

Mix-networks with Restricted Routes

George Danezis

University of Cambridge, Computer Laboratory,
William Gates Building, 15 JJ Thomson Avenue,
Cambridge CB3 0FD, United Kingdom.
`George.Danezis@cl.cam.ac.uk`

Abstract. We present a mix network topology that is based on sparse expander graphs, with each mix only communicating with a few neighbouring others. We analyse the anonymity such networks provide, and compare it with fully connected mix networks and mix cascades. We prove that such a topology is efficient since it only requires the route length of messages to be relatively small in comparison with the number of mixes to achieve maximal anonymity. Additionally mixes can resist intersection attacks while their batch size, that is directly linked to the latency of the network, remains constant. A worked example of a network is also presented to illustrate how these results can be applied to create secure mix networks in practise.

Keywords: mix networks, mix cascades, traffic analysis, anonymity.

1 Introduction

Mix networks were introduced by Chaum [4] as a technique to provide anonymous communications. The messages to be anonymized are relayed by a sequence of trusted intermediaries, called mixes, to make the task of tracing them through the network as difficult as possible. Nested layers of encryption and strict length restrictions are additionally used to make the inputs of each mix node bitwise unlinkable to its outputs.

Further research into mix networks, has been divided between real time systems, primarily for web browsing, such as onion routing [12], webmixes [1] or the freedom network [3], and non real-time systems such as babel [13], mixmaster [17] and the newer mixminion [6]. Other issues have been the trade off between real time guarantees and anonymity properties, proper metrics to quantify anonymity [22, 7], and the importance of cover traffic to maintain anonymity.

In this paper we present and discuss some proposals about the topology that mix networks might assume. These are on one hand a fully connected graph, on the other a mix cascade. We then discuss the advantages and disadvantages of a restricted network topology, that can be modeled as a network corresponding to a sparse constant degree graph, and analyze it using existing work on expander graphs. Finally we compare the anonymity and other properties provided by this new topology against the more traditional ones.

We prove that such restricted networks scale well in the number of mix nodes. The route length necessary to provide maximal anonymity grows only logarithmically in the number of nodes in the network, and the total amount of genuine traffic required to protect the network against traffic analysis and intersection attacks grows linearly with the number of mix nodes.

The paper is organized in the following fashion: in section 2 we present previous work on mix network topologies namely fully connected mix networks and cascades. We then proceed in section 3 to introduce a new network topology based on expander graphs and give a brief summary of its advantages. In section 4 we introduce a framework for analyzing mix networks and definitions of possible attacks. We then in section 5 analyze the properties of the new expander graph topology introduced, and in particular the route lengths necessary, the volumes of traffic to resist intersection attacks and its resilience to corrupt nodes. In section 6 we compare the new topology with an analysis of fully connected mix networks and cascades. We then illustrate our results, and how they can be used in practice by creating and studying an example network in section 7. Finally we present some possible avenues for future work in section 8.

2 Previous Work

Some work has already been done on the topology of mix networks. The network topology influences how clients choose the path their messages take through the mix network. The original proposal by Chaum [4] assumes a fully connected graph, while mix cascades [2] force a particular sequence of mixes to be used. The freedom network [3] only allows restricted routes to be used, for performance reasons but without any published analysis about what repercussions on anonymity such restrictions on the network might have. In [2] Berthold *et al* briefly introduces the possibility of having mix networks that are sparse, but then as we will see, focuses on describing the benefits of mix cascades.

2.1 General Mix Networks

In [4] David Chaum introduces mix networks as a collection of nodes that relay messages between each other from their original senders to their final recipients. Throughout the paper there is an implicit assumption that all nodes can communicate with all other nodes, therefore forming a fully connected network. Clients choose the path their messages take by selecting a sequence of nodes at random, from the set of all existing mix nodes. Mixmaster [17] which follows quite closely Chaum's original proposals, also allows clients to choose any route through the network, using reliability statistics [16] to select nodes. Other proposals concerning route selection use reputation metrics [8] to quantify how reliable mixes in a network are.

2.2 Mix Cascades

While the fully connected nature of the mix networks seemed to improve the anonymity provided, Berthold *et al* [2] found that they can be vulnerable against very powerful adversaries, such as those who control all the mix nodes except one. In particular if only one of the mixes in the network is honest the anonymity of the messages going through it will most likely be compromised. Attackers can perform intersection attacks, while the probability that all nodes in a short path are compromised is very high. Additionally if two or more messages follow the same route, attacks are trivial to perform.

As a solution a *cascade* of mixes is proposed. Users of the network are not free to choose which route to take, but are all forced to route their messages through a predefined sequence of mixes. Further work has been done to improve the reliability of networks of cascades against corrupt nodes using reputation [9]. The obvious drawbacks of cascades are the small anonymity sets they provide in the general case, and the fact that they do not scale well to handle heavy load. Cascades are also vulnerable to denial of service attacks, since disabling one node in the cascade will stop the functioning of the whole system. Some solutions are proposed to solve the problem that active attacks could pose, but require user authentication to work properly [2].

