

Relations Among Privacy Notions

Jens-Matthias Bohli
University of Sussex, UK

Andreas Pashalidis
Katholieke Universiteit Leuven, Belgium

December 26, 2009

Abstract

This paper presents a hierarchy of privacy notions that covers multiple anonymity and unlinkability variants. The underlying definitions, which are based on the idea of indistinguishability between two worlds, provide new insights into the relation between, and the fundamental structure of, different privacy notions. We furthermore place previous privacy definitions concerning group signature, anonymous communication, and secret voting systems in the context of our hierarchy; this renders these traditionally disconnected notions comparable.

1 Introduction

With the growing number of services and information offered in the digital world, the number of situations where there is a need to hide the correspondence between individual data items and the people that cause their appearance, is also increasing. A variety of privacy protecting systems address this need; anonymous communication systems, for example, hide how transmitted messages correspond to their senders (and their recipients); group signatures hide the identity of the signer of a given message, and secret voting schemes hide the identity of the voter who cast any given ballot. In general, a system is said to ‘provide privacy’ if it hides, perhaps to an extent, the correspondence between its users and the data items it produces.

What *exactly* it means for any given privacy protecting system to provide privacy naturally varies between system types. The privacy definition for group signatures [5], for example, differs from the one for anonymous credentials [11]. Similarly, privacy for voting schemes [1] is defined differently from privacy in the setting of anonymous communication [27]. Despite efforts for a consistent terminology [35], formal treatments *seem* to define privacy in an inconsistent and sometimes even contradictory manner; while, for example, some authors assert that ‘anonymity and unlinkability are technically the same property’ [5], others show that, although related, they are, in fact, distinct [27], and others insist that they are independent [28]. Unfortunately, it is not only the terminology that is used inconsistently; due to the discrepancies between the formal models, the resulting privacy notions themselves turn out to be incomparable. It remains unclear whether or not it is possible to construct a *single* formal framework in which privacy notions pertaining to *different* system types can be defined in a consistent and comparable manner.

Our Contributions: This work can be seen as first step towards a formal framework that aims to define multiple privacy notions in an application-agnostic manner. By so doing, it provides new insights into the inner structure of privacy notions. Starting from a generic system model that potentially hides the correspondence between digital elements and the users that cause their appearance, we systematically analyse different degrees to which this correspondence may be hidden, and place the resulting privacy notions into a well-characterised hierarchy. Furthermore, we examine the classes of ‘online’ and ‘shuffling’ systems, and show why only some privacy notions apply to these classes. Finally, we place existing definitions for group signature, anonymous communication, and secret voting systems in the context of our framework. This enables us, on the one hand, to understand the relationship between, and to compare, these traditionally disconnected privacy notions. On the other hand, the framework highlights a largely unexplored space of theoretically possible notions some of which may be of practical interest.

An extended abstract of this work appeared in [8]. The present version improves the model description, enriches the hierarchy by introducing new privacy notions, proves new relations between the notions, and discusses the role of shuffling in privacy systems. As a new application, this version examines secret voting schemes in the context of the framework.

Related Work: The framework introduced in [27] has certain commonalities with the framework introduced in this paper; both define, for example, a hierarchy of privacy notions based on the principle that an adversary may break any privacy notion *except* the one of interest, and both follow the idea of left-or-right security introduced in [4]. However, the framework in [27] appears to be specific to anonymous communication systems. Moreover, the hierarchy of privacy notions defined in this paper is richer; when mapped to anonymous communication systems – we do this in section 4.2 – notions beyond those considered in [27] arise. The framework in [28] also has certain commonalities with ours; both support, for example, the specification of privacy notions against adversaries with partial knowledge about a function. However, in contrast to the framework introduced in this paper, the one in [28] does not consider probabilistic adversaries, and is only applied to anonymous communication systems. Moreover, it is unclear how its privacy definitions map to existing and established application-specific ones.

