

Modelling Unlinkability

Sandra Steinbrecher¹ and Stefan Köpsell²

¹ Freie Universität Berlin, Institut für Informatik, Takustr. 9, D-14195 Berlin, Germany, <steinbrecher@acm.org>

² Technische Universität Dresden, Fakultät Informatik, D-01062 Dresden, Germany, <sk13@inf.tu-dresden.de>

Abstract. While there have been made several proposals to define and measure anonymity (e.g., with information theory, formal languages and logics) unlinkability has not been modelled generally and formally. In contrast to anonymity unlinkability is not restricted to persons. In fact the unlinkability of arbitrary items can be measured. In this paper we try to formalise the notion of unlinkability, give a refinement of anonymity definitions based on this formalisation and show the impact of unlinkability on anonymity. We choose information theory as a method to describe unlinkability because it allows an easy probabilistic description. As an illustration for our formalisation we describe its meaning for communication systems.

1 Introduction

Every human being sometimes has the desire to act anonymously. Outreach clinics and doctors are visited by many human beings but in some cases visitors do not want others to get to know about their visit. It is quite obvious that someone visiting a doctor or an outreach clinic might do this for a limited number of reasons. For a doctor a patient's reason is linkable to his visitor while for outsiders it should be unlinkable. Obviously a visitor remains anonymous against outsiders regarding a specific visit's reason if his visit or the reason for his visit is unlinkable to him. Sometimes even the visit might indicate the reason, and if only the reason but not the visit is unlinkable to a human being his anonymity is endangered.

Naturally everyone only can be anonymous within a group of human beings that might be in the same situation, especially might have executed a specific action (e.g., visited a doctor).

In contrast to anonymity the notion of (un)linkability is not restricted to human beings and their actions, actions also might be linkable to each other or not. This might endanger a human being's anonymity. One specific action might be unlinkable to a human being but a succession of actions might only be executed by a specific human being and so each single action becomes linkable to the human being. When Clayton et al. studied technical attacks on an electronic student dating service [7] they found out that none of these possible attacks had been executed but some users tried to make 'social' attacks: They asked others

for some of their habits or actions. With only a few of these pieces of information linkable to each other it was quite easy to break a user's anonymity.

If a user chooses an unfavourable anonymity group even the links between (some of) his actions might indicate that these actions are (probably) his.

If there is a group of users and a number of actions executed by these users but the concrete links between users and actions are unknown to an attacker the exclusion of users from the group by linking all of his actions to him might reduce the other users' anonymity regarding the remaining actions.

The electronic society that has been built during the recent years gives many human beings a fallacious feeling of acting anonymously. But it becomes even more difficult to act as anonymously as in the society of the real world. Linkable information about a human being that might decrease someone's anonymity can often be collected quite easily in the electronic world.

While someone visiting shops might be anonymous in the real world he might not in the electronic world. For example in the latter he might use non-anonymous payment methods or non-anonymous web surfing. So human beings might become afraid of becoming 'transparent beings' and want to measure (or even better to determine themselves) the degree of anonymity they have in certain situations.

Recently there have been made several attempts to define and formalise the notion of anonymity and unlinkability. Most of the models for anonymity are only formulated for communication scenarios. We give an overview of previous approaches and extend the information theoretic approach to arbitrary scenarios in section 2.

To measure anonymity in real world situations it is necessary to measure the linkability between arbitrary items (e.g., actions, pieces of information, and human beings). The notion of unlinkability and untraceability is well-known in electronic commerce. In section 3 we give an overview of the notion underlying known concepts in electronic commerce and more general scenarios. Based on a general notion for unlinkability we present a formalisation of (un)linkability of arbitrary items and related attacker goals to break unlinkability.

Finally in section 4 is studied which influence these definitions have on a more general measurement of anonymity. In contrast to previous approaches our definitions are not restricted to one specific action but consider a set of actions linkable to a set of actors.

Our formalisation of unlinkability is illustrated by its application on communication systems in several examples. In these examples we assume communication systems to be systems with a set of users who may execute two actions: They may send or receive messages within the system. The users make use of anonymising services (e.g., Anonymizer [1], Crowds [15], Onion-Routing [14], Web mixes [2]) to reach sender and recipient anonymity as well as unlinkability of messages and users. If they do not use such a service users and messages become linkable. We abstract from the internal structure of concrete anonymising services but concentrate on the formalisation of the unlinkability and anonymity levels they are able to provide. We assume every message to be sent/received by

exactly one user. In real world scenarios users might send or receive messages with the same content (for example in the case of web surfing), but we assume these messages to be still technically distinguishable by their internal structure. We further neglect that in real world scenarios one human being might act under the names of multiple users. Every user name involved in the system will be counted as one user and every message sent by a user will be counted as one message.

2 Anonymity

A subject only can be anonymous within a group of other subjects. In [12] the following suggestion to standardise the definition of anonymity is given:

'Anonymity is the state of being not identifiable within a set of subjects, the anonymity set.'