3 Mix Networks Based on Expander Graphs

We propose a mix network with a network topology based upon sparse, constant degree graphs, where users may only choose routes following this topology. Furthermore each link out of a node should be chosen according to a predefined probability distribution. Therefore selecting a path in the network can be approximated as a random walk on the corresponding weighted graph. We will show that this network provides some of the advantages of cascades, while being more scalable. We will provide theoretical anonymity bounds using the proposed metric for anonymity based on entropy in [22], and define under which conditions the network provides anonymity close to the theoretical limit. Minimum traffic bounds to prevent the network being vulnerable to traffic analysis and intersection attacks are also calculated.

The topology that we propose for the mix network is based on expander graphs. Expanders are a special family of graphs with the following properties: a D -regular graph is a $(\mathcal{K}, \mathcal{A})$ -expander if for every subset \mathcal{S} of vertexes of \mathcal{G} , if $|\mathcal{S}| \leq \mathcal{K}$, then $|N(\mathcal{S})| > \mathcal{A}|\mathcal{S}|$ where $|\mathcal{S}|$ is the number of vertexes in \mathcal{S} and $|N(\mathcal{S})|$ is the number of nodes sharing an edge (neighbouring) with a vertex in \mathcal{S} . In plain language it means that any random subset of nodes will have “many” different neighbouring nodes. In practise expanders have a relatively small and constant degree D in relation to the number of edges of the graph, and a large expansion factor A , that is proportional to the number of “neighbours”. A good introduction to expander graphs and their applications can be found in [15].

A relevant result is that most bipartite graphs with degree at least three provide good expansion properties, which means that a topology based on a

random bipartite graph with each mix node having three fixed neighbors would be an expander with high probability [20]. Therefore such networks can be constructed by brute force, or by using the surveyed or proposed methods in [21]. The families of expanders with explicit constructions presented have a constant, but large, degree but also an arbitrary large number of nodes, which makes them practical for large networks.

The first question that comes to mind is quantifying the anonymity that such networks provide in comparison to fully connected networks. In a fully connected network a message coming out of the network has a probability of originating initially from a particular node proportional to the input load of the node. As we will see a random walk through the expander graph will converge toward the same probability after a number of steps proportional to $\mathcal{O}(\log N)$ where N is the number of nodes in the network [10]. This represents the *a-priori* knowledge of an adversary that only knows the topology of the graph, but no details about the actual traffic going through it.

Intersection attacks presented in [2] rely on the fact that messages using the same sequence of nodes will only occur in a fully connected network with a very small probability. Since a mix network based on a small constant degree graph only provides a limited choice of routes for messages to take, a node can wait so that enough messages are accumulated before sending them, to make sure that all its neighbors always receive messages. Because there is only a linear number of routes to fill with dummy traffic, only order $\mathcal{O}(DN)$ messages are required where N is the number of nodes and D the degree of the graph. This strategy is more efficient than filling all the $\mathcal{O}(N^2)$ links in a fully connected graph, since adding more nodes only requires the total traffic to increase linearly in order to maintain the network's resistance to traffic analysis.

Before we move on to prove the properties described above, as we will do in section 5, we will first introduce a way of quantifying anonymity and some definitions about the attacks that can be performed on mix networks.

4 A Framework for Analyzing Mix Networks

In order to compare fully connected mix networks, mix cascades and restricted routing there is a need to have a way of quantifying not only their security but also their efficiency. Efficiency can be measured following the usual paradigms of communication networks, namely the latency of messages and the load on mix servers. On the other hand security, and in particular anonymity, does not have a well established way of being measured.

In order to quantify the anonymity provided by a network of mixes we will use the metric proposed by Serjantov and Danezis [22]. We will consider the sender anonymity set of a message as the entropy of the probability distribution describing the likelihood particular participants were senders of the message. As expected the anonymity of messages increases as the number of potential senders increases. Given a certain number of participants, the anonymity of a message is also maximized when all participants have an equal probability of having been

the sender of a message. Recipient anonymity can be quantified in an equivalent fashion, so we will only present the analysis of sender anonymous properties.

A message m_e exits the mix network at time t_e from node n_e . The network is made out of N mix nodes, n_1 to n_N . Messages m_{ij} are injected at node n_i at time t_j . The task of an attacker is to link the message m_e with a message m_{ij} .

We consider the probability distribution

$$\begin{aligned} p_{ij} &= \Pr[m_e \text{ is } m_{ij}] \\ &= \Pr[m_e \text{ is } m_{ij} | m_e \text{ inserted at } n_i] \times \Pr[m_e \text{ inserted at } n_i] \end{aligned} \quad (1)$$

that describes how likely the input messages in the network are to have been message m_e . We can express this probability as the probability that a node n_i was used to inject a message, multiplied by the probability a particular message m_{ij} injected at this node is m_e . The entropy of the probability distribution p_i is the effective sender anonymity set of the message. Because of the strong additive property of entropy we can calculate this entropy as:

$$\begin{aligned} \mathcal{A} &= H(p_{ij}) \\ &= H(\Pr[m_e \text{ inserted at } n_i]) \\ &+ \underbrace{\sum_{x \in 1 \dots N} \Pr[m_e \text{ inserted at } n_x]}_{\text{traffic analysis attacks}} \times \underbrace{H(\Pr[m_e \text{ is } m_{ij} | m_e \text{ inserted at } n_x])}_{\text{traffic confirmation attacks}} \end{aligned} \quad (2)$$

