Towards measuring anonymity

Claudia Díaz, Stefaan Seys, Joris Claessens, and Bart Preneel

K.U.Leuven ESAT-COSIC
Kasteelpark Arenberg 10, B-3001 Leuven-Heverlee, Belgium
claudia.diaz@esat.kuleuven.ac.be
http://www.esat.kuleuven.ac.be/cosic/

Abstract. This paper introduces an information theoretic model that allows to quantify the *degree* of anonymity provided by schemes for anonymous connections. It considers attackers that obtain probabilistic information about users. The degree is based on the probabilities an attacker, after observing the system, assigns to the different users of the system as being the originators of a message. As a proof of concept, the model is applied to some existing systems. The model is shown to be very useful for evaluating the level of privacy a system provides under various attack scenarios, for measuring the amount of information an attacker gets with a particular attack and for comparing different systems amongst each other.

1 Introduction

In today's expanding on-line world, there is an increasing concern about the protection of anonymity and privacy in electronic services. In the past, many technical solutions have been proposed that hide a user's identity in various applications and services. Anonymity is an important issue in electronic payments, electronic voting, electronic auctions, but also for email and web browsing.

A distinction can be made between connection anonymity and data anonymity. Data anonymity is about filtering any identifying information out of the data that is exchanged in a particular application. Connection anonymity is about hiding the identities of source and destination during the actual data transfer. The model presented in this paper focuses on the level of connection anonymity a system can provide, and does not indicate any level of data anonymity.

Information theory has proven to be a useful tool to measure the amount of information (for an introduction, see Cover and Thomas [4]). We try to measure the information obtained by the attacker. In this paper, a model is proposed, based on Shannon's definition of entropy [11], that allows to quantify the degree of anonymity of an electronic system. This degree will be dependent on the power of the attacker. The model is shown to be very useful to evaluate the anonymity a system provides under different circumstances, to compare different systems, and to understand how a system can be improved.

Appeared in Proceedings of PET 2002, April 14-15, 2002, San Francisco,
In Hannes Federath (Ed.), *Designing Privacy Enhancing Technologies*, Lecture Notes in Computer Science, 2002.

1.1 Related work

To our knowledge, there have been several attempts to quantify the degree of anonymity of a user provided by an anonymous connection system.

Reiter and Rubin [9] define *the degree of anonymity* as $1 - p$, where p is the probability assigned to a particular user by the attacker. We believe that this degree is useful to get an idea of the anonymity provided by the system to the user who is in the worst case, but it does not give information on how distinguishable the user is *within* the *anonymity set*. For a system with a large number of possible senders the user who is in the worst case may have an assigned probability that is less than $1/2$ but still be distinguishable by the attacker because the rest of the users have very low associated probabilities.

Berthold *et al.* [2] define the degree of anonymity as $A = log_2(N)$, where N is the number of users of the system. This degree only depends on the number of users of the system, and does not take into account the information the attacker may obtain by observing the system. Therefore, it is not useful to measure the robustness of the system towards attacks. The degree we propose in this paper measures the information the attacker gets, taking into account the whole set of users and the probabilistic information the attacker obtains about them.

Wright *et al.* analyze the degradation of anonymous protocols in [12]. They assume that there is a recurring connection between the sender of a message an the receiver.

An anonymity measurement model similar to the one proposed in this paper has been independently proposed by Serjantov and Danezis in [10]. The main difference between the two models is that their system does not normalize the degree in order to get a value relative to the anonymity level of the ideal system for the same number of users.

1.2 Outline of the paper

This paper is organized as follows: Section 2 describes the system and attack model; the actual measurement model is then proposed in Section 3. As a proof of concept, this model is applied to some existing systems in Section 4. Finally, our conclusions and some open problems are presented.

2 System model

In this paper we focus on systems that provide anonymity through mixes. The system model we consider, thus consists of the following entities:

Senders. These are users who send (or have the ability to send) *messages* to recipients. These messages can be emails, queries to a database, requests of web pages, or any other stream of data. The senders can be grouped into the *set of senders*, that is also called the *anonymity set*. These are the entities of the system whose anonymity we want to protect.

Appeared in Proceedings of PET 2002, April 14-15, 2002, San Francisco,
In Hannes Federath (Ed.), *Designing Privacy Enhancing Technologies*, Lecture Notes in Computer Science, 2002.

During the attack, we consider the number of senders constant, and senders behaving as independent, identical Poisson processes. This is a standard assumption for modeling the behavior of users making phone calls [5]. This means that all users send, in average, the same amount of messages, and the interval of time between one message and the next one follows an exponential distribution.

