

Anonymity v.s. Information Leakage in Anonymity Systems *

Ye Zhu Riccardo Bettati
Department of Computer Science
Texas A&M University
College Station TX 77843-3112, USA
zhuye@tamu.edu bettati@cs.tamu.edu

Abstract

Measures for anonymity in systems must be on one hand simple and concise, and on the other hand reflect the realities of real systems. Such systems are heterogeneous, as are the ways they are used, the deployed anonymity measures, and finally the possible attack methods. Implementation quality and topologies of the anonymity measures must be considered as well. We therefore propose a new measure for the anonymity degree, which takes into account possible heterogeneity. We model the effectiveness of single mixes or of mix networks in terms of information leakage and measure it in terms of covert channel capacity. The relationship between the anonymity degree and information leakage is described, and an example is shown.

Keywords: Anonymity, Mix Networks, Covert Channels

1 Introduction

This paper studies the relationship between the anonymity degree and information leakage from an anonymity network.

Since Chaum [2] proposed the mix network, researchers have developed various anonymity systems for different applications. Examples include Crowds [22] for anonymous web transaction, Freenet [4] for distributed anonymous information storage and retrieval, Onion Router [12] for anonymous routing, and Tarzan [11] for p2p networking.

How to quantify the anonymity provided by a whole anonymity system? Researchers proposed various definitions to quantify anonymity, such as *anonymity set size* [15], *effective anonymity set size* [23], and *entropy-based anonymity degree* [8]. While the metrics led to an increasingly better understanding of anonymity, they tend to focus on the anonymity of a *single* message under a *single* anonymity attack. In practice however, metrics are needed that take into account realities of today's use of networks: Communication settings in real systems range from single messages, to message groups, to streams and FTP transfers. In addition, sophisticated attacks can resort to a variety

of techniques to break anonymity: flow correlation attacks [27], intersection attacks [7], trickle attacks [24], and so on.

A measure for the anonymity degree should satisfy a number of requirements: First, the anonymity degree should capture the *quality* of an anonymity system. It has been shown for example that information theoretical means, such as entropy, are more accurate for comparing anonymity systems than, say, anonymity sets. Second, the anonymity degree should take into account the topology of the network, or that of any overlay defined by the anonymity system. The topology influences how much information can be gathered by an attacker, and thus has an impact on the system anonymity degree. For example, a system of fully-connected nodes will have a different anonymity degree from a chain of nodes. Third, the anonymity degree, as measure of the effectiveness of the anonymity system should be independent of the number of users. While a large number of users clearly contributes to anonymity, this not necessary reflects on the quality of the anonymity system. Finally, the anonymity measure must be independent of the threat model, as attackers may use a variety of attack techniques, or combinations thereof, to break the anonymity.

Since the goal of anonymity attacks is to infer the communication relations in a system despite countermeasures, it is natural to model such attacks as covert channels, and interest has focused on the interdependence of anonymity and covert channels [20]. The designer of an anonymity system generally faces the question of how much information may leak from the anonymity network given the unavoidable imperfectness of the anonymity network and how this may affect the anonymity degree. The imperfectness of an anonymity system will result in the information leaking from the system. This information leakage can be evaluated in form of a covert channel.

The major contributions of our study are summarized as follows: First, we propose an anonymity degree to quantify the anonymity provided by *an anonymity network*. This definition generalizes the information theoretic definitions previously proposed in [23, 8]. Then, we propose a new class of covert channels, which we call *anonymity-based covert channels*. We formally prove how to establish covert channels of maximum capacity over a single mix based on anonymity attacks on the mix. Finally, we use anonymity-

*This work is supported in part by the Texas Information Technology and Telecommunication Task Force.

based covert channels to assess the performance of mix networks. We show how the capacity of anonymity-based covert channels can be used to provide simple descriptions of non-perfect mix networks, and can be used to formulate bounds on the provided anonymity.

The rest of the paper is organized as follows: Section 2 reviews the related work. Section 3 describes the proposed anonymity degree and the relationship with other entropy-based anonymity degree definitions. In section 4, we define the anonymity based covert channel. Section 5, Section 6, and Section 7 present the relationship between the covert channel capacity and anonymity degree for a single-mix case and mix-network case. We conclude this paper and discuss the future work in Section 9.

2 Related Work

Chaum [2] pioneered the idea of anonymity in 1981. Since then, researchers have applied the idea to different applications, such as message-based email and flow-based low-latency communications, and they have invented new defense techniques as more attacks have been proposed. For anonymous email applications, Chaum proposed to use relay servers, called mixes, that re-route messages. Messages are encrypted to prevent their tracking by simple payload inspection.

Helsingius [14] implemented the first Internet anonymous remailer, which is a single application proxy and replaces the original email's source address with the remailer's address. Gülcü and Tsudik [13] developed a relatively complete anonymous email system, called Babel. Cottrell [17] developed Mixmaster, which counters a global passive attack by using message padding. It counters trickle and flood attacks [13, 24] by using a pool batching strategy. Danezis, Dingleline and Mathewson [6] developed Mixminion. Mixminion's design considers a large set of attacks that researchers have found [1, 24]. The authors suggest a list of research topics for future study. Tor [9], the second-generation onion router, is developed for circuit-based low-latency anonymous communication recently. It can provide perfect forward secrecy.

To evaluate the effectiveness of such anonymity systems under anonymity attacks, a number of different anonymity degree definitions have been proposed: The anonymity degree proposed in [22] is defined as the probability of not being identified by the attacker. It focuses on each user and does not capture the anonymity of the whole system. Berthold et al. [1] propose an anonymity degree based on the number of the users of an anonymity system. There is an ongoing debate about what the role of the number of users is in providing anonymity. Intuitively, the larger the crowd, the easier it is for an individual to hide in it. In practice, however, attacks proceed by isolating users or groups of users that are more likely to be participants in a communication. This was first considered in the *anonymity set*, introduced in [3]. The anonymity set describes the set of *suspected* senders or receivers of a message. The size of the anonymity set is used in [15] as the anonymity degree.

A big step forward was done by Serjantov and Danezis

[23] and by Diaz *et al.* [8] by proposing anonymity measures that consider probability distributions in the anonymity set. Both measures are based on entropy and can differentiate two anonymity sets that have identical sizes, but different distributions. The measure in [8] normalizes the anonymity degree to discount for the anonymity set size.

A number of efforts have studied the relation between covert channels and anonymity systems. Moskowitz *et al.* [19] focus on the covert channel over a mix-firewall between two enclaves. The covert channel in this case is established by the channel receiver determining whether an anonymized sender is transmitting packets. Newman *et al.* [21] focus on the covert channel over a timed mix. The authors in [20] make a series of excellent observations about the relation between covert channels and anonymity systems. They illustrate this relation by describing the linkage between the lack of complete anonymity (quasi-anonymity) and the covert communication over different types of mixes and propose to use of this covert channel capacity as a metric for anonymity.

