

A Method for Degradation of Anonymity on Mix Systems for E-Mail and Surfing the WWW

Lexi Pimenidis*

Abstract:

We'll show in this paper, that anonymizing networks that provide access to external services are subject to a traffic analysis that can strongly degrade the anonymity provided to single participants. This analysis can be used as an effective measure to preprocess anonymized network streams and build a step stone for further, more elaborated, attacks.

1 Introduction

Encryption hides the content of a conversation, but it doesn't hide the identity of the participating users. One of the first proposals to protect this information was in Chaum's paper [Cha81], where he proposes mixes to protect this data. Thus the users are anonymous, *within the respective anonymity set* [PK05], i.e. the group that participated in the single round of a mix. The grade of protection is generally considered to be linked to the size of this anonymity set.

In this paper we're going to show how the size of an anonymity set is decreased, if requests are relayed to external entities and some of these recipients answer to the anonymous request. In general, this traffic analysis will always be possible, if an anonymity system is used for hiding communication to external services that generate responses directly to requests and are not participants in the network, e.g. web-servers or email communication.

To show the extent of our new traffic analysis, we build a generic model of anonymity systems, show in which cases a successful degradation is possible, how it is done and conduct simulations to provide quantitative results (and prove the theory).

2 Related Works

There is a number of related publications that also measure the level of protection which is provided by an anonymity infrastructure. To the extent of our knowledge there is no publication that regards the grade of linkability that results by answers to anonymous requests.

The results of [KAP02], and [MD04] measure the provided level of anonymity by the

*Mr Pimenidis is funded by the European Commission's 6th Framework Program.

amount of observations, an attacker needs to break a system. They show that repeated communication will disclose a user's peers, even if perfect unlinkability is provided in each single round.

Diaz et al use the size of the anonymity set as a measure for the quality of the provided anonymity in their work [DSD04]. We use the same measure of anonymity in this work, but we do not focus on mix designs, so the results of their work and this are not comparable.

An information theoretical evaluation is done in [DSCP02], [SD02], and [SK03]. Here the grade of anonymity is calculated by probability distributions. While [DSCP02] provides numerical results, the other two papers stick with developing formulas.

3 Models and Methods

We assume that the anonymizing network either consists of a single mix, or can be modeled by a single mix that works in distinct rounds. The attacker is assumed to be an observer of the system, esp. of a user called Alice. She has m distinct peers and uses the anonymity technique to hide her requests in an *anonymity set* of size at most b . Alice runs no server, or gets contacted without asking for it in the first hand.


Alice participates in any round of the system with a probability called α . If Alice sends a message to one of her peers, the message is encoded and relayed through the system such that it reaches the peer. By nature of the anonymity system, a global observer will not be able to unambiguously link Alice's request to any outgoing request of the system. Instead, he will be able to bring down the number of possible peers into a set of size b .

The peer itself answers to the request with a probability ρ and we assume that the answers' delays are distributed according to an exponential distribution with the mean value of δ rounds of the system. We also assume that answers are always linkable to the requests that are forwarded out of the anonymity system. The latter is always the case if Alice surfs the web, and most often in email correspondence. Note that encrypting content doesn't change this.

Degradation of Anonymity

First, let's observe a single round of the system: if Alice sends only a single message using the anonymity system, the recipient is hidden in a set of b recipients. Out of these, $0 \leq b_0 \leq b$ recipients do not send an answer at all, or the answer is not linkable to a request. These recipients remain in the anonymity set. The other $b_1 = b - b_0$ answers, namely from peers P_1, P_2, \dots, P_{b_1} , are returned to the anonymity system that returns them to the original senders. However, since the answers are not necessarily simultaneously, e.g. if they are emails, they are processed in different anonymity sets, namely $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_{b_1}$. For all \mathcal{A}_i , where Alice is not a recipient of any message of the anonymity set \mathcal{A}_i , we know that P_i 's answer was not an answer to Alice request. These P_i can be *excluded* from the anonymity set.

Thus, the following self-evident Lemma is valid:


 Figure 1: Results of the first simulation for the parameters α , δ , and b

Lemma 1: *Given a user U sends a message μ that is hidden in a set of b other messages in an anonymity set $\mu \in \mathcal{A}^*$. These messages are directed to the users $\mathcal{R} = \{R_1, R_2 \dots R_b\}$. Assume, that $\mathcal{P} = \{P_1, P_2, \dots P_{b_1}\} \subseteq \mathcal{R}$ send linkable answers to the requests from \mathcal{A}^* and these answers are returned to the original senders in the anonymity sets $\mathcal{A}_1, \dots, \mathcal{A}_{b_1}$. It may be possible that $\mathcal{A}_i = \mathcal{A}_j$ for some $i \neq j$. In this case: the effective anonymity set \mathcal{A}_μ of the message μ consists of*

$$\mathcal{A}_\mu = \{\mathcal{R} - \mathcal{P}\} \cup \{P_j : P_j \in \mathcal{P}, \text{ and } U \in \mathcal{A}_j\} \quad (1)$$

$$\iff \mathcal{A}_\mu = \mathcal{R} - \{P_j : P_j \in \mathcal{P}, U \notin \mathcal{A}_j\} \quad (2)$$

Note that $\mathcal{A}_\mu \subseteq \mathcal{R}$ and thus $|\mathcal{A}_\mu| \leq b$.