In real world scenarios a subject's anonymity usually is related to an action. Then the anonymity set is formed by all actors who might have executed the action. The notion further given to measure the anonymity of a subject within such a set is that 'anonymity is the stronger, the larger the respective anonymity set is and the more evenly distributed' the action's execution 'of the subjects within that set is.' I.e., not only the size of the respective anonymity set determines the anonymity of a certain subject but also how likely a subject of the anonymity set might have executed the action.

Usually subjects cannot have the same anonymity against every possible participant and outsider who might be an attacker on the subject's anonymity. Depending on the attacker's knowledge the above set of possible subjects and the likelihood with which they have caused an action varies. For a specific attacker's view anonymity only can decrease. Thus the definition of anonymity in [12] is an analog to the definition of 'perfect secrecy' by Claude E. Shannon [18] as the authors indicate.

There have been made several proposals to describe anonymity with formal languages and logics. Syverson and Stubblebine describe anonymity properties in epistemic language based on group principals [21]. Their description includes the information that should be protected and the nature of the protection (degree of anonymity). This approach is demonstrated with the simple example of an anonymous proxy [1] which removes identifying information of the person requesting a website through it. In [16] a process algebraic formalisation in the modelling language CSP is given. This approach is illustrated with the example of Dining Cryptographers, the DC-Net [4]. Both papers follow the possibilistic approach, that is both metrics only consider the size of the anonymity set not the probability distribution on it.

2.1 Anonymity in communication systems

For communication systems three types of anonymity can be distinguished [12]: 'Sender/recipient anonymity as the properties that a particular message is not

linkable to any sender/recipient and that to a particular sender/recipient, no message is linkable.’ Relationship anonymity as the property that it is unlinkable who communicates with whom.

In open environments like the Internet a user is a member of the anonymity set if the probability that he initiated the action is non-zero [11]. But the size of the anonymity set is not sufficient to measure a user’s anonymity as already outlined.

Reiter and Rubin [15] introduce a degree of anonymity which they view as ‘informal continuum’. Within this continuum they define six degrees that might be reached with ‘absolute privacy’ as best and ‘provably exposed’ as worst case. They use these degrees to describe the anonymity their anonymising service Crowds provides. Shmatikov formalises their model by describing Crowds with Markov Chains and expressing the anonymity degrees as temporal probabilistic logic formulas [19].

Based on a mathematical abstraction describing the partial knowledge of a function (the so-called function view [10]) Hughes and Shmatikov develop a modular approach to specify anonymity properties. In this model equivalence relations are used to describe an attacker’s inability to distinguish between system configurations (observational equivalence). This approach can be used independent of the underlying algebra or logic. Unfortunately in this approach probabilism is not included.

Information theoretic models can help to precise the above notion of ‘the more evenly distributed’ by assigning probability distributions to anonymity sets [9, 17, 8]. These models compare the optimal situation (where every subject in the anonymity set might have executed the action with the same probability) with the situation where the subjects might be assigned different probabilities because of additional information.

In [9, 17] only the connection level is considered. In particular they analyse mix-based systems consisting of senders and recipients of messages, mixes used to send/receive these messages anonymously and possible attackers (especially Crowds [15] and Onion-Routing [14] in [9]).

The anonymity on the data level of a communication system is studied in [8]. Their scenario is web surfing of users who are grouped in subsets with different profiles. Every user in a group has the same profile, so when visiting a web site and using this profile this group is his anonymity set. The user remains anonymous but he profits from getting services related to his profile as well as the server profits from placing advertisement fitting the profile.

2.2 Anonymity for arbitrary actions

Because we want to develop and study a general model of unlinkability between arbitrary items (not restricted to communication systems) we unify and extend the definitions presented in [9, 17, 8] slightly in this section.

Let A be a non-empty set of actions of arbitrary size and $U = \{u_1, \dots, u_n\}$ be a set of subjects (the anonymity set regarding a specific action $a \in A$) of size n . Every subject $u_i \in U$ with $i \in \{1, \dots, n\}$ executes a with probability $p_i > 0$.

Example 1. In communication systems the set A is defined as $A = \{\text{sending, receiving}\} \times \{\text{messages}\}_{i \in I}$ with I an index set to enumerate the number of possible messages. In the case of web surfing it simply holds $A = \{\text{requesting}\} \times \{\text{website}_i\}_{i \in I}$. According to [11] in open communication systems like the Internet only subjects whose probability that they have executed the action is non-zero are members of the anonymity set U .

Ideally before the execution of action a every u_i will execute it with the a priori probability $\frac{1}{n}$ for a possible attacker's view on the system. This is the basis against which the a posteriori probabilities the attacker assigns to the users is compared in [9, 17, 8]. The attacker model depends on the concrete application and its requirements. Attackers might get the opportunity to perform several attacks during the execution of the action by which they might get additional information which helps to change the probability distribution on the anonymity set. On the connection level possible attacks are traffic analysis or timing attacks.