The work presented in this paper takes a system-level view of covert channels and anonymity, and differs from previous work, such as [19, 20, 21], in two ways: First, we assume that the existence of various sources of information leakage in the elements (mixes, batchers, padders, ...) of an anonymity system are a reality that system designers and operators have to deal with. Some of the resulting covert channels can be identified and either measured or analyzed using techniques described in [19, 21]¹ In addition, any cautious anonymity system designer or operator must assume that even mixes presumed to be perfect are not so, even if the particular weakness is not known *a priori*. In this paper, we use covert channel capacity as a generic measure to model weaknesses (known or unknown) in the anonymity system infrastructure. This gives a tool for designers or operators to uniformly describe both known weaknesses (i.e. results of attacks), or merely suspected ones, and to analyze their effect on the anonymity provided by the system. Second, the anonymity degree of the mix network is a result of system-level effects: changes in the user population or application mix affect the anonymity provided. So do topology of the anonymity system and routing preferences within the system. As a result, there is no one-to-one mapping from the anonymity degree to covert channel capacities of elements in a mix network and *vice versa*. In this paper, we investigate the relationship between anonymity degree and covert channel capacity in terms of what effect one has on the other.

3 Anonymity Degree

A number of attacks have been described recently that give raise to reasonably high capacity channels on mixes. Several attacks to simple mixes lend themselves to an accurate analysis of the exploited covert channels, such as in [19, 20, 21]. For other attacks the covert channel capacity can be merely estimated, using statistical means. Examples

¹Statistical techniques can be used as well, as we describe in Section 3.

are intersection attacks [7], timing attacks [16], Danezis’s attack on continuous mixes [5], and the flow correlation attack [27]. The timing attack [16] uses cross-correlation to match flows given the packet timestamps of the flow. Danezis’s attack on continuous mix [5] uses likelihood ratios to detect a flow in aggregate traffic. The flow correlation attack [27] employs statistical methods to detect TCP flows in aggregate traffic. The flow correlation attack can achieve high detection rates for all the mixes described in [24] and for continuous mixes.

3.1 Attack Model

We model a single mix (Figure 1) as a communication node that connects m senders $S = (s_1, s_2, s_3, \dots, s_m)$ to n receivers $R = (r_1, r_2, r_3, \dots, r_n)$. Every Sender s_i may communicate to every Receiver r_j . We say that a *communication* exists between s_i and r_j whenever s_i communicates to r_j . A communication between s_i and r_j is denoted by the term $[s_i, r_j]$. It can consist of a single packet being sent, or of an established flow.

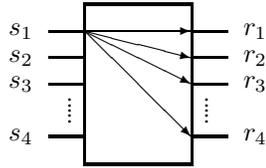


Figure 1. Model of a Mix

We model an *attack* to such a node in terms of its effectiveness in determining who is talking to whom: the set of probabilities $p([s_u, r_v]_s | [s_i, r_j]_a)$ denotes the probability that Communication $[s_u, r_v]_s$ is suspected, given that communication $[s_i, r_j]$ is actually taking place. In other words, a probability $p([s_u, r_v]_s | [s_i, r_j]_a)$ denotes the probability of erroneously suspecting s_u sending to r_v when in actuality s_i is sending to r_j . This model allows for an accurate description of many different attacks, as the probability $p([\cdot, \cdot]_s | [\cdot, \cdot]_a)$ can be defined based on the observation of single packets, a number of packets, a flow or a session, depending on the particular attack method used. For example, the passive attack described in [25] determines a flow successfully when the flow is alone on a link. So the probability $p([s_i, r_j]_s | [s_i, r_j]_a)$ of correctly identifying communication $[s_i, r_j]$ is equal to the chance that the flow is alone on the output link from the mix to Receiver r_j . Alternatively, Danezis’s attack on the continuous mix, the probability $p([s_i, r_j]_s | [s_i, r_j]_a)$ is the probability that the likelihood of the hypothesis assuming that the flow of interest is going through the link between the mix and Receiver r_j is greater than any other hypothesis assuming that the flow of interest is going to any other receiver. Finally, for the flow correlation attack, the probability of $p([s_i, r_j]_s | [s_i, r_j]_a)$ is equal to the probability that the mutual information between the flow of interest and the aggregate traffic on the link between the mix and Receiver r_j is larger than the mutual information between the flow of interest and the aggregate traffic on any other outgoing link.

We note that the attacker may use different attack methods to estimate the probability $p([s_u, r_v]_s | [s_i, r_j]_a)$ for different communications on different mixes, or even on the same mix.

The model above describes attacks on sender-receiver anonymity, where both sender and receiver are anonymous. It can be easily extended to sender anonymity or receiver anonymity, that is, cases where the sender only or the receiver only are anonymous, respectively. For example, we can describe the results of a sender-anonymity attack in terms of $p([s_u, *]_s | [s_i, *]_a)$ or just $p([s_u]_s | [s_i]_a)$. To keep the following discussion simple and general, we will focus on sender-receiver anonymity, with the understanding that sender anonymity or receiver anonymity can be modeled just as well.

3.2 Proposed Anonymity Degree

We define a new measure, D , for the anonymity degree based on the following rationale: Let the random variable $[S, R]_a$ indicate the *actual* sender and receiver pair, and the random variable $[S, R]_s$ in turn indicate the *suspected* sender and receiver pair. If the attack identifies the communicating pairs with high accuracy, then the dependence between the two random variables $[S, R]_a$ and $[S, R]_s$ will be high.

In general, the dependence of two random variables can be measured using the *mutual information* of the two random variables. The mutual information $I(X; Y)$ of two random variables X and Y is a function of the entropies of X and Y as follows:

$$I(X; Y) = H(X) - H(X|Y). \quad (1)$$

Therefore, the effectiveness of the attack can be described in terms of the mutual information $I([S, R]_a; [S, R]_s)$.

To give a more figurative interpretation of mutual information as measure of the attack effectiveness, we use an analogy to communication channels: Mutual information is typically used to describe the amount of information sent across a channel from a sender X to a receiver Y where $H(X)$ is the information at the input of the channel and $H(X|Y)$ describes the information attenuation caused by noise on the channel. (See [18] for an easy-to-read introduction to the information theory used in this context.) This gives an intuition of why mutual information describes the effectiveness of an anonymity attack: Let $[S, R]_a$ be the random variable that describes the actual sender and receiver pair. Let the attacker’s estimate of $[S, R]_a$ through observation of the system, i.e. the attack, be $[S, R]_s$. The information carried through the observation channel provided by the attack is therefore $I([S, R]_a; [S, R]_s)$. The higher this carried information, the more accurate the anonymity attack. Using the textbook definition for entropy, the effectiveness of an anonymity attack can be described as follows:

$$\begin{aligned} I([S, R]_a; [S, R]_s) &= H([S, R]_a) - H([S, R]_a | [S, R]_s) \\ &= \sum_{[s, r]_a, [s, r]_s} p([s, r]_a, [s, r]_s) \log \frac{p([s, r]_s | [s, r]_a)}{p([s, r]_s)}. \end{aligned} \quad (2)$$

In Equation (2), we let $p([s, r]_a, [s, r]_s) = p([s, r]_a)p([s, r]_s|[s, r]_a)$ and $p([s, r]_s) = \sum_{[s, r]_a} p([s, r]_a, [s, r]_s)$. We let $p([s, r]_a)$ denote the *a priori* probability of s communicating to r , typically derived from the expected traffic from s to r .

