

Modeling identity-related properties and their privacy strength^{*}

Meilof Veeningen, Benne de Weger, and Nicola Zannone

Eindhoven University of Technology
{m.veeningen, b.m.m.d.weger, n.zannone}@tue.nl

Abstract. In the last years several attempts to define identity-related properties such as identifiability, pseudonymity and anonymity have been made to analyze the privacy offered by information systems and protocols. However, these definitions are generally incomparable, making it difficult to generalize the results of their analysis. In this paper, we propose a novel framework for formalizing and comparing identity-related properties. The framework employs the notions of detectability, associability and provability to assess the knowledge of an adversary. We show how these notions can be used to specify well-known identity-related properties and classify them with respect to their logical relations and privacy strength. We also demonstrate that the proposed framework is able to capture and compare several existing definitions of identity-related properties.

1 Introduction

With the growth of Internet usage, companies and institutions are gathering more and more personally identifiable information about the individuals in their target groups, such as customers or citizens. Also, more and more people voluntarily publish this information, e.g. on social network websites. Individuals will be associated with an identifier such as a customer number, Social Security Number or nickname, that uniquely identifies them within the service domain. Privacy-sensitive data (name, address, age, sex, all kinds of personal preferences, etc.) will be associated with this identifier.

People are often asked to provide personally identifiable information with the alleged reason of enabling personalization of services, while there may be other, hidden, reasons for this personal data collection, as it often also enables profiling of individual persons or target groups, e.g. for marketing reasons, or for (credit) risk management.

Although this trend promises advantages, it also threatens the privacy of the persons involved, because of the easiness of overuse of personally identifiable information in ways that may not be in line with the reason why the person handed over this information in the first place, or may even be against his will. Before subscribing to a service, people may want to know the privacy policies and practices of the service provider, in particular they may want to know which information is personally identifiable, i.e. can be linked to the originating individual, and who is able to establish those links.

Privacy, defined as *the right to control one's own personally identifiable information*, is studied extensively in the literature. Several identity-related properties such as

^{*} This work is partially funded by the Dutch Sentinel Mobile IDM project (#10522).

pseudonymity, anonymity and identifiability are used in the literature to assess the privacy offered by information systems together with techniques for their verification. Unfortunately, there is no universal, widely accepted definition for these properties. Sometimes, the properties are defined informally and lack a precise semantics [1], making it difficult to reason about them. Elsewhere, they are formally characterized using different foundation frameworks (e.g., process algebra [2,3], modal logic [4], information theory [5], etc.) and are specific to a certain application domain. These differences make it difficult to compare the existing definitions and generalize the results of such studies.

In this paper, we propose a formal framework that allows the specification and comparison of identity-related properties. In particular, we make these contributions:

- Using the ideas of [1] as a starting point, we formally model the personally identifiable information in an information system, and the (partial) knowledge actors can have about this information. In this way, we are able to precisely and unambiguously define many identity-related properties. We also include non-repudiation and deniability properties that are not found in other similar works;
- By analyzing our model, we obtain the logical relations between our properties. As a result, one can easily and formally see which combinations of properties are possible and which are not;
- By introducing the concept of “privacy strength” we order the defined identity-related properties so that one can compare them in terms of the privacy they offer;
- We formalize the identity-related properties presented in [1,6,7] to demonstrate the expressiveness of our model. In particular, we show that our framework provides a means for comparing existing definitions of identity-related properties and consequently for generalizing the results of such studies.

This paper is structured as follows. First, we introduce the concepts of personal information model (§2) and actor view (§3). Then, we show how these concepts are used to formalize identity-related properties (§4). Next, we analyze the properties in terms of logical relations and privacy offered (§5), and compare the results with existing taxonomies (§6). Finally, we present final remarks and directions for future work (§7).

2 Personal Information Models

This section introduces the concept of *personal information (PI) model*, which is used to represent the context of a computer system containing personally identifiable information. A PI model consists of *items of interest* within the system (i.e., data items, identifiers, and entities) and the relations between them. A *data item* is a piece of information that we are interested in from a privacy perspective: information belonging to an *entity* (i.e., a natural person [8]) that might be regarded as sensitive. A precise characterization of data items depends on the application domain. For example, in a communication setting, data items (“hello”, sent) and (“hello”, recv) represent the actions of sending and receiving message “hello” over a network; in a database setting, entries in the database such as a person’s age, address, etc. are data items.

Data items, however, are not directly coupled with entities; rather, they are coupled with *identifiers*, which represent an attribute of an entity (e.g., social security number)




Fig. 1. $\mathcal{D} = \{d_1, d_2\}$; $\mathcal{I} = \{i_1, i_2\}$; $\mathcal{E} = \{e\}$. The θ and ψ functions indicate to which (single) identifier a data item belongs, and to which (single) entity an identifier belongs; the “related” relation \leftrightarrow is the minimal equivalence relation containing ψ and θ .

that can be used to uniquely identify an entity within the system. An identifier, in turn, is coupled with the corresponding entity. Distinguishing identifiers from entities makes it possible to model situations in which the coupling of a data item to an identifier is known, but the coupling to the entity is not; and to model situations in which an entity makes use of multiple identifiers (i.e., pseudonyms) within the same system.

A data item in our model is assumed to belong to one unique identifier, which in turn belongs to one unique entity. This way, when one speaks e.g. about the “anonymity” of a data item, it is immediately clear whose privacy is protected. Given the sets of data items \mathcal{D} , identifiers \mathcal{I} and entities \mathcal{E} , we can see these links as functions $\theta : \mathcal{D} \rightarrow \mathcal{I}$ and $\psi : \mathcal{I} \rightarrow \mathcal{E}$ (Figure 1(a)). We introduce the relation *related*, denoted by \leftrightarrow , to indicate that two items of interest are related to each other. Essentially, \leftrightarrow is the minimal equivalence relation on data items, identifiers and entities such that a data item d is related to an identifier i (represented as $d \leftrightarrow i$) if $\theta(d) = i$ and an identifier i is related to an entity e (represented as $i \leftrightarrow e$) if $\psi(i) = e$ (Figure 1(b)).

Note that according to the definition of \leftrightarrow , a data item is related to *any* identifier of an entity, and not just to the one given by θ . In fact, the relation \leftrightarrow defines equivalence classes which consist of exactly one entity and all identifiers and data items related to it. This representation makes it possible to infer all the information related to an entity and better evaluate the level of entity’s privacy.