For a random variable X let $p_i = P_a(X = u_i)$ denote the attacker's a posteriori probability that given an action a , X takes the value u_i (or u_i executed the action a). Naturally $\sum_{i=1}^n p_i = 1$.

Entropy can be used as a measure to describe the degree of anonymity the system provides against a specific attacker. The attacker's a posteriori entropy is

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

Serjantov and Danezis [17] define the a posteriori entropy to be the effective size of the anonymity probability distribution (p_1, \dots, p_n) .

Obviously the maximum entropy is

$$\max(H(X)) = \log_2(n).$$

The information the attacker has learned is $(\max(H(X)) - H(X))$. Diaz et al. normalise this information [9] and define

Definition 1 (Degree of anonymity). *The degree of anonymity a system provides is*

$$d(U) := 1 - \frac{\max(H(X)) - H(X)}{\max(H(X))} = \frac{H(X)}{\max(H(X))}. \quad (1)$$

The normalisation has the effect that only the probability distribution not the size of the anonymity set is measured in the degree of anonymity. According to definition 1 both an anonymity set U_i with $i > 0$ subjects and a posteriori probabilities $p_1 = \dots = p_i = \frac{1}{i}$ and an anonymity set U_j with $j > i$ subjects and a posteriori probabilities $q_1 = \dots = q_i = \frac{1}{j}$ have degree $d(U_i) = d(U_j) = 1$.

The advantage of the normalisation is the the finite range $[0, 1]$ the degree lies within. The degree's maximum/minimum is reached if

$$d(U) = 0 \quad \Leftrightarrow \quad \exists i \in \{1, \dots, n\} : p_i = 1,$$

$$d(U) = 1 \quad \Leftrightarrow \quad \forall i \in \{1, \dots, n\} : p_i = \frac{1}{n}.$$

Note both degree and size of an anonymity set have to be given to describe the anonymity a system provides. An alternative is the definition of effective size of an anonymity set [17] above.

3 Unlinkability

The notion of anonymity (regarding a specific action) is usually restricted to users, while the notion of unlinkability is applicable to arbitrary items within a given system. For example in communication systems a sender of a message might not be linkable to that message but two messages sent by the same user might be linkable to each other. In [12] the following definition is given:

'Unlinkability of two or more items (e.g., subjects, messages, events, actions,...) means that within this system, these items are no more and no less related than they are related concerning the a priori knowledge.'

Or to say it inversely, items are linkable if they are more or less related than they are without any knowledge of the system. With full knowledge of the system, items within the system are either related to each other or not. Note this description of linkability does describe the inverse of unlinkability but not exactly its usual notion because it includes 'less related'.

An attacker on unlinkability of items within the system a priori knows the items within the system while his knowledge about their relation depends on the concrete system and the concrete attacker. Ideally his view of the system only contains the items.

After the attacker got time to observe/influence the system his knowledge might have increased. A passive attacker only observes the system. Whether he also has the opportunity to become an active attacker and execute several types of attacks influencing the system depends on the concrete system.

The above notion of unlinkability implies that the attacker is successful if his a posteriori probability that items are related has increased or decreased in comparison to his a priori probability. This means not only related items should be protected against detecting this but also unrelated items should a posteriori be related just as much as a priori.

Note the notion of linkability used for electronic payment systems is slightly less restrictive:

'The privacy requirement for the users is that payments made by users should not be linkable (informally, linkability means that the a posteriori probability of matching is nonnegligibly greater than the a priori probability) to withdrawals, even when banks cooperate with all the shops (untraceability). Untraceability guarantees that users remain anonymous, since their identity is only linked to withdrawals.' [3].

Known anonymous cash systems follow this notion.

Digital pseudonyms introduced by Chaum [6] guarantee unlinkability of their use to the corresponding user to make him untraceable. But all transactions

executed under the same pseudonym are linkable to each other. If users want to use different pseudonyms for different purposes he should be the only one who is able to link this pseudonyms. The use of credentials [5] enables him to transform statements made about one of his pseudonyms to statements about another one of his pseudonyms while the pseudonyms are still unlinkable to each other for everyone except himself.

In [20] a protocol for unlinkable serial transactions usable in electronic commerce is presented. The tokens (or credentials) used in the protocol fulfill the requirements of users and vendors: both fraud (sharing or abusing tokens) and unlinkability of users' transactions are guaranteed.

In this section we give a formalisation of the above notion of unlinkability. We start with a simple system model for unlinkability within one set in 3.1 and then extend this model to a model for unlinkability between sets which tries to meet real world conditions slightly better in section 3.2. Section 3.3 gives an overview of possible attacker models. If an attacker only learns the numbers of linkable items within a set his a posteriori probabilities of unlinkability will have increased in comparison to his a priori probabilities as we will finally show in section 3.4.

