

# Perfect Matching Disclosure Attacks

Carmela Troncoso, Benedikt Gierlichs, Bart Preneel, and Ingrid Verbauwhede

K.U. Leuven, ESAT/SCD-COSIC, IBBT  
Kasteelpark Arenberg 10, B-3001 Leuven-Heverlee, Belgium  
`firstname.lastname@esat.kuleuven.be`

**Abstract.** Traffic analysis is the best known approach to uncover relationships amongst users of anonymous communication systems, such as mix networks. Surprisingly, all previously published techniques require very specific user behavior to break the anonymity provided by mixes. At the same time, it is also well known that none of the considered user models reflects realistic behavior which casts some doubt on previous work with respect to real-life scenarios. We first present a user behavior model that, to the best of our knowledge, is the least restrictive scheme considered so far. Second, we develop the Perfect Matching Disclosure Attack, an efficient attack based on graph theory that operates without any assumption on user behavior. The attack is highly effective when de-anonymizing mixing rounds because it considers all users in a round at once, rather than single users iteratively. Furthermore, the extracted sender-receiver relationships can be used to enhance user profile estimations. We extensively study the effectiveness and efficiency of our attack and previous work when de-anonymizing users communicating through a threshold mix. Empirical results show the advantage of our proposal. We also show how the attack can be refined and adapted to different scenarios including pool mixes, and how precision can be traded in for speed, which might be desirable in certain cases.

## 1 Introduction

Traffic analysis exploits traffic data to infer information about observed communications. It is the most powerful known attack against anonymous networks. More precisely, Disclosure (or Intersection) attacks use the fact that users' communication patterns are repetitive to uncover communication relationships between them [1,4].

Previous work on Disclosure Attacks [1,4,5] considers a very simplistic model, where users send messages to a fixed set of contacts through a threshold mix. Users choose amongst their communication partners with uniform probability and the effectiveness of these attacks strongly relies on this model. In this paper we present a new attack, the Perfect Matching Disclosure Attack, that requires no assumption on the users' behavior in order to reveal their relationships. Besides its capability to uncover relations amongst users, *i.e.* their sending profiles, in an arbitrary scenario, we demonstrate the strength of our attack in

de-anonymizing individual messages, *i.e.* finding the links between messages arriving to the network and messages leaving it. Our method’s advantage stems from the fact that it considers all users in a round at once, rather than single users iteratively. This approach is likely to de-anonymize a large fraction of the set correctly in scenarios where a per user approach fails with high probability.

We analyze and compare the Statistical Disclosure Attack (SDA) and the Perfect Matching Disclosure Attack (PMDA) empirically in two scenarios. In both scenarios, we chose a simple threshold mix as communication channel such that we can focus on presenting our techniques. With respect to a simple user behavior model we observe that the SDA and the PMDA perform very similar. In a generic user model the PMDA outperforms the SDA for a limited increase of computational cost. Simulation results show that our method is more accurate when linking senders and receivers of de-anonymized messages and that it allows to derive better estimations of users’ profiles. We also propose the Normalized SDA, a trade-off between precision and speed, which yields results nearly as good as the PMDA with a running time slightly higher than the one of the original SDA.

This paper is organized as follows. Section 2 provides an overview of the state-of-the-art of attacks on mix networks. We explain the system model and our models for user behavior in Sect. 3. Section 4 describes the mathematical background for our attack and its application to a threshold mix. In Sect. 5 we show how our attack and the SDA can be applied in practice. An evaluation of both methods is presented in Sect. 6. We explain in Sect. 7 how to construct enhanced user profiles while Sect. 8 deals with further improvements and variants of the PMDA. Finally, we pose some open questions and conclude in Sect. 9.

## 2 Related Work

Mixes were proposed by David Chaum [3]. Chaum’s proposal consists of a router that receives a number of messages of fixed length, performs some cryptographic operations on them changing their appearance and outputs the result in a random order. This ensures that linking inputs and outputs based on timing information is impossible. Mixes can be combined in networks, such that even if a mix is compromised, the user’s anonymity is guaranteed. They are widely used in the literature to implement anonymous email [6,17] or e-voting protocols [11,14].

Although mix networks provide good anonymity, they are vulnerable to long-term traffic analysis attacks. An attacker who observes a mix network can collect what is called traffic data: the identities of the messages’ senders and receivers, together with the timing of these events. The family of Disclosure Attacks [1,12] aims at identifying users’ communication patterns. It is assumed that the participants communicate through a threshold mix. This mix collects a certain number of messages per round, and outputs them in a random order. Applying the Disclosure Attack, an adversary observing the mix, *i.e.* senders and receivers per round, over enough time can uncover the set of Alice’s friends. Nevertheless, the Disclosure Attack is very expensive, as it relies on solving an NP-problem and

is only feasible for very small systems. A more efficient approach to obtain the exact solution, the Hitting Set Attack was proposed in [13].

Danezis presents a different efficient approach, the Statistical Disclosure Attack [4], which reveals the most likely set of Alice’s friends using statistical methods and approximations assuming the same model as in [1]. The attack model was extended to include anonymous replies in [5] and to consider a pool mix instead of a threshold mix in [7]. More complex models are analyzed and tested by simulation in [16]. Besides discovering a user’s set of friends, Danezis proposes in [4] to use Alice’s sending profile derived in the attack to individually trace each of the messages she sends to the network.

An approach to measure anonymity has been developed independently by Edman et al. and was recently published in [9]. It is to some extent related to our work as it applies the same fundamental notions of graph theory and optimization problems. However, the goals of their and our work are different. Edman et al. argue that anonymity metrics reflecting the perspective of a single user have certain defects and define a metric that, as they claim, benchmarks the system as a whole. However, they also make clear that their metric is only supposed to complement entropy based metrics and that it can not express the degree of anonymity provided to a single user.

By contrast, we look at mix networks and the anonymity they provide from an adversarial point of view. We derive a robust attack that does not rely on an assumption about the user behavior and focus on pinpointing the success probability of an adversary.

Although the work of Edman et al. is supposed to support system designers while our approach clearly reflects the adversarial side, both works have, to some extent, a common conclusion: whether it is to measure anonymity or to derive strong attack methodologies — considering the perspective of a single user is not good enough. At the same time the works are separate. Their metric is not self-contained and cannot express some necessary aspects of anonymity. Therefore, it can only complement previously derived information-theoretic metrics. We seek to put our proposal into context, empirically rate it against previous work, and show that it is superior in relevant and generic scenarios.

### 3 System and User Models

In this section we introduce our notation to describe anonymous channels and propose a new generic user model.

