Towards an Information Theoretic Metric for Anonymity

Andrei Serjantov and George Danezis

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

Abstract. In this paper we look closely at the popular metric of anonymity, the *anonymity set*, and point out a number of problems associated with it. We then propose an alternative information theoretic measure of anonymity which takes into account the probabilities of users sending and receiving the messages and show how to calculate it for a message in a standard mix-based anonymity system. We also use our metric to compare a pool mix to a traditional threshold mix, which was impossible using anonymity sets. We also show how the maximum route length restriction which exists in some fielded anonymity systems can lead to the attacker performing more powerful traffic analysis. Finally, we discuss open problems and future work on anonymity measurements.

1 Introduction

Remaining anonymous has been an unsolved problem ever since Captain Nemo. Yet in some situations we would like to provide guarantees of a person remaining anonymous. However, the meaning of this, both on the internet and in real life, is somewhat elusive. One can never remain truly anonymous, but relative anonymity can be achieved. For example, walking through a crowd of people does not allow a bystander to track your movements (though be sure that your clothes do not stand out too much). We would like to express anonymity properties in the virtual world in a similar fashion, yet this is more difficult. The users would like to know whether they can be identified (or rather the probability of being identified). Similarly, they would like to have a metric to compare different ways of achieving anonymity: what makes you more difficult to track in London — walking through a crowd or riding randomly on the underground for a few hours?

In this paper, we choose to abstract away from the application level issues of anonymous communication such as preventing the attacker from embedding URLs pointing to the attacker's webpage in messages in the hope that the victim's browser opens them automatically. Instead, we focus on examining ways of analysing the anonymity of a messages going through mix-based anonymity systems [Cha81] in which all network communication is observable by the attacker.

In such a system, the sender, instead of passing the message directly to the recipient, forwards it via a number of *mixes*. Each mix waits for n messages to arrive before decrypting and forwarding them in a random order, thus hiding the correspondence between incoming and outgoing messages.

Perhaps the most intuitive way of measuring the anonymity of a message M in a mix system is to just count the number of messages M has been mixed with while passing through the system. However, as pointed out in [Cot94] and [GT96], this is not enough as all the other messages could, for instance, come from a single known sender. Indeed, the attacker may mount the so called $n-1$ attack based on this observation by sending $n-1$ of their own messages to each of the mixes on M's path. In this case, the receiver of M ceases to be anonymous.

Another popular measure of anonymity is the notion of anonymity set. In the rest of this section we look at how anonymity sets have previously been defined in the literature and what systems they have been used in.

1.1 Dining Cryptographers' Networks

The notion of anonymity set was introduced by Chaum in [Cha88] in order to model the security of Dining Cryptographers' (DC) networks. The size of the anonymity set reflects the fact that even though a participant in a Dining Cryptographers' network may not be directly identifiable, the set of other participants that he or she may be confused with, can be large or small, depending on the attacker's knowledge of particular keys. The anonymity set is defined as the set of participants who could have sent a particular message, as seen by a global observer who has also compromised a set of nodes. Chaum argues that its size is a good indicator of how good the anonymity provided by the network really is. In the worst case, the size of the anonymity set is 1, which means that no anonymity is provided to the participant. In the best case, it is the size of the network, which means that any participant could have sent the message.

1.2 Stop and Go Mixes

In [KEB98] Kesdogan *et al.* also use sets as the measure of anonymity. Furthermore, they define the anonymity set of users as those who had a non-zero probability of having the role $\mathcal{R}$ (sender or recipient) for a particular message. The size of the set is then used as the metric of anonymity. Furthermore, *deterministic anonymity* is defined as the property of an algorithm which always yields anonymity sets of size greater than 1.

The authors also state that it is necessary to protect users of anonymity systems against the $n-1$ attack described earlier and propose two different ways doing so: the Stop-and-Go-mixes and a scheme for mix cascades[1]. Stop-and-Go are a variety of mixes that, instead of waiting for a particular number of messages to arrive, flush them according to some delay which is included in the message. They protect against the $n-1$ attack by discarding the messages if they are received outside the specified time frame. Thus, the attacker cannot delay messages which is required to mount the $n-1$ attack.

[1] An anonymity system based on mix cascades is one where all the senders send all their messages through one particular sequence of mixes.

1.3 Standard terminology

In an effort to standardise the terminology used in anonymity and pseudonymity research publications and clarify different concepts, Pfitzmann and Köhntopp [PK00] define anonymity itself as:

> "Anonymity is the state of being not identifiable within a set of subjects, the *anonymity set*."