3.1 Unlinkability within one set

Let $A = \{a_1, \dots, a_n\}$ be the set of items within a given system. For someone with full knowledge of the system some items of this set are related while others are not. We consider a notion of 'is related' that forms an equivalence relation $\sim_{r(A)}$ on the set A . Then by this relation A is split in l ($1 \leq l \leq n$) equivalence classes A_1, \dots, A_l with $\forall i, j \in \{1, \dots, l\}, i \neq j: A_i \cap A_j = \emptyset$ and $A_1 \cup \dots \cup A_l = A$. Items are related to each other iff they belong to the same equivalence class.

Example 2 (Communication system). A could be a set of messages sent. All messages sent by the same sender are related to each other for him but should not for an attacker. But not all relations on A are equivalence relations. Obviously the relation 'sent by the same sender' is one. But the relation 'not sent by the same sender' is no equivalence relation because this relation is neither reflexive nor transitive.

In the following we use this equivalence relation $\sim_{r(A)}$ instead of the notion 'is related' to describe unlinkability of items. An attacker on unlinkability of items within one set knows A . A priori he should not know the structure of $\sim_{r(A)}$ but by observing and attacking the system he might learn more about it.

The following example shows the notion 'a priori' here is slightly different to real world scenarios:

Example 3 (Communication system). By knowing A a priori the attacker even has an advantage in comparison to real world scenarios where the messages that will be sent usually are not known to an attacker beforehand. But by assuming this knowledge the difference between the ideal situation (the attacker learns nothing) and the imperfect situation (the attacker learns something) can be

measured more easily. This assumption is similar to the assumption that the set of possible senders a priori is known to an attacker in open environments. But in real world scenarios he will not have learned the set of senders before they have sent their messages.

Unlinkability of two items within one set For a random variable X let $P(a_i \sim_{r(A)} a_j) := P(X = (a_i \sim_{r(A)} a_j))$ denote the attacker's a posteriori probability that given two items a_i and a_j , X takes the value $(a_i \sim_{r(A)} a_j)$ (or a_i and a_j are related). And $P(a_i \not\sim_{r(A)} a_j)$ denotes the analog probability that a_i and a_j are not related. Quite clearly it holds:

$$P(a_i \sim_{r(A)} a_j) + P(a_i \not\sim_{r(A)} a_j) = 1 \quad \forall a_i, a_j \in A. \quad (2)$$

As for the measurement of anonymity we use the attacker's entropy to measure two items' unlinkability. Let $H(i, j) := H(X)$.

Definition 2 (Degree of unlinkability). *The degree of (i, j) -unlinkability $d(i, j)$ describing the unlinkability of two items $a_i, a_j \in A$ a system provides is*

$$\begin{aligned} d(i, j) &:= H(i, j) \\ &= -P(a_i \sim_{r(A)} a_j) \cdot \log_2(P(a_i \sim_{r(A)} a_j)) \\ &\quad - P(a_i \not\sim_{r(A)} a_j) \cdot \log_2(P(a_i \not\sim_{r(A)} a_j)). \end{aligned}$$

Obviously it holds $0 \leq d(i, j) \leq 1$ and the minimum/maximum is reached if

$$\begin{aligned} d(i, j) = 0 &\Leftrightarrow (P(a_i \sim_{r(A)} a_j) = 1 \quad \vee \quad P(a_i \sim_{r(A)} a_j) = 0) \\ d(i, j) = 1 &\Leftrightarrow P(a_i \sim_{r(A)} a_j) = P(a_i \not\sim_{r(A)} a_j) = \frac{1}{2}. \end{aligned}$$

The latter (the maximum of unlinkability) is the worst case for an attacker: From his view $\forall a_i, a_j \in A$ ($i \neq j$) the probability that this pair of items is related is equal to the probability it is not.

Unlinkability of arbitrary many items Let $2 < k \leq n$ and $\{a_{i_1}, \dots, a_{i_k}\}$ be a subset of the set A . Then we define $\sim_{r(\{a_{i_1}, \dots, a_{i_k}\})}$ to be an equivalence relation on $\{a_{i_1}, \dots, a_{i_k}\}$. By this relation $\{a_{i_1}, \dots, a_{i_k}\}$ is split in equivalence classes. For a random variable X let

$$P\left(\left(\sim_{r(A)} \mid_{\{a_{i_1}, \dots, a_{i_k}\}}\right) = \left(\sim_{r(A)}\right)\right) := P\left(X = \left(\left(\sim_{r(A)} \mid_{\{a_{i_1}, \dots, a_{i_k}\}}\right) = \left(\sim_{r(A)}\right)\right)\right).$$

denote the attacker's a posteriori probability that the equivalence relation $\sim_{r(A)}$ restricted to $\{a_{i_1}, \dots, a_{i_k}\}$ is the same relation as $\sim_{r(\{a_{i_1}, \dots, a_{i_k}\})}$. This describes the probability that the distribution of the elements a_{i_1}, \dots, a_{i_k} on equivalence classes in $\{a_{i_1}, \dots, a_{i_k}\}$ is the same as in A . Items that are unlinkable in $\{a_{i_1}, \dots, a_{i_k}\}$ are unlinkable in A as well.

