

Practical Traffic Analysis: Extending and Resisting Statistical Disclosure

Nick Mathewson and Roger Dingledine

The Free Haven Project
{nickm,arma}@freehaven.net

Abstract. We extend earlier research on mounting and resisting passive long-term end-to-end traffic analysis attacks against anonymous message systems, by describing how an eavesdropper can learn sender-receiver connections even when the substrate is a network of pool mixes, the attacker is non-global, and senders have complex behavior or generate padding messages. Additionally, we describe how an attacker can use information about message distinguishability to speed the attack. We simulate our attacks for a variety of scenarios, focusing on the amount of information needed to link senders to their recipients. In each scenario, we show that the intersection attack is slowed but still succeeds against a steady-state mix network. We find that the attack takes an impractical amount of time when message delivery times are highly variable; when the attacker can observe very little of the network; and when users pad consistently and the adversary does not know how the network behaves in their absence.

1 Introduction

Mix networks aim to allow senders to anonymously deliver messages to recipients. One of the strongest attacks against current deployable designs is the *long-term intersection attack*. In this attack, a passive eavesdropper observes a large volume of network traffic and notices that certain recipients are more likely to receive messages after particular senders have transmitted messages. That is, if a sender (call her Alice) maintains a fairly consistent pattern of recipients over time, the attacker can deduce Alice’s recipients.

Researchers have theorized that these attacks should be extremely effective in many real-world contexts, but so far it has been difficult to reason about when these attacks would be successful and how long they would take.

Here we extend a version of the long-term intersection attack called the statistical disclosure attack [13] to work in real-world circumstances. Specifically, whereas the original model for this attack makes strong assumptions about sender behavior and only works against a single batch mix, we show how an attacker can learn Alice’s regular recipients even when:

- Alice chooses non-uniformly among her communication partners, and can send multiple messages in a single batch.

- The attacker lacks *a priori* knowledge of the network’s average behavior when Alice is not sending messages.
- Mixes use a different batching algorithm, such as Mixmaster’s dynamic-pool algorithm [25, 30] or its generalization [17].
- Alice uses a mix network (of any topology, with synchronous or asynchronous batching) to relay her messages through a succession of mixes, instead of using just a single mix.
- Alice disguises when she is sending real messages by sending padding traffic to be dropped by mix nodes in the network.
- The attacker can only view a subset of the messages entering and leaving the network (so long as this subset includes some messages from Alice and some messages to Alice’s recipients).
- The cover traffic generated by other senders changes slowly over time. (We do not address this case completely.)

Each deviation from the original model reduces the rate at which the attacker learns Alice’s recipients, and increases the amount of traffic he must observe.

Additionally, we show how an attacker can exploit additional knowledge, such as distinguishability between messages, to speed these attacks. For example, an attacker who sees message contents can take into account whether messages are written in the same language or signed by the same pseudonym, and thereby partition messages into different classes and analyze the classes independently.

The attacks in this paper fail to work when:

- Alice’s behavior is not consistent over time. If Alice does not produce enough traffic with the same recipients, the attacker cannot learn her behavior.
- The attacker cannot observe how the network behaves in Alice’s absence. If Alice always sends the same number of messages, in every round, forever, a passive attacker cannot learn who receives messages in Alice’s absence.
- The attacker cannot tell when Alice is originating messages.

We begin in Section 2 by presenting a brief background overview on mix networks, traffic analysis, the disclosure attack, and the statistical disclosure attack. In Section 3 we present our enhancements to the statistical disclosure attack. We present simulated experimental results in Section 4, and close in Section 5 with recommendations for resisting this class of attacks, implications for mix network design, and a set of open questions for future work.

2 Previous work

Chaum [10] proposed hiding the correspondence between sender and recipient by wrapping messages in layers of public-key cryptography, and relaying them through a path composed of *mixes*. Each mix in turn decrypts, delays, and re-orders messages, before relaying them toward their destinations. Because some mixes might be controlled by an adversary, Alice can direct her messages through a sequence or ‘chain’ of mixes in a network, so that no single mix can link her to her recipient.

Many subsequent designs have been proposed, including Babel [21], Mixmaster [25], and Mixminion [14]. We will not address the differences between these systems in any detail: from the point of view of a long-term intersection attack, the internals of the network are irrelevant so long as the attacker can observe messages entering and leaving the network, and can guess when a message entering the network is likely to leave.

Another class of anonymity designs aims to provide low-latency connections for web browsing and other interactive services [6, 9, 18, 28]. We do not address these systems here because short-term timing and packet counting attacks seem sufficient against them [31].

Attacks against mix networks aim to reduce the anonymity of users by linking anonymous senders with the messages they send, by linking anonymous recipients with the messages they receive, or by linking anonymous messages with one another. For detailed lists of attacks, consult [2, 27]. Attackers can trace messages through the network by observing network traffic, compromising mixes, compromising keys, delaying messages so they stand out from other traffic, or altering messages in transit. They can learn a given message’s destination by flooding the network with messages, replaying multiple copies of a message, or shaping traffic to isolate a target message from other unknown traffic [30]. Attackers can discourage users from using honest mixes by making them unreliable [2, 19]. They can analyze intercepted message text to look for commonalities between otherwise unlinked senders [26].