Definition 1. A personal information model or PI model is a tuple $(\mathcal{E}, \mathcal{I}, \mathcal{D}, \leftrightarrow)$ where:

1. \mathcal{E} , \mathcal{I} and \mathcal{D} are disjoint finite sets of entities, identifiers and data items, respectively;
2. \leftrightarrow is an equivalence relation on $\mathcal{E} \cup \mathcal{I} \cup \mathcal{D}$, called the related relation, such that every data item is related to at least one identifier, and every identifier is related to exactly one entity.

Note that this definition implies that any data item or identifier is related to one unique entity, as indicated above. Note, also, the disjointness of \mathcal{E} , \mathcal{I} and \mathcal{D} : this is to avoid ambiguity in the definitions of our properties later on. For the sake of simplicity, we use \mathcal{O} to denote the set of all items of interest in the system (i.e., $\mathcal{O} = \mathcal{E} \cup \mathcal{I} \cup \mathcal{D}$) when there is no need to distinguish between the different types. Moreover, we use the symbol \nleftrightarrow to indicate that two items of interest are not related to each other.

An obvious domain where identity-related properties are relevant is communication networks where various parties send messages to each other. The following example shows how our model applies to this domain.




Fig. 2. The two consistent PI models from Example 2 (a), and a 's views from Example 5 (b).





Fig. 3. Combined actor view (Example 3)

Example 1. Consider a communication network such as the Internet. The IP address used in communication can be modeled as the identifier of the person accessing the network. A person may use several IP addresses to access the network (e.g., the IP addresses corresponding to his home Internet connection, to his mobile phone, etc.). In our model, these IP addresses corresponds to different identifiers of the same entity.

A message m transmitted over the network is then represented as two data items: (m, sent) and (m, recv) . This representation makes it possible to identify the sender and the receiver of a specific message.

Different systems may share common entities, identifiers, or data items. In particular, a data item in one system can be used as an identifier in another system. In case of overlap, if different models $(\mathcal{E}_i, \mathcal{I}_i, \mathcal{D}_i, \leftrightarrow_i)$ are to give a consistent view of reality, they must correspond on the overlapping part. Formally, let \leftrightarrow be the minimal equivalence relation on $\bigcup \mathcal{O}_i$ such that $\leftrightarrow \supset \bigcup \leftrightarrow_i$; then the restriction of \leftrightarrow to \mathcal{O}_i is equal to \leftrightarrow_i .

Example 2. Let $M_1 = (\mathcal{E}_1, \mathcal{I}_1, \mathcal{D}_1, \leftrightarrow_1)$ and $M_2 = (\mathcal{E}_2, \mathcal{I}_2, \mathcal{D}_2, \leftrightarrow_2)$ be two PI models where M_1 represents a social network site and M_2 an e-mail system. Let e be an entity that occurs in both systems and i be its identifier on the social network. Now, the e-mail address d_1 of e is mentioned in his profile on the social networking site, so it occurs as a data item in M_1 . However, d_1 is also e 's identifier in M_2 , linked to an e-mail message d_2 . The \leftrightarrow relation as defined above now relates all items $\{e, i, d_1, d_2\}$. The relation \leftrightarrow_1 defines the equivalence class $\{e, i, d_1\}$ and \leftrightarrow_2 the equivalence class $\{e, d_1, d_2\}$. Clearly, M_1 and M_2 are consistent (see Figure 2(a)).

3 Actor View

The PI model describes the “full reality”, that is, all the items of interest within the system and the relations between them. Different actors may see different parts of the system; their *views* model the partial knowledge they have about the reality at a given moment in time. In this paper, \mathcal{A} is used to denote the set of actors. There is no inherent relation between this set and the set of entities defined in the system. Actors may be normal users of a system, or they may be attackers that want to collect as much information as possible – it is in terms of these latter kinds of attackers that identity-related properties are usually defined [1,6,7]. Note that we only consider knowledge of a given

actor at a static moment in time. Thus, the concept of view is not bound to a certain kind of attacker model in which an attacker is static/active or inside/outside the system.

An actor may or may not know about the existence of certain data items, identifiers or entities. Also, even if the actor knows about the existence of two items, he may not be aware of whether they are related. Let us indicate how these aspects are formalized in the actor's view. First, the view includes the set of items of interest *detectable* by the actor. In other words, the actor “can sufficiently distinguish whether [these items] exist or not” [1]. This means that the actor not only “sees” that the item exists, but also that it is real (for instance, in cases where dummy traffic is added to a communication network [9], this dummy traffic is not part of the view). Next, an actor a can observe relations between items of interest within the system. This is captured by the *associability* relation (denoted by \leftrightarrow_a): if x, y are related in the system, then $x \leftrightarrow_a y$ means that a can “sufficiently distinguish whether these items of interest are related” [1]. Finally, if $x \leftrightarrow_a y$, then it is possible that a is even able to convince other actors that the items are indeed related: for this we introduce a stronger version of \leftrightarrow_a called the *provability* relation, denoted by \rightsquigarrow_a . Note that providing a formal definition of detectability, associability, and provability is out of the scope of this paper. We refer to Section 7 for a discussion on existing works that can be used to formalize these notions.

Actors can combine their knowledge about the PI model. One application is federated identity management [10], in which service providers share information with partners and deliver services across multiple domains. To capture this case, we define the view of sets of actors, indicating the knowledge they would collectively have if they would (hypothetically) come together and share their information about the PI model.

Definition 2. Let $M = (\mathcal{E}, \mathcal{I}, \mathcal{D}, \leftrightarrow)$ be a PI model and $A \subset \mathcal{A}$ be a set of actors. The view of A on M is a tuple $(\mathcal{E}_A, \mathcal{I}_A, \mathcal{D}_A, \leftrightarrow_A, \rightsquigarrow_A)$ such that:

1. $\mathcal{E}_A \subset \mathcal{E}, \mathcal{I}_A \subset \mathcal{I}, \mathcal{D}_A \subset \mathcal{D}$ are the items of interest detectable by A ; their union is denoted $\mathcal{O}_A = \mathcal{E}_A \cup \mathcal{I}_A \cup \mathcal{D}_A$
2. $\leftrightarrow_A \subset (\leftrightarrow \cap (\mathcal{O}_A \times \mathcal{O}_A))$ is an equivalence relation: the associability relation of A
3. $\rightsquigarrow_A \subset \leftrightarrow_A$ is an equivalence relation: the provability relation of A
4. If $A = A_1 \cup A_2$ for non-empty $A_1, A_2 \subset \mathcal{A}$, then $\mathcal{E}_A \supset \mathcal{E}_{A_1} \cup \mathcal{E}_{A_2}; \mathcal{I}_A \supset \mathcal{I}_{A_1} \cup \mathcal{I}_{A_2}; \mathcal{D}_A \supset \mathcal{D}_{A_1} \cup \mathcal{D}_{A_2}; \leftrightarrow_A \supset (\leftrightarrow_{A_1} \cup \leftrightarrow_{A_2}); \rightsquigarrow_A \supset (\rightsquigarrow_{A_1} \cup \rightsquigarrow_{A_2})$.