Figure 1 illustrates this for the small example of $|A| = 11$ with $|A_1| = 3$, $|A_2| = 5$, $|A_3| = 2$, $|A_4| = 1$, the equivalence relation 'having the same filling' and a subset containing 2 items from A_1 and 4 items from A_2 . Items having the same filling in A have so in its subset.




Fig. 1. Example: Same distribution on equivalence classes in A and its subset

Let I_k be an index set enumerating all possible equivalence relations on $\{a_{i_1}, \dots, a_{i_k}\}$. It holds $|I_k| = 2^{k-1}$. The sum over I_k of the above probabilities has to add up to 1 as it does for two items in formula (2):

$$\sum_{j \in I_k} P\left((\sim_{r_j(A)} |_{\{a_{i_1}, \dots, a_{i_k}\}}) = (\sim_{r(A)})\right) = 1. \quad (3)$$

Let $H(i_1, \dots, i_k) := H(X)$. We generalise the degree of unlinkability for two items to arbitrary items:

Definition 3 (Degree of unlinkability). *Let $2 < k \leq n$. The degree of (i_1, \dots, i_k) -unlinkability $d(i_1, \dots, i_k)$ describing the unlinkability of k items $a_{i_1}, \dots, a_{i_k} \in A$ a system provides is*

$$\begin{aligned} d(i_1, \dots, i_k) &:= H(i_1, \dots, i_k) \\ &= - \sum_{j \in I_k} \frac{1}{|I_k|} \left[P\left((\sim_{r_j(A)} |_{\{a_{i_1}, \dots, a_{i_k}\}}) = (\sim_{r(A)})\right) \right. \\ &\quad \left. \cdot \log_2 \left(P\left((\sim_{r_j(A)} |_{\{a_{i_1}, \dots, a_{i_k}\}}) = (\sim_{r(A)})\right) \right) \right]. \end{aligned}$$

Obviously it holds $0 \leq d(i_1, \dots, i_k) \leq 1$ and the minimum/maximum is reached if

$$d(i_1, \dots, i_k) = 0 \quad \Leftrightarrow \quad \exists j \in I_k : P\left((\sim_{r_j(A)} |_{\{a_{i_1}, \dots, a_{i_k}\}}) = (\sim_{r(A)})\right) = 1,$$

$$d(i_1, \dots, i_k) = 1 \quad \Leftrightarrow \quad \forall j \in I_k : P\left((\sim_{r_j(A)} |_{\{a_{i_1}, \dots, a_{i_k}\}}) = (\sim_{r(A)})\right) = \frac{1}{|I_k|}$$

For someone with full knowledge of the system, the set A is split in its equivalence classes uniquely, so is every $\{a_{i_1}, \dots, a_{i_k}\}$. It holds $d(1, \dots, n) = 0$.

3.2 Unlinkability between sets

The example of communication systems shows that the model of unlinkability within one set might not be sufficient to describe communication systems. It does not include the important (un)linkability between specific senders and specific messages. In section 4 we will outline how this model will help to refine anonymity. Now we extend the set of items by the set of all items sending messages, usually the set of senders. Then the set of items consists of items sending messages and being messages. But no item can send and be a message. Example 2 can be refined as follows:

Example 4 (Communication system). Let $A = A_s \cup A_m$ with A_s the set of senders and A_m the set of messages within the system. Then the relation 'being sent by the same sender or being this sender' forms an equivalence relation on A with equivalence classes consisting of one item (a sender) of A_s and arbitrary many items of A_m .

Whenever A is composed of sets with different types of items it seems to make sense to extend the model of unlinkability within one set to a model for unlinkability between sets.

Let both $A = \{a_1, \dots, a_k\}$ be a set of items (e.g. actions) and $U = \{u_1, \dots, u_n\}$ (e.g. users) within a given system. For someone with full knowledge of the system every item in U is related to at least one item in set A and every item in A is related to exactly one item in U . It follows that $|A| \geq |U|$.

The notion 'is related' forms a relation $\sim_{r(U,A)}$ between these sets that can be described as an enumeration of pairs.

Concurrently through $\sim_{r(U,A)}$ on A an equivalence relation $\sim_{r(A)}$ is defined as 'is related to the same item in U '. This equivalence relation has the same features than the one described in section 3.1. By this relation A is split in l ($1 \leq l \leq n$) equivalence classes A_1, \dots, A_l with $A_i \cap A_j = \emptyset \forall i, j \in \{1, \dots, l\}$, $i \neq j$ and $A_1 \cup \dots \cup A_l = A$. Items are related to each other iff they belong to the same equivalence class.

Figure 2 illustrates this for the small example of A a set of boxes with $|A| = 11$ and U a set of circles with $|U| = 4$. By 'has the same filling' a relation between A and U is defined as well as it defines an equivalence relation on A .