Recipients. These are the entities that receive the messages from the senders. Recipients can be *active* (if they send back answers to the senders) or *passive* (if they do not react to the received message). Depending on the system there is a large variety of recipients. Some examples are web servers, databases, email accounts or bulletin boards where users can post their messages. The attacker may use the reply messages to gain information.

Mixes. These are the nodes that are typically present in solutions for anonymous connections. They take messages as input, and output them so that the correlation with the corresponding input messages is hidden. There are many different ways to implement a mix; if more than a single mix is used (which is usually done in order to achieve better security), there are several methods to route the message through a chain of mixes; a summary can be found in [2, 7]. In some of the systems, e.g., Crowds, the nodes do not have mixing properties as the ones described by Chaum [3]. In these cases the actual properties of the intermediate nodes will be mentioned.

Note that in some systems the intersection between the different sets might be non-empty (e.g., a sender could be at the same time a recipient or a mix).

Examples of systems that provide anonymous connections are Crowds [9] and Onion Routing [8]. The proposed measurement model is shown to be suitable for these systems. It is however generally applicable to any kind of system.

2.1 Attack model

The degree of anonymity depends on the probabilities that the users have sent a particular message; these probabilities are assigned by the attacker. The degree is therefore measured *with respect to* a particular attack: the results obtained for a system are no longer valid if the attack model changes. Concrete assumptions about the attacker have to be clearly specified when measuring the degree of anonymity.

We briefly describe the attacker properties we consider:

- *Internal-External:* An internal attacker controls one or several entities that are part of the system (e.g., the attacker can prevent the entity from sending messages, or he may have access to the internal information of the entity); an external attacker can only compromise communication channels (e.g., he can eavesdrop or tamper with messages).
- *Passive-Active:* A passive attacker only listens to the communication or reads internal information; an active attacker is able to add, remove and modify messages or adapt internal information.

Appeared in Proceedings of PET 2002, April 14-15, 2002, San Francisco,
In Hannes Federath (Ed.), *Designing Privacy Enhancing Technologies*, Lecture Notes in Computer Science, 2002.

– *Local-Global:* A global attacker has access to the whole communication system, while a local attacker can only control part of the resources.

Different combinations of the previous properties are possible, for instance a global passive external attacker is able to listen to all the channels, while a local internal active attacker can control, for example, a particular mix, but is unable to get any other information.

In our model, an attacker will carry out a *probabilistic attack*. It has been pointed out by Raymond in [7] that these attacks have not been thoroughly addressed so far. With such an attack, the adversary obtains probabilistic information of the form *with probability p, A is the sender of the message.*

3 Proposed measurement model

First of all, we should give a precise definition of *anonymity*. In this paper we adopt the definition given by Pfitzmann and Köhntopp in [6]. Anonymity is *the state of being not identifiable within a set of subjects, the anonymity set.* A sender is identifiable when we get information that can be linked to him, e.g., the IP address of the machine the sender is using.

In this paper we only consider *sender anonymity*. This means that for a particular message the attacker wants to find out which subject in the *anonymity set* is the originator of the message. The *anonymity set* in this case is defined as *the set of honest*[1] *users who might send a message.* It is clear that the minimum size of the anonymity set is 2 (if there is only one user in the anonymity set it is not possible to protect his identity).

Our definition for the degree of anonymity is based on probabilities: after observing the system, an attacker will assign to each user a probability of being the sender.

3.1 Degree of anonymity provided by the system

According to the previous definitions, in a system with N users, the maximum degree of anonymity is achieved when an attacker sees all subjects in the anonymity set as equally probable of being the originator of a message. Therefore, in our model the degree of anonymity depends on the distribution of probabilities and not on the size of the anonymity set, in contrast with previous work [1, 2]. This way, we are able to measure the quality of the system with respect to the anonymity it provides, independently from the number of users who are actually using it. Nevertheless, note that the size of the *anonymity set* is used to calculate the distribution of probabilities, given that the sum of all probabilities must be 1.

The proposed model compares the information obtained by the attacker after observing the system against the optimal situation, in which all honest users

[1] Users controlled by the attacker are not considered as part of the anonymity set, even if they are not aware of this control.

Appeared in Proceedings of PET 2002, April 14-15, 2002, San Francisco,
In Hannes Federath (Ed.), *Designing Privacy Enhancing Technologies*, Lecture Notes in Computer Science, 2002.

seem to be equally probable as being the originator of the message, that is, in a system with N users, the situation where the attacker sees all users as being the originator with probability $1/N$.