Consider a set of users  $U$  of cardinality  $u$ . We define the sending profile of a user  $x \in U$ , say  $x$  is Alice, as the probability distribution  $\mathbb{P}_{Alice}$  of the same size. A given element of the distribution expresses the probability that Alice sends a message to a given user  $y \in U$ , say  $y$  is Bob. So for example,  $\mathbb{P}_{Alice}(Bob)$  is the probability that Alice sends a message with Bob. The distribution as a whole describes Alice’s sending behavior with respect to the entire population (including herself). For completeness, we note that  $\sum_y \mathbb{P}_x(y) = 1$  for all  $x$ . As done in previous work [5], we model the sending rate of each individual user

$x \in U$  as a Poisson distribution with parameter  $\lambda_x$ . Further, we use the following notion of friendship: we say  $y$  is a friend of  $x$ , if  $x$  sends a message to  $y$  with non-zero probability. That is, if  $\mathbb{P}_x(y) > 0$ .

We consider two types of populations. The first one,  $U_0$ , is a simple and very restrictive user behavior model. Gradually relaxing assumptions on the number of users' friends and the user sending behavior, we construct a series of populations  $U_1$  to  $U_5$ . The latter is the most generic model considered in the literature so far to the best of our knowledge and the second population we deal with in this work. We define the models as follows:

- $U_0$ : a single user, Alice, has  $k$  randomly selected friends; her sending behavior toward her friends is uniform;  $\mathbb{P}_{Alice}$  contains  $k$  times the value  $\frac{1}{k}$  and  $u - k$  times the value zero; all other user profiles contain  $u$  times  $\frac{1}{u}$ ;
- $U_5$ : every user  $x$  has an individual number  $k_x$  of friends that is chosen at random; the sending probabilities toward the friends are randomly chosen from a uniform distribution and normalized such that  $\sum_y \mathbb{P}_x(y) = 1$  for all  $x$ ;

The anonymous channel, used by both populations, is modeled as a threshold mix. The mix's sole parameter is the threshold  $t$  which defines the number of messages in a round.

### 3.1 Comparison with Previous Models

The original Disclosure Attack and its first sequels [1,4,13] use a model that is almost equivalent to our model with population  $U_0$ . The sole difference is that, in their model, Alice sends exactly one message per round in which she participates, contrary to our model where this limitation does not exist.

Mathewson and Dingledine introduce in [16] a more complex model. First, Alice is allowed to send more than one message per round in which she participates and second, all the participants have a set of friends. Nevertheless, their behavior toward them is still uniform. In some of their experiments they go a step further and let Alice, but not the rest of the users, choose with non-uniform probability amongst her friends, thus obtaining a model a bit closer to our  $U_5$ . A recently published attack, the Two-Sided Statistical Disclosure Attack [5], is tested under  $U_0$  traffic and in a variant where all users have the same number of friends to which they send with uniform probability. Both models permit several messages of Alice per round in which she participates. The main drawback of the aforementioned models is their narrowness. With the proposed model  $U_5$  we aim at covering a wider range of scenarios, including previous work.


In particular,  $U_5$  requires no assumption about the number of users that have friends, the number of friends they have, and the sending behavior toward their friends.

## 4 Mathematical Background

In this section we recapitulate the required basic notions of graph theory and introduce our optimization problem. Then we show how we model a threshold

mix using these notions and in particular bipartite graphs. Next, we explain how maximum weighted bipartite matchings can be used to efficiently de-anonymize users communicating through a threshold mix. For further reading about graph theory in an anonymity context we refer the interested reader to [10].

A graph  $G = (N, E)$  consists of a set of nodes  $N$  and a set of edges  $E$ . Without loss of generality we assume  $N \neq \emptyset$ . A bipartite graph  $G = (S \cup R, E)$  is a graph whose nodes can be divided into two distinct sets  $S$  and  $R$  such that every edge in  $E$  connects one node in  $S$  and one node in  $R$ . In other words, there exists no edge between nodes from the same set. In this paper we focus on sets  $S$  and  $R$  of equal and finite cardinality  $t > 1$ . A set of edges  $M \subseteq E$  is called a matching in the bipartite graph  $G$  if no node in  $G$  is incident to more than one edge. A perfect matching additionally requires that every node is incident to exactly one edge. In a weighted bipartite graph  $G$  each edge  $e_i \in E$  is associated with a weight  $w_i$ . Figure 1 illustrates the definitions.



**Fig. 1.** Bipartite graphs: matching, perfect matching, and perfect matching with weights

A maximum weighted bipartite matching is defined as a perfect matching for which the sum of the weights  $w_i$  associated with the edges in the matching has a maximal value, *i.e.* the perfect matching  $M$  maximizes  $\sum_i w_i |e_i \in M$ . If the graph is not complete bipartite, *i.e.* edges which would not violate the requirements of a bipartite graph are missing due to other restrictions, one usually inserts the missing edges with an associated weight of zero. In the rest of this work we focus on maximum weighted bipartite matchings and assume completeness of the graph.

In the literature, finding such matchings is often called the assignment problem. Usually it is assumed that i) the distinct sets of nodes are of equal and finite size and ii) the total weight of the assignment (or matching) is equal to the sum of the weights associated to the edges in the assignment. In this case one deals with a *linear* assignment problem. Algorithms to solve linear assignment problems include the Hungarian algorithm [15] with complexity  $\mathcal{O}(N^2E)$  which can be optimized to  $\mathcal{O}(N^2 \log(N) + NE)$ , the Bellman-Ford algorithm [2]  $\mathcal{O}(N^2E)$  and the Dijkstra algorithm [8]  $\mathcal{O}(N^2 \log(N) + NE)$ .

#### 4.1 The Optimization Problem

Let  $S$  and  $R$  be sets of nodes of cardinality  $t$  in a complete bipartite graph  $G = (S \cup R, E)$ . We define an assignment  $M$  as a perfect matching on  $G$ . Let

$P'$  be a  $t \times t$  matrix containing weights  $w_{s,r}$ , representing probabilities, for all possible edges  $e_{s,r}$  in  $G$ . Applying Bayes theorem, the conditional a posteriori probability  $p(M|S, R)$  can be computed as

$$p(M|S, R) = \frac{p(S, R|M) \cdot p(M)}{p(S, R)}.$$

Given an assignment  $M$ , the sets of nodes  $S$  and  $R$  are implicitly fixed and thus  $p(S, R|M) = 1$ . It follows that  $p(M|S, R) = p(M)/p(S, R)$ . Since the sets  $S$  and  $R$  are given in the condition,  $p(S, R)$  is a constant term and independent of a considered assignment  $M$ . Therefore, the assignment  $M$  maximizing  $p(M)$  also maximizes  $p(M|S, R)$ .