We write $\not\leftrightarrow_A$ for the complement of \leftrightarrow_A and $\not\rightsquigarrow_A$ for the complement of \rightsquigarrow_A . The above definition states that (1) only “real” items can be detected; (2) the associability relation is an equivalence relation and a set of actors A can only associate items that are related in the system and detectable by A ; (3) the provability relation is an equivalence relation and an actor can prove that two items of interest are associated only if he can associate them; (4) two sets of actors combined have at least as much knowledge as the separate sets. Note that \leftrightarrow_A contains at least the transitive closure of \leftrightarrow_{A_1} and \leftrightarrow_{A_2} , but possibly even more (and similarly for \rightsquigarrow_A). So, a group may be able to deduce more than what follows directly from the models of the single actors (or subsets of actors) forming the group (see Example 3).

Example 3. Consider a research group with only one Italian postdoc e , with i as employee number. Suppose actor a knows the nationality $d_1 = \text{Italian}$ of i , and actor b

knows the position $d_2 = \text{postdoc}$ of i . Data items d_1 and d_2 alone are not sufficient to uniquely identify e , as several postdocs and several Italians may work in the research group. However, the set $\{a, b\}$ can link i to e (Figure 3).

Example 4. Consider the scenario of Example 1. Detecting an IP address means being able to say whether the IP address is in use in the network. Detectability of messages means knowing that a given message is a real message sent over the network regardless of its sender or recipient. We assume that if a message m is detectable, then data items $d_1 = (m, \text{sent})$, $d_2 = (m, \text{recv})$ are detectable. However, the associability of the message to the sender and to recipient are independent from each other.

The views of a set of actors on two overlapping PI models must be consistent. Similarly to what we described earlier, given views $(\mathcal{E}_{i,A}, \mathcal{I}_{i,A}, \mathcal{D}_{i,A}, \leftrightarrow_{i,A}, \leftrightarrow_{i,A})$ of actors $A \subset \mathcal{A}$ on system models $(\mathcal{E}_i, \mathcal{I}_i, \mathcal{D}_i, \leftrightarrow_i)$, we say that the $\leftrightarrow_{i,A}$ should be restrictions to $\mathcal{O}_{i,A} \times \mathcal{O}_{i,A}$ of \leftrightarrow_A . The same holds for the provability relations $\leftrightarrow_{i,A}$.

Example 5. Let us revisit Example 2. Now suppose that an actor $a \in \mathcal{A}$ is able to see the e-mail address of e on e 's profile, so $i \leftrightarrow_{1,a} d_1$. Also, in his e-mail application he sees that the e-mail address belongs to e , so in the view of the e-mail system, $d_1 \leftrightarrow_{2,a} e$. Then the consistency requirement states that also in a 's view on M_1 , the e-mail address (as data item) is associable to e : $d_1 \leftrightarrow_{1,a} e$, and thus also $i \leftrightarrow_{1,a} e$ (Figure 2(b)).

4 Identity-Related Properties

This section presents a formalization of well known identity-related properties (e.g. [1]) in terms of actor view. These properties (Table 1) are defined from the perspective of data items by describing their detectability and the amount of associability that data items have to other items of interest.

- The *detectability properties* (D, UD) indicate the occurrence of the data item in the detectability set of the actor view;
- The *identifiability properties* (I, PI, CI) indicate that the actor can associate the data item to an entity, identifier, or both — note that the identifier that a data item d is associated with does not necessarily need to be $\theta(d)$: it can also be any other identifier of $\psi(\theta(d))$ observable by the actor;
- The *anonymity properties* (A, PA, CA) indicate that such a link can not be made;
- The *linkability properties* (L, UL) indicate the ability to associate the data item with other data items;
- The *non-repudiability properties* (EN, IN, CN) indicate that a link to an identifier, entity, or both, can be proved;
- The *deniability properties* (ED, ID, CD) indicate that such a link can not be proved.

The following observations follow directly from our model:

Proposition 1. *Let $M = (\mathcal{E}, \mathcal{I}, \mathcal{D}, \leftrightarrow)$ be a PI model, $A \subset \mathcal{A}$ a set of actors, and $d_1, d_2 \in \mathcal{D}$ s.t. $d_1 \leftrightarrow d_2$. If d_1 and d_2 are identifiable, then they are linkable. If d_1 and d_2 are pseudo-identifiable to the same i , then they are linkable.*

(Un)detectability, identifiability, anonymity, (un)linkability, entity-non-repudiability and entity-deniability can be similarly defined as properties of identifiers; we do not present them in the paper due to lack of space.