We can now formulate the *Anonymity Degree* D as a function of the attack effectiveness as follows:

$$D = 1 - \frac{I([S, R]_a; [S, R]_s)}{\log(m \cdot n)}. \quad (3)$$

Since $I([S, R]_a; [S, R]_s) \leq H([S, R]_a) \leq \log(m \cdot n)$, we use $\log(m \cdot n)$ to normalize the anonymity degree into the range of $[0, 1]$ in Equation (3). Alternatively, one could choose $H([S, R]_a)$ as normalization factor. However the latter depends on a *a priori* probability of communication between each pair of sender and receiver. The impact of this *a priori* probability been taken into account by the term $p([s, r]_a)$ in Equation (2).

The equality $I([S, R]_a; [S, R]_s) = H([S, R]_a)$ holds when perfect identification is achieved, that is, $p([s_i, r_j]_s|[s_i, r_j]_a) = 1$ for each pair of sender and receiver. This corresponds to the situation where anonymity is totally broken, in which case the anonymity degree measure D is zero

3.3 Relationship to Previous Anonymity Degree Definitions

The anonymity degree definition D is a generalization of the entropy-based definitions proposed in [23, 8]. In fact, we can rewrite the attack effectiveness $I([S, R]_a; [S, R]_s)$ as

$$\begin{aligned} I([S, R]_a; [S, R]_s) &= H([S, R]_s) - H([S, R]_s|[S, R]_a) \\ &= H([S, R]_s) \\ &- \sum_{[s, r]_a} p([s, r]_a)H([S, R]_s|[S, R]_a = [s, r]_a) \end{aligned} \quad (4)$$

In Equation (4), the term $H([S, R]_s|[S, R]_a = [s, r]_a)$ represents the conditional entropy of the suspected sender-receiver pair distribution given the communication $[s, r]$. This corresponds to the anonymity degree definition described in [23] and also to the core of the anonymity degree defined in [8].

In our mutual-information based anonymity degree, the entropy-based degree is included by averaging according to $p([s, r]_a)$, the *a priori* probability of traffic between each pair. In comparison with entropy-based definitions above, our proposed definition describes the anonymity provided by a network of mixes.

4 Anonymity-Based Covert Channels

Less-than-perfect anonymity systems give raise to a form of covert channel that is exploited by anonymity attacks. We call this form of covert channel *anonymity-based* covert channel. The input symbols of this type of covert channel are the *actual* sender-receiver pairs $[s, r]_a$, and the channel output symbols are the *suspected* sender-receiver pairs

$[s, r]_s$. The channel transition probability $p([s, r]_s|[s, r]_a)$ (i.e. the probability that $[s, r]_s$ is suspected as communication given that $[s, r]_a$ is the actual communication) describes the result of the anonymity attack.

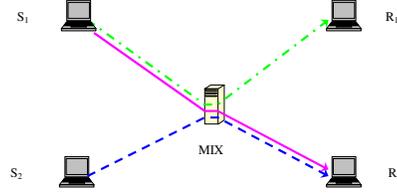


Figure 2. Single-Mix Scenario

We use the simple scenario shown in Figure 2 as an example. We assume that the attacker can collect data at the output ports of the mix as well as some additional information about incoming traffic from the senders. The details on how this information is collected and evaluated depend on the particular attack. See Section 3.1 for examples. Given sufficient collected data, the attacker can detect individual communications, such as $[s_2, r_2]$, with some non-negligible probability, despite the anonymity preserving count-measures in the mix.

The fact that the attacker is able to gain information about communications indicates that a covert channel of the following form exists: A covert channel sender can send a symbol by establishing a communication from some Sender s_2 to Receiver r_1 and send another symbol by establishing a communication from Sender s_2 to another Receiver, r_2 . The covert channel receiver can use the anonymity attack to detect the flow's direction and then make the decision. The channel model is as shown in Figure 3. For sake of simplicity, in this example we limit the covert channel sender to establishing communications from Sender s_2 . Allowing communications from Sender s_1 increases the set of input symbols accordingly.

We compute the capacity of the (anonymity-based) covert channel in textbook fashion by maximizing the mutual information over all input symbol distributions:

$$\begin{aligned} C &= \max_{p([s_2, r]_a)} I([s_2, R]_a; [s_2, R]_s) \\ &= \max_{p([s_2, r]_a)} \sum_{i=1}^2 \sum_{j=1}^2 (p([s_2, r_i]_a, [s_2, r_j]_s) \\ &\cdot \log \frac{p([s_2, r_j]_a, [s_2, r_i]_s)}{p([s_2, r_i]_a)p([s_2, r_j]_s)}). \end{aligned} \quad (5)$$

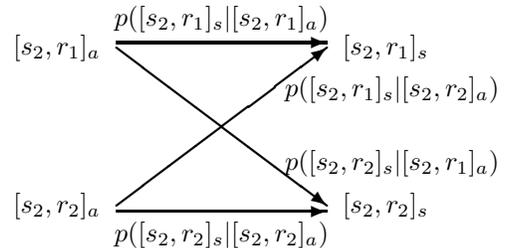


Figure 3. Anonymity-Based Covert Channel Model

The covert channels previously proposed in the context of mix networks [20, 19, 21] are not anonymity-based in the sense described above, as the signal is not received across the channel as the result of an anonymity attack. Rather, they describe information leakage in low-level mechanisms that are used to realize mixes, such as batching mechanisms in [20, 19]. These covert channels are then exploited by the anonymity attacks, which in turn can be used to establish the type of anonymity-based covert channels described in this paper.

5 Single-Mix Case

In a mix with a single Sender s_1 , a covert-channel sender can establish a covert channel by having s_1 communicate with any combination of j among the n receivers. For this covert channel, the set of input symbols is $\{[s_1, r_k]_a : 1 \leq k \leq n\}$ and the set of output symbols is $\{[s_u, r_v]_s : 1 \leq u \leq m, 1 \leq v \leq n\}$. We can include all communications into the set of output symbols because the improbability of any particular communication being declared as suspected by a particular attack can be appropriately reflected by a zero transition probability.