In order to further refine the concept of anonymity and anonymity set and in an attempt to find a metric for the quality of the anonymity provided they continue:

> "Anonymity is the stronger, the larger the respective anonymity set is and the more evenly distributed the sending or receiving, respectively, of the subjects within that set is."

The concept of "even distribution" of the sending or receiving of members of the set identifies a new requirement for judging the quality of the anonymity provided by a particular system. It is not obvious anymore that the size is a very good indicator, since different members may be more or less likely to be the sender or receiver because of their respective communication patterns.

2 Difficulties with Anonymity Set Size

The attacks against DC networks presented in [Cha88] can only result in partitions of the network in which all the participants are still equally likely to have sent or received a particular message. Therefore the size of the anonymity set is a good metric of the quality of the anonymity offered to the remaining participants. In the Stop-and-Go system [KEB98] definition, the authors realise that different senders may not have been equally likely to have sent a particular message, but choose to ignore it. We note, however, that in the case they are dealing with (mix cascades in a system where each mix verifies the identities of all the senders), all senders have equal probability of having sent (received) the message. In the standardisation attempt [PK00], we see that there is an attempt to state, and take into account this fact in the notion of anonymity, yet a formal definition is still lacking.

We have come to the conclusion that the potentially different probabilities of different members of the anonymity set actually having sent or received the message are unwisely ignored in the literature. Yet they can give a lot of extra information to the attacker.

2.1 The Pool Mix

To further emphasise the dangers of using sets and their cardinalities to assess and compare anonymity systems, we note that some systems have very strong "anonymity set" properties. We take the scenario in which the anonymity set of a message passing through a mix includes (at least) the senders of all the messages which have *ever* passed through that mix.

This turns out to be the case for the "pool mix" introduced by Cottrell in [Cot94]. This mix always stores a pool of n messages (see Figure 1). When incoming N messages have accumulated in its buffer, it picks n randomly out of the $n+N$ it has, and stores them, forwarding the other ones in the usual fashion. Thus, there is always a small probability of any message which has ever been through the mix not having left it. Therefore, the sender of every message should be included in the anonymity set (we defer the formal derivation of this fact until Section 5). At this point we must consider the anonymity provided by this system. Does it really give us very strong anonymity guarantees or is measuring anonymity using sets inappropriate in this case? Our intuition suggests the latter, [2] especially as the anonymity set seems to be independent of the size of the pool, n.

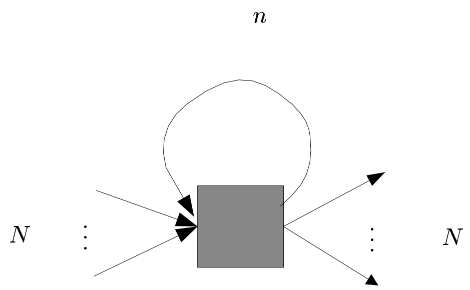

Fig. 1. A Pool Mix

2.2 Knowledge Vulnerability

Yet another reason for being sceptical of the use of anonymity sets is the vulnerability of this metric against an attacker's additional knowledge. Consider the arrangement of mixes in Figure 2. The small squares in the diagram represent senders, labelled with their name. The bigger boxes are mixes, with threshold of 2. Some of the receivers are labelled with their sender anonymity sets.

Notice that if the attacker somehow establishes the fact that, for instance, A is communicating with R, he can derive the fact that S received a message from E. Indeed, to expose the link $E \to S$, all the attacker needs to know is that

[2] A side remark is in order here. In a practical implementation of such a mix, one would, of course, put an upper limit on the time a message can remain on the mix with a policy such as: "All messages should be forwarded on within 24 hours + K mix flushes of arrival".

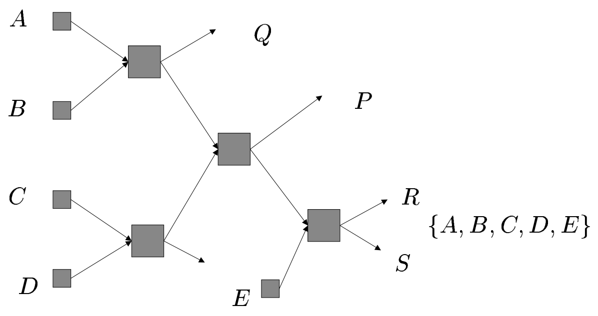

Fig. 2. Vulnerability of Anonymity Sets

one of A, B, C, D is communicating to R. And yet this is in no way reflected in S's sender anonymity set (although E's receiver anonymity set, as expected, contains just R and S).