Property of $d \in \mathcal{D}$	Condition
detectability (D)	$d \in \mathcal{D}_A$
undetectability (UD)	$d \notin \mathcal{D}_A$
identifiability (I)	$\exists e \in \mathcal{E}_A$ s.t. $d \leftrightarrow_A e$; i.e. $d \leftrightarrow_A \psi(\theta(d))$
pseudo-identifiability (PI)	$\exists i \in \mathcal{I}_A$ s.t. $d \leftrightarrow_A i$
complete identifiability (CI)	$\exists e \in \mathcal{E}_A$ s.t. $d \leftrightarrow_A e$ and $\exists i \in \mathcal{I}_A$ s.t. $d \leftrightarrow_A i$;
anonymity (A)	$d \notin \mathcal{D}_A$, or $\forall e \in \mathcal{E}_A$ $d \not\leftrightarrow_A e$; i.e. $d \notin \mathcal{D}_A \vee d \not\leftrightarrow_A \psi(\theta(d))$
pseudonymity (PA)	$\forall e \in \mathcal{E}_A$ $d \not\leftrightarrow_A e$ and $\exists i \in \mathcal{I}_A$ s.t. $d \leftrightarrow_A i$
complete anonymity (CA)	$d \notin \mathcal{D}_A$, or $\forall e \in \mathcal{E}_A$ $d \not\leftrightarrow_A e$ and $\forall i \in \mathcal{I}_A$ $d \not\leftrightarrow_A i$
linkability (L)	$\exists d' \in \mathcal{D}_A$ s.t. $d \leftrightarrow_A d'$
unlinkability (UL)	$\forall d' \in \mathcal{D}_A$ $d \not\leftrightarrow_A d'$
entity-non-repudiability (EN)	$\exists e \in \mathcal{E}_A$ s.t. $d \rightsquigarrow_A e$; i.e. $d \rightsquigarrow_A \psi(\theta(d))$
identifier-non-repudiability (IN)	$\exists i \in \mathcal{I}_A$ s.t. $d \rightsquigarrow_A i$
complete non-repudiability (CN)	$\exists e \in \mathcal{E}_A$ s.t. $d \rightsquigarrow_A e$ and $\exists i \in \mathcal{I}_A$ s.t. $d \rightsquigarrow_A i$
entity-deniability (ED)	$d \notin \mathcal{D}_A$, or $\forall e \in \mathcal{E}_A$ $d \not\rightsquigarrow_A e$; i.e. $d \notin \mathcal{D}_A \vee d \not\rightsquigarrow_A \psi(\theta(d))$
identifier-deniability (ID)	$d \notin \mathcal{D}_A$, or $\forall i \in \mathcal{I}_A$ $d \not\rightsquigarrow_A i$
complete deniability (CD)	$d \notin \mathcal{D}_A$, or $\forall e \in \mathcal{E}_A$ $d \not\rightsquigarrow_A e$ and $\forall i \in \mathcal{I}_A$ $d \not\rightsquigarrow_A i$

Table 1. Let $(\mathcal{E}, \mathcal{I}, \mathcal{D}, \leftrightarrow)$, $A \subset \mathcal{A}$, $M_A = (\mathcal{E}_A, \mathcal{I}_A, \mathcal{D}_A, \leftrightarrow_A, \rightsquigarrow_A)$ be a PI model, a set of actors and the view of that set of actors, respectively. Let $d \in \mathcal{D}$. The table shows the conditions under which identity-related properties hold for d w.r.t. A .

5 Taxonomies of Identity-Related Properties

This section presents an analysis of the relationships between the properties defined in Section 4. We introduce the concept of “privacy strength” of a data item, which represents the information that an actor has about the data item. This concept is then used to compare the properties in terms of their logical relations and of privacy offered.

5.1 Strength of Privacy

The properties introduced in Section 4 are defined in terms of several aspects of a data item: its detectability, its associability to an entity, and its associability to an identifier. We use these three aspects of privacy to define the “privacy strength” of a data item:

Definition 3. Let $(\mathcal{E}, \mathcal{I}, \mathcal{D}, \leftrightarrow)$ be a PI model and $A \subset \mathcal{A}$ be a set of actors. Let $d \in \mathcal{D}$. The privacy strength σ of d w.r.t. A is the tuple $(\delta, \epsilon, \iota)$, where:

- $\delta = 1$ if $d \in \mathcal{D}_A$; $\delta = 0$ otherwise;
- $\epsilon = 2$ if there is $e \in \mathcal{E}_A$ such that $d \rightsquigarrow_A e$; $\epsilon = 1$ if not, but there is $e \in \mathcal{E}_A$ such that $d \leftrightarrow_A e$; $\epsilon = 0$ otherwise;
- $\iota = 2$ if there is $i \in \mathcal{I}_A$ such that $d \rightsquigarrow_A i$; $\iota = 1$ if not, but there is $i \in \mathcal{I}_A$ such that $d \leftrightarrow_A i$; $\iota = 0$ otherwise.

Note that the \leftrightarrow_A and \rightsquigarrow_A relations only have meaning if $d \in \mathcal{D}_A$, so the *admissible privacy strengths* are those in $S = \{(0, 0, 0)\} \cup (\{1\} \times \{0, 1, 2\} \times \{0, 1, 2\})$.

We can now rephrase our properties as conditions of the privacy strength of a data item. Each identity-related property then corresponds to a subset $S' \subset S$ of admissible

privacy property	condition	privacy property	condition
undetectability (UD)	$\delta = 0$	detectability (D)	$\delta = 1$
anonymity (A)	$\epsilon = 0$	identifiability (I)	$\epsilon \geq 1$
pseudonymity (PA)	$\epsilon = 0, \iota \geq 1$	pseudo-identifiability (PI)	$\iota \geq 1$
complete anonymity (CA)	$\epsilon = \iota = 0$	complete identifiability (CI)	$\epsilon \geq 1, \iota \geq 1$
entity-deniability (ED)	$\epsilon < 2$	entity-non-repudiability (EN)	$\epsilon = 2$
identifier-deniability (ID)	$\iota < 2$	identifier-non-repudiability (IN)	$\iota = 2$
complete deniability (CD)	$\epsilon < 2, \iota < 2$	complete non-repudiability (CN)	$\epsilon = 2, \iota = 2$

Table 2. Let $(\mathcal{E}, \mathcal{I}, \mathcal{D}, \leftrightarrow)$ be a PI model and $A \subset \mathcal{A}$ a set of actors. The table shows the conditions on the privacy strength $\sigma = (\delta, \epsilon, \iota)$ of a data item $c \in \mathcal{D}$ corresponding to the given identity-related properties.

privacy strengths. The correspondence is shown in Table 2. Note that we do not consider (un)linkability as it is not about associating the data item with identifiers and entities.

5.2 Logical Relations between Properties

The formulation of identity-related properties in terms of their privacy strength enables us to analyze the logical relations between them (i.e., whether they overlap, are mutually exclusive, or one implies the other). The following result follows directly from Table 2:

Proposition 2. *Let $(\mathcal{E}, \mathcal{I}, \mathcal{D}, \leftrightarrow)$ be a PI model, and $A \subset \mathcal{A}$ a set of actors. Then Table 3 shows the logical relations between the identity-related properties for a data item $d \in \mathcal{D}$ with respect to A .*

Table 3 makes it possible to identify, for example, which properties it is impossible to achieve at the same time. For instance, having both identifiability and pseudonymity is not possible, but one can have identifiability with pseudo-identifiability. It also shows that anonymity automatically guarantees entity-deniability.

Although the taxonomy in Table 3 exactly indicates the logical relations between properties, it does not say which ones are more desirable from a privacy standpoint. In the next section, we propose an ordering of properties with respect to privacy.