Other related work includes the literature on *measuring* privacy (e.g. [3, 13, 14, 16, 17, 15, 18, 19, 22, 26, 32, 37, 38, 40]). The proposed metrics appear, however, to pertain to particular privacy notions, if not system types. Multiple, sometimes inconsistent metrics for the same notion have also been proposed. While, for example, the metric in [18], proposed for the anonymity in the setting of anonymous communication systems, focuses on the relationship between incoming and outgoing messages, the metric proposed in [24] focuses on the relationship between senders and receivers. Similarly, the metric for unlinkability proposed in [22, 39] does not take into account the skewness of the adversary’s view on possible solutions, while the metric proposed in [21] does. The only work we are aware of that places multiple privacy notions into a single framework [34], does not relate the metrics it defines to privacy definitions from the cryptographic literature. It is important to note that most privacy metrics cited above are probabilistic. That is, they measure the degree to which a system provides privacy. In contrast to this, the privacy definitions in this paper are ‘all-or-nothing’; a system either provides or does not provide a given privacy notion.

While the most popular adversarial model in the anonymous communication literature is perhaps that of the ‘global passive’ adversary (see, for example, [29, 33, 36]), in this paper we consider an adaptive adversary that may corrupt users. This is in line with definitions from group signatures [5], anonymous credentials [11], and some of the literature on anonymous communication (e.g. [7, 19]). Moreover, our privacy definitions classify systems as either succeeding, or failing to provide a given privacy notion; while this is in contrast with some works on anonymous communication that consider ‘soft’, probabilistic measures (see, for example, [19, 26]), it, too, is in line with works on group signatures and some of the literature on anonymous communication systems (see, for example, [24]).

Outline: The rest of this paper is organised as follows. The next section introduces our notation and formal model, and Section 3 presents the hierarchy of privacy notions and examines its structure. Section 3.3, in particular, examines ‘online’ systems and shows why only some privacy notions apply in such systems, and section 3.4 examines the conditions under which shuffling is beneficial for a system, and shows that a particular class of shuffling systems (which we call ‘stateless’) always provides at least a certain degree of privacy. Section 4 examines group signature, anonymous communication, and secret voting systems in the context of the hierarchy. Section 5 concludes.

2 Preliminaries

This section introduces our notation and formal model. In particular, the next section introduces the class of systems that are considered in this paper, Section 2.2 introduces the different privacy notions considered, and Section 2.3 describes the adversarial model.

2.1 System model

In this paper, we consider systems that may be *sequentially invoked* a finite number of times and that, for each invocation, produce an element $e \in \{0, 1\}^*$. It is required that each invocation is uniquely associated with a user and with an input parameter $\alpha \in A$, where A is a system-specific parameter space, and where the value of α may influence the behaviour of the system. It is furthermore required that each user is identified by means of a unique identifier from an identifier space \mathcal{U} ; it is required that this space is of cardinality at least polynomial in a given security parameter, or infinite if no such parameter is given.

We assume that the system, denoted by Φ^A in the sequel, produces its output in batches of (potentially varying sizes). That is, it is assumed that, on input a batch of invocations $(u_1, \alpha_1), (u_2, \alpha_2), \dots, (u_c, \alpha_c) \in (\mathcal{U} \times A)^c$, Φ^A outputs a sequence $((e_1, \dots, e_c), \beta)$, where the sequence (e_1, \dots, e_c) contains the elements that Φ^A produced as a result of the invocations. The order in which the elements appear in this sequence is determined by the system, and may differ from the order of the invocations. In particular, $e_{\pi(i)}$ is the element that Φ^A produces for the invocation (u_i, α_i) , for some potentially secret Φ^A -specific permutation π . Finally, $\beta \in \{0, 1\}^*$ denotes some additional information that Φ^A outputs and that pertains to a batch as a whole, i.e. that is not associated with any specific invocation.