It is also clear that not all senders in this arrangement are equally vulnerable to this, as is the fact that other arrangements of mixes may be less so. Although we have highlighted the attack here by using mixes with threshold of 2, it is clear that the principle can be used in general to cut down the size of the anonymity set.

3 Entropy

We have now discussed several separate and, in our view, important issues with using anonymity sets and their cardinalities for measuring anonymity. We have also demonstrated that there is a clear need to reason about information contained in probability distributions. One could therefore borrow mathematical tools from Information Theory [Sha48]. The concept of *entropy* was first introduced to quantify the uncertainty one has before an experiment. We now proceed to define our anonymous communication model and the metrics that use entropy to describe its quality. The model is very close to the one described in [KEB98].

Definition 1. *Given a model of the attacker and a finite set of all users Ψ, let $r \in \mathcal{R}$ be a role for the user ($\mathcal{R} = \{sender, recipient\}$) with respect to a message $\mathcal{M}$. Let $\mathcal{U}$ be the attacker's a-posteriori probability distribution of users $u \in \Psi$ having the role r with respect to $\mathcal{M}$.*

In the model above we do not have an anonymity set but an r *anonymity probability distribution* $\mathcal{U}$. For the mathematically inclined, $\mathcal{U} : \Psi \times \mathcal{R} \to [0, 1]$ s.t. $\sum_{u \in \Psi} \mathcal{U}(u, r) = 1$. In other words, given a message M, we have a probability distribution of its possible senders and receivers, as viewed by the attacker. $\mathcal{U}$ may assign zero probability to some users which means that they cannot possibly have had the role r for the particular message $\mathcal{M}$. For instance, if the message we are considering was seen by the attacker as having arrived at Q, then $\mathcal{U}(receiver, Q) = 1$ and $\forall S \neq Q \;\; \mathcal{U}(receiver, S) = 0$.[3]. If all the users that are not assigned a zero probability have an equal probability assigned to them, as in the case of a DC network under attack, then the size of the set could be used to describe the anonymity. The interesting case is when users are assigned different, non zero probabilities.

Definition 2. *We define the effective size $\mathcal{S}$ of an r anonymity probability distribution $\mathcal{U}$ to be equal to the entropy of the distribution. In other words*

$$\mathcal{S} = -\sum_{u \in \Psi} p_u \log_2(p_u)$$

where $p_u = \mathcal{U}(u, r)$.

One could interpret this effective size as the number of bits of additional information that the attacker needs in order to definitely identify the user u with role r for the particular message $\mathcal{M}$. It is trivial to show that if one user is assigned a probability of 1 then the effective size of is 0 bits, which means that the attacker already has enough information to identify the user.

There are some further observations:

- It is always the case that $0 \leq \mathcal{S} \leq \log_2 |\Psi|$.
- If $\mathcal{S} = 0$ the communication channel is not providing any anonymity.
- If for all possible attacker models, $\mathcal{S} = \log_2 |\Psi|$ the communication channel provides perfect $\mathcal{R}$ anonymity.

We now go on to show how to derive the discrete probability distribution required to calculate the information theoretic metric of anonymity presented above.

3.1 Calculating the Anonymity Probability Distribution

We now show how to calculate the sender anonymity probability distribution for a particular message passing through a mix system with the standard threshold mixes. We assume that we have the ability to distinguish between the different senders using the system. This assumption is discussed in Section 6. To analyse a run of the system (we leave this notion informal), we have to have knowledge of all of the messages which have been sent during the run. (This includes mix-user, user-mix and mix-mix messages and is consistent with the model of the attacker who sees all the network communications, but has not compromised any mixes.) The analysis attaches a sender anonymity probability distribution to every message. The starting state is illustrated in Figure 3a.

[3] Alternatively, we may choose to view the sender/receiver anonymity probability distribution for a message $\mathcal{M}$ as an extension of the underlying sender/receiver anonymity set to a set of pairs of users with their associated (non-zero) probabilities of sending or receiving it.