An assignment  $M$  is a perfect matching on  $G$ , thus  $p(M)$  is the joint probability of the individual edges  $e_{s,r} \in M$ . Assuming that the edges  $e_{s,r} \in M$  are independent, the joint probability  $p(M)$  is the product of the individual edge probabilities

$$p(M) = \prod_{e_{s,r} \in M} w_{s,r}.$$

## 4.2 Mapping to a Threshold Mix

The  $t$  messages sent during one round of the mix form the set  $S$ . Each node  $s \in S$  is labeled with the sender's identity  $sen(s)$ . That is, two messages from one sender are represented by two different nodes with the same label (note that a node does not represent a specific user, but a message sent by a specific user). Equivalently, the  $t$  messages received during one round form the set  $R$  where each node  $r$  is labeled with the receiver's identity  $rec(r)$ . An edge  $e_{s,r}$  in this graph always connects a sent message  $s$  with a received message  $r$ , implying that these two messages are the same ( $s = r$ ) and therefore exhibiting the link between sender and receiver. The nodes  $S \cup R$  and the edges  $E$  form the complete bipartite graph  $G = (S \cup R, E)$ . A perfect matching  $M$  on  $G$  links all  $t$  sent and received messages.

The weights  $w_{s,r}$  associated with the edges  $e_{s,r} \in E$  are derived from user profiles  $\mathbb{P}_x$ . We discuss how to estimate these user profiles and practical issues in a separate section. Recall that for each user  $x$ ,  $\mathbb{P}_x$  describes the sending behavior toward the entire population but, for a given round, only those  $x$  and elements of  $\mathbb{P}_x$  associated with senders and receivers in the round are of interest. Therefore we derive the  $t \times t$  matrix  $P'(s, r) := \mathbb{P}_{sen(s)}(rec(r))$ ,  $s \in S, r \in R$ .

In the bipartite graph  $G = (S \cup R, E)$ , an edge  $e_{s,r}$  between a message  $s \in S$  sent by user  $sen(s)$  and a message  $r \in R$  received by user  $rec(r)$  is associated with  $w_{s,r} = P'(s, r)$ . Note that *a priori* the graph is complete bipartite as every sent message can be linked to every received message. If the user profiles exclude certain individuals from the list of possible communication partners due to  $\mathbb{P}_{sen(s)}(rec(r)) = 0$ , the relation is represented by an edge of weight zero.

In our model, all senders send with the same sending rate such that all combinations of senders  $sen(S)$  are equally likely to be observed. Each sender chooses

the recipient(s) of her message(s) independently of the choice(s) of all other senders. Further, if a user sends multiple messages, the receivers of these messages are also chosen independently. Therefore, we can model the case that a user sends two (or more) messages by considering her two (or more) distinct senders with identical profiles that each send one message to independently chosen receivers.

Given a round observation, which consists of multisets of senders  $sen(S)$  and receivers  $rec(R)$ , the probability of each assignment  $M$  is  $\prod_{e_{s,r} \in M} w_{s,r}$ . The assignment  $M$  maximizing  $p(M)$  also maximizes  $p(M|S, R)$ .

## 5 Attack Description

In this section we describe the profiling step and the de-anonymization step of the Statistical Disclosure Attack and the improved de-anonymization step of the Perfect Matching Disclosure Attack.

An attacker deploying a Disclosure Attack observes the system during  $\rho$  rounds, collecting the identity of the senders and receivers in each of them. We denote  $sen(S_i)$  the set of the senders of the  $t$  messages arriving to the mix in round  $i$  and  $rec(R_i)$  the set of the corresponding receivers. We denote the whole set of  $\rho$  round observations as the trace  $T = (S_i, R_i), 1 \leq i \leq \rho$ . We note that both  $sen(S_i)$  and  $rec(R_i)$  are multisets and may contain repeated elements, meaning that users can send (or receive) more than one message in each round.

### 5.1 Profiling with the Statistical Disclosure Attack

The SDA, as presented by Danezis in [4], focuses on revealing the *likely* set of friends of a target user, Alice. It was proposed for a scenario very close to our  $U_0$  scenario, where Alice is the only user in the system that has a set of friends ( $\mathbb{P}_{Alice}$  contains  $k$  positions with value  $1/k$  corresponding to her  $k$  friends), and the rest of the population choose their recipients uniformly amongst all the users ( $\mathbb{P}_{sen(s)}(rec(r)) = \frac{1}{u}$  for all  $s \in S, r \in R, sen(s) \neq Alice$ ). The sole difference of Danezis' model with respect to our definition of  $U_0$  is that in his model Alice sends exactly one message per round in which she participates.

In each round where Alice is sending a message, an attacker deploying the SDA considers the probability distribution  $O$  of the potential recipients of this message as a combination of the profiles of all the participating senders

$$O = \frac{1}{t}\mathbb{P}_{Alice} + \frac{t-1}{t}\mathbb{P}_x, x \in sen(S_i) \setminus \{Alice\}. \quad (1)$$

For a sufficient number  $i$  of observed rounds, the law of large numbers allows to estimate Alice's profile from the empirical mean over the observed rounds:

$$\bar{O} = \frac{1}{t} \sum_i O_i \approx \frac{\mathbb{P}_{Alice} + (t-1)\mathbb{P}_x}{t} \Rightarrow \tilde{\mathbb{P}}_{Alice} \approx t \frac{\sum_i O_i}{t} - (t-1)\mathbb{P}_x. \quad (2)$$

Using the round observations contained in  $T$  as input to this method, the attacker estimates the profiles of all the users in the system. We denote the estimated profile of user  $x$  obtained in this phase  $\tilde{\mathbb{P}}_{x,SDA}$ , for each user  $x$  in the population, and we denote the whole set of these profiles as  $\tilde{\mathbb{P}}_{SDA}$ .

## 5.2 De-anonymization with the Statistical Disclosure Attack

As suggested in [4,5], the estimated profile can be used to rank the potential receivers of a message from Alice according to the likelihood that Alice would send to them. The most likely receiver  $rec(r)$  of her message in a round  $i$  can thus be easily identified as

$$rec(r) = \operatorname{argmax}_{rec(r)} \tilde{\mathbb{P}}_{Alice,SDA}(rec(r)), r \in R_i. \quad (3)$$

When de-anonymizing the receivers of several messages in one round, the most obvious, though naïve approach is to repeat this procedure for each individual sent message. Figure 2 depicts the entire de-anonymization process, where the box marked as SDA profiling represents the profiling step described in the previous section, and the output  $D_{SDA}$  is the de-anonymization result of the attack.



Fig. 2. De-anonymization with the Statistical Disclosure Attack

## 5.3 De-anonymization with the Perfect Matching Disclosure Attack

In a nutshell, our idea is to link all messages sent and received during one round such that each message is linked and the joint probability of all links is maximized. Thus, we aim at finding a maximum weighted bipartite matching on the underlying graph, which in terms of algorithmic computer science is an assignment problem. We denote the space of all perfect matchings on the graph  $G$  by  $\mathcal{M}$  and require that an eligible set of edges belongs to this space, *i.e.* it must be a perfect matching  $M \in \mathcal{M}$ .