Therefore $\sum_{j=1}^n \binom{n}{j}$ different covert channels can be established. Similarly, if the covert channel sender has control over multiple senders, there are at least $\sum_{i=1}^m \binom{m}{i} \sum_{j=1}^n \binom{n}{j}$ different covert channels that can be established. Which of these $\sum_{i=1}^m \binom{m}{i} \sum_{j=1}^n \binom{n}{j}$ covert channels has the maximum capacity?

LEMMA 1. For a single sender s_i on a single mix, maximum covert channel capacity is achieved when s_i communicates to all receivers.

Proof: By having s_i communicate to all receivers, the covert channel sender can send all the possible symbols $[s_i, r_j]_a, 1 \leq j \leq n$. We call this covert channel x . Without loss of generality, we assume another covert channel y is established by communicating only to a subset of receivers, $r_1, r_2, \dots, r_l, 1 \leq l < n$.

By definition, the capacity of channel x is the maximal mutual information over the distributions $p([s_i, r_1]_a), p([s_i, r_2]_a), \dots, p([s_i, r_n]_a)$, where $\sum_{j=1}^n p([s_i, r_j]_a) = 1$, that is:

$$C_x = \max_{\substack{p([s_i, r_1]_a), p([s_i, r_2]_a), \\ \dots, p([s_i, r_n]_a)}} I([S, R]_a; [S, R]_s) \quad (6)$$

If Sender s_i does not send to Receiver r_j , the probability $p([s_i, r_j]_a)$ is zero. By constraining some of the probabilities to zero, the maximum value of the capacity does not increase.

$$\begin{aligned} C_x &\geq \max_{\substack{p([s_i, r_1]_a), p([s_i, r_2]_a), \dots, \\ p([s_i, r_l]_a), \underbrace{0, \dots, 0}_{n-l}}} I([S, R]_a; [S, R]_s) \\ &= \max_{\substack{p([s_i, r_1]_a), p([s_i, r_2]_a), \\ \dots, p([s_i, r_l]_a)}} I([S, R]_a; [S, R]_s) = C_y \end{aligned}$$

Hence, the capacity of Channel x communicating to all receivers is larger or equal to the capacity of all other covert channels that communicating to only a subset of receivers. ■

THEOREM 1. For a single mix, the maximum covert-channel capacity is achieved when the covert channel sender controls all the Senders s_1, s_2, \dots, s_m , and the input symbols of the corresponding channel include all the possible pairs $[s_i, r_j]_a$.

The proof of Theorem 1 follows the same approach as the proof of Lemma 2.

From Theorem 1, we can get the following corollary.

COROLLARY 1. For the single-mix model shown in Figure 1, the maximum covert-channel capacity is $C = \max_{p([s, r]_a)} I([S, R]_a; [S, R]_s)$.

From Corollary 1 and Equation (3), we get the relationship between the quality of a single mix (i.e. the capacity of any covert channel that allows information to leak from the mix) and the anonymity degree. (Note that this relationship is trivial for the single-mix case. However, we make use of this result in the analysis of networks of mixes.)

LEMMA 2. Given a single mix with a possible maximum information leakage that is upper-bounded by C_{upper} , the anonymity degree of the single mix is lower-bounded by $1 - \frac{C_{upper}}{\log(m \cdot n)}$. Similarly, given that the anonymity degree provided by a single mix is upper-bounded by D_{upper} , the maximum information leakage of the mix is lower-bounded by $(1 - D_{upper}) \log(m \cdot n)$.

Proof: If the covert channel capacity is upper-bounded by C_{upper} ,

$$\begin{aligned} D &= 1 - \frac{I([S, R]_a; [S, R]_s)}{\log(m \cdot n)} \\ &\geq 1 - \frac{C}{\log(m \cdot n)} \\ &\geq 1 - \frac{C_{upper}}{\log(m \cdot n)}. \end{aligned}$$

If the anonymity degree is upper-bounded by D_{upper} ,

$$\begin{aligned} C &= \max(I([S, R]_a; [S, R]_s)) \\ &\geq I([S, R]_a; [S, R]_s) \\ &= (1 - D) \log(m \cdot n) \\ &\geq (1 - D_{upper}) \log(m \cdot n). \end{aligned}$$

Lemma 2 describes how the design and implementation quality of a mix affects effectiveness. In the following sections, we will describe this relation for the case of mix networks. ■

6 Mix Network Case

6.1 Anonymity Degree of a Mix Network

We generalize the anonymity degree for a single mix defined in Equation (3) to the network case by observing that the effectiveness of a mix network can be represented similarly to that of a ‘‘super mix’’. Let R_M and S_M represent the set of senders and receivers of the super mix, respectively. The anonymity degree of the super mix (and of the mix network) is

$$D = 1 - \frac{I([S_M, R_M]_a; [S_M, R_M]_s)}{\log(m \cdot n)} \quad (7)$$

where, similarly to the single-mix case,

$$\begin{aligned} I([S_M, R_M]_a; [S_M, R_M]_s) = & \sum_{[s_i, r_j]_a, [s_u, r_v]_s} (p([s_i, r_j]_a, [s_u, r_v]_s) \\ & \cdot \log \frac{p([s_u, r_v]_s | [s_i, r_j]_a)}{p([s_u, r_v]_s)}). \end{aligned} \quad (8)$$

$I([S_M, R_M]_a; [S_M, R_M]_s)$ is determined by $p([s_i, r_j]_a)$ and $p([s_u, r_v]_s | [s_i, r_j]_a)$, where probability $p([s_i, r_j]_a)$ is the proportion of traffic between s_i and r_j , and the probability $p([s_u, r_v]_s | [s_i, r_j]_a)$ is determined by the results of the anonymity attack at one or more mixes in the mix network. In the following sections, we describe how to make use of the single-mix attack result to describe the effectiveness of a mix network.

6.2 Effectiveness of Single-Mix vs. Super Mix

In the following, we use the term $p_h([s_u, r_v]_s | [s_i, r_j]_a)$ to represent the transition probabilities that are the result of some anonymity attack on Mix M_h , and $p([s_u, r_v]_s | [s_i, r_j]_a)$ to represent the end-to-end transition probability for the super mix. Without loss of generality, we assume in the following that the super mix transition probability we are interested in is $p([s_u, r_v]_s | [s_i, r_j]_a)$. The process to determine the relationship between $p_h([s_u, r_v]_s | [s_i, r_j]_a)$ and $p([s_u, r_v]_s | [s_i, r_j]_a)$ can be divided into three steps.

Step 1: Find the set P_{uv} of all the possible paths between s_u and r_v . Clearly

$$\begin{aligned} p([s_u, r_v]_s | [s_i, r_j]_a) = & \sum_{P_a \in P_{uv}} p([s_u, r_v]_{s, P_a} | [s_i, r_j]_a) \end{aligned} \quad (9)$$

where $p([s_u, r_v]_{s, P_a} | [s_i, r_j]_a)$ denotes the probability of suspecting communication $[s_i, r_j]_a$ to be communication $[s_u, r_v]_s$ over Path P_a . Note that the actual communication between s_i and r_j takes only one path, which we call Path P_0 .