In this model Example 4 can be refined as follows:

Example 5 (Communication system). Let $U = A_s$ and $A = A_m$. If two items $a \in A$ and $u \in U$ are related to each other by $\sim_{r(U,A)}$ then u sent a . If $u \sim_{r(U,A)} a_1$ and $u \sim_{r(U,A)} a_2$ for $a_1, a_2 \in A$ it holds $a_1 \sim_{r(U,A)} a_2$ and these messages are sent by the same sender u and related to each other within A by $\sim_{r(A)}$.

An attacker on unlinkability between sets knows A and U . A priori he should not know the structures of $\sim_{r(U,A)}$ and $\sim_{r(A)}$ but by observing and attacking the system he might learn more about it.




Fig. 2. Example:

Unlinkability of two items Accordingly to formula (2) it holds

$$P(u_i \sim_{r(U,A)} a_j) + P(u_i \not\sim_{r(U,A)} a_j) = 1 \quad \forall u_i \in U, a_j \in A.$$

As within one set we define the degree of (u_i, a_j) -unlinkability as

$$d(u_i, a_j) = H(u_i, a_j).$$

Optimally $\forall (u_i, a_j) \in U \times A$ the probability that this pair is related to each other is equal to the probability it is not.

The above definitions for unlinkability between two sets can be extended to definitions for unlinkability between arbitrary many sets.

3.3 Attacks on Unlinkability

So far we gave only lifeless definitions of unlinkability without exact consideration of attacker goals. These will be given in this section because different attackers might have quite different goals. If an attacker wants to observe a specific victim his attacks might not violate other members of the victim's anonymity set. For example an insurance company trying to find out its (potential) customer's state of health might not collect data about other user's because storing huge amounts of useless data might be expensive.

Security considerations usually distinguish between existential and selective break of a system. We adopt these to the notion of unlinkability introduced:

1. **Existential break:** There exist any two items for which the attacker's a priori probability that they are related to each other is unequal to the a posteriori probability.
2. **Selective break:** The attacker is allowed to choose the items which unlinkability should in- or decrease.
 - (a) **Chosen subset of items:** The attacker may choose a subset of at least two items. For these items his a priori probability that they are related to each other is unequal to the a posteriori probability.
 - (b) **Chosen Item:** The attacker chooses one item. For this item there exist other items for which the a posteriori probabilities they are related to this specific item are unequal to the corresponding a priori probabilities.

The worst case for unlinkability within one set is that the chosen subset equals A and all a posteriori probabilities either are 0 or 1 in the selective break. For unlinkability between sets accordingly the subset would be $U \cup A$. In authentication or encryption systems existential breaks sometimes are neglected because the attacker success might be no problem for real world applications, e.g., a senseless message with a correct signature does not endanger the system's security. In systems guaranteeing unlinkability linkability between items not selected by the attacker might influence the linkability of items he has selected. In [9, 17] several examples are given where anonymity of a specific item is decreased because of this effect. Attackers on unlinkability typically reach their goal by excluding other items to be linkable to the items they are interested in.

3.4 The relation guaranteeing unlinkability

While anonymity regarding a specific action depends on the probability distribution on the anonymity set which a priori is uniform, (un)linkability depends on the equivalence classes induced by $\sim_{r(A)}$ or $\sim_{r(A)}$ on the set A .

The attacker's knowledge about the structure of the relation $\sim_{r(A)}$ on the given set A of items influence his probability distribution of unlinkability. For instance if the sizes of the equivalence classes are publicly known then optimal a priori probabilities cannot be reached as a posteriori probabilities.

Example 6 (Communication system). If an attacker gets to know how many messages every sender sends in the scenario of Example 1 he knows the size of every equivalence class, i.e. $\forall i \in \{1, \dots, l\} |A_i|$ becomes known to the attacker.

The structure of the equivalence classes has an impact on the a posteriori probabilities even in an existential break. The probability that t items a_{i_1}, \dots, a_{i_t} chosen arbitrarily from A lie in the same equivalence class A_v with $v \in \{1, \dots, l\}$ is

$$P(a_{i_1} \sim_{r(A)} \dots \sim_{r(A)} a_{i_t}) = \frac{\sum_{v=1}^l \binom{|A_v|}{t}}{\binom{n}{t}}.$$

with $\binom{n}{t} := 0$ for $n < t$.

For the special case $t = 2$ this leads to

$$P(a_{i_1} \sim_{r(A)} a_{i_2}) = \frac{(\sum_{v=1}^l |A_v|^2) - n}{n^2 - n}$$

Accordingly a_{i_1} and a_{i_2} lie in different equivalence classes with probability

$$\begin{aligned} P(a_i \not\sim_{r(A)} a_j) &= 1 - \frac{(\sum_{v=1}^l |A_v|^2) - n}{n^2 - n} \\ &= \frac{n^2 - \sum_{v=1}^l |A_v|^2}{n^2 - n} \end{aligned}$$

Theorem 1. Let A be a set of size $|A| > 1$ and $\sim_{r(A)}$ be an equivalence relation on it. If the sizes of the equivalence classes A is split into is known it cannot be reached that all pairs of items a_{i_1} and a_{i_2} chosen arbitrarily from A have degree of unlinkability $d(i_1, i_2) = 1$.

