

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

Lexi Pimenidis

## Abstract:

Some anonymizing networks, like e.g. JAP, relay messages to external services, i.e. allow communication to entities that are not participants of the network. In general this is the case if networks provide access to the world wide web (like e.g. JAP, and Tor) or allow email messaging with pseudonymous return addresses (like e.g. Mixmaster).

Even if answers are indistinguishable from requests in the anonymizing network, they can severely damage the amount of anonymity provided by a system. This is due to the fact that outside the anonymous network answers are linkable to their request, take an individual amount of time to be generated, and thus are returned to the original sender in different anonymity sets.

We will show in this paper that anonymizing networks which provide access to external services are subject to a traffic analysis that can strongly degrade the anonymity provided to their participants. The degradation is based on the fact that replies to messages are linkable outside the anonymizing network. This analysis can be used as an effective measure to preprocess all kinds of anonymized communication and build a step stone for further, more elaborated, attacks.

## 1 Introduction

Typically today's networks do not hide sender/recipient relationships at all, since all data packets or email-messages traversing the Internet usually contain the sender and the recipient address. While encryption hides the content of a conversation it does not protect the user against traffic analysis, i.e. does not hide the identity of a user's peers. One of the first proposals to protect this information was by Chaum's seminal paper [Cha81]. Chaum proposed "mixes" to solve this problem.

A basic mix collects a batch of messages, changes their order and appearance and relays them to their final destination. This provides untraceability in each single round and an observer is unable to correlate senders and recipients. Thus the users are anonymous, *within the respective anonymity set* [PH06], i.e. the group that participated in the single round of a mix. Thus, it is the nature of a mix that a third party can only observe sets of communicating users without being able to trivially link single entities as peers.

Yet, there are a lot of improvements proposed to the basic mix design, as e.g. Stop-and-Go-Mixes [KEB98], Pool-Mixes, and more. An overview is given in [SDS02]. This work will focus on the basic mix design in an *open environment*. That is, users are free to join and leave the system at any time and contact any available service, especially those that are external to the network. To this end, software at the user side encrypts the request, e.g. a

HTTP-request, with the public key of the mix, and sends it there at any arbitrary time. The mix collects a set of requests, decrypts them, reorders them, and forwards them to their final destination – keeping internal state of the sender/recipient relationship in order to correctly relay responses.

The common strength of anonymity techniques is to hide a user's communication in a set of other communications, thus building the so called anonymity set. The grade of protection is generally considered to be linked to the size of this anonymity set, i.e. the larger the set and evenly distributed the elements, the better the protection against traffic analysis. The degree of protection can be measured by the amount uncertainty of linking incoming to outgoing messages. This value, the entropy, is measured in *bits* and quantifies the protection against traffic analysis: the higher the uncertainty, the better the protection.

Our novel traffic analysis makes use of the messages' linkability outside the network in order to reduce the size of a user's anonymity set. We use the fact that requests take an individual amount of time in order to be processed and answered, in addition to the properties of open environments, i.e. a setup, where the network does not control the number of messages per user and time interval, or the like. Thus, it happens frequently that responses to requests will be returned in *different* anonymity sets. If this happens, we can intersect the potential recipients of the response with the potential senders of the request. Users which are not in the intersection, can then be removed from the set of potential senders, thus effectively reducing the size of the anonymity set and the degrading the level of protection.

There are a number of works that give measures on the quality of protection in these systems, like [DSCP02] and [SD02] and the extension of their work in [SK03], which gives an information theoretical methodology. Measurements of protection in real systems are given in [KAP02], [Dan03] and [MD04]. Alas, none of these works take into account that some client requests may receive answers. In fact, we will show that if an anonymizing network allows requests to external services and answers to be received, the effective size of the anonymity set is typically reduced.

## 1.1 Contributions

In this paper we are 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 kind of answers will always occur if an anonymity system is used for hiding communication to external services that generate responses directly to requests (e.g. the www) and are not participants in the network.