An attacker might attempt to reduce the anonymity by subjecting the network to traffic analysis in order to reduce the uncertainty of $\Pr[m_e \text{ inserted at } n_i]$. We shall therefore name $\mathcal{A}_{\text{network}} = H(\Pr[m_e \text{ inserted at } n_i])$, the anonymity provided by the network. This quantifies how effectively the traffic injected to or ejected from particular nodes is mixed with traffic from other nodes. Given a particular threat model if no technique is available for the attacker to reduce the uncertainty of $\mathcal{A}_{\text{network}}$ beyond her a-priori knowledge, we can say that *the network is resistant to traffic analysis in respect to that particular threat model*.

The attacker can also try to reduce the anonymity set of the message by reducing the uncertainty of the probability distribution describing the traffic introduced at the mix nodes, $\Pr[m_e \text{ is } m_{ij} | m_e \text{ inserted at } n_x]$. The attacker can do this by using additional information about m_e and m_{ji} , like the times t_e the message comes out of the network or t_j the time it was injected in the network. She can also do this by flooding nodes, or stopping messages arriving to the initial nodes. It is worth noting that a network might protect users from traffic analysis, but still provide inadequate anonymity because of such side information leaked by messages as they enter and exit the mix network. Side information is not limited to time, but can also be associated with the protocol or mechanism used, client type, unique identifiers or route length indications observed at the edges of the mix network. Attacks that try to link messages using such side information leaked at the edges of the network, instead of tracing the message through the network, are called *traffic confirmation attacks* [24].

In analyzing and comparing the anonymity provided by networks with restricted routes we will limit ourselves into considering the traffic analysis resistance since it depends heavily on the topology while traffic confirmation attacks depend on the particular mix batching and flushing strategy individual nodes use. Having defined a way of quantifying the anonymity provided by the network, we will study in the next section, the route length necessary to archive maximal anonymity in expander graph based mix networks and the volumes of traffic necessary to avoid traffic analysis attacks.

5 Anonymity Analysis of Restricted Network Topologies

In a fully connected mix network it is intuitive that a message that comes out of a mix node, after a number of hops, could have originated from any node in the network with a probability proportional to its input load. Since users chose their initial nodes at random, or taking into account in the case of mixmaster reliability statistics [16], we can say that the probability the messages originated from an initial node is equal to the probability a client has to choose this node as an entry point to the network. The same probability distribution is often used to determine the intermediate and final node of the anonymous path. This observation allows us to compute $\mathcal{A}_{\text{network}}$ for fully connected networks, using the probability distribution describing the selection of the entry node.

For a graph that is not fully connected we need to calculate what the probability is that a message that is present in a node after a number of mixing steps has originated from a particular initial node. This requires us to trace the message backwards in the network. If the graph is not directed the likelihood a message was injected at a particular node is equal to the probability a random walk starting at the final node finishes on a particular node after a certain number of hops.

Therefore, we consider the network as a graph and the act of selecting a path through it as a random walk, and we model the route selection procedure and actual communication as a Markov process. In practice some anonymous route selection algorithms exclude nodes from being present on the path of a message twice, which violates the memoryless property of the Markov process. Despite this if we assume that a Markov process is still a good approximation to the route selection process, after an infinite number of steps the probability a message is present on a particular node should be equal to the stationary probability distribution π of the process. Therefore the maximum anonymity provided by the network will be equal to its entropy, $\mathcal{A}_{\text{network}} = H(\pi)$

For reasons of efficiency we need to find how quickly the probability distribution $q^{(t)}$ describing where a message is after a number of random steps t , converges to the stationary probability π of the Markov process. A smaller t would allow us to minimize the number of hops messages need to be relayed for, therefore reducing the latency and increasing the reliability of the network. Motwani and Raghavan [18] provide a theoretical bound on how quickly a random walk on a graph converges to the stationary probability. If π_i is the stationary

distribution of a random walk on a graph G and $q^{(t)}$ the probability distribution after t number of steps starting from any initial distribution $q^{(0)}$. We define $\Delta(t)$ as the relative point wise distance as follows:

$$\Delta(t) = \max_i \frac{|q_i^{(t)} - \pi_i|}{\pi_i} \quad (3)$$

It can be shown [18] that this distance after a number of random steps t is bound by the number of nodes in the network N and the second eigenvalue λ_2 of the transition matrix corresponding to the graph of the network:

$$\Delta(t) \leq \frac{\sqrt{N}(\lambda_2)^t}{\min_i \pi_i} \quad (4)$$

Therefore the distance decreases exponentially as the number of step t , for which the message travels through the networks, increases linearly.