Given the trace  $T$  of round observations, the adversary first estimates simple user profiles  $\tilde{\mathbb{P}}_{SDA}$  as described in 5.1. Then she uses these profiles to de-anonymize mixing rounds, see Fig. 3. For a round  $i$ , she derives the  $t \times t$  matrix  $P'(s, r) := \tilde{\mathbb{P}}_{sen(s),SDA}(rec(r))$ ,  $s \in S_i, r \in R_i$ . The joint probability of all  $t$  links in an assignment  $M$  is  $p_{joint} = \prod_{e_{s,r} \in M} P'(s, r)$ .

As derived in Sect. 4.1, the assignment  $M$  that maximizes  $p_{joint}$  is the adversary's best guess. Note that maximizing  $p_{joint}$  does not fit the definition of a linear assignment problem because a maximum weight bipartite matching is achieved by maximizing the *sum* of edge weights in a perfect matching. In order to model our problem as a linear assignment problem one more step has to be




taken. To linearize the problem, we replace each element of the matrix  $P'(s, r)$  with its logarithmic value  $\log_{10}(P'(s, r))$  before associating it to the edge  $e_{s,r}$  linking message  $s$  to message  $r$ .

It is well known that the logarithm is a monotonically ascending function, if the basis is greater than or equal to one. Thus maximizing  $\log_{10}(p_{joint})$  is equivalent to maximizing  $p_{joint}$ . The advantage is that  $\log_{10}(p_{joint})$  can be calculated as a sum

$$\log_{10}(p_{joint}) = \log_{10}\left(\prod_{e_{s,r} \in M} P'(s, r)\right) = \sum_{e_{s,r} \in M} \log_{10}(P'(s, r)). \quad (4)$$

Having each edge associated with a log-probability, the assignment problem is linearized and can be solved efficiently. Using as input the matrix  $P'$ , a suitable algorithm to solve linear assignment problems outputs the most likely sender-receiver combination for all  $t$  messages in the round as the perfect matching  $M \in \mathcal{M}$ . It is a maximum weighted bipartite matching on the graph  $G = (S \cup R, E)$  and maximizes  $p_{joint}$  for this round. We summarize the approach for one round:

1. sent messages are nodes in  $S_i$  and marked with their senders' identities
2. received messages are nodes in  $R_i$  and marked with their receivers' identities
3. derive the  $t \times t$  matrix:  $P'(s, r) := \mathbb{P}_{sen(s), SDA}(rec(r))$ ,  $s \in S_i, r \in R_i$
4. replace  $P'(\cdot, \cdot)$  with  $\log_{10}(P'(\cdot, \cdot))$
5. solving the linear assignment problem yields the maximum weight bipartite matching  $M$ .



**Fig. 3.** De-anonymization with the Perfect Matching Disclosure Attack

In order to implement the attack, a subtle detail needs to be considered. Taking into account that, before applying the logarithm,  $0 \leq P'(\cdot, \cdot) \leq 1$  and that  $\log(0)$  is not defined, we need to define  $\log(0) = -\infty$  in order for the algorithm to maximize the joint probability. Note however, that i) this case is rarely encountered in practical scenarios unless one has access to very precise user profiles and that ii) replacing 0 with  $-\infty$  solely prevents numerical errors and has no influence on the output  $M$  of the matching algorithm. Further, some implementations of algorithms for linear assignment problems aim at *minimizing* the sum of the edge weights (*e.g.* costs) in a perfect matching. However, a linear maximization problem can be turned into a linear minimization problem by substituting  $P'(\cdot, \cdot) = -P'(\cdot, \cdot)$ .

## 6 Empirical Evaluation of De-anonymization Techniques

In order to evaluate the performance of the Perfect Matching Disclosure Attack, we deploy it in different scenarios and compare it to the original Statistical

Disclosure Attack. Our goal is to study the impact of system parameters on the effectiveness and viability of both attacks.

### 6.1 Experimental Settings

Our experiments are carried out on populations  $U$  of size  $u = 1000$  users that send messages through a threshold mix with threshold  $t = 100$ , ensuring that a considerable fraction of the users participate in each mixing round. Every user  $x \in U$  chooses her recipients according to her profile  $\mathbb{P}_x$ , which depends on the considered user behavior model (see Sect. 3), and initiates communications with the same frequency  $\lambda$ . We note that, given that the attacks need full rounds of mixing, the choice of this parameter’s value is arbitrary. As long as all users send messages to the network with equal rate, their frequency of appearance as senders does not depend on the precise sending rate. Although real users are expected to send messages with different frequencies, we chose to fix this parameter in order to create a scenario that allows us to clearly illustrate our techniques.

We study how the number of rounds observed by the attacker affects the performance of the PMDA and the SDA. Both from the adversarial and the designer’s points of view, this consists of exploring the effectiveness, efficiency, and scalability of the attacks. For the purpose of our studies we have generated 100 000 mixing rounds. An experiment consists of 1) estimating all user profiles  $\tilde{\mathbb{P}}_{SDA}$  from  $\rho$  round observations, 2) de-anonymizing 5000 rounds with the SDA, and 3) de-anonymizing the same 5000 rounds with the PMDA (except when  $\rho = 1000$ , when we only de-anonymized 1000 rounds). Table 1 summarizes the parameters and their values in the experiments.

**Table 1.** Parameters of the experiments ( $N=1000$ ,  $t=100$ ),  $\mu$  is average number of messages used to profile one user,  $\gamma$  is average number of de-anonymization trials per user

Param \ $\rho$	1k	5k	10k	25k	50k	100k	Population	N° of friends $k$	Profile
$\mu$	100	500	1000	2500	5000	10000	$U_0$	{5, 25, 50}	Uniform
$\gamma$	100	500	500	500	500	500	$U_5$	random [5, 50]	Non-uniform

### 6.2 Results

In this section we present the results of our experiments. To measure the effectiveness of the attacks we define two metrics, the *individual success rate* and the *round success rate*. The former expresses the accuracy of the attack when de-anonymizing the receiver of a message from a particular sender, *i.e.* successfully linking a specific sender to a receiver. It is computed by counting how many messages sent by each user in the population have been correctly de-anonymized during the attack, then deriving the success rate per sender by dividing by the number of messages sent by this user. The latter shows the percentage of links

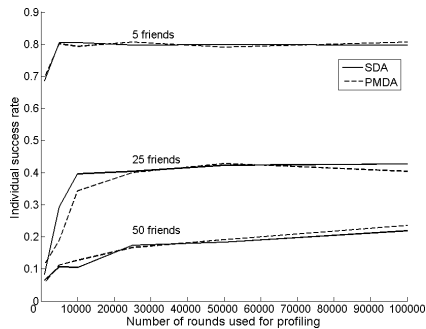
correctly de-anonymized per round. We calculate it as the average number of relations successfully identified per round. Both metrics are computed over all 5000 (1000) rounds.