To show the extent of our new traffic analysis, we

1. build a generic model of anonymity systems,
2. show in which cases a successful degradation is possible and how it is done,
3. conduct simulations to provide results.

## 1.2 Roadmap

In Section 2, we discuss related works to this topic and show where this work differs and contributes. The models and methods used in the remainder of this paper are introduced in Section 3 and brought into practice in Section 4. The results are discussed in Section 5 and the paper ends with conclusions in Section 6.

## 2 Related Works

In this section we are giving a short overview over related publications that also measure the level of protection which is provided by an anonymity infrastructure.

The results of [KAP02], [Dan03], and [MD04] measure the provided level of anonymity by the amount of observations, an attacker needs to break a system, i.e. to discover the peers of a single user. They show that repeated communication will disclose a user's peers with high probability, even if perfect unlinkability is provided in each single round. The number of observations typically depends on the number of peers of a user, the size of the provided anonymity sets and the total number of possible peer partners.

However, since these systems only take into account a total break of anonymity and do not measure partial information gain, we cannot compare our results to theirs.

Diaz et al use the size of the anonymity set as a measure for the quality of the provided anonymity in their work [DSD04].

Information theoretical evaluations are done in [DSCP02], [SD02], and [SK03]. Here the grade of anonymity depends on probability distributions that govern the flow of an anonymity system and the result in its output. While [DSCP02] provides some quantitative results similar to the papers in the previous paragraphs, the other two papers stick with developing formulas.

We will use the same measure of anonymity in this work, since we consider the information theoretical approach to be the most general.

Finally, the publications [SS96], [SS99], and [HS03] give formal definitions and logics to define and measure anonymity. All three publications remain on the formal side of the problem and do not give numbers or estimations of protection levels in real or simulated systems.

To the extent of our knowledge the only publication that regards the grade of linkability that results by answers to anonymous requests is [DDT07]. In difference to this paper, the authors are taking into account only replies to messages within an anonymizing network. However, in networks like e.g. Tor and JAP, the answer to a request is send outside the network. The authors also have chosen not to use information theoretical measures to display the impact of their attack, thus the results are not comparable to our results. Similarities include that we also use Poisson processes to simulate user behaviour.

### 3 Models and Methods

This section introduces the model we are working with and describes the method we use to degrade the previously assumed level of anonymity.

Instead of considering real implementations of anonymizing networks, we will focus our work on an abstract model of anonymizing networks. Besides making theoretic evaluation possible, this allows adaption of the results to all implementations of anonymization technique that meet the given preconditions.

#### 3.1 Model

Our considered model of an anonymity technique is depicted in figure 1. 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. This is in analogy to [PH06], where anonymity systems are depicted as protecting users by providing *anonymity sets*. Thus, we consider this choice to be one of the most general models and applicable to most systems<sup>1</sup>.

The mix collects a number of messages each turn, decrypts them, reorders them, and finally relays them to their destination. We consider the attacker to be a passive observer of all incoming and outgoing messages. Without loss of generality, we assume that the attacker wants to learn the peer partners of a user called Alice, and that the implementation of the anonymizing network is robust against any active capabilities of the attacker, such that he must rely on observing the network for more reconnaissance.

Without loss of generality, *Alice* uses an anonymity technique that hides her requests in an *anonymity set* of size at most  $b$ . We also assume that at any point in time there are enough other honest users to provide Alice being disclosed by trivial means. There is no special assumption we make about the remaining users, other than that they provide Alice with a constant stream of non-trivial cover traffic. We also assume that Alice runs no server, or gets contacted without asking for it in the first place.

Alice chooses on her own whether or not to participate in any round of the anonymous infrastructure. The probability of her contributing to a round of the system is called  $\alpha$ . The other users of the system act accordingly. If Alice sends a message to one of her peers, the message is encoded and relayed through the system so that it finally reaches the peer. By the nature of the anonymity system, a passive global observer will not be able to unambiguously link Alice's request to any outgoing request of the system. Instead, by definition, he will be able to bring down the number of possible peers to a set of size  $b$ .