We take the case of the attacker performing "pure" traffic analysis. In other words, he does not have any *a-priori* knowledge about the senders and receivers and the possible communications between them. [4] The attacker's assumption arising from this is that a message, having arrived at a mix, was equally likely to have been forwarded to all of the possible "next hops", independent of what that next hop could be.

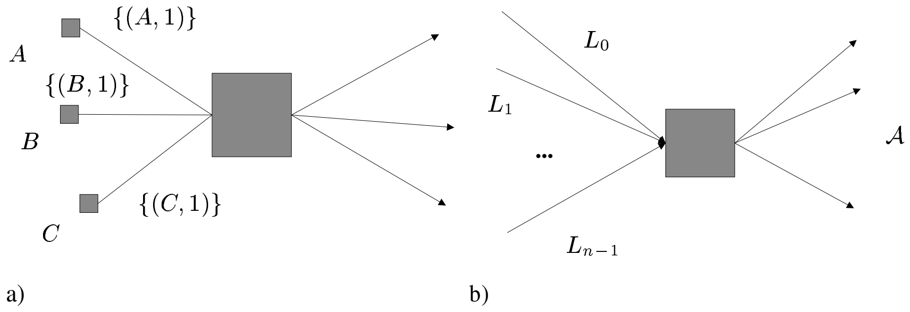

Fig. 3. a) The start of the analysis. b) Deriving the anonymity probability distribution of messages going through a mix.

For a general mix with n incoming messages with anonymity probability distributions $L_0 \ldots L_{n-1}$, which we view as sets of pairs, we observe that the anonymity probability distributions of all of the messages coming out of the mix, are going to be the same. This distribution $\mathcal{A}$ is defined as follows:

$(x,p) \in \mathcal{A}$ iff $\exists i.(x,p') \in L_i$ and

$$p = \frac{\sum_{i.(x,p_j)\in L_i} p_j}{n}$$

Thus, the anonymity probability distribution on each of the outgoing arrows on Figure 3a is $\{(A,\frac{1}{3}),(B,\frac{1}{3}),(C,\frac{1}{3})\}$.

In the next section we will discuss how we can calculate the effective anonymity size of systems composed of other mixes.

3.2 Composing mix systems

Given some arrangement of individual mixes connected together, it is possible to calculate the maximum effective anonymity size of the system from the effective anonymity size of its components. The assumption necessary to do this is that the inputs to the different entry points of this system originate from distinct

[4] This is a simplification. In practice, the attacker analysing email can choose to assign lower probabilities to, for example, potential Greek senders of an email in Russian which arrived in Novosibirsk.

users. In practice this assumption is very difficult to satisfy, but at least we can get an upper bound on how well a "perfect" system would perform.

Assume that there are l mixes each with effective sender anonymity size $S_i, 0 < i \leq l$. Each of these mixes sends some messages to a mix we shall call sec. The probability a message going into sec originated from mix i is $p_i, 0 < i \leq l, \sum_i p_i = 1$.

Using our definitions it is clear that $S_{sec} = \sum_{0<i\leq l} p_i \log(p_i)$ is the effective anonymity size of this second mix.

The effective sender anonymity size of messages going through the system described above is $\sum_{0<i\leq l} \sum_{0<j\leq f(i)} p_j p_i \log(p_j p_i)$ which simplifies to

$$S_{total} = S_{sec} + \sum_{0<i\leq l} p_i S_i$$

where $f(i)$ is the number of inputs that mix i takes and $p_j, 0 < j \leq f(i)$ is the probability corresponding to the j^{th} input of i.

Using this rule we can calculate the effective sender anonymity set size of mix systems or other anonymous communication channels using the effective sizes of their components and information about how they are interconnected. A similar approach can be used to calculate the effective recipient anonymity set size.

In the next section, we look at how knowledge about the system available to the attacker can be used to perform a better anonymity analysis.

4 Route length

Having included probabilities in the model and demonstrated that they can give the attacker more information about the system than just anonymity sets, we note that the standard attacks aimed at reducing the size of the anonymity set will now have the effect of narrowing the anonymity probability distribution. If we consider this distribution as a set of pairs (of a sender and its respective non-zero probability of having sent the message), then narrowing the probability distribution is the process of deriving that some senders have zero probability of sending the message and can therefore be safely excluded from the set.

We now look at an attack which not only has the ability to narrow the probability distribution, but also to alter it in such a way as to reduce the entropy of the anonymity probability distributions without affecting the underlying anonymity set.