Remark 1. *The system output being generated in batches models the behaviour of certain privacy-protecting systems that do not generate an output immediately after each invocation, but rather collect several inputs before producing some output. Mix networks [12] and secret voting schemes [1], for example, operate in this way: mix networks can provide privacy only if they forward multiple messages at a time, and secret voting schemes require multiple votes for different candidates to be cast before the tally is published in order to provide privacy. However, some privacy protecting systems, for example group signatures [5], do not exhibit this behaviour, i.e. have batch size equal to one. These systems are examined in detail in Section 3.3.*

2.2 Privacy model

Let n denote the number of times the system is invoked in a given period. The correspondence between (the serial numbers of) the elements that occur in this period and the set of its users is modelled as a function $f \in \mathfrak{F}$, where $\mathfrak{F} = \{f : \{1, 2, \dots, n\} \rightarrow \mathcal{U}\}$ is the space of functions that map the serial number of each output element to the (identifier of the) user it corresponds to. The privacy notions considered in this paper describe potentially different degrees to which f remains hidden from an adversary. The adversary's goal is to identify f , or some 'interesting property' of f , possibly with respect to some subset of elements, through interaction with, or observation of, the system Φ^A . We consider the following properties of f with respect to a subset $I \subseteq \{1, 2, \dots, n\}$ of element serial numbers, which may be of interest to an adversary.

$U_{f,I} = \{f(i) : i \in I\} \subset \mathcal{U}$ denotes the *participant set*, i.e. the set of user identifiers that are associated with the elements in I .

$Q_{f,I} = \{(u, \#u_{f,I}) : u \in U_{f,I}\}$, where $\#u_{f,I} = |\{i \in I : f(i) = u\}| \in \{1, 2, \dots, |U_{f,I}|\}$, denotes *usage frequency set*, i.e. the collection of records that indicate how many elements correspond to each participant from I 's participant set.

$P_{f,I} = \{I'_1, I'_2, \dots, I'_{|U_{f,I}|}\} \vdash I$ denotes the *linking relation*, i.e. the partition of I that is induced by f . That is, $P_{f,I}$ denotes the partition that divides I into non-overlapping subsets such that, for all $i, i' \in I'_j$, $f(i) = f(i')$. Note that $\bigcup_j I'_j = I$.

$C_{f,I} = \{(1, c_1), (2, c_2), \dots, (|I|, c_{|I|})\}$ where, for all $i \in \{1, 2, \dots, |I|\}$, $c_i = |\{I' \in P_{f,I} : |I'| = i\}|$ denotes the *cardinalities of equivalence classes* as induced by the linking relation. Thus, $C_{f,I}$ is the multiset of equivalence class sizes with respect to the linking relation $P_{f,I}$, i.e. $C_{f,I}$ reveals how many elements correspond to each participant from I 's participant set. Note that, in contrast to $Q_{f,I}$, $C_{f,I}$ does not reveal the identifiers of the participants and how they correspond to equivalence class sizes and, in contrast to $P_{f,I}$, $C_{f,I}$ does not reveal how to partition the elements into equivalence classes (of the revealed sizes).

In the remainder of this paper, omission of the modifier I implies that the property under consideration refers to the entire invocations of the system, i.e. that $I = \{1, 2, \dots, n\}$. Given the above properties, and based on the principle that the adversary should be allowed to break any privacy notion *except* the one of interest, we informally derive the following privacy notions. These notions are further formalised in Section 3.