Proof: For $n \notin \{0, 1\}$ the following requirement holds:

$$\begin{aligned}
d(i_1, i_2) = 1 &\Leftrightarrow P(a_{i_1} \sim_{r(A)} a_{i_2}) = P(a_{i_1} \not\sim_{r(A)} a_{i_2}) \\
&\Leftrightarrow \frac{(\sum_{v=1}^l |A_v|^2) - n}{n^2 - n} = \frac{n^2 - \sum_{v=1}^l |A_v|^2}{n^2 - n} \\
&\Leftrightarrow 2 \cdot \sum_{v=1}^l |A_v|^2 - \left(\sum_{v=1}^l |A_v| \right)^2 + \sum_{v=1}^l |A_v| = 0 \\
&\Leftrightarrow \sum_{v=1}^l (|A_v| (2|A_v| - n + 1)) = 0.
\end{aligned}$$

Either all l summands have to equal 0 or the summands have to add up to 0. Because no equivalence class is empty ($\forall i. |A_i| \neq 0$) and $n > 0$ the v -th addend of the sum is

1. > 0 iff $|A_v| > \frac{n-1}{2}$
2. $= 0$ iff $|A_v| = \frac{n-1}{2}$
3. < 0 iff $|A_v| < \frac{n-1}{2}$.

It follows that there exist at most two summands > 0 :

- If the v_1 -th and the v_2 -th summand are > 0 it holds $|A_{v_1}| = |A_{v_2}| = \frac{n}{2}$. But it follows $|A_{v_1}| + |A_{v_2}| = |A|$ and the requirement above cannot be fulfilled.
- If only the v_1 -th summand is > 0 it holds $|A_{v_1}| \geq \frac{n}{2}$ and

$$\sum_{v=1, v \neq v_2}^l (|A_v| (2|A_v| - n + 1)) = -|A_{v_1}| (2|A_{v_1}| - n + 1) = 0.$$

And as with l summands it follows either all remaining $l - 1$ summands have to equal 0 or add up to 0. And this can be repeated till $l = 1$ where it will not be fulfilled.

- If no summand is > 0 all have to equal 0.
 \Rightarrow All $|A_v|$ have to equal $\frac{n-1}{2}$ and this is impossible.
- \Rightarrow There exists no equivalence relation on arbitrary sets A with $|A| > 1$ that guarantees $d(i, j) = 1 \forall i, j \in \{1, n\}$.

4 Anonymity in terms of unlinkability

In terms of unlinkability anonymity in communication systems is defined as 'the properties that a particular message is not related to any sender (recipient) and that to a particular sender (recipient), no message is related.' [12]. The formalisation of this notion of anonymity was given in [9, 17] and extended in section 2 to arbitrary actions. This definition is indicated in [12] by anonymity of an item as 'it is not related to any identifier, and the anonymity of an identifier as not being related to any item of interest'.

Please note 'unlinkability is a sufficient condition of anonymity, but it is not a necessary condition' as outlined in [12].

In contrast to previous approaches in real-world scenarios not only the anonymity of an actor within its anonymity set regarding one specific action has to be measured but the unlinkability of a subset of actions and an actor has to be measured for all possible subsets of actions and all actors within a given system. Our definitions from section 3 are suited for this more general scenario.

Recall the definitions for unlinkability between sets from section 3.2. By $\sim_{r(A)}$ the set A is split in l equivalence classes A_1, \dots, A_l . This means every item u_i in U is described uniquely by a subset $A_i \subseteq A$. If (a part of) this unique description becomes known to an attacker and the unlinkability of the items in A_i decreases the item u_i 's anonymity decreases.

In communications systems both connection and data level have to be considered to give a measurement of anonymity. The basic approach on the connection level is traffic analysis (overview in [13]). The data level is more difficult to analyze:

Example 7 (Communication system). Specific users may have specific interests depending on common personal characteristics e.g., their age, sex, job, religion. These characteristics are often available to the public. And typically additional information is available to an attacker because he usually knows his victim or might influence him [7].

This involves the fact that every item $a_j \in A$ is related to only one item $u_i \in U$. For communication systems this would mean:

Example 8 (Communication system). Our definition assumes users not to send exactly the same messages. In mix-based systems [6] this is realistic because this is forbidden to prevent replay attacks. Nevertheless users might send similar contents. But this similarity and the uncertainty about a user's unique description will hopefully prevent an attacker from decreasing the unlinkability of a certain subset to 0 and especially from decreasing a user's anonymity to 0.

Here we come to a point where we only might estimate anonymity and unlinkability because an exact measurement for a single user would assume him knowing how much an attacker knows about his unique description and how much his description varies from other user's description. And to end with an example for this:

Example 9 (Web surfing). A user using a unusual combination of operating system and browser and requesting contents not typically for his anonymity group will have a quite low unlinkability degree to his set of web requests. A user should consider this fact when choosing his anonymity group e.g. choosing users of the same age and sex. But because users want to be as anonymous as possible even against members of the same anonymity group this claim might be senseless beneath the fact that the anonymity group might give a single user the fallacious feeling of being anonymous (flooding attack or a social variant of it).