As suggested in [BPS00], route length is important and some arrangements of mixes are more vulnerable to route length based attacks than others. Here, we demonstrate that maximum route length should be taken into account when calculating anonymity sets. Note that, of course, the maximum route length in a traditional mix-based anonymity system exists and is known to the attacker[5]. Several mix systems have been designed to remove the maximum route length

[5] The reason for this is standard, as follows: All the messages in a mix-based system have to have the same size, otherwise an attacker could trace particular messages.

constraint, for instance via tunnelling in Onion Routing [STRL00] or Hybrid mixes [OA00], but it exists in fielded systems such as Mixmaster [MC00] (maximum route length of 20) and so can be used by the attacker.

It may also be possible to obtain relevant information by compromising a mix. Some mix systems will allow a mix to infer the number of mixes a message has already passed through and therefore the maximum number of messages it may go through before reaching the destination. Such information would strengthen our attack, so care needs to be taken to design mix systems (such as Mixmaster [Cot94]) which do not give it away.

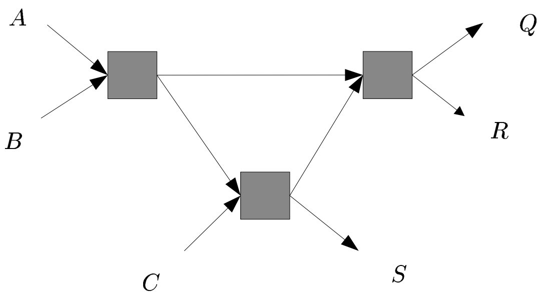

Fig. 4. Using maximum route length to reduce anonymity sets

We illustrate the problem by example. Consider the situation in Figure 4, where each arrow represents a message observed by the attacker. Now let us suppose that the maximum route length is 2, i.e. any message can pass through no more than 2 mixes. The arrow from the bottom to the rightmost mix could only have been the message from C as otherwise this message, coming from A or B would have gone through 3 mixes. From this, we infer that C was not communicating to S, which makes S's sender anonymity set $\{A, B\}$. Of course, without taking maximum route length into account, this anonymity set would have been $\{A, B, C\}$.

We now illustrate how the same fact can alter the sender anonymity probability distribution of a particular receiver and therefore reduce its entropy.

Here we use the same arrangement of mixes, but look at a different receiver, Q. The anonymity probability distribution worked out using the algorithm (without the route length constraint) in the previous section is shown in Figure 5. If the attacker knows that the maximum route length is 2, the arrow from mix 2 to mix 3 has the sender probability distribution of: $\{(C, 1)\}$ and thus the probability distribution at Q (or R) is $\{(A, \frac{1}{4}), (B, \frac{1}{4}), (C, \frac{1}{2})\}$. This reduced the entropy from 1.5613 down to 1.5. Compare this with the entropy of 1.585 for a uniform

Yet each message (when leaving the sender) has to include inside it all the addresses of all the servers it will be forwarded via. Thus, there is a limit on the number of the mixes a message can pass through, and it is known to the attacker.

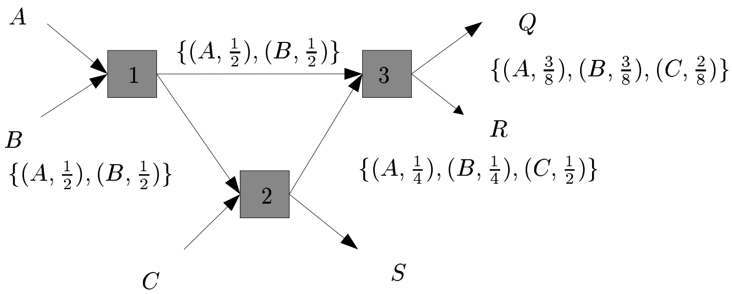

Fig. 5. Using maximum route length to alter anonymity probability distribution

distribution. It is also worth noting that an attack which eliminated one of the possible senders, would reduce the entropy to at most 1 bit and an attack which would expose a single host as the sender of a particular message to Q — to 0 bits. Thus, our metric is capable of not only comparing the effectiveness of systems, but also the power of different attacks. A similar idea has been proposed by [DSCP02]. However, comparing "the effectiveness" of different attacks using this method in general is beyond the scope of this paper and is the subject of future work.

5 Analysis of the Pool Mix

Recall from Section 2.1 that a pool mix stores n messages and receives N messages every round. It then puts together the stored and received messages and outputs N of them (chosen randomly). The remaining n messages are stored for the next round. On round zero the mix stores n "dummy messages" that look to outsiders as if they were real, but were created by the mix itself.