- Strong anonymity, denoted SA: A system that provides SA does not enable the adversary to learn any information about how elements correspond to users, i.e., it does not leak any information about f .
- Strong unlinkability with participation hiding, denoted SUP: A system that provides SUP does not leak any information about f beyond the number of participants $|U_f|$. In particular, it does not enable the adversary to learn any information about (a) the participant set U_f beyond its size, and (b) the linking relation P_f beyond the number of equivalence classes it contains.
- Strong unlinkability with usage hiding, denoted SUU: A system that provides SUU does not leak any information about f beyond the participant set U_f . In particular, it does not enable the adversary to learn any information about (a) the usage frequency set Q_f beyond the participant identifiers that appear in it, and (b) the linking relation P_f beyond the number of equivalence classes it contains.
- Weak unlinkability with participation hiding, denoted WUP: A system that provides WUP does not leak any information about f beyond the equivalence class sizes C_f . In particular, it does not enable the adversary to learn any information about (a) the participant set U_f beyond its size (which is the number of non-empty equivalence classes), and (b) the linking relation P_f beyond the sizes of its equivalence classes.
- Weak unlinkability with usage hiding, denoted WUU: A system that provides WUU does not leak any information about f beyond the participant set U_f and equivalence class sizes C_f . In particular, it does not enable the adversary to learn any information about (a) the usage frequency set Q_f beyond the participant identifiers that appear in it, and (b) the linking relation P_f beyond the sizes of its equivalence classes.
- Weak unlinkability, denoted WU: A system that provides WU does not leak any information about f beyond the usage frequency set Q_f . In particular, it does not enable the adversary to learn any information about the linking relation P_f beyond the sizes of its equivalence classes.
- Pseudonymity, denoted PS: A system that provides PS does not leak any information about f beyond the linking relation P_f . In particular, it does not enable the adversary to learn any information about the participant set U_f beyond what it learns from P_f (i.e. its size). This notion is called ‘pseudonymity’ because the adversary may be able to assign, to each equivalence class in P_f , a unique label, or ‘pseudonym’.
- Anonymity, denoted AN: A system that provides AN does not leak any information about f beyond the linking relation P_f and the participant set U_f . Intuitively, a system that provides AN may enable the adversary to divide all elements into non-overlapping groups, and also determine the set of participants they correspond to, but does not enable it to determine which group corresponds to which participant.
- Weak anonymity, denoted WA: A system that provides WA does not leak any information about f beyond the linking relation P_f and the usage frequency set Q_f . Similarly to AN, WA requires that the system hides the correspondence between element groups and participants. However, since knowledge of Q_f may enable the adversary to at least partially establish this correspondence, systems that provide WA (but not AN) hide less information about it than systems that provide AN (which do not reveal any information about it). In the worst case, namely the case where all equivalence classes in P_f are of *different* sizes, the adversary may be able to unambiguously determine f in its entirety, since each user in Q_f would have a unique frequency by which he can be mapped to the correct equivalence class in P_f .

Remark 2. *A system that provides any of the ‘strong unlinkability’ variants may leak the number of equivalence classes induced by the linking relation. A system that provides any of the ‘weak unlinkability’ variants may not only leak the number, but also the sizes of the equivalence classes induced by the linking relation; this information typically allows an adversary to establish the linking relation to a significantly higher degree than knowledge of merely the number of equivalence classes allows. The notions ‘pseudonymity’, ‘anonymity’, and ‘weak anonymity’ apply to systems that do not provide any unlinkability guarantee, i.e. that may enable the adversary to unambiguously establish the linking relation. Moreover, systems that provide PS or a notion with ‘participation hiding’ do not leak information about the participant identifiers beyond their number. Systems that provide, on the other hand, AN or a notion with ‘usage hiding’ may reveal the participant identifiers, but do not leak information about how often each participant invokes the system.*

2.3 Adversarial model

This section specifies the adversarial model considered in this paper. The adversary, denoted by \mathcal{A} in the sequel, adaptively controls the usage of Φ^A . Its interaction with Φ^A is modelled via queries in an experiment that a challenger arranges for \mathcal{A} . During this experiment, \mathcal{A} may (adaptively) corrupt users via the `corrupt` query; this models insider attacks.

Remark 3. *Our system model does not feature a `reveal` oracle as known in models for key establishment [6], which allows \mathcal{A} to query session information pertaining to a particular invocation. Such an oracle may appear desirable in order to model situations in which \mathcal{A} obtains context information that affects the privacy with respect to a particular invocation. Due to the fact that our adversarial model is based on left-or-right indistinguishability, however, a `reveal` oracle is not needed as \mathcal{A} knows and even chooses itself all user behaviour. The system’s internal randomness (if any), including the randomness possibly used to compute the potentially secret permutation, is considered to be a system property that faces no danger to be revealed by careless users.*

At the beginning of the experiment, the user identifier space \mathcal{U} and, if necessary, a security parameter $k \in \mathbb{N}$, are fixed and Φ^A is set up. The experiment, depicted in Figure 1, starts with the challenger selecting a bit $b \in \{0, 1\}$ uniformly at random, and by setting the initial value of the input counter `ic` and the batch counter `bc` to zero. The challenger then offers the following interfaces to \mathcal{A} , through which the system can be controlled.