After observing the system for a while, an attacker may assign some probabilities to each sender as being the originator of a message, based on the information the system is leaking, by means of traffic analysis, timing attacks, message length attacks or more sophisticated attacks.

For a given distribution of probabilities, the concept of entropy in information theory provides a measure of the information contained in that distribution [4]. We use entropy as a tool to calculate the degree of anonymity achieved by the users of a system towards a particular attacker. The entropy of the system after the attack is compared against the maximum entropy (for the same number of users). This way we get an idea of how much information the attacker has gained, or, in other words, we compare how distinguishable the sender is within the set of possible senders after the attack.

Lex X be the discrete random variable with probability mass function $p_i = Pr(X = i)$, where i represents each possible value that X may take. In this case, each i corresponds to an element of the anonymity set (a sender). We denote by $H(X)$ the entropy of the system after the attack has taken place. For each sender belonging to the senders set of size N, the attacker assigns a probability p_i. $H(X)$ can be calculated as:

$$H(X) = -\sum_{i=1}^{N} p_i \log_2(p_i) \ .$$

Let H_M be the maximum entropy of the system we want to measure, for the actual size of the anonymity set:

$$H_M = \log_2(N) \ ,$$

where N is the number of honest senders (size of the anonymity set).

The information the attacker has learned with the attack can be expressed as $H_M - H(X)$. We divide by H_M to normalize the value. We then define the **degree of anonymity** provided by the system as:

$$d = 1 - \frac{H_M - H(X)}{H_M} = \frac{H(X)}{H_M} \ .$$

For the particular case of one user we assume d to be zero.

This degree of anonymity provided by the system quantifies the amount of information the system is leaking. If in a particular system a user or a small group of users are shown as originators with a high probability with respect to the others, this system is not providing a high degree of anonymity.[2]

It follows immediately that $0 \leq d \leq 1$:

[2] On the other hand, note that any system with equiprobable distribution will provide a degree of anonymity of one, therefore a system with two senders will have $d = 1$ if both of them are assigned probability $1/2$. This is because the definition of anonymity we are using is independent of the number of senders.

Appeared in Proceedings of PET 2002, April 14-15, 2002, San Francisco,
In Hannes Federath (Ed.), *Designing Privacy Enhancing Technologies*, Lecture Notes in Computer Science, 2002.

- $d = 0$ when a user appears as being the originator of a message with probability 1.
- $d = 1$ when all users appear as being the originator with the same probability ($p_i = 1/N$).

4 Measuring the degree of anonymity provided by some systems

In this section we apply our proposed measurement model in order to analyze the degree of anonymity provided by some existing systems, in particular Crowds and Onion Routing.

4.1 A simple example: mix based email.

As a first example, let us consider the system shown in Fig. 1. Here we have a system that provides anonymous email with 10 potential senders, a mix network and a recipient. The attacker wants to find out which of the senders sent an email to this particular recipient. By means of timing attacks and traffic analysis, the attacker assigns a certain probability to each user as being the sender. The aim of this example is to give an idea on the values of the degree of anonymity for different distributions of probabilities.

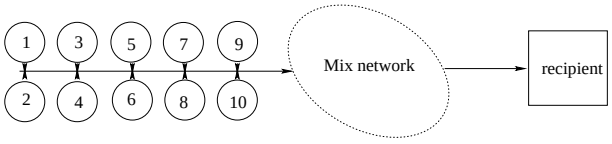

Fig. 1. A simple example of a mix based email system

Active attack. We first consider an active internal attacker who is able to control eight of the senders (that means that these eight users have to be excluded from the anonymity set). He is also able to perform traffic analysis in the whole mix network and assign probabilities to the two remaining senders. Let p be the probability assigned to *user 1* and $1 - p$ the probability assigned to *user 2*.

The distribution of probabilities is:

$$p_1 = p \ ; \quad p_2 = 1 - p \ ,$$

and the maximum entropy for two honest users is:

$$H_M = log_2(2) = 1 \ .$$

Appeared in Proceedings of PET 2002, April 14-15, 2002, San Francisco,
In Hannes Federath (Ed.), *Designing Privacy Enhancing Technologies*, Lecture Notes in Computer Science, 2002.