Step 2: Determine the probability of suspecting an actual communication over Path P_0 to be the communication over another path P_a . Depending on how Path P_a and Path P_0

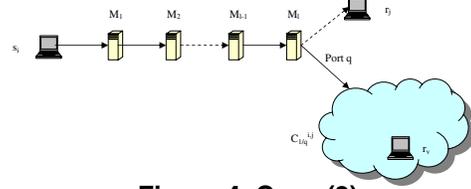


Figure 4. Case (2)

overlap, we distinguish three situations: (i) There is only one segment where the two paths overlap. (ii) The two paths share multiple segments. (iii) There is no overlap between the two paths. Since there is no overlap in Situation (iii), the probability of suspecting a communication over path P_0 to be the communication over path P_a is zero. Hence, we only need to further pursue Situation (i) and Situation (ii).

Situation (i) can be divided into four sub-cases:

Case (1): P_0 and P_a are identical. This implies that $s_u = s_i$ and $r_v = r_j$. In this case, the probability of suspecting correctly is the product of the probabilities of locally suspecting correctly at all mixes along Path P_0 . If we denote the mixes on Path P_0 to be M_1, M_2, \dots, M_l , then

$$\begin{aligned} p([s_i, r_j]_{s, P_0} | [s_i, r_j]_a) = & p_1([s_i, M_2]_s | [s_i, M_2]_a) \\ & \cdot \left(\prod_{d=2}^{l-1} p_d([M_{d-1}, M_{d+1}]_s | [M_{d-1}, M_{d+1}]_a) \right) \\ & \cdot p_l([M_{l-1}, r_j]_s | [M_{l-1}, r_j]_a). \end{aligned} \quad (10)$$

This follows directly from the fact that correct guesses at each mix on the path cause the attacker to correctly suspect the actual path.

Case (2): P_0 and P_a share the same path from s_i through the first Mix M_1 to some Mix M_l , and then diverge due to an error at Mix M_l . This is illustrated in Figure 4 where, in order to emphasize the path P_0 and P_a , other possible connections among the mixes and other possible mixes are ignored. The fact that P_0 and P_a share the same path from s_i means that s_i is correctly suspected, i.e. $s_u = s_i$.

In this subcase, the probability of erroneously suspecting some receiver r_v other than r_j is the result of correctly identifying the path up to some Mix M_{l-1} , and then making a mistake at Mix M_l . Once an error has been made, the remaining mixes on the path to any erroneously suspected Receiver r_v are not on Path P_0 . According to the attack model described in Section 3, no differentiation can be made between r_v and any other receiver that can be reached after making an error at Mix M_l . We therefore aggregate all receivers that can be reached after an error at Mix M_l into what we call a *cloud* of receivers. We denote by $C_{l/q}^{ij}$ the cloud that is a result of an error at Mix M_l , where communication $[s_i, r_j]_a$ is incorrectly identified because Port q was erroneously selected instead of the port taken by $[s_i, r_j]_a$. For the example in Figure 4, the probability of suspecting receiver to be inside Cloud $C_{l/q}^{ij}$ is

$$p([s_i, C_{l/q}^{ij}]_s | [s_i, r_j]_a) = p_1([s_i, M_2]_s | [s_i, M_2]_a)$$

Constraints: The optimization problem is subject to the following three sets of constraints:

[Constraint Set 1:] The sum of all path-independent transition probabilities to all the end nodes in a group of clouds is identical to the sum of path-dependent end-to-end transition probabilities to the clouds in the group. For simplicity of notation, we formulate this for the special case of a correctly suspected Sender s_i . The extension to the general case is cumbersome, but straightforward. Let $GR_v^{i,j}$ be the smallest set of receiver clouds that contain r_v and all receivers in $GR_v^{i,j}$.

$$\begin{aligned} \forall r_v : & \sum_{r_w \in GR_v^{i,j}} p([s_i, r_w]_s | [s_i, r_j]_a) \\ & = \sum_{RC_{l/q}^{i,j} \in GR_v^{i,j}} p([s_i, RC_{l/q}^{i,j}]_{s,P_b} | [s_i, r_j]_a) \end{aligned} \quad (14)$$

[Constraint Set 2:] The sum of all path-independent transition probabilities to a sub-group of receivers is larger than the sum of the path-dependent end-to-end transition probabilities to the clouds which only contain the receivers in the sub-group. It is true because one receiver in the sub-group may be contained in another cloud which contains the receivers not in the sub-group. Let R_{sub} be a subset of the set R of all receivers. Define $H_{R_{sub}}^{i,j}$ to be the set of all clouds that contain *only* receivers in R_{sub} . For the simple case of a correctly suspected Sender s_i :

$$\begin{aligned} \forall R_{sub} : & \sum_{r_v \in R_{sub}} p([s_i, r_v]_s | [s_i, r_j]_a) \\ & \geq \sum_{RC_{l/q}^{i,j} \in H_{R_{sub}}^{i,j}} p([s_i, RC_{l/q}^{i,j}]_{s,P_b} | [s_i, r_j]_a). \end{aligned} \quad (15)$$

[Constraint Set 3:] The sum of all path-independent transition probabilities to a sub-group of receivers is less than the sum of the path-dependent end-to-end transition probabilities to the clouds which have at least one receiver in the sub-group. It is true because these clouds may have other receivers which are not in the sub-group. Let R_{sub} be a subset of the set R of all receivers. Define $I_{R_{sub}}^{i,j}$ to be the set of all clouds that contains *at least one* of the receivers in R_{sub} . We can conclude:

$$\begin{aligned} \forall R_{sub} : & \sum_{r_v \in R_{sub}} p([s_i, r_v]_s | [s_i, r_j]_a) \\ & \leq \sum_{RC_{l/q}^{i,j} \in I_{R_{sub}}^{i,j}} p([s_i, RC_{l/q}^{i,j}]_{s,P_b} | [s_i, r_j]_a). \end{aligned} \quad (16)$$

[Constraint set 4:] The end-to-end transition probabilities for all suspects for all actual communications sum up to 1:

$$\forall i, j \sum_{s_u, r_v} p([s_u, r_v]_s | [s_i, r_j]_a) = 1. \quad (17)$$

The solution of this optimization problem is the set of the end-to-end transition probabilities of the super mix that minimize the anonymity degree of the mix network.

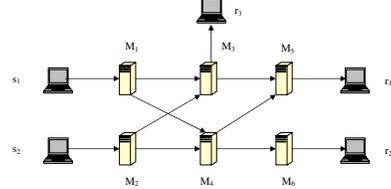


Figure 7. A Small Example

6.3 A Small Example

We use the example mix network displayed in Figure 7 to illustrate how to compute end-to-end transition probabilities as described in Step 2 of Section 6.2.