2.1 The long-term intersection attack

Even if we foil all the above attacks, an adversary can mount a *long-term intersection attack* by correlating times when senders and receivers are active [8].

A variety of countermeasures make intersection attacks harder. Kesdogan’s stop-and-go mixes [23] provide probabilistic anonymity by letting users specify message latencies, thereby broadening the range of time when messages might leave the mix network. Similarly, batching strategies [30] as in Mixmaster and Mixminion use message pools to spread out the possible exit times for messages.

Rather than expanding the set of messages that might have been sent by a suspect sender, other designs expand the set of senders that might have sent a target message. A sender who also runs a node in the mix network can conceal whether a given message originated at her node or was relayed from another node [5, 20, 29]. But even with this approach, the adversary can observe whether certain traffic patterns are present when a user is online (possibly sending) and absent when a user is offline (certainly not sending) [33, 34].

A sender can also conceal whether she is currently active by consistently sending decoy (dummy) traffic. Pipenet [11] conceals traffic patterns by constant padding on every link. Unfortunately, a single user can shut down this network simply by not sending. Berthold and Langos aim to increase the difficulty of intersection attacks with a scheme for preparing plausible dummy traffic and having other nodes send it on Alice’s behalf when she is offline [7], but their design has many practical problems.

Finally, note that while the adversary can perform this long-term intersection attack entirely passively, active attacks (such as blending attacks [30] against a suspected sender) can help him reduce the set of suspects at each round.

2.2 The disclosure attack

In 2002, Kesdogan, Agrawal, and Penz presented the disclosure attack [22], an intersection attack against a single sender on a single batch mix.

The disclosure attack assumes a global passive eavesdropper interested in learning the recipients of a single sender Alice. It assumes that Alice sends messages to m recipients; that Alice sends a single message (recipient chosen at random from m) in each batch of b messages; and that the recipients of the other $b - 1$ messages are chosen at random from the set of N possible recipients.

The attacker observes the messages leaving the mix and constructs sets R_i of recipients receiving messages in batch i . The attacker then performs an NP-complete computation to identify m mutually disjoint recipient sets R_i , so that each of Alice’s recipients is necessarily contained in exactly one of the sets. Intersecting these sets with subsequent recipient sets reveals Alice’s recipients.

2.3 The statistical disclosure attack

In 2003, Danezis presented the statistical disclosure attack [13], which makes the same operational assumptions as the original disclosure attack but is far easier to implement in terms of storage, speed, and algorithmic complexity. Unlike its predecessor, statistical disclosure only reveals *likely* recipients; it does not identify Alice’s recipients with certainty.

In the statistical disclosure attack, we model Alice’s behavior as an unknown vector \vec{v} whose elements correspond to the probability of Alice sending a message to each of the N recipients in the system. The elements of \vec{v} corresponding to Alice’s m recipients will be $1/m$; the other $N - m$ elements of \vec{v} will be 0. We model the behavior of the covering “background” traffic sent by other users as a known vector \vec{u} each of whose N elements is $1/N$.

The attacker derives from each output round i an observation vector \vec{o}_i , each of whose elements corresponds to the probability of Alice’s having sent a message to each particular recipient in that round. That is, in a round i where Alice has sent a message, each element of \vec{o}_i is $1/b$ if it corresponds to a recipient who has received a message, and 0 if it does not. Taking the arithmetic mean \vec{O} of a large set of these observation vectors gives (by the law of large numbers):

$$\vec{O} = \frac{1}{t} \sum_{i=1}^t \vec{o}_i \approx \frac{\vec{v} + (b-1)\vec{u}}{b}$$

From this, the attacker estimates Alice’s behavior:

$$\vec{v} \approx b \frac{\sum_{i=1}^t \vec{o}_i}{t} - (b-1)\vec{u}$$

Danezis also derives a precondition that the attack will only succeed when $m < \frac{N}{b-1}$, and calculates the expected number of rounds to succeed (with 95% confidence for security parameter $l = 2$ and 99% confidence for $l = 3$) [12]:

$$t > \left[m \cdot l \left(\sqrt{\frac{N-1}{N}(b-1)} + \sqrt{\frac{N-1}{N^2}(b-1) + \frac{m-1}{m}} \right) \right]^2$$

3 Extending the statistical disclosure attack

3.1 Broadening the attack

Here we examine ways to extend Danezis’s statistical disclosure attack to systems more closely resembling real-world mix networks. We will simulate the time and information requirements for several of these attacks in Section 4 below.

Complex senders, unknown background traffic: First, we relax the requirements related to sender behavior. We allow Alice to choose among her recipients with non-uniform probability, and to send multiple messages in a single batch. We also remove the assumption that the attacker has full knowledge of the distribution \vec{u} of cover traffic sent by users other than Alice.

To model Alice’s varying number of messages, we use a probability function P_m such that in every round Alice sends n messages with probability $P_m(n)$. We still use a behavior vector \vec{v} to represent the probability of Alice sending to each recipient, but we no longer require Alice’s recipients to have a uniform $1/m$ probability. Alice’s expected contribution to each round is thus $\vec{v} \sum_{n=0}^{\infty} n P_m(n)$.