In Fig. 2a we show the variation of the degree of anonymity with respect to p. As we could expect from the definitions, we see that d reaches the maximum value ($d = 1$) when both users are equiprobable ($p = 1/2$). Indeed, in this case the attacker has not gained any information about which of the two active users is the real sender of the message by analyzing the traffic in the mix network. The minimum level ($d = 0$) is reached when the attacker can assign probability one to one of the users ($p = 0$ or $p = 1$).

This simple example can be useful to get an idea on the minimum degree of anonymity that is still adequate. Roughly, we suggest that the system should provide a degree $d \geq 0.8$. This corresponds to $p = 0.25$ for one user and $p = 0.75$ for the other. In the following examples, we will again look at the probability distributions that correspond to this value of the degree, in order to compare the different systems. Nevertheless, the *minimum acceptable degree* for a particular system may depend on the anonymity requirements for that system, and we believe that such a minimum cannot be suggested before intensively testing the model.

Passive attack. We now consider a passive global external attacker who is able to analyze the traffic in the whole system, but who does not control any of the entities (the anonymity set is, therefore, composed by 10 users). The maximum entropy for this system is:

$$H_M = log_2(10) \ .$$

The attacker comes to the following distribution:

$$p_i = \frac{p}{3} \ , \quad 1 \leq i \leq 3 \ ; \quad p_i = \frac{1-p}{7} \ , \quad 4 \leq i \leq 10 \ .$$

In this case we have two groups of users, one with three users and the other one with seven. Users belonging to the same group are seen by the attacker as having the same probability.

In Fig. 2b we can see the variation of d with the parameter p. The maximum degree $d = 1$ is achieved for the equiprobable distribution ($p = 0.3$). In this case d does not drop to zero because in the worst case, the attacker sees three users as possible senders with probability $p = 1/3$, and therefore he cannot identify a single user as the sender of the message. The reference value of $d = 0.8$ is reached when three of the users are assigned probability $p_i = 0.25$, and the remaining seven users are assigned probability $p_i = 0.036$.

4.2 Crowds

Overview of the system. Crowds [9] is designed to provide anonymity to users who want to access web pages. To achieve this goal, the designers introduce the notion of "blending into a crowd": users are grouped into a set, and they forward requests within this set before the request is sent to the web server. The web server cannot know from which member the request originated, since it gets the request from a random member of the crowd, that is forwarding the message

Appeared in Proceedings of PET 2002, April 14-15, 2002, San Francisco,
In Hannes Federath (Ed.), *Designing Privacy Enhancing Technologies*, Lecture Notes in Computer Science, 2002.

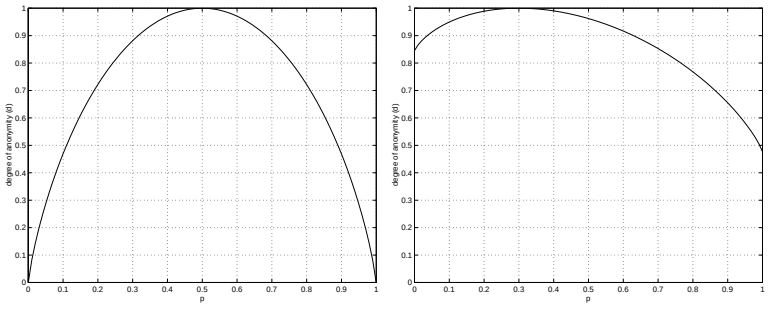

(a) (b)

Fig. 2. Degree of anonymity for a simple example

on behalf of the real originator. The users (members of the crowd) are called *jondos*.

The system works as follows: when a *jondo* wants to request a web page it sends the request to a second (randomly chosen) *jondo*. This *jondo* will, with probability p_f, forward the request to a third *jondo* (again, randomly chosen), and will, with probability $(1-p_f)$ submit it to the server. Each *jondo* in the path (except for the first one) chooses to forward or submit the request *independently* from the decisions of the predecessors in the path.

Communication between *jondos* is encrypted using symmetric techniques, and the final request to the server is sent in clear text. Every *jondo* can observe the contents of the message (and thus the address of the target server), but it cannot know whether the predecessor is the originator of the message or whether he is just forwarding a message received by another member.

Note that for this system the *mixes* are the *jondos*, and they do not have some of the expected characteristics. In particular, they do not make any effort to hide the correlation between incoming and outgoing messages.

Attacker. In this paper we calculate the degree of anonymity provided by Crowds with respect to collaborating crowd members, that is, a *set of corrupted jondos that collaborate in order to disclose the identity of the jondo that originated the request*. The assumptions made on the attacker are:

- *Internal*: The attacker controls some of the entities that are part of the system.
- *Passive*: The corrupted *jondos* can listen to communication. Although they have the ability to add or delete messages, they will not gain extra information about the identity of the originator by doing so.
- *Local*: We assume that the attacker controls a limited set of *jondos*, and he cannot perform any traffic analysis on the rest of the system.