First, we calculate the sender anonymity set and the sender anonymity set size. Denote the anonymity sets associated with the N messages arriving at round i, $K_{i,1} \ldots K_{i,N}$ and let $\hat{K}_i = K_{i,1} \cup \ldots \cup K_{i,N}$. Now the sender anonymity set of the outgoing messages after round k will be the union of the anonymity sets of the messages stored in the mix (all of which are the same and are equal to the anonymity sets of all the messages which left the mix at the previous round) and the messages which arrived from the network.

$$A_0 = \{mix\}$$
$$A_i = S_{i-1} \cup \hat{K}_i$$

Now, assume that all of the messages arrive at the pool mix directly from senders, which are all different from each other. Formally, $\forall i,j,k,l.(i \neq j \vee k \neq l) \Rightarrow K_{i,j} \neq K_{k,l}$. This implies that the size of the set after round k is

$$|A_k| = N \times k + 1$$

and for $k \to \infty$

$$\lim_{k\to\infty} |A| \to \infty$$

It is clear that the size of A_k does not provide us with a useful insight on how well this mix performs. In particular, it does not capture the fact that senders of past rounds have smaller probabilities of being the senders of messages that come out of the mix at the last round. Finally, this metric does not allow us to compare the pool mix with other mixing architectures, including conventional threshold mixes.

We therefore compute the effective size, based on the entropy, of the sender anonymity set.

The probability that a message which comes out of the mix at round k was introduced by a sender in the mix at round $0 < x \leq k$ is

$$p_{round_x} = \frac{N}{N+n}\left(\frac{n}{N+n}\right)^{k-x}$$

$$p_{round_0} = \left(\frac{n}{N+n}\right)^{k}$$

Definition 3. *Now, assume that each message was coming directly from a sender and all senders only send one message. Note that after round* 0*, the only sender involved is the mix itself. The effective size of the sender anonymity set for round* k *is*

$$E_k = -\sum_{x=1}^{k}\left(\frac{N}{N+n}\left(\frac{n}{N+n}\right)^{k-x}\log\left(\frac{1}{N+n}\left(\frac{n}{N+n}\right)^{k-x}\right)\right) - \left(\frac{n}{N+n}\right)^{k}\log\left(\frac{n}{N+n}\right)^{k}$$

After a large number of rounds ($k \to \infty$) *the above expression of the effective size converges towards*

$$\lim_{k\to\infty} E = \left(1+\frac{n}{N}\right)\log\left(N+n\right) - \frac{n}{N}\log n$$

The effective size of the set provides us with useful information about how the mix is performing. As one would expect if there is no pool then the effective sender anonymity set is the same as for a threshold mix architecture with N inputs.

Example 1. When there is no pool ($n = 0$) the effective anonymity set size is

$$\lim_{k\to\infty} E = \log N$$

Example 2. When only one message is fed back to the mix ($n = 1$)

$$\lim_{k\to\infty} E = \left(1 + \frac{1}{N}\right) \log\left(N+1\right)$$

So a mix of this type that takes $N = 100$ inputs will have an effective size of $\lim_{k\to\infty} E = 6.725$. This is equivalent to a threshold mix that takes ≈ 106 inputs.

Example 3. A pool mix with $N = 100$ inputs out of which $n = 10$ are fed back will have an effective size of $\lim_{k\to\infty} E = 7.129$. That is equivalent to a threshold mix with $N = 2^{7.127} \approx 140$ inputs.

The additional anonymity that the pool mix provides is not "for free" since the average latency of the messages increases from 1 round to $1 + \frac{n}{N}$ rounds with a variance of $\frac{n \times (N+n)^2}{N^3}$.

6 Discussion

Let us now examine the scenarios in which our analysis may be useful and demonstrate that one would not be able to use other well-known attacks to compromise anonymity.

The new entropy measure of anonymity is useful in analysing mix networks, rather than cascades or DC nets where there is no possibility of members of anonymity sets having different probabilities of taking on particular roles. The route length techniques are applicable in mix network systems which have a maximum route length constraint such as Mixmaster [MC00].

It is worth mentioning that a similar information theoretic metric was independently proposed in [DSCP02] and used to compare different anonymity systems. Here we concentrate on using it for analysing mix systems and show how it can be used to express new attacks.