It is important to note that we consider a message as de-anonymized correctly if and only if the attack has identified the receiver of that message correctly. Note that this does not necessarily require to match a sent message to the correct received message. We apply a hard yes/no metric on whether the identity of the matched recipient is correct. This is a more rigorous criterion than the one used in [4,5] where the rank in the sorted probability distribution of potential receivers is taken into account.

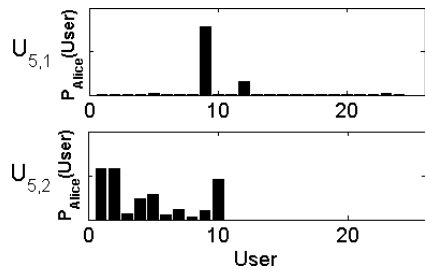
**Population  $U_0$ .** We test both attacks in three  $U_0$  populations where Alice has  $k$  friends. We look at the influence of the number  $\rho$  of rounds used in the profiling step on the success rates of the attacks.

Figure 4 illustrates the individual success rates. As all users except for Alice send uniformly to the entire population, no information can be inferred about them. Therefore the results refer only to Alice’s messages and we only consider her individual success rate. We see that the PMDA does not get any advantage in these scenarios, and both attacks score similarly. On the one hand this is due to the lack of information that the rest of senders in the round provide. Since their profile is uniform, they give no hints about who Alice is *not* sending to. On the other hand, Alice chooses uniformly amongst her friends. Therefore, if two or more of her friends appear in the set  $rec(R)$ , the best the algorithm can do is choose randomly amongst them. This last problem also affects the SDA’s effectiveness. One can observe in the graph that, the smaller the number of friends (thus the smaller the probability that this difficulty appears) the higher the success rate of both attacks. The graph shows that in some cases the PMDA performs slightly worse/better than the SDA, but these small differences have no statistical significance.

As expected, increasing the number  $\rho$  of rounds to profile users increases the likelihood of successful attacks. It is remarkable, however, that this rate does



**Fig. 4.** Individual success rate in a  $U_0$  population



**Fig. 5.** Alice’s profile in a  $U_{5,1}$  and  $U_{5,2}$  scenario

not increase constantly. When the number of Alice’s friends is small ( $k = 5$ ), not much improvement is achieved by increasing the number of profiling rounds above 10 000. Nevertheless, having more rounds helps the attacker when the number of friends increases, as more rounds, in which Alice participates, are needed to observe her sending messages to all of her friends.


**Population  $U_5$ .** For our second set of experiments, we prepared two bench tests,  $U_{5,1}$  and  $U_{5,2}$ , where users had a complex behavior corresponding to  $U_5$  populations. In both cases each user had a random number of friends chosen uniformly from  $[5, 50]$ . However, the scenarios differ in the way the sending probabilities are distributed amongst these friends. Users corresponding to the  $U_{5,1}$  example have a set of contacts where there are one or two very good friends (which they choose as recipients in more than 60% of the cases) and the rest have small probability of being chosen. The users forming the population for the second test,  $U_{5,2}$ , do not have strong preferences about their contacts, still, their distribution is non-uniform. Figure 5 depicts Alice’s profile in the  $U_{5,1}$  and  $U_{5,2}$  scenarios.

Contrary to the  $U_0$  case, where the SDA and the PMDA performed similarly, the PMDA achieves higher de-anonymization success rates when applied to a  $U_5$  scenario. Figure 6 shows the percentage of users participating in the communication for which the attacks obtain a certain individual success rate in both  $U_5$  scenarios. We represent different values for the number  $\rho$  of rounds used for profiling with different line styles.


We see that the PMDA outperforms the SDA in both experiments, but there is a significant difference between them. With respect to the  $U_{5,1}$  case (on the left side of the figure) and  $\rho = 10\,000$  one can observe that the SDA achieves an average individual success rate of 71.5% while the PMDA scores an average individual success rate of 96.04% and de-anonymizes more than 90% of the messages correctly for 99.6% of the users. With respect to the  $U_{5,2}$  case (right side of the figure) and  $\rho = 10\,000$  the SDA achieves an average individual success rate of 26% while the PMDA scores 55.35%.

Figure 7 presents the round success rates of the SDA and the PMDA. Like in the individual success rate, our attack outperforms the SDA. In the  $U_{5,1}$  case (left), the SDA has a high rate (71.5% in average) of round de-anonymization, whichever is the number of rounds observed. However, the PMDA improves this result de-anonymizing in average 96.05% of the messages in each round when 10 000 rounds have been used for profiling and correctly de-anonymizing the full set of links in 17.22% of the cases. The success of both attacks diminishes when the user’s sending patterns tend to be more uniform toward their friends (case  $U_{5,2}$ , right). For the same number of  $\rho = 10\,000$  observed rounds the SDA achieves an average round success rate of 25.6% and the PMDA 55.3%.

It is important to note the influence of the number of rounds observed by the attacker on the success rates of the attacks. Increasing the number of observations makes both attacks more accurate. However, there are notable differences in the effect of this increase depending on the type of population attacked as well as on the attack itself. Analyzing a higher number of rounds provides more



**Fig. 6.** Individual success rate attacking  $U_5$  populations ( $U_{5,1}$  left,  $U_{5,2}$  right)



**Fig. 7.** Round success rate attacking  $U_5$  populations ( $U_{5,1}$  left,  $U_{5,2}$  right)

information, a fact exploited by the PMDA. On the contrary, the SDA's simple decision algorithm takes little advantage of this extra information and we see that almost no improvement is achieved by observing more than 5000 rounds. Moreover, when the attacks are carried out in a  $U_{5,1}$  scenario, the users' profiles have a low entropy, thus the strong friends are early identified and no additional information is extracted from new round observations.

### 6.3 Scalability of the Attacks

We evaluate the efficiency of both attacks in terms of time. We implemented both attacks in the high-level interpreted language of a commercial numerical computing environment without any optimizations. In our implementation of the PMDA we use the Hungarian algorithm [15] to solve the linear assignment problem of finding the most likely perfect matching between inputs and outputs of the mix. We show in Table 2 the time it takes to carry out all the operations depicted in Fig. 2 for the original SDA and in Fig. 3 for the PMDA in  $U_{5,2}$  scenarios with mix thresholds 100, 500 and 1000. In all cases the profiles  $\tilde{\mathbb{P}}_{SDA}$

have been derived from  $\rho = 50\,000$  rounds and have been used to de-anonymize 5000 rounds (*i.e.*, find the recipients for all  $s$  in  $S_i$ ,  $1 \leq i \leq 5000$ ). The code for scenarios with threshold 100 and 500 rounds was executed on a machine with a processor running at 2.8 GHz and 512 KB cache and for the threshold 1000 scenario we used a machine with a processor running at 2.2 GHz and 1 MB cache. We include the success rates for  $t = 100$  to illustrate the trade off between accuracy and speed.

