Information Theory and Anonymity

Claudia Díaz, Joris Claessens, Stefaan Seys, 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. There is an increasing concern about the protection of anonymity and privacy in electronic services. Many web sites store and process personal data of users in order to provide a better service. Users prefer not to give away personal information if this can be avoided. This paper introduces a model, based on Shannon's definition of entropy, with which the degree of anonymity of a set of users grouped in several subsets can be quantified in an objective way. This model is applied to a practical case in which users are associated to profiles, and the degree of anonymity achieved by those users, is measured under different conditions.

1 Introduction

In today's expanding on-line world, there is an increasing concern about the protection of anonymity and privacy in electronic services. Information theory [2] has proven to be a useful tool to measure the amount of information. In previous work [3], we proposed a model to quantify the degree of anonymity in an electronic system. We focused on connection anonymity, i.e., hiding the identities of source and destination during communication. In this paper, we introduce a similar model, also based on Shannon's definition of entropy [7], with which the degree of anonymity of a set of users grouped in several subsets can be quantified in an objective way. We here consider anonymity at the level of the data that is exchanged in a system. The model is in particularly suited for measuring the anonymity of users when they are associated with profiles (e.g., targeted advertising). If there are many different profiles, users will be easily distinguishable but less anonymous; if there are few different profiles, users cannot be targeted anymore on an individual basis, but are more anonymous. The model is intended to find a good compromise between anonymity and usefulness of such systems.

This paper is organized as follows: Sect. 2 describes the system model we consider; the actual measurement model is then proposed in Sect. 3. A practical case is studied in Sect. 4. Finally, our conclusions and future work are dicussed.

Appeared in Proceedings of the 23rd Symposium on Information Theory in the Benelux, May 29–31, 2002, Louvain la Neuve, Belgium

2 System model

We consider a number of users, who are distributed into groups. Each user is identified as a member of a group when he generates requests (e.g., requests for a web page). We explain the model below into more detail, and we set the notation.

Users. Let N be the number of users, u_1, ... , u_N.

Groups. The users are distributed into groups. All the users that belong to a group are indistinguishable. Let M be the number of groups, g_1, ... , g_M. We assume that the number of users, N, is much larger than the number of groups. Users may belong to different groups at different moments. Each group g_i contains N_i users. We can see that:

$$N = \sum_{i=1}^{M} N_i \ .$$

Requests. Each user u_j generates r_j requests. We call R the total number of requests produced by the set of users in a certain amount of time. To identify a particular request, we use the notation R_i, (R_1, ... , R_R):

$$R = \sum_{j=1}^{N} r_j \ .$$

Rg_k denotes the number of requests that belong to the same group g_k.

Connection level. We assume that all users are connected through a mix network [1] to achieve anonymity at the connection level. This mix network works for example like Onion Routing [6]. The attacker is not able to perform traffic analysis within the mix network, but he may be able to analyze the input and the output of the mix network. We consider real time applications. For this reason, the delay of the request cannot be too big. The attacker will know that a request that comes out of the mix network has been recently generated. In this paper, we focus on the anonymity at the data level, and do not analyze the level of anonymity at the connection level, i.e., we assume the mix network to perfectly hide the link between source and destination of a particular request.

Appeared in Proceedings of the 23rd Symposium on Information Theory in the Benelux, May 29–31, 2002, Louvain la Neuve, Belgium

3 Proposed measurement model

A measurement model has been proposed in previous work [3] to measure the anonymity at the connection level. Here, we introduce a model with which the degree of anonymity of a set of users grouped in several subsets, as defined above, can be quantified in an objective way.

3.1 Definition of anonymity

First of all, we should give a precise definition of *anonymity*. In this paper we adopt the definition given by Pfitzmann in [5]. Anonymity is *the state of being not identifiable within a set of subjects, the anonymity set.*

3.2 Definition of the degree of anonymity

According to the previous definition, in a system with N active users, the maximum degree of anonymity is achieved when an attacker sees all users equally probable as being the originator of a request. In our case, this situation is achieved when there is only one group that contains all the users, and the number of requests generated by each user is the same.

Therefore, in our model the degree of anonymity depends on the distribution of probabilities and not on the number of users. 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 number of active users should be large enough in comparison with the number of groups, in order to ensure that there are no groups that contain a small number of users. If a group contains a single active user, this user is no longer anonymous.

The proposed model compares the information obtained by the attacker after observing the system against the optimal situation from the anonymity point of view, in which all users seem to be equally probable as being the originator of the message, that is, in a system with N users, the situation where all users belong to the same group and make the same number of requests.

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 he has stored. For a given distribution of probabilities, the concept of entropy in information theory provides a measure of the information contained in that distribution. We will use entropy as a tool to calculate the degree of anonymity achieved by the users.

The entropy of the system after the attack will be compared against the maximum entropy (for the same number of users). In 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 will correspond to a user u_i. We denote by $H(X)$ the entropy of the system after the attack has taken place. For each user u_i, the attacker will assign 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 number of users:

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

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

The information the attacker has learned with the attack about a particular request can be calculated 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 for a particular request R_j as:

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

For the particular case of one possible sender we assume d_j to be zero. It follows immediately that $0 \leq d_j \leq 1$:

- $d_j = 0$ when a user appears as being the originator of a request with probability 1.
- $d_j = 1$ when all users appear as being the originator with the same probability.

3.3 Average degree of anonymity

The proposed model allows us to calculate the degree of anonymity obtained for each request. Given that during the attack R requests have been produced, we define the **average degree of anonymity** as:

$$d = \frac{\sum_{j=1}^{R} d_j}{R} \ .$$

This gives an accurate idea on the degree of anonymity provided by the system for request on average.