To mount the attack, the attacker first obtains an estimate of the background distribution \vec{u} by observing a large number t' of batches to which Alice has *not* contributed any messages.¹ For each such batch i , the attacker constructs a vector \vec{u}_i , whose elements are $1/b$ for recipients that have received a message in that batch, and 0 for recipients that have not. The attacker then estimates the background distribution \vec{u} as:

$$\vec{u} \approx \bar{U} = \frac{1}{t'} \sum_{i=1}^{t'} \vec{u}_i$$

The attacker then observes, for each round i in which Alice *does* send a message, the number of messages m_i sent by Alice, and computes observations \vec{o}_i as before. Taking the arithmetic mean of these \vec{o}_i gives us

$$\bar{O} = \frac{1}{t} \sum_{i=1}^t \vec{o}_i \approx \frac{\bar{m} \cdot \vec{v} + (b - \bar{m})\bar{U}}{b} \text{ where } \bar{m} = \frac{1}{t} \sum m_i$$

¹ The attack can still proceed if few such Alice-free batches exist, so long as Alice contributes more to some batches than to others. Specifically, the approach described below (against pool mixes and mix networks) can exploit differences between low-Alice and high-Alice batches to infer background behavior.

From this, the attacker estimates Alice’s behavior as

$$\vec{v} \approx \frac{1}{\bar{m}} [b \cdot \bar{O} - (b - \bar{m})\bar{U}]$$

Attacking pool mixes and mix networks: Most designs have already abandoned fixed-batch mixes in favor of other algorithms that better hide the relation between senders and recipients. Such algorithms include timed dynamic-pool mixes, generalized mixes, and randomized versions of each [17, 30]. Rather than reordering and relaying all messages whenever a fixed number b arrive, these algorithms store received messages in a *pool*, and at fixed intervals relay a *fraction* of the pooled messages based on the pool’s current size.

When attacking such a mix, the attacker no longer knows for certain which batches contain messages from Alice. Instead, the attacker can only estimate, for each batch of output messages, the probability that the batch includes one or more of Alice’s messages.

Following Díaz and Serjantov’s approach in [17], we treat these mixing algorithms as follows: a mix relays a number of messages at the end of each round, depending on how many messages it is currently storing. All messages in the mix’s pool at the end of a round have an equal probability of being included in that round’s batch. Thus, we can characterize the mix’s pooling algorithm as a probability function $P_{\text{MIX}}(b|s)$ —the probability that the mix relays b messages when it has s messages in the pool.

We denote by $P_R^i(r)$ the probability that a message arriving in round i leaves the mix r rounds later. We assume that the attacker has a fair estimate of P_R .² Now, when Alice sends a message in round i , the attacker observes round i through some later round $i + k$, choosing k so that $\sum_{j=k+1}^{\infty} P_R^i(j)$ is negligible. The attacker then uses P_R to compute \bar{O}_w , the mean of the observations from these rounds, weighted by the expected number of messages from Alice exiting in each round:

$$\bar{O}_w = \sum_i \sum_{r=0}^k P_R^i(r) \cdot m_i \cdot o_{i+r} \approx \frac{\bar{m} \cdot \vec{v} + (\bar{n} - \bar{m}) \vec{u}}{\bar{n}}$$

To solve for Alice’s behavior \vec{v} , the attacker now needs an estimate for the background \vec{u} . The attacker gets this by averaging observations \vec{u}_i from batches with a negligible probability of including messages from Alice. Such batches, however, are not essential: If the attacker chooses a set of \vec{u}_i such that each round contains (on average) a small number $\delta_a > 0$ of messages from Alice, averaging them gives:

$$\bar{U}' \approx \frac{\delta_a}{\bar{n}} \vec{v} + \frac{1 - \delta_a}{\bar{n}} \vec{u}$$

² The attacker can estimate P_R by sending test messages through the mix, or by counting the messages entering and leaving the mix and deducing the pool size.

and the attacker can solve again for \vec{v} in the earlier equation for $\overline{O_w}$, now using

$$\vec{u} \approx \frac{1}{1 - \delta_a} [\vec{n} \cdot \overline{U'} - \delta_a \cdot \vec{v}]$$

Senders can also direct their messages through multi-hop paths in a network of mixes. While using a mix network increases the effort needed to observe all messages leaving the system, it has no additional effect on intersection attacks beyond changing the system’s delaying characteristics. Assume (for simplicity) that all mixes have the same delay distribution P_R , and that Alice chooses paths of length ℓ_0 . The chance of a message being delayed by a further d rounds is now

$$P'_R(\ell_0 + d) = \binom{\ell_0 + d - 1}{d} (1 - P_D)^{\ell_0} P_D^d$$

Danezis has independently extended statistical disclosure to pool mixes [12]; Danezis and Serjantov have analyzed it in detail [15].

Dummy traffic: Alice can also reduce the impact of traffic analysis by periodically sending messages that are dropped inside³ the network.