Degree of anonymity. Figure 3 shows an example of a crowds system. In this

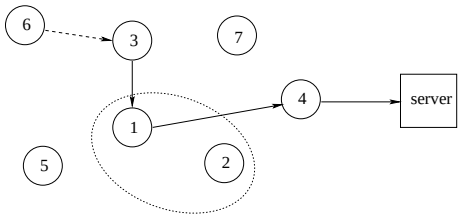

Fig. 3. Example of a Crowds system with 7 *jondos*

example the *jondos* 1 and 2 are controlled by the attacker, i.e., they are *collaborating crowd members*. A non-collaborating *jondo* creates a path that includes at least one corrupted *jondo*[3]. The attacker wants to know which of the non-collaborating *jondos* is the real originator of the message.

Generally, let N be the number of members of the crowd, C the number of collaborators, p_f the probability of forwarding and p_i the probability assigned by the attacker to the *jondo* i of having sent the message. The *jondos* under the control of the attacker can be excluded from the anonymity set. The maximum entropy H_M, taking into account that the size of the anonymity set is $N-C$, is equal to:

$$H_M = \log_2 (N - C) \ .$$

From [9] we know that, under this attack model, the probability assigned to the predecessor of the first collaborating *jondo* in the path (let this *jondo* be number *C+1*) equals:

$$p_{C+1} = \frac{N - p_f(N - C - 1)}{N} = 1 - p_f \frac{N - C - 1}{N} \ .$$

The probabilities assigned to the collaborating *jondos* remain zero, and assuming that the attacker does not have any extra information about the rest of non-collaborators, the probabilities assigned to those members are:

$$p_i = \frac{1 - p_{C+1}}{N - C - 1} = \frac{p_f}{N} \ , \quad C + 2 \leq i \leq N \ .$$

Therefore, the entropy of the system after the attack will be:

$$H(X) = \frac{N - p_f(N - C - 1)}{N} \log_2 \left[\frac{N}{N - p_f(N - C - 1)} \right] + p_f \frac{N - C - 1}{N} \log_2 \left[\frac{N}{p_f} \right] \ .$$

The degree of anonymity provided by this system is a function of N, C and p_f. In order to show the variation of d with respect to these three parameters

[3] If the path does not go through a collaborating *jondo* the attacker cannot get any information.

Appeared in Proceedings of PET 2002, April 14-15, 2002, San Francisco,
In Hannes Federath (Ed.), *Designing Privacy Enhancing Technologies*, Lecture Notes in Computer Science, 2002.

we chose $p_f = 0.5$ and $p_f = 0.75$, and $N = 5$ (Fig. 4a), $N = 20$ (Fig. 4b) and $N = 100$ (Fig. 4c). The degree d is represented in each figure as a function of the number of collaborating *jondos* C. The minimum value of C is 1 (if $C = 0$ there is no attacker), and the maximum value of C is $N - 1$ (if $C = N$ there is no user to attack). For the case $C = N - 1$ we obtain $d = 0$ because the collaborating *jondos* know that the real sender is the remaining non-collaborating *jondo*. We