It is clear that the quick rate of convergence of the probability distribution is dependent on the second eigenvalue being small. An extensive body of research has concentrated on linking the value of the second eigenvalue to expansion properties of graphs, to show that good expanders exhibit small second eigenvalues (see [18] for details). There is a fundamental limit of how quickly a network can mix that depends on the degree D of the graph:

$$\lambda_2 \geq \frac{2\sqrt{D-1}}{D} \quad (5)$$

The results above assure us that a mix network with a topology corresponding to a good expander graph would mix well, in a number of steps logarithmic in its size, $\mathcal{O}(\log N)$. This means that in this small number of steps a message that enters the network will leave the network at a node selected with probability approaching the probability after an infinite number of steps, namely the stationary probability distribution π . Furthermore its degree could be bound in order to allow for links to be padded with cover traffic, to protect against intersection attacks or traffic analysis attacks, as we will study in the next section.

In fact the methods described above can be used to calculate the theoretical probability that a messages that comes out at a mode n_e of the network has been injected at another node n_i . In theory the a-priori knowledge of the attacker, concerning where a message was injected, corresponds to the probability distributions after the random walk on the graph representing the network. It also depends on the way that initial nodes are being chosen by clients, using reliability statistics, or other information. As the number of intermediaries grows this distribution converges towards being the same for all initial choices of entry nodes. Therefore as the number of hops grow a network based on expander graphs offers *uniform anonymity*, which means that the anonymity provided by the network is the same regardless of the node used to inject a message. In the next section we will study how much traffic is needed to make the network resistant to traffic analysis, in other words an actual observation of its running will not give the attacker any additional information beyond the theoretical calculations presented above.

A number of ways can be employed in order to find an expander graph that would represent a good anonymous communication network. If the number of nodes is small it can be done by brute force, until a graph is found with a second eigenvalue that approaches the limit described above. Explicit constructions employing Ramanujan graphs [11] or zig zag products [21] can also be employed to construct the network. Standard graphs such as multi dimensional hyper-cubes also exhibit properties that could be suitable.

5.1 Protection Against Intersection Attacks

An advantage of mix cascades, as argued in [2], is that they are not susceptible to intersection attacks. Such attacks use the fact that many messages are sent using the same path, to perform traffic analysis attacks and follow the messages through the network. The authors note that, if every time a message is sent by the user under surveillance, the set of possible destinations of every mix is intersected with the set of possible destinations of previous messages, then the actual path of the message will become apparent. This is due to the very small probability the same, even partial, route is used by different messages. Since in mix cascades all messages use the same sequence of intermediary nodes, such an attack does not yield any results. Of course traffic confirmation is always possible, by observing all the edges of the network, and find correlations between initial senders and final recipients. Such attacks will always be possible if the network does not provide full unobservability [19], or other specific countermeasures.

In a mix network with restricted routes and small degree, such as one based on expander graphs described in the previous section, the potential for intersection attacks described above, can be greatly reduced. This is done by making sure that all the links from a node to its neighbors are used in a flushing cycle. This is possible in practice since the number of these links in the network is relatively small, and does not grow proportionally to $\mathcal{O}(N^2)$ as for fully connected networks. Making sure that all links are used is sufficient to foil the simplest intersection attacks, that use the intersection of sets of potential senders to trace messages [14]. Traffic analysis is still possible if the probability a messages has used a link is skewed. Therefore we need enough genuine traffic to be mixed together for the observable load on the network links to be proportional to the theoretical probability distribution described by the transition matrix.

Using a threshold mix as an example we will calculate how much traffic is needed for no link of a node to be left empty. We assume that clients select the next hop from a particular node using a probability distribution p_n , where n is the number of N_i neighboring nodes. Then the probability that the link to a node is left empty in a batch of b messages is:

$$\Pr[\exists i.N_i \text{ empty}] < \Pr[N_1 \text{ empty}] + \dots + \Pr[N_n \text{ empty}] \quad (6)$$

$$\Pr[\exists i.N_i \text{ empty}] < \sum_{\forall N_i} (1 - p_i)^b \quad (7)$$

As the size of the batch of messages to be processed grows, the probability that a link is empty decreases exponentially, making simple intersection attacks infeasible. It is important to note that the same effect can be achieved by adding dummy traffic on the links that are not used. Again the amount of dummy traffic in the network will only grow linearly with the number of nodes in the network.

In order to avoid the attacker gaining any more information than the theoretical anonymity, which is the entropy of the stationary probability distribution on the nodes E_{π_i} , the actual flows of messages on the links should be as close as possible to the matrix describing the network topology. As described above each node receives a number of messages b , some of which will be output on a particular link i according to a binomial distribution, with probability p_i . We can require the number of messages that are actually transmitted not to diverge on a particular round or time period by more than a small percentage f from the average mean. As the number of messages b received by the mix increases the probability that X the number of messages transmitted on the link i , is close to the expected mean bp_i increases:

$$\Pr[(1-f)bp_i \leq X \leq (1+f)bp_i] = 1 - 2\Phi\left(-fk^{\frac{1}{2}}\sqrt{\frac{p_i}{1-p_i}}\right) \quad (8)$$

Where Φ is the cumulative probability distribution of a normal random variable, with mean zero and variance one. We can require f to be arbitrary small, like .05, by mixing more messages together in a threshold mix [23]. Expressing the above formula to calculate f makes it clear that the deviation from the mean expected traffic gets smaller proportionally to the inverse square root of the number of messages processed. This result can then be used in conjunction with p_{min} , the probability associated with the link that is least likely to be used in the network or mix, to derive how much genuine traffic would be necessary in a node to protect against traffic analysis.