We focus on communication $[s_1, r_1]$. Suppose the actual communication takes the route $P_0: s_1 \rightarrow M_1 \rightarrow M_3 \rightarrow M_5 \rightarrow r_1$. In this case, the probability of (erroneously) suspecting communications $[s_1, r_3]$ is computed as follows:

$$\begin{aligned} p([s_1, r_3]_s | [s_1, r_1]_a) & = p_1([s_1, M_3]_s | [s_1, M_3]_a) \\ & \cdot p_3([M_1, r_3]_s | [M_1, M_3]_a). \end{aligned} \quad (18)$$

This computation is simple, since there is only one path from s_1 to r_3 .

The situation of (correctly) suspecting communication $[s_1, r_1]_a$ is more complicated, because two paths can be taken. One is $P_0: s_1 \rightarrow M_1 \rightarrow M_3 \rightarrow M_5 \rightarrow r_1$, the other is $P_1: s_1 \rightarrow M_1 \rightarrow M_4 \rightarrow M_5 \rightarrow r_1$. Clearly, we have

$$\begin{aligned} p([s_1, r_1]_{s,P_0} | [s_1, r_1]_a) & = p_1([s_1, M_3]_s | [s_1, M_3]_a) \\ & \cdot p_3([M_1, M_5]_s | [M_1, M_5]_a) \\ & \cdot p_5([M_3, r_1]_s | [M_3, M_1]_a) \end{aligned} \quad (19)$$

of suspecting $[s_1, r_1]$ over Path P_0 .

For path P_1 , we can not get express $p([s_1, r_1]_{s,P_1} | [s_1, r_1]_a)$ directly in terms of anonymity attack result at mixes, because the wrong guess at Mix M_1 will possibly lead to two receivers, r_1 and r_2 . So we have to aggregate Receiver r_1 and r_2 in receiver cloud $C_{1/q}^{1,1}$, where q denotes the wrongly selected output port at Mix M_1 . So what we can get is

$$\begin{aligned} p([s_1, C_{1/q}^{1,1}]_s | [s_1, r_1]_a) & = \\ p_1([s_1, M_4]_s | [s_1, M_1]_a), \end{aligned} \quad (20)$$

where the erroneous selection of Port q on Mix M_1 leads to the suspected Path $s_1 \rightarrow M_1 \rightarrow M_4$. Clearly both Receiver r_1 and Receiver r_2 can be reached after selecting Port q on Mix M_1 .

In turn, by following Equation (14), we can get

$$\begin{aligned} p([s_1, r_1]_s | [s_1, r_1]_a) & + p([s_1, r_2]_s | [s_1, r_1]_a) = \\ p_1([s_1, M_4]_s | [s_1, M_1]_a) & + p_1([s_1, M_3]_s | [s_1, M_3]_a) \\ \cdot p_3([M_1, M_5]_s | [M_1, M_5]_a) & \\ \cdot p_5([M_3, r_1]_s | [M_3, M_1]_a) & . \end{aligned} \quad (21)$$

After repeating this for all possible sender-receiver pairs, expressions for the end-to-end transition can be formulated, and the optimization described in Step 3 of Section 6.2 can be used to determine the anonymity degree of the network.

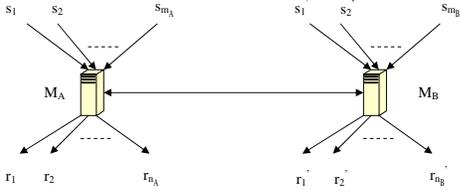


Figure 8. Mix Network of Two Mixes

7 Covert Channel Capacity vs. Anonymity Degree in Mix Networks

The analysis of the effectiveness of anonymity networks is rendered difficult for two reasons, among others: First, attacks on such networks are typically out-of-the-box attacks (for example none of the intersection attacks, trickle attacks, or others target measures taken by the mix network). Second, it is unknown where and how traffic information is collected. Is the attack targeting individual mixes or clusters of mixes? Is the information collected on a per-mix or a per-link basis?

In this section we describe how the anonymity in mix networks can be systematically analyzed and bounded based on estimates of either per-mix weakness (using local covert channels) or the entire mix network (using network-wide covert channels). For this purpose, we investigate the relation between the covert channel capacity of a mix network and the anonymity provided by the network.

7.1 Upper Bound on the Covert Channel Capacity in Mix Networks

Let the mix network have K mixes. For Mix M_h , we use S_h and R_h to represent the set of senders and receivers of Mix M_h respectively. Any anonymity attack on Mix M_h will lead to a set of probabilities of the form $p_h([s_u, r_v]_s | [s_i, r_j]_a)$ with s_u and s_i in S_h and r_v and r_j in R_h .

In a mix network, there are various ways to establish covert channels. For example, in the mix network shown in Figure 8, there are at least two ways to establish the covert channels using the two mixes M_A and M_B . One way is to establish one covert channel on M_A and M_B separately. Alternatively, one can establish a covert channel on the super mix containing both M_A and M_B . We assume each mix can only be contained in one covert channel as before. In the following, we use the notation $cc(\mathcal{M})$ to denote the covert channel that can be established over the set of the mixes \mathcal{M} . If we denote the capacities of $cc(\{M_A\})$ and $cc(\{M_B\})$ to be C_A and C_B , respectively, then the sum of the covert channel capacity clearly is $C_A + C_B$. We have the following lemma:

LEMMA 3. The capacity of $cc(\{M_A, M_B\})$ will be no greater than $C_A + C_B$.

The proof is constructive in nature and can be found in [26]. Extending the two mixes case in Lemma 3, we can get the following Lemma.

LEMMA 4. For two mixes connected with more than one links, the capacity of the covert channel built on the super mix, $cc(\{M_A, M_B\})$ will be no greater than $C_A + C_B$.

The proof is similar to that of Lemma 3. Instead of only one path between M_A and M_B , there are more than one paths between M_A and M_B . But it will not affect the use of the inequalities employed in the proof of Lemma 3.

THEOREM 2. In a mix network of K mixes, the sum of the capacities of all the covert channels in the mix network

will be no greater than $\sum_{h=1}^K C_h$.

Proof: This theorem can be proved by induction on K mixes with the help of Lemma 4, as any set $K + 1$ mixes can be partitioned into a supermix of K mixes and a single mix.

7.2 Relationship

Similarly to the single-mix case in Section 5, we are interested in how bounds on the achievable anonymity degree are affected by the covert channel capacity of the system, and *vice versa*. For example, it is obvious that an upper bound on the anonymity degree will result in a lower bound on the total covert-channel capacity, following the observation that anonymity attacks are more effective in less anonymous mixes.