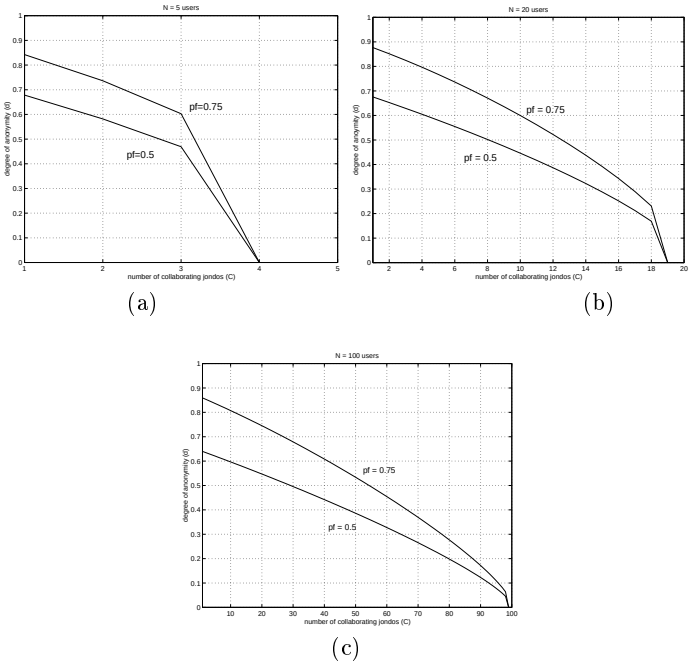

Fig. 4. Degree of anonymity for Crowds

can deduce from the figures that d decreases with the number of collaborating *jondos* and increases with p_f. The variation of d is very similar for systems with different number of users. Regarding the tolerated number of collaborating *jondos* to obtain $d \geq 0.8$, we observe that for $p_f = 0.5$ the system does not tolerate any corrupted *jondo*; for $p_f = 0.75$ the system tolerates: for $N = 5$ users, $C \leq 1$, for $N = 20$ users, $C \leq 4$, and for $N = 100$ users, $C \leq 11$.

In [9] a degree of anonymity is defined as $(1 - p_{sender})$, where p_{sender} is the probability assigned by the attacker to a particular user as being the sender. This measure gives an idea of the degree of anonymity provided by the system for a particular user, and it is complementary with the degree proposed in this paper. It is interesting to compare the results obtained by Reiter and Rubin in [9] with the ones obtained in this paper (for the same attack model): they consider that the worst acceptable case is the situation where one of the *jondos* is seen by the attacker as the sender with probability 1/2. Therefore, they come

Appeared in Proceedings of PET 2002, April 14-15, 2002, San Francisco,
In Hannes Federath (Ed.), *Designing Privacy Enhancing Technologies*, Lecture Notes in Computer Science, 2002.

to the conclusion that, for $p_f = 0.75$, the maximum number of collaborating *jondos* the system can tolerate is $C \leq N/3 - 1$. For the chosen examples we obtain: for $N = 5$ users, $C = 0$, for $N = 20$ users, $C \leq 5$, and for $N = 100$ users, $C \leq 32$.

Degree of anonymity from the point of view of the sender. We have calculated the degree of anonymity of a user who sends a message that goes through a corrupted *jondo*, but this only happens with probability C/N each time the message is forwarded to another *jondo*. We have to take into account that the first *jondo* always forwards the message to a randomly chosen *jondo* of the crowd, and subsequent *jondos* forward with probability p_f to another *jondo*, *independently from previous decisions*. The probability p_H of a message going only through honest *jondos* is:

$$p_H = \frac{N-C}{N}(1-p_f)\sum_{i=0}^{\infty}\left(\frac{N-C}{N}p_f\right)^i = 1 - \frac{C}{N - p_f(N-C)} \ .$$

If a message does not go through any collaborating *jondo*, the attacker will assign all honest senders the same probability, $p_i = 1/(N-C)$, and the degree of anonymity will be $d = 1$ (the maximum degree is achieved because the attacker cannot distinguish the sender from the rest of honest users). Some further discussion about the implications of this fact can be found in the Appendix A.

4.3 Onion Routing

Overview of the system. Onion Routing [8] is a solution for application-independent anonymous connections. The network consists of a number of *onion routers*. They have the functionality of ordinary routers, combined with mixing properties. Data is sent through a path of onion routers, which is determined by an *onion*.

An *onion* is a layered encrypted data structure, that is sent to an onion router. It defines the route of an anonymous connection. It contains the next hop information, key seed material for generating the symmetric keys that will be used by the onion router during the actual routing of the data, and an embedded onion that is sent to the next onion router.

The data is encrypted multiple times using the symmetric keys that were distributed to all the onion routers on the path. It is carried by small data cells containing the appropriate anonymous connection identifier. Each onion router removes/adds a layer of encryption (using the symmetric keys, generated from the key seed material in the onion) depending on the direction of the data (forwards/backwards).

Attack model. Several attack models have been described by Reed, Syverson and Goldschlag in [8]. In this example we consider an attacker who is able to narrow down the set of possible paths. The attacker obtains, as a result of the attack, a subset of the *anonymity set* that contains the possible senders. We do

not make any assumption on the attacker, but that he does not control any user of the system. We make abstraction of the attack, but, in order to illustrate the example, it could be carried out performing a *brute force attack*, starting from the recipient and following all the possible reverse paths to the senders. Another alternative is that the attacker controls some of the onion routers, and he is able to eliminate a group of users from the anonymity set.