Another way of calculating the appropriate threshold value for a threshold mix would be to calculate the number of rounds necessary to perform the intersection attack. The techniques used to do this are related to the statistical disclosures attacks described in [5, 14]. The attacker performs a hypothesis test on each of the links, with H_0 representing the hypothesis that the stream of messages under surveillance are output on the link under observation, and H_1 representing the hypothesis the messages are output on another link. In case H_0 is true the volume of messages on the observed link follows a probability distribution $Y_0 = k + X_{b-1}$ otherwise it follows a probability distribution $Y_1 = X_{b-1}$, where b is the threshold of the mix, k the number of mixing rounds, and p_i the probability the link is used by messages. X_{b-1} is the random variable following the binomial distribution with probability p_i after $b-1$ trials. The mean and standard deviation of these distributions are:

$$\mu_{Y_0} = k + k(b-1)p_i \quad \sigma_{Y_0}^2 = k(b-1)p_i(1-p_i) \quad (9)$$

$$\mu_{Y_1} = k(b-1)p_i \quad \sigma_{Y_1}^2 = k(b-1)p_i(1-p_i) \quad (10)$$

In order to be able to accept or reject hypothesis H_0 we will require the observed volume of traffic to be within a distance of a few standard deviations σ_{Y_0} from the mean μ_{Y_0} of while also at a minimum distance of a few standard deviations σ_{Y_1} from μ_{Y_1} to avoid false positives. The number l of standard deviation depends on the degree of confidence required. The minimum number of mixing rounds k that need to be observed by an attacker to confirm or reject the hypothesis can therefore be calculated by:

$$\mu_{Y_0} - l\sigma_{Y_0} > \mu_{Y_1} + l\sigma_{Y_1} \quad (11)$$

$$k > 4l^2 \frac{p_i}{1 - p_i} (b - 1) \quad (12)$$

For values of $l = 1$ we get a confidence of 68%, for $l = 2$, 95% and for $l = 3$, 99%. The above formula is true both for general mix networks and for mix networks with restricted routes. We can require the value of rounds k to be greater than one $k > 1$, with $l = 0.6745$ for the attacker to have only a confidence of 50% in order to frustrate traffic analysis of messages that are not part of a stream that follows the same route.

5.2 Corrupt Nodes

An important factor that has to be taken into account when judging an anonymous network, is how robust it is to corrupt nodes. In particular one has to assess the likelihood that all the nodes that have been selected to be on the path of a message are corrupt nodes. For the topology presented this amounts to determining the probability $p_{l/c}$ that l nodes selected by a random walk on the expander graph, might include $c \leq l$ corrupt nodes. Gillman provides an upper bound for this probability [10], that is dependent on the expansion properties of the graph, and the ‘‘probability mass’’ of the corrupt nodes.

If the matrix representing the graph of the mix network has a second eigenvalue λ_2 then define $\epsilon = 1 - \lambda_2$. Assume that the set C of nodes is corrupt. Then define π_c as the probability mass represented by this corrupt set, $\pi_c = \sum_{i \in C} \pi_i$ where π is the stationary probability distribution of the random walk on the graph. After a number of steps l the probability that a walk has only been performed on corrupt nodes is:

$$\Pr[t_c = l] \leq \left(1 + \frac{(1 - \pi_c)\epsilon}{10}\right) e^{-l \frac{(1 - \pi_c)^2 \epsilon}{20}} \quad (13)$$

The probability that a path is totally controlled by corrupt nodes therefore depends on the amount of traffic processed by the corrupt nodes, and the mixing properties of the graph, but decreases exponentially as the route length increases. Assuming a particular threat model the route length can therefore be increased until that threat is very improbable. In practice the constant factors of the bound are too large to lead to values of the route length that are practical in real systems. Therefore despite the initially encouraging results, for even modest π_c

other methods might have to be used to determine and minimize the probability that the full route is composed of corrupt nodes.

6 Comparing Sparse Networks with other Topologies

We have studied in the previous sections some properties of sparse mix networks namely their necessary route length, batch size or volume of traffic necessary to provide nearly maximal anonymity. We shall next compare these properties with previously introduced topologies.

Sparse networks based on expander graphs scale well by providing maximal network anonymity for a route length l proportional to $\mathcal{O}(\log N)$. Furthermore they can be made resistant to traffic analysis and intersection attacks using a constant volume of traffic per node, depending on the degree D of the network. By (12) we observe that if the route selection algorithm is uniform, then the batch size b of nodes can be $b < \frac{1}{4l^2}k(D-1) + 1$ which is independent of the number of nodes in the network.

