starts defining $\mu = 1$ as the initial probability “budget” available. Then we recursively

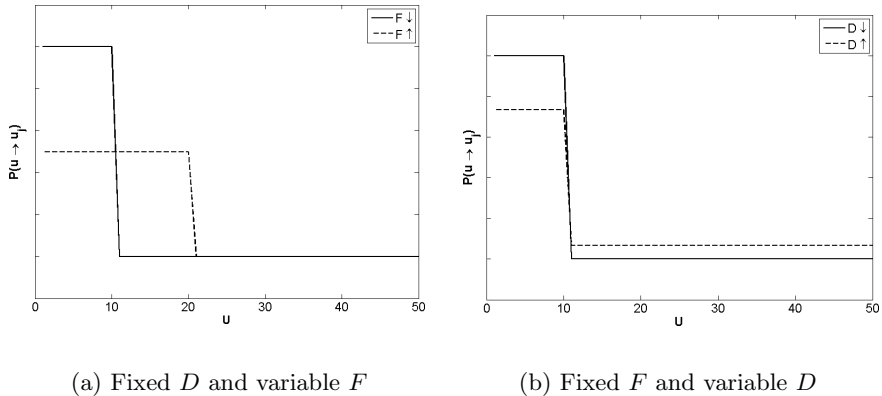


Fig. 7. Variation of P_T (Step) profiles with F and D

assign to each friend a probability p chosen at uniformly at random from the interval $[0, \mu]$, and update the value $\mu = \mu - p$ describing the remaining budget. We repeat the procedure until only one friend is left, to whom we assign the remaining probability μ .

P_T : Step. users with these profiles send messages to the whole population. Nevertheless, they choose their friends as recipients more frequently than non-friends. For user u , the probability assigned in her profile to each of her friends u_f is $\Pr_f = \frac{1/F + D/N}{1+D}$, while the probability assigned to each non-friend u_{nf} is $\Pr_{nf} = \frac{D/N}{1+D}$. F is the cardinality of the set of friends, and the influence of its variation in the profile can be seen in Fig. 7(a). The parameter D influences the relation, in terms of probability, between friends and non-friends. As D increases, the sending profile becomes more uniform in all N potential recipients, diminishing the difference between friends and non-friends, as shown in Fig. 7(b). For $D = 0$, users never send to non-friends, and profiles are uniform on the whole population for $D = \infty$.

B Erroneous Profiles

We simulate the adversary’s imprecise information as follows. For each user $u \in U$ we take her true profile $P(u \rightarrow U)$, generated as explained in Appendix A, and we create a set of “erroneous profiles”, $P_{attacker,i}(u \rightarrow U)$, by applying one of the following transformations:

Tail: we consider that if the adversary does not have accurate knowledge of u ’s profile, she will rather not exclude any user as potential contact of u (note that a similar profile shape is obtained after applying statistical disclosure attacks, with friends getting higher probabilities and non friends getting lower—but not zero—probabilities). We model this by distributing 20% of

the total probability to the set nf_u , and we subtract this probability uniformly from the set of friends f_u , so that the new profile probabilities add up to one. We use this profile as basis when introducing the following errors.

Gaussian: we create two sets of profiles with this method, where we first add Gaussian noise to each element in the profile and then normalize. The noise samples come from two normal distributions $N(\mu_i, \sigma_i^2)$ with $\mu_1 = 0.01$, $\mu_2 = 0.05$, and $\sigma_1 = \sigma_2 = 0.3$.

Eliminate: this error emulates situations where the attacker misses one or more friends of u in her approximation of the profile, and considers them as non-friends. As explained before, becoming a non-friend does not discard a user as potential receiver of u , but it reduces her probability in u 's profile. In our experiments we eliminate an increasing number of friends until only one remains. Each time a friend is eliminated the probabilities of the remaining friends are increased to compensate.

Swap: this error models the case where the attacker not only misses some friends, but wrongly considers non-friends as likely recipients. This effect is modeled by swapping (between one and all) the elements of f_u with elements of nf_u , i.e. when a friend is erased from the set of contacts, a non friend takes his place..

Uniform: this error simulates the case where the attacker has no knowledge about the social network, and thus considers all profiles as uniform over all population.