Although this padding can slow or stop the attacker (as discussed below in Section 4), the change in the attack is trivial: Alice’s behavior vector \vec{v} no longer adds to 1, since there is now a chance that a message from Alice will not reach any recipient. Aside from this, the attack can proceed as before, so long as Alice sends more messages (including dummies) in some rounds than in others.

Partial observation: Until now, we have required that the attacker, as a global passive adversary, observe all the messages entering and leaving the system (at least, all the messages sent by Alice, and all the messages reaching Alice’s recipients). This is not so difficult as it might seem: to be a “global” adversary against Alice, an attacker need only eavesdrop upon Alice, and upon the mixes that deliver messages to recipients. (Typically, not all mixes do so. For example, only about one third of current Mixminion servers support delivery.)

A non-global attacker’s characteristics depend on which parts of the network he can observe. If the attacker eavesdrops on a fraction of the *mixes*, he receives a sample⁴ of the messages entering or leaving the system. If such an attacker can see some messages from Alice and some messages to her recipients, he can guess Alice’s recipients, but will require more rounds of observation.

Alternatively, an attacker who eavesdrops on a fraction of the *users* receives *all* messages sent to or from those users but no messages sent to or from other

³ Alice might also send dummy traffic to ordinary recipients. This approach has its problems: how should Alice generate cover texts, or get the list of all possible recipients? In any case, it is unclear whether Alice can obscure her true recipients without sending equal volumes of mail to all of her non-recipients as well, which is impractical.

⁴ But possibly a biased sample, depending on Alice’s path selection algorithm.

users. So long as one of these users is Alice, the network (to such an attacker) is as if the messages sent by Alice to unobserved recipients were dummy messages. Now the attack converges only against observed recipients: the attacker learns which of observed recipients get messages from Alice, and which do not.

Time-variant background traffic: If Alice’s behavior changes completely and radically over time, long-term intersection attacks cannot proceed: the attacker cannot make enough observations of any version or subset of Alice’s behavior to converge on a \bar{v} for any of them.

On the other hand, if Alice’s behavior \vec{v} remains consistent while the behavior of the background traffic \vec{u} changes slowly, the attacker still has some hope. Rather than estimating a single \bar{U} from rounds to which Alice does not contribute, the attacker estimates a series of successive \bar{U}_i values based on the average behavior of the network during comparatively shorter durations of time. The attacker observes \vec{o}_i and computes the average of $\vec{o}_i - \bar{U}_i$, as before. Now,

$$\vec{v} \propto \frac{1}{t} \sum_{i=1}^t \vec{o}_i - \bar{U}_i$$

So if an attacker can get good local estimates to \vec{u} , the intersection attack proceeds as before.

Attacking recipients: Finally, we note that an attacker can find recipients as well as senders by using slightly more storage and the same computational cost.

Suppose the attacker wants to know who is sending anonymous messages to a given recipient Bob. The analysis remains the same: the attacker compares sender behavior in rounds from which Bob probably receives messages with behavior in rounds from which Bob probably doesn’t receive messages. The only complication is that the attacker cannot tell in advance when Bob will receive a message. Therefore, the attacker must remember a window of recent observations at all times, such that if Bob later receives a message, the chance is negligible that the message was sent before the first round in the window.

3.2 Strengthening the attack

Section 3.1 extended the original statistical disclosure attack to link senders and recipients in a broader range of circumstances. In this section, we discuss ways to reduce the required amount of traffic by incorporating additional information.

Partitioning messages: The attack is simplified if some output messages are *linkable*—that is, if they are likelier to originate from the same sender than are two randomly chosen messages. We consider a special case of linkability, in which we can *partition* messages into separate classes such that messages in the same class are likelier to have the same sender than messages chosen at random.

For example, in a typical pseudonym service, each sender has one or more pseudonyms and each delivered message is associated with a pseudonym. To link senders and recipients, an attacker only needs to link senders to their pseudonyms. He can do so by treating pseudonyms as virtual message destinations: instead of collecting observations \vec{o}_i of recipients who receive messages in round i , the attacker now collects observations \vec{o}_i of linkable classes (e.g. pseudonyms) that receive in round i . Since two distinct senders don't produce messages in the same linkability class, the elements of Alice's \vec{v} and the background \vec{u} are now disjoint, and thus easier for the attacker to separate.

It's also possible that the partitioning may not be complete: sometimes many senders will send messages in the same class. For example, two binary documents written with the same version of MS Word are more likely to be written by the same sender than two messages selected at random.⁵

To exploit these scenarios, the attacker chooses a set of c partitioning classes (such as languages or patterns of use), and assigns each observed output message a probability of belonging to each class. Instead of collecting observation vectors with elements corresponding to recipients, the attacker now collects observation vectors whose elements correspond to number of messages received by each (recipient, class) tuple. (If a message might belong to multiple classes, the attacker sets the corresponding element of each possible class to the probability of the message's being in that class.) The attack proceeds as before, but messages that fall in different classes no longer provide cover for one another.

Exploiting *a priori* suspicion: Finally, the attacker may have reason to believe that some messages are more likely to have been sent by the target user than others. For example, if we believe that Alice studies psychology but not astrophysics, then we will naturally suspect that a message about psychology is more likely to come from Alice than is a message about astrophysics. Similarly, if users have different views of the network, then an attacker will suspect messages exiting from mixes Alice probably doesn't know about less than other messages.