6.1 Mix Cascades

Given our definitions it is clear that a mix cascade is resistant to traffic analysis, since observing the network traffic does not provide an attacker with more information than she originally had about the correspondence of input to output nodes. This is the case because there is no uncertainty about the node where all messages were inserted, since there is only one. The fact that $\mathcal{A}_{\text{network}} = 0$ does not mean that the network does not provide any anonymity to messages, but simply that all the anonymity provided to the messages originates conceptually from the single $H(\Pr[m_e \text{ is } m_{ij} | m_e \text{ inserted at } n_x])$ component of (2).

This absolute protection against traffic analysis comes at a very high cost. The anonymity provided is reduced to the volume of messages that can be processed by the node with least throughput in the cascade. The latency of the messages is also large, since each message has to be processed by all nodes in the cascade.

Despite the inefficiencies presented above mix cascades are a valuable design. They are resistant to very powerful adversaries, that control all nodes but one. They also highlight the advantages of implementing topologies that can be analyzed, in order to understand their anonymity properties.

6.2 Mix Networks

General mix networks are distinct from sparse, constant degree, mix networks because senders of anonymous messages are allowed to follow arbitrary routes through them. This sometime is misinterpreted as meaning that matrix corresponding to the mix network is fully connected. Indeed an attacker that has no additional knowledge of the network, beyond the way routes are selected, has no other way of attributing probabilities linking output messages to input nodes,

other than by using a random walk on this fully connected graph, for a number of steps corresponding to the route length.

An attacker that can observe the traffic in the network, on the other hand, can get much better results. If we assume that the number of nodes is larger than the threshold of the mixes, some links remain unused in each mix round. Furthermore even if the threshold is comparable to the number of mixes, the volume of messages sent will give the attacker a different probability distribution from the theoretical one described by the route selection distribution. Therefore an attacker can use the additional information, extracted from these observation to trace messages more effectively.

We will denote the graph used for the route selection through the network as G . This graph has N nodes, the number of mixes, that are all connected with each other by weighted edges. The weights correspond to the probability that a node is selected as the next mix in the path, and can be uniform if the selection is performed at random, or it can be based on reliability statistics or reputation metrics. Given a column vector v describing where a message was injected, the probability P a messages comes out of the network at a particular nodes after l steps, can be calculated to be $P_l = G^l v$ This is the *a-priori* information that an attacker has about the correspondence between input and output nodes, even before any traffic analysis has been performed.

As the attacker observes the network, for round i it can deduce a matrix G_i with the mixes as the vertexes, and the traffic load that was observed between them during round i as the weights on the edges. It is worth observing that G_i is closely related to G in the sense that the selection of routes for any round is performed using G , but is sparse if the threshold of the mixes is lower than the number of nodes. In fact the weights on the edges follow the same probability distribution as for G , but are going to be different subject the the variance of the multinomial distribution, and the threshold of the mixes. An adversary that observes the actual traffic patterns in the networks will therefore be able to have more accurate information about where the messages injected are going, by calculating the probability distribution $P'_l = G_l \dots G_2 G_1 v$.

The relation of G_i the graph of the traffic observed at round i with the graph G used to route messages, is crucial in understanding the anonymity that generic mix networks provide. The smaller the difference between G_i and G the more resistant the network will be to traffic analysis. In order for G_i to be close to G there needs to be enough traffic to make the mean load on all the links proportional to the probabilities of the route selection, as described for sparse topologies in section 4.2. In general one would expect $\lim_{l \rightarrow \infty} H(P'_l) = \lim_{l \rightarrow \infty} H(P_l)$, but also $\forall l, H(P_l) \leq H(P'_l)$, where $H(\cdot)$ denotes the entropy of a distribution.

For G_i to be therefore a good approximation of G it is necessary each round to fill all links with traffic volumes proportional to the values on the edges of G . This requires the volumes of traffic handled by the network to be proportional to $\mathcal{O}(N^2)$ as the number of nodes N in the network grows. The batch size that needs to be handled by each node therefore grows proportionally in the size

of the network $b < \frac{k}{4t^2}(N - 1) + 1$, as described by (12). The increased batch sizes also has a repercussion on the latency of messages that travel through the network since mixes will wait for more messages before they operate.

Mix networks that do not use a deterministic threshold mixing strategy, where the first batch of messages to go in are also the first batch of messages to go out, can also be analyzed in a similar fashion by redefining G_i . It would then need to represent the probability distributions leading to the effective anonymity sets of messages instead of the volumes of traffic in the network.

7 An Example Network: Putting It All Together

In order to illustrate how all the results presented on mix networks based on expander graphs fit together, we will present an example network and analyze it. We will proceed to calculate the route length necessary for it to provide uniform anonymity, the amount of real traffic that should be present in each node for it to be resistant to traffic analysis and intersection attacks.

7.1 Selecting a Good Topology

We aim to create a network with $N = 400$ mix nodes, each with $D = 40$ neighbors. The neighbor of a mix both sends and receives messages from the mix and therefore we can represent this network as an undirected graph. Furthermore we will assume that senders will choose their path across the network using a random walk on the graph, with equal weights on all the edges. Therefore the probability that a message follows a particular link, given that it is already on a node is equal to $p_n = p_{min} = \frac{1}{40}$.