The peer itself decides whether or not to answer to the request. We call the probability for an answer  $\rho$  and assume that the delays are distributed according to a probability distribution with the mean value of  $\delta$  rounds of the system. We also assume that answers are always linkable to the requests that are forwarded out of the anonymity system.

Linkability is directly possible for web surfing, and email. While surfing the web, be it

---

<sup>1</sup>A discussion on this topic can be found in Section 5.

Figure 1: The model used in this publication.

by HTTP or HTTPS, the answer is sent in the same TCP-connection as the request. Thus linking them is trivial for the attacker. Emails most often contain a reference to the original email's message ID in the header field `In-reply-to`. Encrypting emails with PGP does not hide header information, so linking is easy here, too. If an given answer is not linkable to a request, we consider it the same case as if there were no answer at all<sup>2</sup>.

For example JAP [BFK00] is a system, where the traffic analysis as proposed in the next section would be applicable. Since JAP uses fixed cascades, an attacker only needs to be present at two network lines; in case of the default cascade, it would only be necessary to listen in on a single line.  $\rho$  would be nearly 1 in this case, because most web servers in the WWW respond to queries even if the requested document is not existing. JAP comes preconfigured with  $b = 10$  but, due to the mix cascade, the effective value will be larger. It should be safe to consider  $b < 50$  in this case. Without internal access to logs or traffic dumps of JAP it is difficult to give appropriate values on  $\delta$  but it can be assumed to be  $\delta > 10$ .

An overview of the variables and the assumptions are listed in figure 2. We will now show how linkable answers can reduce the size of the anonymity set.

### 3.2 Degradation of Anonymity

In this section we will show in detail, how requests to external service providers lead to a degradation of the anonymity set. We start with a simple example (see also Figure 3), before we do a general analysis.

---

<sup>2</sup>Note however that all common Internet protocols require a reply to be unambiguously linked to its request, so we weren't able to give an intuitive counter example here.

	Description	No.	Assumption
$\alpha$	probability that Alice contributes to a round of the system	1	Some answers are linkable to their requests
$\delta$	average delay of a peer's answer in rounds	2	Alice runs no server
$\rho$	probability that peers answer to requests	3	Answers to requests are sometimes processed in different anonymity sets
$b$	the maximum size of anonymity sets	4	The attacker can observe incoming and outgoing messages

Figure 2: Variables and Central Assumptions used in this publication

Alice, Bob, and Charles use some anonymizing network to communicate to Xavier, Yuriko, and Zoe. In the first round Alice, Bob, and Charles each send one message; these get relayed to Xavier, Yuriko, and Zoe. Each of these three recipients could be a peer from Alice. We call these three the *original anonymity set* of Alice's communication in turn one. In order to allow the reader to trace the answers to these three messages, we coloured them in figure 3.

In the next round, there is an answer from Xavier to the message he received in the first round, and two more messages from Bob and Zoe. The recipients of these three messages are Bob, Charles, and Yuriko. Since Alice was not contained in this recipient set, we know that Alice's peer in the first round was *not* Xavier, since the answer to the request was not returned to Alice.

In the last round, there is are messages from Alice and Bob, and Yuriko sends an answer to the message she received in turn one. The recipients of these three messages are Alice, Charles, and Xavier. Since Alice in a possible recipient to Yuriko's answer to the request from turn one, we can not exclude Yuriko from the original anonymity set.

If communication stops at this point, we also can not exclude Zoe from the original anonymity set as well, since it could have been the case that Alice never received an answer to her request from turn one. The message she received in turn three could have been from Bob, Yuriko, or herself.

We thus showed in this example how the anonymity set was reduced from three to two.

In the general case, 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 the 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.

The same method can be trivially extended to multiple rounds of the system. We will display this in the following example. The flow of messages is also depicted in figure 3.

### 3.3 Mathematical Analysis

We will use this section to develop a calculus for determining the amount of information gain for an adversary deploying the aforementioned traffic analysis.