7 Conclusion

We have demonstrated serious problems with using the notion of anonymity set for measuring anonymity of mix-based systems. In particular, we exhibited the pool mix as an illustration of the fact that we cannot always use anonymity sets to compare the effectiveness of different mix flushing algorithms.

We have also proposed an information-theoretic metric based on the idea of anonymity probability distributions. We showed how to calculate them and used the metric to compare the pool mix to more traditional mixes.

We must note, however, that our new metric does not really deal with the knowledge vulnerability problem discussed in Section 2.2. We feel that additional structure to enrich the notion of anonymity sets and enable better analysis of knowledge-based vulnerabilities is needed. However, having introduced probabilities into the model, we want to go on and develop a framework capable of answering questions like "What happens to the anonymity probability distribution of receiver S when the attacker knows that A is communicating to P with probability $\frac{1}{3}$ or R with probability $\frac{2}{3}$?"[6]. This is the subject of future work.

We then showed that care must be taken when calculating anonymity probability distributions as the same attacks as used against the anonymity set metric, would also apply here. In particular, we demonstrated that if maximum route length in a mix system exists, it is known to the attacker and can be used extract additional information and gain knowledge which was impossible to express using anonymity sets.

We feel that more sophisticated probabilistic metrics of anonymity should be developed. Moreover, perhaps, if combined with knowledge of the communication protocols executed by the sender and recipient, they can yield powerful attacks against mix-based systems. Moreover, we feel that in a subject like anonymity, formal reasoning is essential if strong guarantees are to be provided. Yet another direction is relating the above to unlinkability and plausible deniability. All these are subjects of future work.

7.1 Acknowledgements

We gratefully acknowledge the support of EPSRC grant GRN24872 Wide Area Programming and helpful contributions and suggestions from Ross Anderson, Mike Bond, Richard Clayton, Aras Kimis, Ira Moskowitz, Richard Newman, Peter Sewell, Paul Syverson and Keith Wansbrough.

References

[BPS00] O. Berthold, A. Pfitzmann, and R. Standtke. The disadvantages of free MIX routes and how to overcome them. In *Designing Privacy Enhancing Technologies: Proceedings of the International Workshop on the Design Issues in Anonymity and Observability*, LNCS 2009. 2000.

[Cha81] D. Chaum. Untraceable electronic mail, return addresses and digital pseudonyms. *Communications of the A.C.M.*, 24(2):84–88, 1981.

[Cha88] D. Chaum. The dining cryptographers problem: Unconditional sender and recipient untraceability. *Journal of Cryptology: the journal of the International Association for Cryptologic Research*, 1(1):65–75, 1988.

[Cot94] L. Cottrell. Mixmaster and remailer attacks, 1994. `http://www.obscura.com/~loki/remailer/remailer-essay.html`.

[DSCP02] C. Diaz, S. Seys, J. Claessens, and B. Preneel. Towards measuring anonymity. In *Workshop on Privacy Enhancing Technologies*, LNCS 2482. 2002.

[GT96] C. Gulcu and G. Tsudik. Mixing email with Babel. In *1996 Internet Society Symposium on Network and Distributed Sytem Security*, pages 2–16. San Diego, CA, 1996.

[KEB98] D. Kesdogan, J. Egner, and R. Buschkes. Stop-and-go-MIXes providing probabilistic anonymity in an open system. In *Proceedings of the International Information Hiding Workshop*, LNCS 1525. 1998.

[6] The reader may wish to refer back to the Figure 2 to see why this is not at all trivial.

[MC00] U. Moeller and L. Cottrell. *Mixmaster Protocol Version 3*, 2000. `http://www.eskimo.com/~rowdenw/crypt/Mix/draft-moeller-v3-01.txt`.

[OA00] M. Ohkubo and M. Abe. A length-invariant hybrid mix. In T. Okamoto, editor, *Advances in Cryptology – ASIACRYPT*, LNCS 1976, page 178 ff. 2000.

[PK00] A. Pfitzmann and M. 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*, LNCS 2009. 2000.

[Sha48] C. Shannon. The mathematical theory of communication. *Bell Systems Technical Journal*, 30:50–64, 1948.

[STRL00] P. F. Syverson, G. Tsudik, M. G. Reed, and C. E. Landwehr. Towards an analysis of onion routing security. In *Designing Privacy Enhancing Technologies: Proceedings of the International Workshop on the Design Issues in Anonymity and Observability*, LNCS 2009. 2000.