Using a brute force algorithm we create a number of networks and compute their second eigenvalue until a network with good expansion properties is found. After testing less than ten candidates we find a graph with a second eigenvalue $\lambda_2 = 0.3171$, which is close to the theoretical limit given by equation (5) of $\lambda_2 > 0.3122$

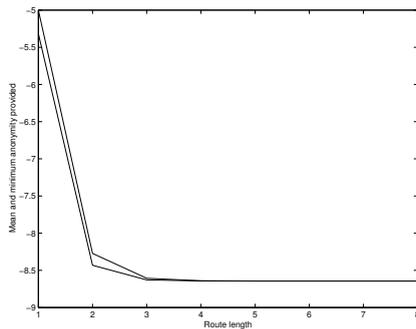
As expected such a graph has $N_l = 16 \cdot 10^3$ links instead of $N^2 = 16 \cdot 10^4$ that a fully connected graph would have. Therefore it is sparse in the sense that only one in ten links are used.

7.2 Mixing Speed

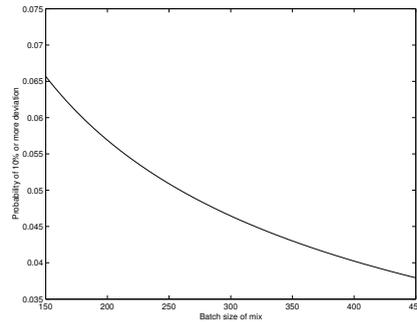
Using the theoretical formula (4) we know that the network will provide nearly uniform anonymity after a number of mixing steps proportional to $\log N$. From the graph we know that the $\min_i \pi_i = \frac{1}{400}$ since the stationary distribution is uniform, and therefore the theoretical anonymity, according to [22], should be equal to $\mathcal{A} = -\log_2 N = -8.6438$.

In theory the relative point wise distance $\Delta(t)$ between the observed $q^{(t)}$ distribution after t steps and the stationary distribution π_i should converge following $\Delta(t) \leq n\sqrt{\pi}\lambda_2^t$. This allows us to calculate that the safe route length is

around six. In practice much tighter bounds can be computed by directly calculating using G^t the distributions $q^{(t)}$ from the available graph G . It is therefore observed that after four steps of the random walk the entropy of $q^{(t)}$ is equal to the theoretical entropy calculated above. Figure a illustrates this by showing how the mean entropy provided to messages entering on any node compares with the minimum entropy that is offered by the network. Their convergence indicates that after four steps the network provides uniform and also maximum anonymity.



(a) Mean and lowest entropy after a random walk



(b) Probability the distribution deviates more than 10%

7.3 Resisting Intersection and Traffic Analysis Attacks

In order to avoid the simplest forms of intersection attacks all the networks links need to be used for every round. The probability a network link is not used is described by equation (7). For this particular network all $p_i = \frac{1}{40}$ where p_i is the probability a link is followed. The probability that any link is left empty for threshold mix with batch size $b = 300$ is therefore, $\Pr[\exists i.N_i \text{ empty}] < 2\%$. Therefore for batches larger than 300 messages the probability a link is left empty is very small.

In order to protect against more sophisticated traffic analysis attacks taking into account statistical differences in the observed distributions from the graph G , we need to calculate the probability this deviation is large (for example larger than 10% as shown in figure b). In practice with a batch size of $b = 300$ as specified above, the attacker would need to observe $k > 4\frac{1}{40-1}(300-1) = 30$ messages in a stream in order to have a confidence of 68% that a particular link was used.

8 Future Work

An area that has not been investigated in depth has been the creation of the graph topology. Since the routes are restricted there is a need to advertise the allowed routes to clients, but also for the mixes to collaboratively decide upon a topology. If a brute force algorithm is used some randomness about the initial seed could be contributed by each mix so that the result is assured not to be biased. If an explicit construction is employed a similar procedure should be used to make sure that the parameters of the network are not set by a minority of potentially corrupt nodes, as discussed in [9].

Besides good mixing properties, expander graphs provide some useful robustness properties against deletion of nodes. A major concern when building a network is the number of nodes an adversary needs to disable necessary to partition the network, or reduce the anonymity it provides. Assessing the impact of removing nodes on the speed of mixing would be a good start for assessing this risk.

Finally strategies for countering active flooding or delaying attacks are necessary. Since the number of neighboring nodes is small, they are more likely than in a fully connected network to all be corrupt and mount active attacks against the surrounded honest nodes.

9 Conclusions

The case has been argued in this paper that sparse networks provide desirable properties against traffic analysis attacks and scale better than fully connected networks or cascades. Some calculations presented, such as the probability a route is fully captured by adversaries, are theoretically appealing but do not provide tight enough bounds to be used in practice, while others are directly applicable for analyzing networks. Maybe tighter bounds could be found by restricting further the topology of the network.