In the situation as depicted in section 3.1 an adversary is able to determine an anonymity set for each act of communication. We denote this set as  $\mathcal{A} = \{u_1, u_2, \dots, u_b\}$ , where each possible peer  $u_i$  has the probability  $p(u_i)$  of being Alice's peer in this turn. This probability distribution is denoted by  $\mathcal{P}$ .

The amount of uncertainty in bits is then determined by

$$I(\mathcal{P}) = - \sum_{i=1}^b p(u_i) \log_2 p(u_i) \quad (1)$$

For each possible peer  $u_x$  with  $p(u_x) > 0$  excluded from the anonymity set, the probability distribution has to be adapted to the change. Without any further knowledge about the system, the context, or anything else, one can only adjust the probability values of the other peers. This changes  $\mathcal{P}$  to  $\mathcal{P}'$ , where

$$p'(u_i) = \begin{cases} 0 & : \text{ if } i = x \\ p(u_i)/(1 - p(u_x)) & : \text{ if } i \neq x \end{cases} \quad (2)$$

From (1) and (2) follows that  $I(\mathcal{P}) > I(\mathcal{P}')$ .

The exact amount of information gain depends on the probability distribution  $\mathcal{P}$  and the system parameters  $\alpha$ ,  $\delta$ ,  $\rho$ , and  $b$ .

## 4 Experiments

We also conducted series of experiments to show the extend of anonymity loss in case of more complicated probability distributions. To this end we simulated the model as given in section 3.1 and 3.3 with a program that created all data with a pseudo random number generator. In this example we had chosen that the delay of messages send was done with the help of a Poisson distribution.

We considered the model as given in section 3.1 and modeled the user behaviour with a random number generator: in each step, there was a probability of  $\alpha$  that Alice communicated to a randomly chosen peer. Answers were generated with a probability of  $\rho$  after a

Poisson distributed delay with an average of  $\delta$  rounds. The complete process was repeated for a period of 1000 rounds of the system and repeated 200 times. In each round we mixed exactly  $b$  messages generated by users and all incoming answers to previous requests.

To show the impact of the variables  $\alpha, \delta, \rho$  and  $b$ , we kept all parameters fixed and varied one along an interval. The default values were  $\alpha = 0.1, \delta = 5, \rho = 1$  and  $b = 10$ .

Plots of the results of the series of parameters  $\alpha, \delta, \rho$  and  $b$  can be seen in figure 4. The uncertainty about Alice's peers without this traffic analysis would have been about  $\log_2(b) \approx 3.3$  in the first three series and  $\log_2(b)$  in the third. Note also that the results are in this case independent on the number of Alice's peer partners or the total number of peer partners.

As can be seen in the first two pictures, it is bad for the provided anonymity if answers are spread over a large time interval and Alice communicates seldom. If this is the case, an attacker can exclude a lot of potential peers because the answers to their queries are often returned in an anonymity set that does not contain Alice.

The third picture shows that avoiding answers, i.e. reducing  $\rho$  helps to reduce the knowledge gain of a potential attacker. Finally, the influence of  $b$  is depicted in the last of the four plots. While there is no (additional) information gain for  $b = 1$  (in this case, there is no anonymity at all), the attacker gains more insights into sender/recipient relationship with a larger batch size. Fortunately, this value is lower than the overall entropy. Thus raising the amount of messages per anonymity set will result in a stronger protection of the system's users.

## 5 Discussion

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 38% of the expected value. The degradation is based nearly exclusively on the fact that answers to requests originated from an anonymizing network are linkable to the requests and is independent to the handling of answers inside the anonymity system. This generalization makes the results deployable to all implementations of mix networks where an attacker is able to observe all incoming and outgoing messages, like e.g. in JAP.

Note that this analysis will most likely not provide a complete break of a system, i.e. it will not 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, or an intersection attack. In [KAPR06] the authors indicate that the difficulty of exact attacks on anonymizing networks might be  $O(b^m)$  where  $m$  is the number of a user's peers. Thus any reduction of  $b$  can severely damage the overall security of anonymity systems because, for an information gain of  $g$  bit, we can estimate an effort decrease for the attacker of the factor  $2^{g \cdot m}$ .