Degree of anonymity. Figure 5 gives an example of an Onion Routing system. There are in total seven users in this system. We assume that the attacker

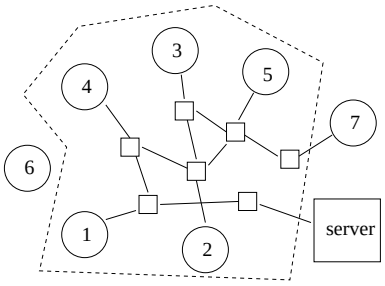

Fig. 5. Example of Onion Routing

managed to exclude users 6 and 7 from the set of possible senders.

Generally, let N be the size of the *anonymity set*; the maximum entropy for N users is:

$$H_M = \log_2(N) \ .$$

The attacker is able to obtain a subset of the *anonymity set* that contains the possible senders. The size of the subset is S ($1 \leq S \leq N$). We assume that the attacker cannot assign different probabilities to the users that belong to this subset:

$$p_i = \frac{1}{S} \ , \quad 1 \leq i \leq S \ ; \quad p_i = 0 \ , \quad S+1 \leq i \leq N \ .$$

Therefore, the entropy after the attack has taken place, and the degree of anonymity are:

$$H(X) = \log_2(S) \ , \quad d = \frac{H(X)}{H_M} = \frac{\log_2(S)}{\log_2(N)} \ .$$

Figure 6 shows the degree of anonymity with respect to S for $N = 5$, $N = 20$ and $N = 100$. Obviously, d increases with S, i.e., when the number of users that the attacker is able to exclude from the *anonymity set* decreases. In order to

Appeared in Proceedings of PET 2002, April 14-15, 2002, San Francisco,
In Hannes Federath (Ed.), *Designing Privacy Enhancing Technologies*, Lecture Notes in Computer Science, 2002.

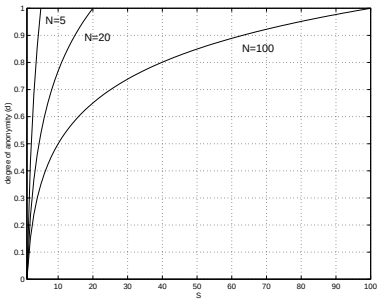

Fig. 6. Degree of anonymity for Onion Routing

obtain $d \geq 0.8$: for $N = 5$ users, we need $S \geq 3$; for $N = 20$ users, we need $S \geq 12$; and for $N = 100$ users, we need $S \geq 40$.

When comparing $N - S$ to the number of collaborating *jondos* C in the Crowds system, it seems that Onion Routing is much more tolerant against 'failing' users/ *jondos* than Crowds. This is because the remaining 'honest' users/*jondos* have equal probability (for this attack model) in the Onion Routing system, while in Crowds there is one *jondo* that has a higher probability than the others.

5 Conclusions and open problems

Several solutions for anonymous communication have been proposed and implemented in the past. However, the problem of how to measure the actual anonymity they provide, has not yet been studied thoroughly. We proposed a general measurement model to quantify the degree of anonymity provided by a system in particular attack circumstances. We applied our model to some existing solutions for anonymous communication. We suggested a intuitive value for the minimum degree of anonymity for a system to provide adequate anonymity. The model showed to be very useful for evaluating a system, and comparing different systems.

In the examples we have chosen, we calculate the degree for a particular message, and we do not take into account the behavior of the system over time. However, the attacker may gain useful information by observing the system for a longer time, and this fact is reflected in the distribution of probabilities. We could apply the model taking into account these changes in the probabilities, and we would obtain information on the evolution of the degree of anonymity with the time.

There are still some open problems. Our model is based on the probabilities an attacker assigns to users; finding this probability distribution in real situations is however not always easy.

Appeared in Proceedings of PET 2002, April 14-15, 2002, San Francisco,
In Hannes Federath (Ed.), *Designing Privacy Enhancing Technologies*, Lecture Notes in Computer Science, 2002.

It would be also interesting to take into account the *a priori* information the attacker may have, and use the model to see the amount of information he has gained with the attack.

The paper only focused on sender anonymity; recipient anonymity can be treated analogously; unlinkability between any sender and any recipient depends on the probability of finding a match.

Finally, the usefulness of our model should be more intensively tested; for example, it would be interesting to measure the effect of dummy traffic in the more advanced anonymous communication solutions, in order to find the right balance between performance and privacy.

Acknowledgments

Claudia Díaz is funded by a research grant of the K.U.Leuven. Joris Claessens and Stefaan Seys are funded by a research grant of the Institute for the Promotion and Innovation by Science and Technology in Flanders (IWT). This work was also partially supported by the IWT STWW project on Anonymity and Privacy in Electronic Services (APES), and by the Concerted Research Action (GOA) Mefisto-2000/06 of the Flemish Government.