The analysis of intersection attacks provides practical bounds to calculate the amount of traffic necessary to defend mixes, but is only applicable to threshold mixes. It is important to generalize it in the future to other mix batching and flushing strategies, such as pool mixes presented in [22, 23]. It also offers a good foundation to decide how many messages in the stream are to travel in the same path. The required volume of data to make the network resistant to some rounds of traffic analysis can also be used as a guide to decide how much cover traffic is to be introduced.

The main contribution of this paper is that it highlights that a middle ground exists between free route mix networks, and extremely restrictive mix cascades. By designing the network carefully, and choosing appropriate topologies, some properties of both can be achieved, such as improved resistance to intersection attacks, along with shorter routes and better scalability.

Acknowledgments: The author would like to thank Andrei Serjantov for the invaluable provocative discussions that greatly improved this work, and Richard Clayton for clearly formulating the problem surrounding intersection attacks.

References

1. Oliver Berthold, Hannes Federrath, and Stefan Köpsell. Web MIXes: A system for anonymous and unobservable Internet access. In *Designing Privacy Enhancing Technologies, LNCS Vol. 2009*, pages 115–129. Springer-Verlag, 2000.
2. Oliver Berthold, Andreas Pfitzmann, and Ronny Standtke. The disadvantages of free MIX routes and how to overcome them. In *Designing Privacy Enhancing Technologies, LNCS Vol. 2009*, pages 30–45. Springer-Verlag, 2000.
3. P. Boucher, I. Goldberg, and A. Shostack. Freedom system 2.0 architecture. <http://www.freedom.net/info/whitepapers/>, December 2000. Zero-Knowledge Systems, Inc.
4. David Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 4(2), February 1982.
5. George Danezis. Statistical disclosure attacks. <http://www.cl.cam.ac.uk/~gd216/StatDisclosure.pdf>, November 2002.
6. George Danezis, Roger Dingledine, and Nick Mathewson. Mixminion: Design of a Type III Anonymous Remailer Protocol. In *IEEE Security and Privacy Symposium*, 2003.
7. Claudia Diaz, Stefaan Seys, Joris Claessens, and Bart Preneel. Towards measuring anonymity. In *Privacy Enhancing Technologies Workshop 2002*, April 2002.
8. Roger Dingledine, Michael J. Freedman, David Hopwood, and David Molnar. A reputation system to increase MIX-net reliability. *Lecture Notes in Computer Science*, 2137:126–141, 2001.
9. Roger Dingledine and Paul Syverson. Reliable MIX Cascade Networks through Reputation. Proceedings of Financial Cryptography 2002.
10. David Gillman. A chernoff bound for random walks on expander graphs. In *IEEE Symposium on Foundations of Computer Science*, pages 680–691, 1993.
11. Yair Glasner. Ramanujan graphs with small girth.
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. Gulcu and G. Tsudik. Mixing E-mail with Babel. In *Network and Distributed Security Symposium - NDSS '96*. IEEE, 1996.
14. Dogan Kesdogan, Dakshi Agrawal, and Stefan Penz. Limits of anonymity in open environments. In *Information Hiding, 5th International Workshop*, Noordwijkerhout, The Netherlands, October 2002. Springer Verlag.
15. Nati Linial and Avi Wigderson. Expander graphs and their applications. Collection of Lecture Notes http://www.math.ias.edu/~avi/TALKS/expander_course.pdf, January 2003.
16. Christian Mock. Mixmaster stats (Austria). <http://www.tahina.priv.at/~cm/stats/mlist2.html>.
17. Ulf Möller and Lance Cottrell. Mixmaster Protocol — Version 2. Unfinished draft, January 2000. <http://www.eskimo.com/~rowdenw/crypt/Mix/draft-moeller-mixmaster2-protocol-00.txt>.

18. Rajeev Motwani and Prabhakar Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.
19. Andreas Pfitzmann and Marit Kohntopp. Anonymity, unobservability and pseudonymity — a proposal for terminology. In *Designing Privacy Enhancing Technologies: Proceedings of the International Workshop on the Design Issues in Anonymity and Observability*, number 2009 in LNCS, pages 1–9. Springer-Verlag, July 2000.
20. M. S. Pinsker. On the complexity of a concentrator. In *Proceedings of the 7th International Teletraffic Conference*, pages 318/1–318/4, Stockholm, 1973.
21. Omer Reingold, Salil P. Vadhan, and Avi Wigderson. Entropy waves, the zig-zag graph product, and new constant-degree expanders and extractors. In *IEEE Symposium on Foundations of Computer Science*, pages 3–13, 2000.
22. Andrei Serjantov and George Danezis. Towards an information theoretic metric for anonymity. In *Privacy Enhancing Technologies Workshop 2002*, San Francisco, CA, May 2002.
23. Andrei Serjantov, Roger Dingledine, and Paul Syverson. From a trickle to a flood: Active attacks on several mix types. In Fabien A.P. Petitcolas, editor, *Information Hiding Workshop*, number 2578 in LNCS, pages 36–52. Springer-Verlag, 2002.
24. P. Syverson, M. Reed, and D. Goldschlag. Private web browsing. *Journal of Computer Security*, 5(3):237–248, 1997.