5 Conclusion

We generalised the definitions for anonymity [9, 17, 8] to arbitrary scenarios, and we gave new definitions for unlinkability based on the notions in [12]. Especially we have shown there exists no equivalence relation on trivial sets that guarantee the best possible unlinkability in an existential break if only the sizes of the equivalence classes have become known to an attacker. Our next task will be to study sub-optimal equivalence classes on given sets. Finally we refined anonymity in terms of unlinkability. Especially we pointed out the limits of measuring anonymity in real world applications. In future work we will try to evaluate the impact of different constructions of users' unique descriptions on the connection and data level on their anonymity within the system.

Acknowledgements We would like to thank Andrei Serjantov, Andreas Pfitzmann and the anonymous reviewers for valuable suggestions and their patience.

References

1. The anonymizer. <http://www.anonymizer.com>.
2. Oliver Berthold, Hannes Federrath, and Stefan Köpsell. Web mixes: A system for anonymous and unobservable internet access. Designing Privacy Enhancing Technologies. Proc. Workshop on Design Issues in Anonymity and Unobservability, LNCS 2009, Springer-Verlag, Heidelberg 2001, pp. 115–129.
3. Stefan Brands. An efficient off-line electronic cash system based on the representation problem. Centrum voor Wiskunde en Informatica, Computer Science/Department of Algorithmics and Architecture, Report CS-R9323, March 1993.
4. David Chaum. The dining cryptographers problem: unconditional sender and recipient untraceability. Journal of Cryptology (1), 1988.
5. David Chaum. Showing credentials without identification - signatures transferred between unconditionally unlinkable pseudonyms. Advances in Cryptology - EUROCRYPT 85, LNCS 219, Springer-Verlag Berlin 1986, pp. 241–244.
6. David Chaum. Untraceable electronic mail, return addresses and digital pseudonyms. Communications of the ACM, 24(2), 1981, pp. 84–88.
7. Richard Clayton, George Danezis, and Markus G. Kuhn. Real world patterns of failure in anonymity systems. Information Hiding 2001, LNCS 2137, Springer-Verlag Berlin 2001, pp. 230–245.

8. Claudia Diaz, Joris Claessens, Stefan Seys, and Bart Preneel. Information theory and anonymity. Proceedings of the 23rd Symposium on Information Theory in the Benelux, May 29–31, 2002, Louvain la Neuve, Belgium.
9. Claudia Diaz, Stefan Seys, Joris Claessens, and Bart Preneel. Towards measuring anonymity. Privacy Enhancing Technologies 2002, LNCS 2482, Springer-Verlag Berlin.
10. D. Hughes and V. Shmatikov. Information hiding, anonymity and privacy: A modular approach. To appear in Journal of Computer Security, 2003.
11. Dogan Kesdogan, Jan Egner, and Roland Büschkes. Stop-and-go-mixes providing probabilistic anonymity in an open system. Information Hiding 1998, LNCS 1525, Springer-Verlag Berlin 1998, pp. 83–98.
12. Marit Köhntopp and Andreas Pfitzmann. Anonymity, unobservability, and pseudonymity - a proposal for terminology. Draft v0.12., June 2001.
13. Jean-Francois Raymond. Traffic analysis: Protocols, attacks, design issues, and open problems. Privacy Enhancing Technologies 2000, LNCS 2009, Springer-Verlag Berlin.
14. 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.
15. M. K. Reiter and A. D. Rubin. Crowds: Anonymity for web transactions. ACM Transactions on Information and System Security 1(1), November 1998, pp. 66–92.
16. Steve Schneider and Abraham Sidiropoulos. CSP and anonymity. ESORICS 1996, LNCS 1146, Springer-Verlag Berlin 1996, pp. 198–218.
17. Andrei Serjantov and George Danezis. Towards an information-theoretic metric for anonymity. Privacy Enhancing Technologies 2002, LNCS 2482, Springer-Verlag Berlin.
18. C. E. Shannon. Communication theory of secrecy systems. The Bell System Technical Journal 28/4 (1949), pp. 656–715.
19. Vitaly Shmatikov. Probabilistic analysis of anonymity. Proc. 15th IEEE Computer Security Foundations Workshop (CSFW) 2002, pp 119–128.
20. Stuart G. Stubblebine, Paul F. Syverson, and David M. Goldschlag. Unlinkable serial transactions: Protocols and applications. ACM Transactions on Information and System Security, Vol. 2, No. 4, Nov.1999, pp. 354–389.
21. Paul F. Syverson and Stuart G. Stubblebine. Group principals and the formalization of anonymity. FM'99 – Formal Methods, Vol. I, LNCS 1708,, Springer-Verlag 1999pp. 814–833.