To exploit this knowledge, an attacker can (as suggested in the original statistical disclosure paper) modify the estimated probabilities in \vec{o}_i of Alice having sent each delivered message.

4 Simulation results

In Section 3.1, we repeatedly claim that each complication of the sender or the network forces the attacker to gather more information. But how much?

⁵ Encrypting all messages end-to-end would address most of these attacks, but is difficult in practice. Most recipients do not run anonymity software, and many don't support encrypted email. Thus, many messages still leave today's mix networks in plaintext. Furthermore, today's most popular encryption standards (such as PGP and SMIME) have enough variation for an attacker to tell which implementations could have generated a given message.

To find out, we ran a series of simulations of our attacks, first against the model of the original statistical disclosure attack, then against more sophisticated models. We describe our simulations and present results below.

The original statistical disclosure attack: Our simulation varied the parameters N (the number of recipients), m (the number of Alice’s recipients), and b (the batch size). The simulated “Alice” sends a single message every round to one of her recipients, chosen uniformly at random. The simulated background sends to $b - 1$ additional recipients per round, also chosen uniformly at random. We ran 100 trial attacks for each chosen $\langle N, m, b \rangle$ tuple. Each attack was set to halt when the attacker had correctly identified Alice’s recipients, or when 1,000,000 rounds had passed. (We imposed this cap to keep our simulator from getting stuck on hopeless cases.)

Figure 1 presents the results of our simulation (the low- m curves are at the bottom). As expected, the attack becomes more effective when Alice sends messages to only a few recipients (small m); when there are more recipients to whom Alice does not send (large N); or when batch sizes are small (small b).




Fig. 1. Statistical disclosure model: median rounds to guess all recipients

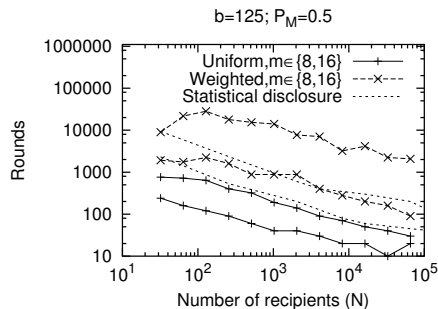


Fig. 2. Unknown background: median rounds to guess all recipients

Complex sender behavior and unknown background traffic: The next simulation examines the consequences of a more complex model for background traffic, and of several related models for Alice’s behavior.

We model the background as a graph of N communicating parties, each of whom communicates with some of the others. We build this graph according to the “scale-free” model [3, 4]. Scale-free networks share the “six degrees of separation property” (for arbitrary values of six) of small-world networks [32], but also mimic the clustering and ‘organic’ growth of real social networks, including citations in journals, co-stars in IMDB, and links in the WWW. For these trial attacks, the background messages were generated by choosing nodes from the graph with probability proportional to their connectedness. This simulates a case where users send messages with equal frequency and choose recipients uniformly from among the people they know.

We simulated trial attacks for different values of N (number of recipients) and m (number of Alice’s recipients). Instead of sending one message per batch, however, Alice now sends messages according to a geometric distribution with parameter P_M (such that Alice sends n messages with probability $P_m(n) = (1 - P_M)P_M^n$). We tried two methods for assigning Alice’s recipients: In the ‘uniform’ model, Alice’s recipients are chosen according to their connectedness (so that Alice, like everyone else, is likelier to know well-known people) but Alice still sends to her recipients with equal probability. In the ‘weighted’ model, not only are Alice’s recipients chosen according to their connectedness, but Alice also sends to them proportionally to their connectedness. We selected these models to examine the attack’s effectiveness against users who behave with the same model as other users’, and against users who mimic the background distribution.

The results are in Figure 2, along with the results for the original statistical disclosure attack as reference. As expected, the attack succeeds easily, and finishes faster against uniform senders than weighted senders for equivalent values of $\langle N, m, b \rangle$. Interestingly, the attack against uniform senders is *faster* than the original statistical disclosure attack—because the background traffic is now clustered about popular recipients, Alice’s recipients stand out more.

Attacking pool mixes and mix network: Pooling slows an attacker by increasing the number of output messages that could correspond to each input message. To simulate an attack against pool mixes and mix networks, we abstract away the actual pooling rule used by the network, and instead assume that the network has reached a steady state, so that each mix retains the messages in its pool with the same probability (P_{delay}) every round. We also assume that all senders choose paths of exactly the same length.

Unlike before, ‘rounds’ are now determined not by a batch mix receiving a fixed number b of messages, but by the passage of a fixed interval of time. Thus, the number of messages sent by the background is no longer a fixed $b - n_a$ (where n_a is the number of messages Alice sends), but now follows a normal distribution with mean $BG = 125$ and standard deviation set arbitrarily to $BG/10$.⁶

To examine the effect of pool parameters, we fixed m at 32 and N at 2^{16} , and had Alice use the ‘uniform’ model discussed above. The results of these simulations are presented in Figure 3. Lines running off the top of the graph represent cases in which fewer than half of the attacks converged upon Alice’s recipients within 1,000,000 rounds, and so no median could be found.