A Extension of the model

In some systems we may get different distributions with a certain probability. For example, in Crowds, there are two cases: the message goes through a corrupted *jondo* with probability p_C, and it goes only through honest *jondos* with probability p_H, where:

$$p_C = \frac{C}{N - p_f(N - C)} \;\; ; \quad p_H = 1 - \frac{C}{N - p_f(N - C)} \;\; .$$

If we want to calculate the degree of anonymity offered by the system taking into account all possibilities, we may combine the obtained degrees as follows:

$$d = \sum_{j=1}^{K} p_j d_j \;\; ,$$

where d_j is the degree obtained under particular circumstances and p_j the probability of occurrence of such circumstances. K is the number of different possibilities.

The degree of anonymity becomes in this case a composite of the degrees obtained for the different cases.

B Alternative solution

It may be the case that, for a particular system, a requirement on the minimum acceptable degree of anonymity is formulated as *users should have at least a*

degree of anonymity equivalent to a system with M users and perfect indistinguishability.

In this case, we could compare the actual entropy of the system against the required one. We should then compare the obtained entropy, $H(X)$ and $log_2(M)$, instead of normalizing by the best the system can do with the number of current users. If $H(X)$ is bigger, then the system is above the minimum; if it is smaller, we may want to use some extra protection in the system, such as dummy traffic.

This might be useful to see if the system is meeting the requirements or not, and to launch an alarm in case the degree of anonymity is lower than the one defined as the *minimum*.

References

1. O. Berthold, H. Federrath and S. Köpsell. Web MIXes: A System for Anonymous and Unobservable Internet Access. In Hannes Federath (Ed.), *Designing Privacy Enhancing Technologies*, Lecture Notes in Computer Science, LNCS 2009, pp. 115–129, Springer-Verlag, 2001.
2. O. Berthold, A. Pfiztmann and R. Standtke. The Disavantages of Free MIX Routes and How to Overcome Them In Hannes Federath (Ed.), *Designing Privacy Enhancing Technologies*, Lecture Notes in Computer Science, LNCS 2009, pp. 30–45, Springer-Verlag, 2001.
3. D. Chaum. Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms. *Communications of the ACM*, Vol. 24, No. 2, pp. 84-88, 1981.
4. T. M. Cover and J. A. Thomas. Elements of Information Theory. *John Wiley & Sons, Inc.*, 1991. ISBN 0-471-06259-6.
5. W. Feller. An Introduction to Probability Theory and its Applications. *John Wiley & Sons, Inc.*, Third edition, 1968.
6. A. Pfitzmann and M. Köhntopp. Anonymity, Unobservability and Pseudonymity – A Proposal for Terminology. In Hannes Federath (Ed.), *Designing Privacy Enhancing Technologies*, Lecture Notes in Computer Science, LNCS 2009, pp. 1–9, Springer-Verlag, 2001.
7. J.-F. Raymond. Traffic Analysis: Protocols, Attacks, Design Issues, and Open Problems. In Hannes Federath (Ed.), *Designing Privacy Enhancing Technologies*, Lecture Notes in Computer Science, LNCS 2009, pp. 10–29, Springer-Verlag, 2001.
8. M. G. Reed, P. F. Syverson and D. Goldschlag. Anonymous Connections and Onion Routing. *IEEE Journal on Selected Areas in Communication*. Special Issue on Copyright and Privacy Protection, 1998.
9. M. K. Reiter and A. D. Rubin. Crowds: Anonymity for Web Transactions. *Communications of the ACM*, Vol. 42, No. 2, pp. 32-48, 1999.
10. A. Serjantov and G. Danezis. Towards an Information Theoretic Metric for Anonymity. In Hannes Federath (Ed.), *Designing Privacy Enhancing Technologies*, Lecture Notes in Computer Science, 2002.
11. C. E. Shannon. A Mathematical Theory Of Communication. Bell System Tech. J., 27:379–423; 623–656, 1948.
12. M. Wright, M. Adler, B. Levine and C. Shields. An Analysis of the Degradation of Anonymous Protocols. In Proceedings of *Network and Distributed System Security Symposium*, February 2002.

Appeared in Proceedings of PET 2002, April 14-15, 2002, San Francisco,
In Hannes Federath (Ed.), *Designing Privacy Enhancing Technologies*, Lecture Notes in Computer Science, 2002.