5.3 A Partial Order on Privacy Strengths

The notion of “privacy strength” makes it possible to compare the privacy of different data items: we say that a data item has (strictly) stronger privacy than another data item if on each of the three aspects above, its privacy is stronger than that of the other item. This gives a partial order on the admissible privacy strengths (Figure 4).

Definition 4. *Let b and c be two data items, and $\sigma_b = (\delta_b, \epsilon_b, \iota_b) \in S$ and $\sigma_c = (\delta_c, \epsilon_c, \iota_c) \in S$ be their privacy strengths. We say that σ_b is stronger than σ_c , denoted by $\sigma_b \succeq \sigma_c$, if $\delta_b \leq \delta_c, \epsilon_b \leq \epsilon_c, \iota_b \leq \iota_c$. Accordingly, we say that b has stronger privacy than c (denoted by $b \succeq c$) if $\sigma_b \succeq \sigma_c$.*

	detect.		identif.			anonymity			non-repud.			deniability		
	D	UD	I	PI	CI	A	PA	CA	EN	IN	CN	ED	ID	CD
Detectability (D)	=	\neg	\Leftarrow	\Leftarrow	\Leftarrow	\emptyset	\Leftarrow	\emptyset	\Leftarrow	\Leftarrow	\Leftarrow	\emptyset	\emptyset	\emptyset
Undetectability (UD)	\neg	=	\emptyset	\emptyset	\emptyset	\uparrow	\emptyset	\uparrow	\emptyset	\emptyset	\emptyset	\uparrow	\uparrow	\uparrow
Identifiability (I)	\uparrow	\emptyset	=	\emptyset	\Leftarrow	\neg	\emptyset	\emptyset	\Leftarrow	\emptyset	\Leftarrow	\emptyset	\emptyset	\emptyset
Pseudo-identifiability (PI)	\uparrow	\emptyset	\emptyset	=	\Leftarrow	\emptyset	\Leftarrow	\emptyset	\emptyset	\Leftarrow	\Leftarrow	\emptyset	\emptyset	\emptyset
Complete identifiability (CI)	\uparrow	\emptyset	\uparrow	\uparrow	=	\emptyset	\emptyset	\emptyset	\emptyset	\emptyset	\Leftarrow	\emptyset	\emptyset	\emptyset
Anonymity (A)	\emptyset	\Leftarrow	\neg	\emptyset	\emptyset	=	\Leftarrow	\Leftarrow	\emptyset	\emptyset	\emptyset	\uparrow	\emptyset	\emptyset
Pseudonymity (PA)	\uparrow	\emptyset	\emptyset	\uparrow	\emptyset	\uparrow	=	\emptyset	\emptyset	\emptyset	\emptyset	\uparrow	\emptyset	\emptyset
Complete anonymity (CA)	\emptyset	\Leftarrow	\emptyset	\emptyset	\emptyset	\uparrow	\emptyset	=	\emptyset	\emptyset	\emptyset	\uparrow	\uparrow	\uparrow
Entity-non-repudiability (EN)	\uparrow	\emptyset	\uparrow	\emptyset	\emptyset	\emptyset	\emptyset	\emptyset	=	\emptyset	\Leftarrow	\neg	\emptyset	\emptyset
Identifier-non-repudiability (IN)	\uparrow	\emptyset	\emptyset	\uparrow	\emptyset	\emptyset	\emptyset	\emptyset	\emptyset	=	\Leftarrow	\emptyset	\neg	\emptyset
Complete non-repudiability (CN)	\uparrow	\emptyset	\uparrow	\uparrow	\uparrow	\emptyset	\emptyset	\emptyset	\uparrow	\uparrow	=	\emptyset	\emptyset	\emptyset
Entity-deniability (ED)	\emptyset	\Leftarrow	\emptyset	\emptyset	\emptyset	\Leftarrow	\Leftarrow	\Leftarrow	\neg	\emptyset	\emptyset	=	\emptyset	\Leftarrow
Identifier-deniability (ID)	\emptyset	\Leftarrow	\emptyset	\emptyset	\emptyset	\emptyset	\emptyset	\Leftarrow	\emptyset	\neg	\emptyset	\emptyset	=	\Leftarrow
Complete deniability (CD)	\emptyset	\Leftarrow	\emptyset	\emptyset	\emptyset	\emptyset	\emptyset	\Leftarrow	\emptyset	\emptyset	\emptyset	\uparrow	\uparrow	=

Table 3. Logical relations between the various identity-related properties. \uparrow : left property implies top property; \Leftarrow : left property implied by top property; \emptyset : mutually exclusive; \neg : each others negation; \emptyset : properties partially overlap.

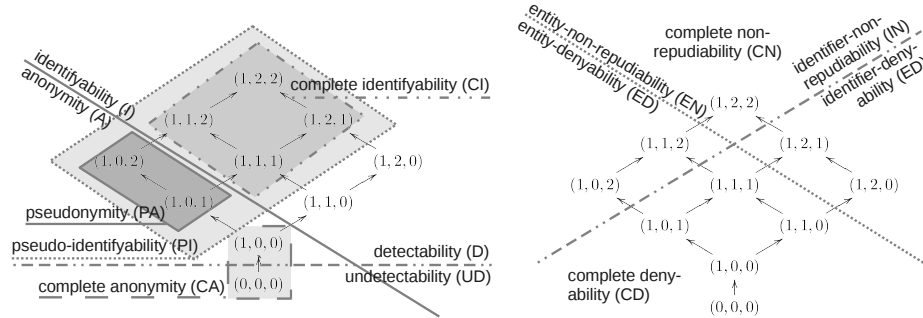


Fig. 4. Partial order on the set S of admissible privacy strengths. In the left figure, for the detectability, identifiability and anonymity properties their subsets $S' \subset S$ are drawn (for the colored areas, the line style of the area corresponds to that of the property name); in the right figure the non-repudiability and deniability properties are indicated.

The definition above together with the formalization of identity-related properties in Table 2 allows us to order properties with respect to their privacy strength.

We consider two natural ways to define when $S_1 \subset S$ offers stronger privacy than $S_2 \subset S$. First, we consider a strict condition under which S_1 should be pairwise stronger than S_2 . In this case, we say that S_1 has *absolutely stronger privacy* than S_2 . We also consider a weaker condition under which S_1 just has *stronger privacy* than S_2 .

Definition 5. Let P_1 and P_2 be properties, and $S_1 \subset S$ and $S_2 \subset S$ their respective sets of privacy strengths.