From these results, we see that increased variability in message delay slows the attack by increasing the number of output messages that may correspond to any input message from Alice, effectively ‘spreading’ each message across several output rounds. More interestingly, pooling is most effective at especially high or

⁶ It’s hard to determine standard deviations for actual message volumes on the deployed remailer network: automatic reliability checkers that send messages to themselves (“pingers”) contribute to a false sense of uniformity, while some users generate volume spikes by sending enormous fragmented files, or maliciously flooding discussion groups and remailer nodes. Neither group blends well with the other senders.




Fig. 3. Pool mixes and mix networks: median rounds to guess all recipients

especially low volumes of traffic from Alice: the ‘spreading’ effect here makes it especially hard for the attacker to discern rounds that contain messages from Alice when she sends few messages, or to discern rounds that don’t contain Alice’s messages when she sends many messages.

The impact of dummy traffic: Several proposals exist for using dummy messages to frustrate traffic analysis. Although several of them have been examined in the context of low-latency systems [24], little work has been done to examine their effectiveness against long-term intersection attacks.

First, we choose to restrict our examination (due to time constraints) to the effects of dummy messages in several cases of the pool-mix/mix network simulation above. Because we are interested in learning how well dummies thwart analysis, we choose cases where, in the absence of dummies, the attacker had little trouble in learning Alice’s recipients.

Our first padding strategy (“independent geometric padding”) is based on the algorithm used in current versions of Mixmaster: Alice generates a random number of dummy messages in each round according to a geometric distribution with parameter P_{junk} , independent of her number of real messages.

This strategy slows the attack, but does not *necessarily* stop it. As shown in Figure 4, independent geometric padding is most helpful when the mix network has a higher variability in message delay to ‘spread’ the padding between rounds. Otherwise, Alice must send far more padding messages to confuse the attacker.

Our second padding strategy (“imperfect threshold padding”) assumes that Alice attempts to implement the otherwise unbreakable threshold padding strategy (always send M messages total in every round, adding dummies up to M and delaying messages after M as necessary), but that she is only sometimes online




Fig. 4. Independent geometric padding: median rounds to guess all recipients




Fig. 5. Imperfect threshold padding: median rounds to guess all recipients

(with probability P_{online}), and cannot send real messages or padding when she is offline. (This will be typical for most real-world users attempting to implement threshold padding in a world of unreliable hardware and network connections.)

Figure 5 shows the result of imperfect threshold padding. As before, Alice benefits most from padding in networks with more variable delays. Interestingly, in the low delay-variability cases (short paths, low P_{delay}), padding does not thwart the attack even when Alice is online 99% of the time.

For our final dummy traffic simulation, we assume that Alice performs threshold padding consistently, but that the attacker has had a chance to acquire a view of the network’s background behavior before Alice first came online.⁷ Here, our goal was to confirm our earlier suspicion that padding helps not by disguising how many messages Alice sends, but by preventing the attacker from learning how the network acts in Alice’s absence.

Figure 6 compares results when Alice uses consistent threshold padding and the attacker knows the background to results when Alice does not pad and the background \vec{u} is unknown. Not only can an attacker who knows the background distribution identify Alice’s recipients with ease, regardless of whether she uses padding, but such an attacker is *not* delayed by increased variability in message delays.

The impact of partial observation: Finally, we examine the degree to which a non-global adversary can mount a statistical disclosure attack.

Clearly, if Alice chooses only from a fixed set of entry and exit mixes as suggested by [34], and the attacker is watching none of her chosen mixes, the attack will fail—and conversely, if the attacker is watching all of her chosen mixes, the attack proceeds as before. For our simulation, therefore, we assume that all senders (including Alice) choose all mixes as entry and exit points with equal

⁷ As usual, we assume that the background traffic patterns are unchanging. If background traffic changes significantly over time, Alice can defeat this attack by joining the network, sending nothing but padding until the network’s background characteristics have changed on their own, and only then beginning to send her messages.




Fig. 6. Full padding, background known: median rounds to guess all recipients




Fig. 7. Partial observation: median rounds to guess all recipients

probability for each message, and that the attacker is watching some fraction f of the mixes. We simulate this by revealing each message entering or leaving the network to the attacker with probability $P_{\text{observe}} = f$. The attacker sees a message when it enters *and* when it exits with probability $(P_{\text{observe}})^2$.

The results in Figure 7 show that the attacker can still implement a long-term intersection attack even when he is only observing part of the network. When most of the network is observed ($P_{\text{observe}} > 70\%$ in our results), the attack is hardly impaired at all. As more of the network is concealed ($.4 < P_{\text{observe}} < .7$) the attack becomes progressively harder. Finally, as P_{observe} approaches 0, the required number of rounds approaches infinity.

5 Conclusions

Our results demonstrate that long-term end-to-end intersection attacks can succeed even in the presence of complicating factors. Here we suggest several open questions for future work, and offer recommendations for mix network designs.

A more realistic model: Our model differs from reality in five major ways.