The approach can be trivially extended to multiple communications of Alice.

4 Experiments

Lemma 1 shows that the size of the *effective* anonymity set, is smaller or equally sized to a system that relays requests to external entities, as in a system where requests to external entities aren't allowed. The question is, whether the set actually is smaller and if it is, how much. Thus, we conducted series of experiments to show the extend of anonymity loss. As a measure of the grade of anonymity we use the average size of a user's anonymity set per communication step. We compare the original value, i.e. the value that would be valid, if there were no answers allowed in the system, to the *effective* value, that shows the reduced size of the anonymity set.

In the first simulations, we considered the model as given in section 3 and modeled the user behaviour with a random number generator. To show the impact of the variables α , δ , ρ , m and b , we kept all parameters fixed and varied one along an interval. The default values were $\alpha = 0.1$, $\delta = 5$, $\rho = 1$, $m = 10$ and $b = 10$. Plots of the results of the series of parameters α , δ , and b can be seen in figure 1.

As can be seen, the lesser Alice communicates, the more often it is possible to exclude peers from the anonymity set. The second picture shows that, the longer the average delay of an answer is, the more $\mathcal{A}_{t,i}$ are different. Finally, the influence of b is depicted in the last of the four plots.

5 Discussion and Conclusion

As was shown in the previous section, the fact that an anonymity system allows requests to external entities and returns answers, can reduce the size of the anonymity set to possibly 20% of the expected value.

Note that this analysis will most likely not provide a complete break of a system, i.e. it won't provide a list of a user's peers. On the other hand this weakness can be used as a preprocessing step to other attacks, like e.g. a disclosure attack, an intersection attack. Since these attacks are often bound to solve NP-complete problems whose computational effort is usually strongly linked to b , reducing this value can severely damage the overall security of anonymity systems.

One solution to circumvent this weakness is to wait for and collect all answers to a single anonymity set in another single anonymity set, and forward them by the time they all arrived. But this solution is prone to denial-of-service attacks, since a single answer missing would result in no answer being returned at all. Additionally it is not always possible for the anonymity system to wait for the answers, e.g. if the messages are emails. Thus, the best way of avoiding this weakness, is including external services into the network and thus making them internal.

References

- [Cha81] David L. Chaum. Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms. *Communications of the ACM*, 24(2):84 – 88, Feb 1981.
- [DSCP02] Claudia Díaz, Stefaan Seys, Joris Claessens, and Bart Preneel. Towards measuring anonymity. In Roger Dingledine and Paul Syverson, editors, *Proceedings of Privacy Enhancing Technologies Workshop (PET 2002)*. Springer-Verlag, LNCS 2482, April 2002.
- [DSD04] Claudia Díaz, Len Sassaman, and Evelyne Dewitte. Comparison between two practical mix designs. In *Proceedings of 9th European Symposium on Research in Computer Security (ESORICS)*, LNCS, France, September 2004.
- [KAP02] Dogan Kesdogan, Dakshi Agrawal, and Stefan Penz. Limits of Anonymity in Open Environments. In *Information Hiding, 5th International Workshop*. Springer Verlag, 2002.
- [MD04] Nick Mathewson and Roger Dingledine. Practical Traffic Analysis: Extending and Resisting Statistical Disclosure. In *Proceedings of Privacy Enhancing Technologies workshop (PET 2004)*, LNCS, May 2004.
- [PK05] Andreas Pfitzmann and Marit Köhntopp. Anonymity, Unobservability, and Pseudonymity: A Proposal for Terminology. Draft, version 0.23, August 2005.
- [SD02] Andrei Serjantov and George Danezis. Towards an Information Theoretic Metric for Anonymity. In Roger Dingledine and Paul Syverson, editors, *Proceedings of Privacy Enhancing Technologies Workshop (PET 2002)*. Springer-Verlag, LNCS 2482, April 2002.
- [SK03] S. Steinbrecher and S. Köpsell. Modelling Unlinkability. *Proceedings of Privacy Enhancing Technologies workshop (PET 2003)*, LNCS, May 2003.