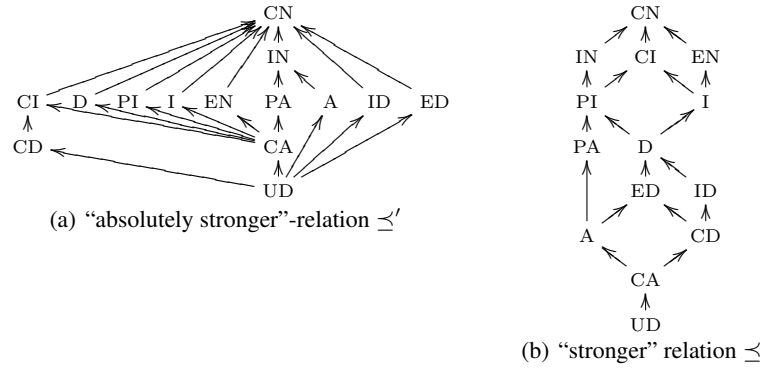


Fig. 5. Partial orders on identity-related properties: arrow means weakening of privacy.

- P_1 has absolutely stronger privacy than P_2 , or $P_1 \preceq' P_2$, if for all $\sigma_1 \in S_1, \sigma_2 \in S_2$ one has $\sigma_1 \preceq \sigma_2$
- P_1 has stronger privacy than P_2 , or $P_1 \preceq P_2$, if for all $\sigma \in S_1$, there exists a $\tau \in S_2$ such that $\sigma \preceq \tau$ and for all $\tau \in S_2$, there exists a $\sigma \in S_1$ such that $\sigma \preceq \tau$.

It is clear that \preceq and \preceq' defined above are relations; however, it may not be immediately clear that they are partial orders (i.e., they are reflexive, anti-symmetric and transitive relations). This does turn out to be true.¹

Proposition 3. *The partial orders \preceq, \preceq' on identity-related properties are as in Fig. 5.*

As seen in the figure, the absolutely stronger relation already gives us several logical conclusions: for instance, we can see that undetectability (UD) is the strongest privacy guarantee, whereas complete non-repudiation (CN) offers the weakest — in fact, no — privacy. However, there are also many things it does *not* allow us to conclude. For instance, anonymity (A) does not offer absolutely stronger privacy than identifiability (I). The reason is that if anonymity (A) holds for a data item d , then it may still be associable to an identifier; if identifiability (I) holds for d' then it may be unassociable to an identifier. Using the less strict “stronger” relation (Figure 5(b)), we *can* make this claim and several similar ones.

5.4 Relating the Taxonomies

Each identity-related property corresponds to an area in Figure 4. Given two identity-related properties, we can compare them on their privacy strength. Loosely speaking, if property A has stronger privacy than property B , then it will lie lower than B in the figure. Also, we can see the logical relation between A and B in the figure: for example, if A and B overlap as properties, then the areas in the figure will overlap.

The comparison of two properties both on strength and on logical relation determines how the areas of the properties relate to each other. For instance, if $A \preceq B$ and

¹ Technically, \preceq, \preceq' give a partial ordering on all subsets $S' \subset S$ that have no gaps, i.e., if $x, z \in S'$ and $x \preceq' y \preceq' z$, then also $y \in S'$. This property holds for all of our properties.

	$A \cup B = \emptyset$	$A \subset B$	$B \subset A$	$A \cup B \neq \emptyset$
$A \preceq B$	B completely above	A is bottom part	B is top part	B partly above
	A (CA \preceq CN)	of B (UD \preceq CA)	of A (IN \preceq CN)	A (A \preceq IN)
A, B incomp.	A, B next to each other (PA, ID)	A in middle of B (PA \subset ED)	B in middle of A (ED \subset PA)	A, B overlap, not one above (A, Cd)

Table 4. Possible situations when comparing two identity-related properties A, B and their interpretations as areas in Figure 4. Between brackets an example of the situation is given. Note that $A = S \setminus B$ is a special case of $A \cup B = \emptyset$.

$A \subset B$, then the area of A in the figure must be the lower part of the area of B . A summary of the many different combinations that can occur is shown in Table 4.

6 Comparing Identity-Related Property Taxonomies

In this section, we discuss several frameworks [1,6,7] that define identity-related properties, and we compare them with our framework. Note that these frameworks mainly focus on a communication setting. We refer to Examples 1 and 4 on how this domain can be represented in our model. To smooth our terminology in this particular context, given a message m we refer to “anonymity of (m, sent) ” simply as “anonymity of the sender of m ”, and similarly for other properties.

The frameworks presented in [1,6,7] do not differentiate between entities and identifiers (although [6] does mention this as a possible extension). Consequently, they cannot capture pseudonymity as a property. Pfitzmann and Hansen [1] do define pseudonymity, but by this they mean the mechanism of using pseudonyms. They then specify different degrees of linkability between the pseudonym and entity that correspond to our notions of identifiability/anonymity of identifiers. Also non-repudiability and deniability properties are not considered in these frameworks.

6.1 Privacy by Data Minimization

In their evolving work [1], Pfitzmann and Hansen present a terminology for properties related to what they call *privacy by data minimization*. From the basic concepts of detectability and linkability of items of interest, they define informally more advanced properties. These (informal) definitions are the basis of our framework, albeit that we use the term associability instead of linkability. This is because we want to distinguish the property of linking two data items (or two identifiers) to one another from the general relation used to associate items of interest with each other.

Table 5 provides a formalization of the identity-related properties of [1] in our model. For instance, sender anonymity can be translated into our model as follows: given a message m , (m, sent) is anonymous from the perspective of adversary a (i.e., *anonymity of sender*). Similarly, recipient anonymity – the dual of sender anonymity – corresponds to *anonymity of recipient*.