First, although real social networks behave more like scale-free networks than like the original disclosure attack’s model, our **models for user behavior** still have room for improvement. Real users do not send messages with a time-invariant geometric distribution: most people’s email habits are based on a 24-hour day, and a 7-day week. Early research on traffic patterns in actual mix networks [16] suggests that this variation is probably significant.

Second, **real user behavior changes over time**. Section 3.2 discusses how an attacker might handle a scenario where the background traffic changes slowly over time, and perhaps a similar approach would also help against a sender whose recipients were not constant. In the absence of a model for time-variant user behavior, however, we have not simulated attacks for these cases.

Third, it seems clear that systems with **message linkability**, such as pseudonymous services, will fall to intersection attacks far faster than anonymizing ser-

vices without linkability. How linkable are messages “in the wild,” how much does this linkability help an attacker, and how can it be mitigated?

Fourth, real attackers are not limited to passive observation. We should generalize our attacks to incorporate information gained by an **active attacker**. Past work on avoiding blending attacks [30] has concentrated on preventing an attacker from being certain of Alice’s recipients—but in fact, an active attack that only reveals slight probabilities could speed up the attacks in this paper.

Fifth, Alice has incentive to **operate a mix**, so an attacker cannot be sure if she is originating messages or just relaying them [1]. Can we treat this relayed traffic (which goes to actual recipients) as equivalent to padding (which goes to no recipients)? Can Alice employ this relayed traffic for a cheaper padding regime, without opening herself up to influence from active attacks?

Other questions for future research: Our analysis has focused on the impact of Alice’s actions on Alice alone. How do Alice’s actions (for example, choice of padding method) affect other users in the system? Are there incentive-compatible strategies that provide good security for all users?

It would be beneficial to find closed-form equations for expected number of rounds required to mount these attacks, as Danezis does for statistical disclosure.

Many of our simulations found “sweet spots” for settings such as mix pool delay, message volume, padding volume, and so on. Identifying those points of optimality in the wild would be of great practical help for users. Systems could perhaps then be designed to adaptively configure their pooling strategies to optimize their users’ anonymity.

Implications for mix network design: First, **high variability** in message delays is essential. By ‘spreading’ the effects of each incoming message over several output rounds, variability in delay increases each message’s anonymity set, and amplifies the effect of padding.

Padding seems to slow traffic analysis, especially when the padding is consistent enough to prevent the attacker from gaining a picture of the network in Alice’s absence. On the other hand, significant padding volumes may be too cumbersome for most users, and perfect consistency (sending padding from the moment a network goes online until it shuts down) is likely impractical.

Users should be educated about the effects of **message volume**: sending infrequently is relatively safe, especially if the user doesn’t repeat the same traffic pattern for long.

Mix networks should take steps to **minimize the proportion of observed messages** that a limited attacker can see entering and exiting the network. Possible approaches include encouraging users to run their own mixes; choosing messages’ entry and exit points to cross geographical and organization boundaries; and (of course) increasing the number of mixes in the network.

Much threat analysis for high-latency mix networks has aimed to provide perfect protection against an eavesdropper watching the entire network. But unless we adopt an unacceptable level of resource demands, it seems that some

highly distinguishable senders will fall quickly, and many ordinary senders will fall more slowly, to long-term intersection attacks. We must stop asking whether our anonymity designs can forever defend every conceivable sender. Instead, we should attempt to quantify the risk: *how long* our designs can defend *which senders* against an adversary who sees *how much*.

Acknowledgments

Thanks go to Gerald Britton, Geoffrey Goodell, Novalis, Pete St. Onge, Peter Palfrader, Alistair Riddoch, and Mike Taylor for letting us run our simulations on their computers; to Peter Palfrader for helping us with information on the properties of the Mixmaster network; and to George Danezis for his comments on drafts of this paper.

References

1. Alessandro Acquisti, Roger Dingledine, and Paul Syverson. On the economics of anonymity. In Rebecca N. Wright, editor, *Financial Cryptography*. Springer-Verlag, LNCS 2742, 2003.
2. Adam Back, Ulf Möller, and Anton Stiglic. Traffic analysis attacks and trade-offs in anonymity providing systems. In Ira S. Moskowitz, editor, *Proceedings of Information Hiding Workshop (IH 2001)*, pages 245–257. Springer-Verlag, LNCS 2137, 2001.
3. Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286:509–512, October 1999.
4. Albert-László Barabási, Réka Albert, and Hawoong Jeong. Mean-field theory for scale-free random networks. *Physica A*, 272:173–187, 2000.
5. Krista Bennett and Christian Grothoff. GAP – practical anonymous networking. In Roger Dingledine, editor, *Proceedings of Privacy Enhancing Technologies workshop (PET 2003)*. Springer-Verlag, LNCS 2760, March 2003.
6. Oliver Berthold, Hannes Federrath, and Stefan Köpsell. Web MIXes: A system for anonymous and unobservable Internet access. In H. Federrath, editor, *Proceedings of Designing Privacy Enhancing Technologies: Workshop on Design Issues in Anonymity and Unobservability*, pages 115–129. Springer-Verlag, LNCS 2009, July 2000.
7. Oliver Berthold and Heinrich Langos. Dummy traffic against long term intersection attacks. In Roger Dingledine and Paul Syverson, editors, *Proceedings of Privacy Enhancing Technologies workshop (PET 2002)*. Springer-Verlag, LNCS 2482, April 2002.
8. Oliver Berthold, Andreas Pfitzmann, and Ronny 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.
9. Philippe Boucher, Adam Shostack, and Ian Goldberg. Freedom systems 2.0 architecture. White paper, Zero Knowledge Systems, Inc., December 2000.
10. David Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 4(2), February 1982.