3.4 Attack models

The degree of anonymity depends on the probabilities of having sent a particular request that the attacker is able to assign to the users. The degree is therefore measured *with respect to* a particular attack: the results obtained 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 consider a very powerful attacker, who can monitor all communication lines of the system and knows the number of active users in the system and the number of groups. The attacker also knows the group of the user that generated a particular request, and the number of requests produced by every user (r_j). The attacker wants to find out the identity of the user that generated a particular request. If there are several users that belong to a group, the attacker is not able to distinguish which member of the group generated the request. The attacker uses all the available information to assign probabilities of being the originator of the request to all the users. The attacker can record all traffic information in the system, and then, for each request, the attacker analyzes the outputs of the system in a period of time which is the maximum delay of the network. He uses this information to assign to every user a probability of having produced a particular request. During an attack, we assume that the number of users in the system, N, is constant, and that the groups are static.

4 Practical case: targeted advertising

In this section we propose a system that provides a tool to allow targeted advertising (banners), while protecting the anonymity of the users.

4.1 Description of the system

There are three different entities in the model, which we describe below.

Banner server. The banner server serves banners to a set of web sites. Each time a user requests a web page within the domain of the banner server, the user contacts the banner server. The user sends his profile to the banner server, which decides which banner to send to the user, depending on the profile.

Web sites. The web sites have a contract with the banner server. When a user requests a page that contains a banner, the web site sends a link to the user, so he accesses the banner server to get the banner.

Users. The users surf the Internet, go to different web sites and request pages. Every time a user requests a page within the domain of our banner server, he will access the banner server (transparently to the user) to get the banner. The user is not identified by the banner server, instead, he sends the profile. This profile will be the same for all the users who have similar interests. He is anonymous but still gets banners that are related to his interests. Users access the Web through a mix network.

4.2 Attacks

We want to protect the user from the banner server, so we assume that the attacker controls the banner server and can also observe the traffic between the users and the mix network.

We consider two cases, in the first one we assume that the attacker, given a user, does not know the group of the user (cannot access the profile), and in the other we assume the attacker knows this information.

4.3 Attack 1

This attacker has access to the information stored in the banner server (relationship profile–web page), and can also see how many requests are generated by each user. The attacker cannot see the profile of each user, so he cannot know the group to which the user belongs.

For each request, the attacker wants to find the user who generated it, so he will take all the users who made a request in a period of time equal to the maximum delay of the network and calculate the probability for each user u_i as follows:

$$p_i = \frac{r_i}{R} \ .$$

Where R is the total number of requests produced in the period of time and r_i is the number of requests produced by user u_i.

To calculate the degree of anonymity obtained for the request R_j we apply the formula:

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

Where H_M is the maximum entropy for a number of users equal to the total number of active users and $H(X)$ is the entropy calculated for the distribution obtained with the p_i.

In this case, the users who are more active just before the request arrives to the banner server appear more likely than the others.

To calculate the average degree of anonymity provided by the system, we calculate d_j for each R_j and then compute the average (d).

4.4 Attack 2

In this case the attacker also knows the group (profile) of each user. Once the attacker gets a request, he will only look at the active users that belong to a particular group.

Let us assume that the request the attacker is analyzing was generated by a user belonging to group g_k. This group contains N_k users who made a request during the attack time.

The attacker will calculate the distribution of probabilities for each user u_i that belongs to g_k as:

$$p_i = \frac{r_i}{Rg_k} \ ,$$

and p_i is zero for the users that belong to other groups. Rg_k denotes the number of requests that arrived close to the request we want to attack and that have associated the group (profile) g_k.

The rest of the analysis is analogous to the previous case. Note that we still compare the obtained entropy with the optimal case (H_M is the same), so the degree of anonymity we will obtain in this case is much less than in the previous one.

5 Conclusions and future work

We proposed a general measurement model to quantify the degree of anonymity provided by a system in particular attack circumstances. We prove the usefulness of information theory in this field of research. This paper provides a starting point to combine customized services (targeted advertising) and anonymity, and proposes a model to measure the actual level of anonymity we can achieve.

In the future, we plan to get some numeric results in order to find a good tradeoff between anonymity and number of profiles (groups). We will also look at other attack models and see how the system behaves under different attack circumstances.

Acknowledgements

This work was supported in part 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. Joris Claessens and Stefaan Seys are funded by a research grant of the Institute for the Promotion of Innovation by Science and Technology in Flanders (IWT).

References

1. D. Chaum. Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms. *Communications of the ACM*, Vol. 24, No. 2, pp. 84-88, 1981.
2. T. M. Cover and J. A. Thomas. Elements of Information Theory. *John Wiley & Sons, Inc.*, 1991. ISBN 0-471-06259-6.
3. C. Diaz, S. Seys, J. Claessens and B. Preneel Towards Measuring Anonymity In Hannes Federath (Ed.), *Designing Privacy Enhancing Technologies*, Lecture Notes in Computer Science, Springer-Verlag, 2002.
4. W. Feller. An Introduction to Probability Theory and its Applications. *John Wiley & Sons, Inc.*, Third edition, 1968.
5. 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.
6. 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.
7. C. E. Shannon. A Mathematical Theory Of Communication. Bell System Tech. J., 27:379–423; 623–656, 1948.

Appeared in Proceedings of the 23rd Symposium on Information Theory in the Benelux, May 29–31, 2002, Louvain la Neuve, Belgium