**Table 2.** Timings of the attacks: estimation of profiles from 50 000 rounds and de-anonymization of 5000 rounds


Attack	$t = 100$		$t = 500$	$t = 1000$
	Time	Success rate, mean (min)	Time	Time
SDA profiling	3.08m	-	38.33m	66.16m
SDA de-anon	10m	25.6% (0.00%)	3.48h	12.91h
PMDA de-anon	10.2m	62.9% (38.8%)	12.9h	4.69days
NSDA de-anon	13.33m	60.2% (33.5%)	4.28h	15.3h

The PMDA de-anonymization is slower than the SDA de-anonymization and the difference grows as the size of the threshold and thus the underlying bipartite graph increases. Nevertheless, it yields higher success rates. In Sect. 8.2 we propose the Normalized Statistical Disclosure Attack (NSDA), that combines accuracy and speed. Table 2 includes the success rate and timings for the operations shown in Fig. 12 inside the dotted line. Note that all of the attacks’ efficiencies would substantially benefit from optimized implementations. Further, the PMDA in particular is suited for parallelization.

## 7 Enhanced Profiling with the Perfect Matching Disclosure Attack

So far we have focused on the PMDA’s de-anonymization capability. In this section, we show how the derived maximum weighted bipartite matchings  $M_i$  can be used to better estimate user profiles.

A better estimation of a profile, say  $\mathbb{P}_{Alice}$ , is built by, instead of considering all possible receivers of her message(s) in a round  $i$  as equally likely, considering the receiver(s) indicated by the matching  $M_i$  as the most likely. Instead of assigning a probability of  $1/t$  to each receiver in  $rec(R_i)$ , the attacker assigns  $z$  to the receiver assigned to Alice’s message(s) by  $M_i$  and  $(1-z)/(t-1)$  to the rest of the elements in  $rec(R_i)$ . This step is marked with “PMDA profiling” in Fig. 8. The choice of the weight  $z$  is not that crucial. It expresses the confidence one has in the perfect matchings  $M_i$ . We experimented with different values for this parameter but observed that the effect on the profile estimation is minor. However, there is one hard bound. Choosing the weight  $z$  such that  $z = (1-z)/(t-1)$  turns the second




**Fig. 8.** Obtaining enhanced profiles with the Perfect Matching Disclosure Attack

profiling step useless as this setting reflects the original SDA, and choosing  $z < (1-z)/(t-1)$  will effectively hide the actual users' relationships. We chose  $z = 0.5$  without a specific motivation. Note that the same  $\rho$  round observations used to construct the simple profile  $\tilde{\mathbb{P}}_{Alice,SDA}$  are reused to estimate the enhanced profile  $\tilde{\mathbb{P}}_{Alice,PMDA}$ .


The same procedure can be applied to the decision  $D_i$  of the de-anonymization phase of the SDA, yielding a more accurate profile than the one estimated by the original SDA and denoted by  $\tilde{\mathbb{P}}_{Alice,eSDA}$ .

For a  $U_0$  scenario where Alice has five friends, Fig. 9 shows the profile  $\mathbb{P}_{Alice}$  we initially generated for Alice, her profile after the PMDA's profiling step, the approximation of her profile derived with the enhanced SDA, and her profile estimated using the original SDA. Figure 10 shows the corresponding set of profiles for a  $U_{5,2}$  scenario. We observe in both cases that the profile estimation  $\tilde{\mathbb{P}}_{Alice,eSDA}$  is more precise than  $\tilde{\mathbb{P}}_{Alice,SDA}$  but not as good as  $\tilde{\mathbb{P}}_{Alice,PMDA}$ .



**Fig. 9.** Alice's profile and estimations (logscale) for  $U_0$ ,  $\rho = 100\,000$ . From left to right:  $\mathbb{P}_{Alice}$ ,  $\tilde{\mathbb{P}}_{Alice,PMDA}$ ,  $\tilde{\mathbb{P}}_{Alice,eSDA}$ , and  $\tilde{\mathbb{P}}_{Alice,SDA}$ .

In the  $U_0$  scenario, all three estimations allow the adversary to easily identify the set of Alice's friends, even if the exact number  $k$  of friends is unknown. However, the enhanced methods increase the contrast between friends and non-friends. In the  $U_{5,2}$  scenario,  $\tilde{\mathbb{P}}_{Alice,SDA}$  does not allow to identify friends, and even worse, there exist non-friends of Alice that have higher probability than some of her friends.  $\tilde{\mathbb{P}}_{Alice,eSDA}$  improves the estimation and allows to identify Alice's best friends (those with high probability in  $\mathbb{P}_{Alice}$ ), but it fails to show more unlikely receivers as for example user 19. In  $\tilde{\mathbb{P}}_{Alice,PMDA}$  the estimation is further improved and all of her friends have higher probabilities than her non-friends.




**Fig. 10.** Alice’s profile and estimations (logscale) for  $U_{5,2}$ ,  $\rho = 100\,000$ . From left to right:  $\mathbb{P}_{Alice}$ ,  $\tilde{\mathbb{P}}_{Alice,PMDA}$ ,  $\tilde{\mathbb{P}}_{Alice,eSDA}$ , and  $\tilde{\mathbb{P}}_{Alice,SDA}$ .

## 8 Extending the Perfect Matching Disclosure Attack

In this section we present variants and extensions of the Perfect Matching Disclosure Attack. The iterated PMDA and the normalized SDA are alternatives with different trade-offs between precision and computational load. We also outline how the PMDA, and any of its variants can be applied to the more realistic scenario of a pool mix.

### 8.1 Iterated PMDA

The profiles  $\tilde{\mathbb{P}}_{PMDA}$  can be used as input to a subsequent PMDA de-anonymization step, yielding the perfect matchings  $M_i^*$  as output, which uncover the actual relations between senders and receivers for each round with an even higher rate of success than the PMDA, particularly in  $U_5$  scenarios. Further, the  $M_i^*$  can be used for a subsequent PMDA profiling step, yielding user profiles that are slightly better than the  $\tilde{\mathbb{P}}_{PMDA}$ . Figure 11 illustrates the chaining for two iterations of the PMDA.



**Fig. 11.** Iterated Perfect Matching Disclosure Attack

In fact, the PMDA can be chained arbitrarily often, each time yielding a (slight) improvement over the outputs of the previous iteration, and asymptotically approaching the optimal result. The concept of this iterated approach is known as expectation maximization.

Note, however, that each additional instance of the PMDA implies an increase of computational cost. Again, it is possible to trade certainty for speed substituting the PMDA de-anonymization step by the SDA de-anonymization step. Table 3 presents de-anonymization success rates of a two-instances PMDA and a two-instances eSDA when applied to  $U_0$  and  $U_5$  scenarios.