11. Wei Dai. PIPENET 1.1. Usenet post, August 1996. <<http://www.eskimo.com/~weidai/pipenet.txt>> First mentioned to the cypherpunks list, Feb. 1995.
12. George Danezis. *Better Anonymous Communications*. PhD thesis, University of Cambridge, December 2003.
13. George Danezis. Statistical disclosure attacks: Traffic confirmation in open environments. In Gritzalis, Vimercati, Samarati, and Katsikas, editors, *Proceedings of Security and Privacy in the Age of Uncertainty, (SEC2003)*, pages 421–426, Athens, May 2003. IFIP TC11, Kluwer.
14. George Danezis, Roger Dingledine, and Nick Mathewson. Mixminion: Design of a type III anonymous remailer protocol. In *2003 IEEE Symposium on Security and Privacy*, pages 2–15. IEEE CS, May 2003.
15. George Danezis and Andrei Serjantov. Statistical disclosure or intersection attacks on anonymity systems. In *Proceedings of Information Hiding Workshop (IH 2004)*. Springer-Verlag, 2004 (forthcoming).
16. Claudia Díaz, Len Sassaman, and Evelyne Deweite. Comparison between two practical mix designs. Forthcoming, 2004.
17. Claudia Díaz and Andrei Serjantov. Generalising mixes. In Roger Dingledine, editor, *Proceedings of the Privacy Enhancing Technologies workshop (PET 2003)*. Springer-Verlag, LNCS 2760, March 2003.
18. Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The Second-Generation Onion Router. In *Proceedings of the 13th USENIX Security Symposium*, August 2004.
19. Roger Dingledine and Paul Syverson. Reliable MIX Cascade Networks through Reputation. In Matt Blaze, editor, *Financial Cryptography*. Springer-Verlag, LNCS 2357, 2002.
20. Michael J. Freedman and Robert 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.
21. Ceki Gülcü and Gene Tsudik. Mixing E-mail with Babel. In *Network and Distributed Security Symposium (NDSS 96)*, pages 2–16. IEEE, February 1996.
22. Dogan Kesdogan, Dakshi Agrawal, and Stefan Penz. Limits of anonymity in open environments. In Fabien Petitcolas, editor, *Proceedings of Information Hiding Workshop (IH 2002)*. Springer-Verlag, LNCS 2578, October 2002.
23. Dogan Kesdogan, Jan Egner, and Roland 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.
24. Brian N. Levine, Michael K. Reiter, Chenxi Wang, and Matthew Wright. Timing attacks in low-latency mix-based systems. In Ari Juels, editor, *Financial Cryptography*. Springer-Verlag, LNCS (forthcoming), 2004.
25. Ulf Möller, Lance Cottrell, Peter Palfrader, and Len Sassaman. Mixmaster Protocol — Version 2. Draft, July 2003.
26. Josyula R. Rao and Pankaj Rohatgi. Can pseudonymity really guarantee privacy? In *Proceedings of the 9th USENIX Security Symposium*, pages 85–96. USENIX, August 2000.
27. J. F. Raymond. Traffic Analysis: Protocols, Attacks, Design Issues, and Open Problems. In H. Federrath, editor, *Designing Privacy Enhancing Technologies: Workshop on Design Issue in Anonymity and Unobservability*, pages 10–29. Springer-Verlag, LNCS 2009, July 2000.
28. Michael G. Reed, Paul F. Syverson, and David M. Goldschlag. Anonymous connections and onion routing. *IEEE Journal on Selected Areas in Communications*, 16(4):482–494, May 1998.

29. Michael Reiter and Aviel Rubin. Crowds: Anonymity for web transactions. *ACM Transactions on Information and System Security*, 1(1), June 1998.
30. Andrei Serjantov, Roger Dingledine, and Paul Syverson. From a trickle to a flood: Active attacks on several mix types. In Fabien Petitcolas, editor, *Proceedings of Information Hiding Workshop (IH 2002)*. Springer-Verlag, LNCS 2578, October 2002.
31. Andrei Serjantov and Peter Sewell. Passive attack analysis for connection-based anonymity systems. In *Computer Security – ESORICS 2003*. Springer-Verlag, LNCS (forthcoming), October 2003.
32. Duncan J. Watts and Steven H. Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393:440–442, June 1998.
33. Matthew Wright, Micah Adler, Brian Neil Levine, and Clay Shields. An analysis of the degradation of anonymous protocols. In *Network and Distributed Security Symposium (NDSS 02)*. IEEE, February 2002.
34. Matthew Wright, Micah Adler, Brian Neil Levine, and Clay Shields. Defending anonymous communication against passive logging attacks. In *IEEE Symposium on Security and Privacy*, pages 28–41. IEEE CS, May 2003.