Although we restricted ourself in this publication to anonymizing systems that are round-based, like e.g. threshold mixes or timed-mixes, the work is extendable to continuous-time mixes. As shown by *Serjantov and Danezis* in [SD02], all of the anonymizing networks

have the common property that they map single transactions to an anonymity set. This anonymity set is either a set of a certain size, or consists of an infinite amount of hosts. Serjantov et al. showed in [SD02] that the latter can be reduced to sets of fixed sizes if all items with a probability of less than some  $1 \gg \epsilon > 0$  are cut off. Thus, all methods map single transactions into anonymity sets of finite size and our analysis is applicable. This is of interest to apply our traffic analysis on implemented systems like e.g. Reliable [Rel04]. For similar reasons, it is possible to apply this analysis not only to single mixes but also to mix networks. A mix network is a set of mixes where messages are routed with onion routing as e.g. described in [GRS96]. Most mix networks can be viewed as a single mix for theoretical evaluation purposes: for every hop of a message that an attacker wants to trace, she puts all outgoing messages of the mix's batch in a set and follows up on *their* paths, until all messages have left the mix network. The resulting set is the anonymity set for the recipient. While this leads to a large, possibly even exponential, increase in the size of an anonymity set, it can be used as one way to abstract mix networks as single mixes.

## 5.1 Protection Methods

To protect against this type of traffic analysis, at least one of the attack's preconditions (see figure 2) have to be mitigated:

**Answers are linkable to their requests** To avoid that answers are linkable to their requests, we have no other choice than to include external services in the network. Otherwise, the communication protocols used will always lead to this condition and thus to the applicability of the proposed attack.

**Alice runs no server** It is infeasible for the average end user to deploy a server that is accessed by random users without the help of dummy traffic. Since the use of dummy traffic puts too heavy load on the network, we do not consider this a good solution or suggest this method to mitigate this problem. In addition, running a public service puts a user to the risk of being vulnerable to security issues of the service, or mistakes made when configuring the software.

**Answers are processed in different anonymity sets** 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. However, this solution is prone to denial-of-service attacks, since a single answer missing would result in no answer being returned to any client at all. Additionally it is not always possible for the anonymity system to wait for the answers depending on the type of messages relayed, e.g. emails.

**Observing incoming and outgoing messages** For simple systems like e.g. JAP, tapping two lines is sufficient for an attacker to hold this requirement. For some cascades, even a single line is sufficient, i.e. in the case of the "Dresden-Dresden" cascade.

But even in more distributed systems, it is not unreasonable for a single entity to keep up with this, as shown in [MZ07].

From the discussion above we can conclude that there seems to be no other practical way of avoiding this weakness, than by either using distributed systems like Tor in order to avoid attackers observing incoming and outgoing messages at the same time, or to include external services into the network and thus making them internal. Thus answers to requests are not distinguishable from requests, or at least not linkable to single requests.

The fact that this attack is very difficult to defeat makes it more dangerous than it might look in the first place.

## 6 Conclusion

We showed that the size of an anonymity set of certain anonymizing networks can be reduced to a fraction of the original size given these networks relay access to external services, e.g. the world wide web or email messaging. Even if answers are processed like normal messages and are thus indistinguishable from requests, they can severely damage the amount of anonymity provided by a system. This analysis can be used as an effective measure to preprocess anonymized network streams and build a step stone for further, more elaborated, attacks.

Future extensions of this work include a mathematical and analytical analysis of this weakness, as well as more detailed user and traffic models.