The upper bound D_{upper} on the anonymity gives rise to a lower bound C_{lower} on the sum of the local channel capacities:

$$C_{lower} = \min\left(\sum_{n=1}^K C_h\right) \quad (22)$$

Equation (22) gives rise to a minimization problem over anonymity attack results $p_h([s_u, r_v]_s | [s_i, r_j]_a)$, with the following three constrains: First, the local *a priori* probabilities for communications at each Mix M_h must sum to one:

$$\sum_{i=1}^{m_h} \sum_{j=1}^{n_h} p_h([s_i, r_j]_a) = 1. \quad (23)$$

Second, the transition probability from each input symbol $[s_i, r_j]_a$ of each mix should sum up to one:

$$\sum_{u=1}^{m_h} \sum_{v=1}^{n_h} p_h([s_u, r_v]_s | [s_i, r_j]_a) = 1, \quad (24)$$

Third, the anonymity of the system, as computed in Section 6.1, should not exceed D_{upper} .

We can solve this constrained optimization problem analytically by using Lagrange multipliers and Kuhn-Tucker conditions. Or we can use numerical methods such as Monte-Carlo.

Similarly, given upper bound C_{upper} on the total covert channel capacity of the mix network, we would like to find out a lower bound D_{lower} for anonymity degree of the mix network.

The objective function becomes

$$D_{lower} = \min\left[1 - \frac{I([S_M, R_M]_a; [S_M, R_M]_s)}{\log(m \cdot n)}\right] \quad (25)$$

Communication	Actual Path
$[s_1, r_1]_a$	$s_1 \rightarrow M_1 \rightarrow M_3 \rightarrow M_5 \rightarrow r_1$
$[s_1, r_2]_a$	$s_1 \rightarrow M_1 \rightarrow M_3 \rightarrow M_6 \rightarrow r_1$
$[s_2, r_1]_a$	$s_1 \rightarrow M_2 \rightarrow M_4 \rightarrow M_5 \rightarrow r_1$
$[s_2, r_2]_a$	$s_2 \rightarrow M_2 \rightarrow M_4 \rightarrow M_6 \rightarrow r_2$

Table 1. Path of the Actual Communications

This optimization problem is over all possible anonymity attack result $p_h([s_u, r_v]_s | [s_i, r_j]_a)$. Constraints (23) and (24) still in this case. The new constraint is

$$C_{upper} \geq \sum_{h=1}^K C_h \quad (26)$$

8 Evaluation

We use the mix network shown in Figure (9) as an example to illustrate the relationships established in the previous section. We choose six mixes because it is not a trivial topology, and both a mix cascade and a stratified network case [10] can be established on the six mixes.

We assume that communications between each sender-receiver pair have the same *a priori* probability (alternatively, the same share of total traffic volume). Since there are two senders and two receivers, we have four sender-receiver pairs. The actual path for Communication $[s_i, r_j]_a$ is shown in Table 1 if the actual path is not specified and the path is possible in the topology. We assume the anonymity attack at each mix is useful, meaning the attack can identify the actual local communication $[s_i, r_j]_a$ with a probability equal or larger than random guess. For our examples, we use adaptive simulated annealing to solve the optimization problem to establish D_{lower} from a known bound on the mix network capacity.

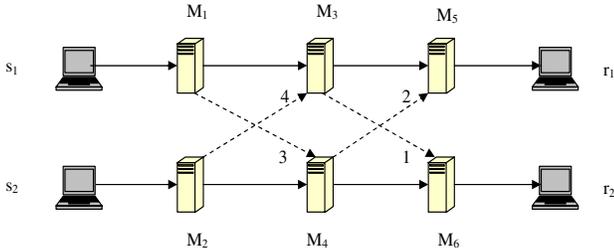


Figure 9. An Example Mix Network

Impact of the Connectivity Obviously the connectivity will affect the anonymity degree in a mix network. In our first set of examples, the base topology contains only the solid lines in Figure 9. Then edges are incrementally added to the base topology in the order of the label assigned to each edge. The average degree of the topologies including base topology are $2, \frac{14}{6}, \frac{16}{6}, 3, \frac{20}{6}$ respectively.

For every mix in the base topology, there is only one input link and one output link. So there is only one sender receiver pair for the mix in the base topology. A channel which has only one input symbol and one output symbol will have capacity zero. So the capacity C_{sum} is zero for base topology.

From Figure 10(a), first we can observe that the lower bound of the anonymity degree decreases with increasing bound on the capacity, just as we expect. In addition, the capacity C_{sum} increases with increasing connectivity. For a given upper bound of the capacity C_{sum} , increasing connectivity will increase the anonymity degree. Third, we can observe that there is large gap between the base topology and the topology of the next higher average degree. This is because adding the edge of label 1 will connect s_1 and r_2 and the Communication $[s_1, r_2]_a$ can be suspected as $[s_1, r_1]_s$. So the initial edge added to the topology can increase the anonymity degree significantly. In comparison, the effect of adding edge with Label 4 is marginal.

Effect of Adding Different Edges In the second set of examples, we use the solid lines and edge with label 1 as base topology. Then we add one more edge 2, 3 or 4 to the base topology. We label the new topology as *A*, *B* and *C* respectively. Clearly these topologies are of the same average degree. From Figure 10(b), we can observe that the anonymity degree increase cause by adding edge with label 3 is smaller than adding the other two edges. This is because adding the other two edges can make Communication $[s_2, r_1]_a$ possible and the Communication $[s_2, r_1]_a$ can be suspected as other communications.

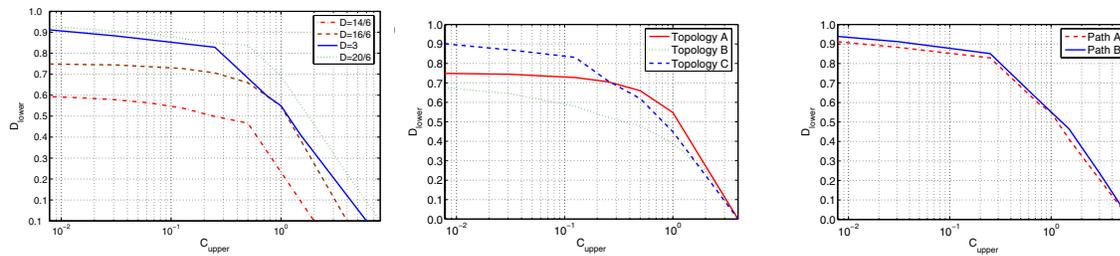
Effect of Path Selection In this set of examples, we focus on the topology containing all the solid and dashed lines except the edge with label 3. We consider two cases. In one case, the actual path for Communication $[s_2, r_1]_a$ follows Path *A* as in Table 1. In the other case, the actual path *B* for Communication $[s_2, r_1]_a$ is $s_1 \rightarrow M_2 \rightarrow M_3 \rightarrow M_5 \rightarrow r_1$.

We can observe going through Mix M_3 will slightly increase the anonymity from Figure 10(c). This is because Mix M_3 has more output and input links than the other mixes. So the communication through Mix M_3 is more easy to hide.