**Table 3.** Individual success rates of two-instances PMDA and two-instances eSDA de-anonymization; all profiles are derived reusing the same set of  $\rho = 10\,000$  rounds, success rates are evaluated from de-anonymization of 5000 rounds

$U_0$	eSDA	PMDA	$U_{5,2}$	eSDA	PMDA
k = 5	80.43	78.67		26.56	60.24
k = 50	10.47	12.15			

## 8.2 Normalized Statistical Disclosure Attack

The Normalized Statistical Disclosure Attack, illustrated in Fig. 12, has a similar structure as the SDA but it additionally constructs the matrix  $P'$  as in the PMDA and it includes a matrix normalization step.



**Fig. 12.** Normalized Statistical Disclosure Attack

We transform  $P'$  into a doubly stochastic transition matrix that, by definition, has the property that each row and each column sums up to one. We use the method proposed by Sinkhorn in [19] in 1964. He showed that an arbitrary positive  $N \times N$  matrix, *i.e.* each element is greater than zero, can be transformed into a doubly stochastic matrix by iterative proportional fitting. This means iteratively normalizing the rows and the columns of the matrix. Sinkhorn also proved that the iteration converges and has a unique solution.

An element of the normalized transition matrix  $P'$  represents the probability of a link between input messages (row) and output messages (column). This ensures that each sent message is received (all rows sum up to 1) and each received message was sent (all columns sum up to 1). The receiver of a given message  $s$  is chosen as the one who maximizes the individual link probability  $P'(s, \cdot)$ .

The normalization step has two important effects on  $P'$  that stem from the fact that the iterative proportional fitting spreads the information contained in each element of  $P'$  over the entire matrix. The first effect is best explained in a noise-free toy example. Consider the matrix  $P'$  before and after normalization

$$P' = \begin{pmatrix} 0.5 & 0.5 & 0 \\ 1 & 0 & 0 \\ 0 & 0.5 & 0.5 \end{pmatrix} \xrightarrow{\text{normalize}} \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

The per sender maximum likelihood decision approach of the SDA achieves 66.66% success rate when assigning receivers to senders based on the original version of  $P'$ .

The normalization process over the matrix  $P'$  implicitly takes interdependencies between the matrix elements in different rows and columns into account and eliminates impossible combinations. In the toy-example, the certainty  $P'(2, 1) = 1$  implies  $P'(1, 1) = 0$ . Hence,  $P'(1, 2)$  becomes 1 to fulfill the doubly-stochastic requirement in the first row. This implies that  $P(3, 2)$  becomes also 0 and hence  $P'(3, 3) = 1$ . Therefore, a per sender maximum likelihood decision approach based on the normalized matrix takes more information into account and leads to the only correct assignment with success rate one.

To explain the second effect, we use a noisy version of the same initial matrix  $P'$  that contains Gaussian noise with standard deviation 0.1

$$P' = \begin{pmatrix} 0.4006 & 0.4208 & 0.1786 \\ 0.7810 & 0.1432 & 0.0757 \\ 0.0997 & 0.4580 & 0.4424 \end{pmatrix} \xrightarrow{\text{normalize}} \begin{pmatrix} 0.2776 & 0.4369 & 0.2856 \\ 0.6673 & 0.1834 & 0.1494 \\ 0.0552 & 0.3798 & 0.5651 \end{pmatrix}.$$

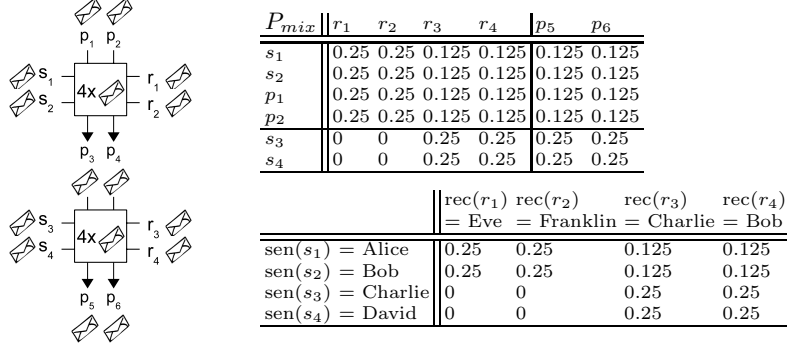
The per sender maximum likelihood decision of the SDA based on the initial  $P'$  leads to the correct assignment for the senders 1 and 2 but to a wrong assignment for sender 3. Based on  $P'$  after the normalization step, also the third assignment is identified correctly. The estimated profiles obtained by an adversary in a realistic scenario contain noise. The normalization step partially eliminates this noise yielding more reliable data.

The combination of these two effects allow the NSDA to de-anonymize messages with a higher success rate than the original SDA. As we show in Table 2, this attack runs faster than the PMDA for  $t = 500$  and  $t = 1000$ , still it achieves a lower success rate. It is a decision of the adversary which method suits her purposes best.

### 8.3 Pool Mix

Finally, we outline how our attack can be applied to a pool mix scenario [17]. Figure 13 depicts a simple example with threshold  $t = 4$  and internal memory of size  $n = 4$ . It also shows the link probabilities between incoming and outgoing messages according to the formula given by Serjantov and Danezis in [18] in the upper table, and the relevant part of users' profiles in the lower table. Such profiles can be derived, for example, by applying the SDA [7]. We observe two rounds of the mix. Initially the mix generates two dummy messages  $p_1$  and  $p_2$  and places them in the pool. After the first round two messages stay in the pool participating in the second round. After the second round, two messages,  $p_5$  and  $p_6$ , remain in the pool.

Before one can apply the PMDA to a pool mix, the scenario needs to be mapped to a bipartite graph. A simple approach for doing so maps each round individually. The set of sent messages in one round is formed by the messages actually sent in this round and the messages that remained in the pool after the previous round. For the first round of our example that is  $S = \{s_1, s_2, p_1, p_2\}$ . Equivalently, the set of received messages is composed of messages that left the mix and the messages remaining in the pool, *i.e.*  $R = \{r_1, r_2, p_3, p_4\}$ . The initial



**Fig. 13.** Left: two rounds of a pool mix scenario; Right: mix probabilities (upper table) and user profiles (lower table)

matrix  $P'$  can then be generated from the mix probabilities given in Fig. 13 on the upper right side, *i.e.*  $P'(s, r) = P_{mix}(s, r)$  for all  $s \in S_i, r \in R_i$ , for each round  $i$ . However, as observed from the experimental results in Sect. 6.2, the uniformity of the entries in this  $P'$  are bad conditions for an attack to operate in.

A better approach is to deal with several observed rounds at once and to compute the probabilities for  $P'$  globally from starting point to end point. In our example both rounds can be combined using  $S = \{s_1, \dots, s_4, p_1, p_2\}$  and  $R = \{r_1, \dots, r_4, p_5, p_6\}$ . Still we do not expect the attacks to perform well due to the same reasons as given above.