## References

- [BFK00] Oliver Berthold, Hannes Federrath, and Stefan Köpsell. Web MIXes: A system for anonymous and unobservable Internet access. In H. Federrath, editor, *Proceedings of Designing Privacy Enhancing Technologies: Workshop on Design Issues in Anonymity and Unobservability*, pages 115–129. Springer-Verlag, LNCS 2009, July 2000.
- [Cha81] David L. Chaum. Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms. *Communications of the ACM*, 24(2):84 – 88, Feb 1981.
- [Dan03] George Danezis. Statistical Disclosure Attacks: Traffic Confirmation in Open Environments. In Gritzalis, Vimercati, Samarati, and Katsikas, editors, *Proceedings of Security and Privacy in the Age of Uncertainty, (SEC2003)*, pages 421–426, Athens, May 2003. IFIP TC11, Kluwer.
- [DDT07] George Danezis, Claudia Diaz, and Carmela Troncoso. Two-sided Statistical Disclosure Attack. In Nikita Borisov and Philippe Golle, editors, *Proceedings of the Seventh Workshop on Privacy Enhancing Technologies (PET 2007)*, Ottawa, Canada, July 2007.
- [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.

- [GRS96] David M. Goldschlag, Michael G. Reed, and Paul F. Syverson. Hiding Routing Information. In R. Anderson, editor, *Proceedings of Information Hiding: First International Workshop*, pages 137–150. Springer-Verlag, LNCS 1174, May 1996.
- [HS03] D. Hughes and V. Shmatikov. Information hiding, anonymity and privacy: A modular approach. To appear in *Journal of Computer Security*, 2003.
- [KAP02] Dogan Kesdogan, Dakshi Agrawal, and Stefan Penz. Limits of Anonymity in Open Environments. In *Information Hiding, 5th International Workshop*. Springer Verlag, 2002.
- [KAPR06] Dogan Kesdogan, Dakshi Agrawal, Vinh Pham, and Dieter Rautenbach. Fundamental Limits on the Anonymity Provided by the MIX Technique. In *The 2006 IEEE Symposium on Security and Privacy, Oakland, California, USA*, May 2006.
- [KEB98] Dogan Kesdogan, Jan Egnér, and Roland Büschkes. Stop-and-Go-Mixes Providing Anonymity in an Open System. In D. Aucsmith, editor, *Information Hiding 98 - Second International Workshop*, pages 83 – 98. Springer Verlag, 1998.
- [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.
- [MZ07] Steven J. Murdoch and Piotr Zielinski. Sampled Traffic Analysis by Internet-Exchange-Level Adversaries. In *Proceedings of the Seventh Workshop on Privacy Enhancing Technologies (PET 2007)*, Ottawa, Canada, June 2007.
- [PH06] Andreas Pfitzmann and Marit Hansen. Anonymity, Unlinkability, Unobservability, Pseudonymity, and Identity Management - A Consolidated Proposal for Terminology. Draft, version 0.28, May 2006.
- [Rel04] Reliable. <http://www.bigfoot.com/potatoware/reli/>, 2004. Remailing software, visited June 2004.
- [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.
- [SDS02] Andrei Serjantov, Roger Dingledine, and Paul Syverson. From a Trickle to a Flood: Active Attacks on Several Mix Types. In Fabien Petitcolas, editor, *Proceedings of Information Hiding Workshop (IH 2002)*. Springer-Verlag, LNCS 2578, October 2002.
- [SK03] S. Steinbrecher and S. Köpsell. Modelling Unlinkability. *Proceedings of Privacy Enhancing Technologies workshop (PET 2003)*, LNCS, May 2003.
- [SS96] Steve Schneider and Abraham Sidiropoulos. CSP and Anonymity. In *ESORICS '96: Proceedings of the 4th European Symposium on Research in Computer Security*, pages 198–218, London, UK, 1996. Springer-Verlag.
- [SS99] P. F. Syverson and S. G. Stubblebine. Group principals and the formalization of anonymity. pages 814 – 833. *FM'99 - Formal Methods, Vol. I*, LNCS 1708, Springer-Verlag, 1999.

Turn	Messages	Remark
1		<p>All three recipients could be Alice's peer, thus the size of the anonymity set for her first message is three.</p>
2		<p>We can exclude Xavier from the anonymity set because Alice did under no circumstances receive the answer to the request, which Xavier received in turn one.</p>
		<p>We cannot exclude Xavier from the anonymity set because</p>

Figure 4: Results of the first simulation for the parameters  $\alpha, \delta, \rho$  and  $b$