9 Summary and Future Work

We propose a new mutual information based anonymity degree. It gives out one number which is between zero and one to indicate the overall effectiveness of a whole mix network. We also gives out a proof on how to achieve maximal covert channel capacity for a single mix based on anonymity attacks on the mix. The relationship between the anonymity degree and anonymity attack based covert channel capacity is derived for both a single mix case and mix network case.

Our work is the first to give out the relationship between anonymity degree and the capacity of anonymity-based covert channel. In the mix network case, these relationship are described in a scenario-oriented fashion. What is needed is a set of rules to map and cluster arbitrary networks into super mixes and clouds. Further research will be on the multicast or broadcast channels in the anonymity network and its relationship with anonymity degree. Finally, we need to extend the work from anonymity-based covert channels to general covert channel in mix networks, such as the non-anonymity based covert channels as in [20, 19, 21] or other formalizations of information leakage based at-



(a) Impact of the Connectivity (b) Effect of Adding Different Edges (c) Effect of Path Selection

Figure 10. Effect of Topology

tacks. Eventually, a conclusion is needed that allows to aggregate attacks and so formulate the level of anonymity provided by systems with less-than-perfect components.

References

- [1] O. Berthold, A. Pfitzmann, and R. Standtke. The disadvantages of free MIX routes and how to overcome them. In H. Federrath, editor, *Proceedings of Designing Privacy Enhancing Technologies: Workshop on Design Issues in Anonymity and Unobservability*, pages 30–45. Springer-Verlag, LNCS 2009, July 2000.
- [2] D. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 4(2), February 1981.
- [3] D. Chaum. The dining cryptographers problem: Unconditional sender and recipient untraceability. *Journal of Cryptology*, 1:65–75, 1988.
- [4] I. Clarke, O. Sandberg, B. Wiley, and T. W. Hong. Freenet: A distributed anonymous information storage and retrieval system. In *Proceedings of Designing Privacy Enhancing Technologies: Workshop on Design Issues in Anonymity and Unobservability*, pages 46–66, July 2000.
- [5] G. Danezis. The traffic analysis of continuous-time mixes. In *Proceedings of Privacy Enhancing Technologies workshop (PET 2004)*, LNCS, May 2004.
- [6] G. Danezis, R. Dingledine, and N. Mathewson. Mixminion: Design of a Type III Anonymous Remailer Protocol. In *Proceedings of the 2003 IEEE Symposium on Security and Privacy*, May 2003.
- [7] G. Danezis and A. Serjantov. Statistical disclosure or intersection attacks on anonymity systems. In *Proceedings of 6th Information Hiding Workshop (IH 2004)*, LNCS, Toronto, May 2004.
- [8] C. Diaz, S. Seys, J. Claessens, and B. Preneel. Towards measuring anonymity. In R. Dingledine and P. Syverson, editors, *Proceedings of Privacy Enhancing Technologies Workshop (PET 2002)*. Springer-Verlag, LNCS 2482, April 2002.
- [9] R. Dingledine, N. Mathewson, and P. Syverson. Tor: The second-generation onion router. In *Proceedings of the 13th USENIX Security Symposium*, August 2004.
- [10] R. Dingledine, V. Shmatikov, and P. Syverson. Synchronous batching: From cascades to free routes. In *Proceedings of Privacy Enhancing Technologies workshop (PET 2004)*, LNCS, May 2004.
- [11] M. J. Freedman and R. Morris. Tarzan: A peer-to-peer anonymizing network layer. In *Proceedings of the 9th ACM Conference on Computer and Communications Security (CCS 2002)*, Washington, DC, November 2002.
- [12] D. Goldschlag, M. Reed, and P. Syverson. Onion routing for anonymous and private internet connections. *Communications of the ACM (USA)*, 42(2):39–41, 1999.
- [13] C. Gülcü and G. Tsudik. Mixing E-mail with Babel. In *Proceedings of the Network and Distributed Security Symposium - NDSS '96*, pages 2–16. IEEE, February 1996.
- [14] J. Helsingius. Press release: Johan Helsingius closes his Internet Remailer. <http://www.penet.fi/press-english.html>, 1996.
- [15] D. Kesdogan, J. Egner, and R. Büschkes. Stop-and-go MIXes: Providing probabilistic anonymity in an open system. In *Proceedings of Information Hiding Workshop (IH 1998)*. Springer-Verlag, LNCS 1525, 1998.
- [16] B. N. Levine, M. K. Reiter, C. Wang, and M. K. Wright. Timing attacks in low-latency mix-based systems. In A. Juels, editor, *Proceedings of Financial Cryptography (FC '04)*. Springer-Verlag, LNCS 3110, February 2004.
- [17] U. Möller and L. Cottrell. Mixmaster Protocol — Version 2. Unfinished draft, January 2000.
- [18] I. S. Moskowitz and M. H. Kang. Covert channels – here to stay? In *Proceedings of COMPASS 1994*. IEEE Press, June 1994.
- [19] I. S. Moskowitz, R. E. Newman, D. P. Crepeau, and A. R. Miller. Covert channels and anonymizing networks. In *Proceedings of the Workshop on Privacy in the Electronic Society (WPES 2003)*, Washington, DC, USA, October 2003.
- [20] I. S. Moskowitz, R. E. Newman, and P. F. Syverson. Quasi-anonymous channels. In *Proceedings of CNIS 2003*, December 2003.
- [21] R. E. Newman, V. R. Nalla, and I. S. Moskowitz. Anonymity and covert channels in simple timed mix-firewalls. In *Proceedings of Privacy Enhancing Technologies workshop (PET 2004)*, May 2004.
- [22] M. Reiter and A. Rubin. Crowds: Anonymity for web transactions. *ACM Transactions on Information and System Security*, 1(1), June 1998.
- [23] A. Serjantov and G. Danezis. Towards an information theoretic metric for anonymity. In R. Dingledine and P. Syverson, editors, *Proceedings of Privacy Enhancing Technologies Workshop (PET 2002)*. Springer-Verlag, LNCS 2482, April 2002.
- [24] A. Serjantov, R. Dingledine, and P. Syverson. From a trickle to a flood: Active attacks on several mix types. In F. Petitcolas, editor, *Proceedings of Information Hiding Workshop (IH 2002)*. Springer-Verlag, LNCS 2578, October 2002.
- [25] A. Serjantov and P. Sewell. Passive attack analysis for connection-based anonymity systems. In *Proceedings of ESORICS 2003*, October 2003.
- [26] Y. Zhu and R. Bettati. Anonymity v.s. information leakage in a anonymity systems. Technical Report 2004-10-3, Texas A&M University, Computer Science Department, October 2004.
- [27] Y. Zhu, X. Fu, B. Graham, R. Bettati, and W. Zhao. On flow correlation attacks and countermeasures in mix networks. In *Proceedings of Privacy Enhancing Technologies workshop (PET 2004)*, May 2004.