The authors also define the notion of unobservability: for a message m sent by s and received by r , it means undetectability of m w.r.t. $\mathcal{A} \setminus \{s, r\}$, sender anonymity w.r.t. r ,

Property from [1]	Interpretation	Property
Detectability	$(m, \text{sent}) \in \mathcal{D}_a \wedge (m, \text{recv}) \in \mathcal{D}_a$	detectability
Undetectability	$(m, \text{sent}) \notin \mathcal{D}_a \wedge (m, \text{recv}) \notin \mathcal{D}_a$	undetectability
Linkability	$o_1 \leftrightarrow_A o_2$ (with $o_1, o_2 \in \mathcal{O}_a$)	associability
Sender Anonymity	$(m, \text{sent}) \not\leftrightarrow_a s$	anonymity of sender
Recipient Anonymity	$(m, \text{recv}) \not\leftrightarrow_a r$	anonymity of recipient
Unobservability	$\forall b \in \mathcal{A} : b \neq s, b \neq r :$	undetectability w.r.t. b ;
(w.r.t. all actors)	$(m, \text{sent}) \notin \mathcal{D}_b \wedge (m, \text{recv}) \notin \mathcal{D}_b \wedge$ $(m, \text{sent}) \not\leftrightarrow_{rs} s \wedge (m, \text{recv}) \not\leftrightarrow_{sr} r$	anonymity of sender w.r.t. r ; anonymity of recv w.r.t. s

Table 5. Interpretation in our model (for a given message m sent by s and received by r w.r.t. actor a) for identity-related notions from [1], and our name for the given property.

and recipient anonymity w.r.t. s . Differently from other properties, unobservability is a global property in the sense that it is defined w.r.t. the set of all actors \mathcal{A} .

6.2 An Indistinguishability-Based Characterization of Anonymous Channels

In [7], several identity-related notions are defined and classified in terms of indistinguishability of message matrices. A message matrix indicates which messages have to be sent by what sender to what recipient in a run of some communication protocol. An adversary gets to choose two message matrices M_1 and M_2 . The communication protocol will then be run in an experiment using either M_1 or M_2 , and the adversary tries to decide which message matrix was used.

Identity-related notions are then defined in terms of restrictions placed on the message matrices (in other words, what information can be freely chosen by the adversary). The strongest anonymity notion, unobservability, holds when there are no restrictions at all on M_1 and M_2 , but still after observing the communication protocol the adversary cannot say which was used. For the weaker notion of unlinkability, the number of messages sent by each sender and the number of messages received by each recipient must be the same for the two message matrices.

Differently from our approach, the properties of [7] are also defined on the basis of the confidentiality of the message content.² Therefore, many of them (i.e., SUL, RUL, UL, SA*, RA*, and SRA) cannot be captured in our model. To enable us to compare these notions, we extend the view M_a with a set $\mathcal{C}_a \subset \mathcal{D}_a$ consisting of the messages whose content is known to the attacker. Table 6 shows a formalization of their identity-related properties in this extended model. From the table, one can clearly reproduce the trivial relations between properties given in Proposition 1 from [7]. In their formalization, the authors also show that using PKI and key-private secure encryption [11], some weak notions can be transformed to stronger notions. The use of encryption aims to guarantee the confidentiality of messages. Under this assumption, the authors prove that SUL, RUL, and UL are all equivalent as well as SA and SA*, and RA and RA*. Table 6 shows that the same conclusions can be drawn in our framework.

² Although confidentiality of the message is an important aspect, we believe that it is independent of the issue of associating a message with its sender and recipient.

Property from [7]	Interpretation	Property
Sender Unlinkability (SUL)	$(m, \text{sent}) \not\leftrightarrow_{as} \vee ((m, \text{recv}) \not\leftrightarrow_{ar} \wedge m \notin \mathcal{C}_a)$	—
Sender Anonymity (SA)	$(m, \text{sent}) \not\leftrightarrow_{as}$	anonymity of sender
Strong Sender Anonymity (SA*)	$(m, \text{sent}) \not\leftrightarrow_{as} \wedge m \notin \mathcal{C}_a$	—
Receiver Unlinkability (RUL)	$(m, \text{recv}) \not\leftrightarrow_{ar} \vee ((m, \text{sent}) \not\leftrightarrow_{as} \wedge m \notin \mathcal{C}_a)$	—
Receiver Anonymity (RA)	$(m, \text{recv}) \not\leftrightarrow_{ar}$	anonymity of recipient
Strong Receiver Anonymity (RA*)	$(m, \text{recv}) \not\leftrightarrow_{ar} \wedge m \notin \mathcal{C}_a$	—
(Sender-Receiver) Unlinkability (UL)	$((m, \text{recv}) \not\leftrightarrow_{ar} \vee (m, \text{sent}) \not\leftrightarrow_{as}) \wedge m \notin \mathcal{C}_a$	—
Sender-Receiver Anonymity (SRA)	$(m, \text{recv}) \not\leftrightarrow_{ar} \wedge (m, \text{sent}) \not\leftrightarrow_{as} \wedge m \notin \mathcal{C}_a$	—
Unobservability (UO)	$(m, \text{sent}) \notin \mathcal{D}_a \wedge (m, \text{recv}) \notin \mathcal{D}_a$	undetectability

Table 6. Interpretation in our model (for a given message m sent by s and received by r w.r.t actor a) for identity-related notions from [7], and our name for the given property.

Pfitzmann and Hansen [1] have also compared their terminology to [7], with different results. However, we believe our comparison matches more closely with the formal model of [7]. For example, they claim that SUL corresponds to sender anonymity. This latter property means that no message can be associated with its sender. However, suppose an actor can associate a message with its sender, but not see the contents of the message or the recipient. Then, in the formalization of [7], SUL still holds. So clearly, although sender anonymity is a sufficient condition for SUL, it is not a necessary one.³

6.3 Information hiding, anonymity and privacy: a modular approach

In [6], identity-related properties are formalized in terms of the amount of knowledge about sender and recipient functions (i.e., functions that link messages to their sender or recipient respectively). These functions can have the following properties:

- *Value opaqueness*: for a given message, the sender/recipient is not known. This corresponds to anonymity of sender and anonymity of recipient, respectively.
- *Image opaqueness*: for a given entity, it is not known if it has sent/received a message. Image opaqueness is a necessary, but not sufficient, condition for detectability.
- *Kernel opaqueness*: given two distinct messages, it is not known that they have been sent/received by the same entity. This corresponds to our notion of unlinkability.

The authors then define several anonymity properties based on these concepts. The formalization of the notions in [6] is shown in Table 7. Sender anonymity and recipient anonymity are defined as value opaqueness of the sender and recipient functions, which is in line with our definitions. For sender untraceability, in addition one should have kernel opaqueness, so this corresponds to anonymity and unlinkability of the data item (m, sent) ; similarly for recipient untraceability.

Blender anonymity is defined as a combination of value opaqueness of the sender and recipient functions. Accordingly, it can be seen as anonymity of both sender and

³ This assumes that seeing a message does not mean seeing the message contents. It is not entirely clear if [1] makes this assumption. If not, then [1] in addition to sender anonymity should also demand relationship anonymity to capture the meaning of SUL.