We propose to additionally combine both sources of information, mix probabilities and user profiles, into  $P'$ . The senders' choices of their recipients and the choice of the mix on which messages to output are independent. Therefore, one computes the joint probability of two choices as the product of the individual probabilities. We derive  $P'$  as

$$P'(s, r) = P_{mix}(s, r) \cdot \tilde{\mathbb{P}}_{sen(s), SDA}(rec(r)), \quad s \in S, r \in R.$$

For completeness we note that the senders of messages which are in the pool at the beginning of the observation and receivers of messages which are in the pool at the end of the observation need to be added to the population. The “virtual” senders are best modeled with a uniform profile while the “virtual” receivers need to be inserted into all senders' profiles. Once the initial matrix  $P'$  has been generated, the PMDA can be applied to this pool mix scenario.

## 9 Conclusions

The main drawback of previously published practical Disclosure Attacks is their susceptibility to changes in the user behavior model. Each of them seems to be optimized for a specific and restricted scenario. Our first contribution is a

more general user behavior model, where the number of users' friends and the distribution of sending probabilities toward them is not restricted.

Our second contribution is the Perfect Matching Disclosure Attack, that achieves a high rate of success when tracing messages sent through a threshold mix in arbitrary scenarios. Its accuracy arises from the fact that it considers information about all senders participating in a round simultaneously, rather than focusing on individual users iteratively. We empirically compare it with previous work in terms of effectiveness and show that our proposal yields better results when de-anonymizing the sender of a given message in a generic scenario.

The second advantage of the PMDA over previous work is its enhanced ability to estimate user profiles. Concerning a very restrictive user behavior model we empirically confirm that the PMDA yields a better separation of friends and non-friends than previous work. With respect to a generic scenario we show that the PMDA reliably identifies users' friends when previously proposed methods fail.

Although the Perfect Matching Disclosure Attack is computationally more expensive than previously proposed and practical methods, our study of its efficiency shows that it is indeed practical. A particular promising property of our proposal is, that it can be parallelized to a high degree. Further, we show how it can be adapted to different scenarios including pool mixes and how it can be refined to achieve even better results. A significantly sped-up variant, the Normalized Statistical Disclosure Attack, yields slightly worse accuracy than the PMDA but is almost as fast as the original SDA.

Although the new user model presented in this work is more generic than previous proposals, it is not as versatile as one would desire and most probably far from real user behavior. More research needs to be performed on the influence of parameters like the users' sending rate or its variance over time on the effectiveness and efficiency of attacks in order to evaluate their impact on real anonymous communications networks.

Perhaps the most closely related work to ours is the approach toward measuring anonymity proposed in [9]. However, their metric is not self-contained and can only complement entropy based metrics. Our work on the other hand aims at pinpointing an adversary's probability of success though it can also be used as a complement for the evaluation of anonymous systems. Nevertheless, both works allow a common conclusion: whether it is to measure anonymity or to derive strong attack methodologies — considering the perspective of a single user is not good enough.

## Acknowledgements

The authors would like to thank Matthew Wright for shepherding this paper and the anonymous Reviewer 1 for his/her insightful comments.

C. Troncoso is funded by a research grant of the Fundacion Barrie de la Maza (Spain). This work was supported in part by the IAP Programme P6/26 BCRYPT of the Belgian State (Belgian Science Policy), by the IWT SBO

ADAPID project, by FWO projects G.0475.05, and G.0300.07, by the Concerted Research Action (GOA) Ambiorics 2005/11 of the Flemish Government, by the European Commission through the IST Programme under Contract IST-2002-507932 ECRYPT NoE, and by the K.U. Leuven-BOF.

The information in this document reflects only the authors' views, is provided as is and no guarantee or warranty is given that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability.

## References

1. Agrawal, D., Kesdogan, D.: Measuring anonymity: The disclosure attack. *IEEE Security & Privacy* 1(6), 27–34 (2003)
2. Bellman, R.: On a routing problem. *Quarterly of Applied Mathematics* 16, 87–90 (1958)
3. Chaum, D.L.: Untraceable electronic mail, return addresses, and digital pseudonyms. *Commun. ACM* 24(2), 84–90 (1981)
4. Danezis, G.: Statistical disclosure attacks: Traffic confirmation in open environments. In: Gritzalis, Vimercati, Samarati, Katsikas (eds.) *Proceedings of Security and Privacy in the Age of Uncertainty (SEC 2003)*, Athens, May 2003, IFIP TC11 pp. 421–426. Kluwer, Dordrecht (2003)
5. Danezis, G., Diaz, C., Troncoso, C.: Two-sided statistical disclosure attack. In: Borisov, N., Golle, P. (eds.) *PET 2007*. LNCS, vol. 4776, p. 15. Springer, Heidelberg (2007)
6. 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 (May 2003)
7. Danezis, G., Serjantov, A.: Statistical disclosure or intersection attacks on anonymity systems. In: Fridrich, J. (ed.) *IH 2004*. LNCS, vol. 3200. Springer, Heidelberg (2004)
8. Dijkstra, E.W.: A note on two problems in connexion with graphs. *Numerische Mathematik* 1, 269–271 (1959)
9. Edman, M., Sivrikaya, F., Yener, B.: A combinatorial approach to measuring anonymity. In: *ISI*, pp. 356–363. IEEE, Los Alamitos (2007)
10. Hughes, D., Shmatikov, V.: Information hiding, anonymity and privacy: A modular approach. *Journal of Computer Security* 12(1), 3–36 (2004)
11. Jakobsson, M., Juels, A., Rivest, R.L.: Making mix nets robust for electronic voting by randomized partial checking. In: *Proceedings of the 11th USENIX Security Symposium* (August 2002)
12. Kesdogan, D., Agrawal, D., Penz, S.: Limits of anonymity in open environments. In: Petitcolas, F.A.P. (ed.) *IH 2002*. LNCS, vol. 2578, pp. 53–69. Springer, Heidelberg (2003)
13. Kesdogan, D., Pimenidis, L.: The hitting set attack on anonymity protocols. In: Fridrich, J.J. (ed.) *IH 2004*. LNCS, vol. 3200, pp. 326–339. Springer, Heidelberg (2004)
14. Kilian, J., Sako, K.: Receipt-free MIX-type voting scheme - a practical solution to the implementation of a voting booth. In: *EUROCRYPT 1995*. Springer, Heidelberg (1995)

15. Kuhn, H.W.: The Hungarian method for the assignment problem. *Naval Research Logistic Quarterly* 2, 83–97 (1955)
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. Möller, U., Cottrell, L., Palfrader, P., Sassaman, L.: Mixmaster Protocol — Version 2. IETF Internet Draft (July 2003)
18. Serjantov, A., Danezis, G.: Towards an information theoretic metric for anonymity. In: Dingleline, R., Syverson, P. (eds.) *PET 2002*. LNCS, vol. 2482. Springer, Heidelberg (2003)
19. Sinkhorn, R.: A relationship between arbitrary positive matrices and doubly stochastic matrices. *The Annals of Mathematical Statistics* 35(2), 876–879 (1964)