Property from [6]	Interpretation	Property
Sender Anonymity	$(m, \text{sent}) \not\leftrightarrow_a s$	anonymity of sender
Sender Untraceability	$(m, \text{sent}) \not\leftrightarrow_a s \wedge \nexists m' : (m', \text{sent}) \leftrightarrow_a (m, \text{sent})$	anonymity of sender + unlinkability of sender
Recipient Anonymity	$(m, \text{recv}) \not\leftrightarrow_a r$	anonymity of recipient
Recipient Untraceability	$(m, \text{recv}) \not\leftrightarrow_a r \wedge \nexists m' : (m', \text{recv}) \leftrightarrow_a (m, \text{recv})$	anonymity of recipient + unlinkability of recipient
Blender Anonymity	$(m, \text{sent}) \not\leftrightarrow_a s \wedge (m, \text{recv}) \not\leftrightarrow_a r$	anon. of sender and recv
Conversation-Agent-2-Unlink.	$(m, \text{sent}) \not\leftrightarrow_a s \vee (m, \text{recv}) \not\leftrightarrow_a r$	anon. of sender or recv

Table 7. Interpretation in our model (for a given message m sent by s and received by r w.r.t. actor a) of identity-related notions from [6], and our name for the given property.

recipient. “Conversation-agent-2-unlinkability”, on the other hand, is value opaqueness of the function giving the sender and recipient at the same time, so it corresponds to either anonymity of sender or anonymity of recipient. Note that [6] defines several variants of anonymity based on the number of possible senders or recipients for a message. Our model abstracts away from this level of detail so that anonymity means that there is “sufficient” unclarity about the sender or recipient.

7 Conclusions and Future Work

In this work, we presented a formal model to represent privacy-sensitive information in information systems and the knowledge that actors have about it. This model can be used to formalize many identity-related notions such as identifiability, pseudonymity and anonymity. This formalization allows one to compare the properties in terms of their logical relations and their privacy strengths. In particular, the obtained taxonomies show which properties cannot be achieved at the same time and which properties are implied by other properties as well as which properties are more desirable from a privacy standpoint. We also demonstrated that our model is able to capture the identity-related properties and reproduce several results presented in [1,6,7]. However, the mapping is not always straightforward because of the different notions these frameworks employ. For instance, [7] considers confidentiality issues in defining identity-related properties. We showed how such issues can be modeled in our framework.

Interestingly, the formalization in [7] enables one to formally express transformations of communication protocols to achieve certain privacy properties. Our model is static and therefore is not able to specify the behavior of communication parties or adversaries. In forthcoming work we intend to extend our model to capture dynamic situations.

Our approach to defining identity-related properties assumes a formal meaning of the notions of detectability, associability, and provability. Several proposed frameworks can be used to formalize these notions. One relevant research stream proposes to model the knowledge of an actor in a probabilistic way [5,12,13,14]. For instance, [5,13,14] consider the knowledge about who has executed some action as a probability distribution over all possible entities; the entropy of this distribution then is a measure of the

“degree of anonymity” of that action. Another research stream models communication systems using process algebra [2,3] or modal logic [4] and, based on such formal specifications, checks the privacy provided by a system. Concerning the formalization of non-repudiation [15], we mention that it is often seen as a consequence of using cryptography, e.g. the unforgeability property of digital signatures. However, as [16, Ch.4] remarks, the meaning of “proving” depends strongly on the specific context of the application. As future work, we plan to investigate how to use these different frameworks for detectability, associability and provability as a foundation for our model. Of particular interest here is how we can capture low-level details of properties (e.g., 10-anonymity is “stronger” than 3-anonymity) without losing the abstract nature of the model.

References

1. Pfitzmann, A., Hansen, M.: A terminology for talking about privacy by data minimization: Anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management. http://dud.inf.tu-dresden.de/Anon_Terminology.shtml (December 2009) v0.32.
2. Schneider, S., Sidiropoulos, A.: CSP and anonymity. In: Proc. of ESORCS'96. LNCS 2482, Springer (1996) 198–218
3. Chatzikokolakis, K.: Probabilistic and Information-Theoretic Approaches to Anonymity. PhD thesis, Laboratoire d'Informatique (LIX), École Polytechnique, Paris (2007)
4. Syverson, P.F., Stubblebine, S.G.: Group principals and the formalization of anonymity. In: Proceedings of FM'99. LNCS 1708, Springer (1999) 814–833
5. Steinbrecher, S., Köpsell, S.: Modelling unlinkability. In: Proc. of PET'03. LNCS 2760, Springer (2003) 32–47
6. Hughes, D., Shmatikov, V.: Information hiding, anonymity and privacy: a modular approach. *Journal of Computer Security* **12**(1) (2004) 3–36
7. Hevia, A., Micciancio, D.: An indistinguishability-based characterization of anonymous channels. In: Proc. of PETS'08. LNCS 5134, Springer (2008) 24–43
8. European Parliament: Directive 95/46/EC on the protection of individuals with regard to the processing of personal data and on the free movement of such data. *Official Journal of the European Communities* **L 281** (23 November 1995) 31–50
9. Diaz, C., Preneel, B.: Taxonomy of mixes and dummy traffic. In: Proc. of I-NetSec'04, Kluwer Academic Publishers (2004) 215–230
10. Camenisch, J., Pfitzmann, B.: Federated Identity Management. In: Security, Privacy, and Trust in Modern Data Management. Springer (2007) 213–238
11. Bellare, M., Boldyreva, R., Pointcheval, D.: Key-privacy in public-key encryption. In: Proc. of ASIACRYPT'01. LNCS 2248, Springer (2001) 566–582
12. Clauß, S.: A Framework for Quantification of Linkability Within a Privacy-Enhancing Identity Management System. In: Emerging Trends in Information and Communication Security. Springer Verlag, LNCS 3995 (2006) 191–205
13. Diaz, C., Seys, S., Claessens, J., Preneel, B.: Towards measuring anonymity. In: Proc. of PET'02. LNCS 2482, Springer (2002) 54–68
14. Serjantov, A., Danezis, G.: Towards an information theoretic metric for anonymity. In: Proc. of PET'02. LNCS 2482, Springer (2002) 41–53
15. Avizienis, A., Laprie, J.C., Randell, B., Landwehr, C.: Basic concepts and taxonomy of dependable and secure computing. *IEEE Trans. Dependable Sec. Comput.* **1** (2004) 11–33
16. Roe, M.: Cryptography and Evidence. PhD thesis, University of Cambridge (1997)