- `input((·, ·), (·, ·))`: on input $((u_0, \alpha_0), (u_1, \alpha_1)) \in (\mathcal{U} \times A)^2$, the challenger increases the counter `ic` by one. We denote by $u_{0,i}$ (resp. $u_{1,i}$, $\alpha_{0,i}$, $\alpha_{1,i}$) the value of u_0 (resp. u_1 , α_0 , α_1) in \mathcal{A} ’s i th `input((u_0, α_0), (u_1, α_1))` query.
- `nextBatch()`: on reception of this query type, the challenger invokes Φ^A on input $(u_{b,i}, \alpha_{b,i})$ for $i = (\sum_{j=1}^{\text{bc}} c_j) + 1, \dots, (\sum_{j=1}^{\text{bc}} c_j) + \text{ic}$ and returns the system’s output. We say that the challenger outputs a batch of size `ic` in this case.¹ The challenger increases the batch counter `bc` by one and stores the batch size $c_{\text{bc}} := \text{ic}$. Subsequently, the input counter `ic` is reset to zero.
- `corrupt(·)`: on input $u \in \mathcal{U}$, the challenger outputs the internal state of the user identified by u . The specification of the information that is returned to \mathcal{A} is specific to Φ^A .

\mathcal{A} may issue a number of queries over these interfaces and, at some point in time, outputs a guess bit $g \in \{0, 1\}$. We say that \mathcal{A} wins the experiment if and only if $g = b$, and its advantage is given by

$$\text{Adv}_{\Phi^A, \mathcal{A}}^{X^x}(k) = \left| \Pr \left[\text{Exp}_{\Phi^A, \mathcal{A}}^{X^x-0}(k) = 0 \right] - \Pr \left[\text{Exp}_{\Phi^A, \mathcal{A}}^{X^x-1}(k) = 0 \right] \right|.$$

¹The specification of β , π , and how α it influences the output of Φ^A , is specific to Φ^A . Moreover, if \mathcal{A} is polynomially bounded, then the length of $\alpha_0, \alpha_1, \beta$ and all e_i must be polynomial in the system’s security parameter.

Experiment $\mathbf{Exp}_{\Phi^A, \mathcal{A}}^{X^x - b}(k)$	// $b \in \{0, 1\}$
	// $X \in \{\text{SA}, \text{SUP}, \text{SUU}, \text{WUP}, \text{WUU}, \text{WU}, \text{PS}, \text{AN}, \text{WA}\}$
	// $x \in \{\emptyset, \circ, +, *\}$
$g \leftarrow \mathcal{A}^{\text{input}((\cdot, \cdot), (\cdot, \cdot)), \text{nextBatch}(b), \text{corrupt}(\cdot)}$	// If \mathcal{A} breaks the restriction induced by
	// the notion X^x , the experiment is aborted.
return $g == b$	

Figure 1: Experiment $\mathbf{Exp}_{\Phi^A, \mathcal{A}}^{X^x - b}(k)$ for defining privacy notions.

In order to describe the course of an experiment we introduce the following notation. Let κ denote the number of `nextBatch` queries \mathcal{A} has issued up to the point in time it outputs g in an $\mathbf{Exp}_{\Phi^A, \mathcal{A}}^{X^x - b}(k)$ experiment. Let π_j denote the permutation applied by Φ^A for the j th batch. Furthermore, let $n = \sum_{j=1}^{\kappa} c_j$ denote the total number of `input` queries that were issued during the experiment. We define the subsets of invocation serial numbers

$$\begin{aligned}
I_1 &= \{1, 2, \dots, c_1\}, \\
I_2 &= \{c_1 + 1, c_1 + 2, \dots, c_1 + c_2\}, \\
&\vdots \\
I_\kappa &= \{c_1 + c_2 + \dots + c_{\kappa-1} + 1, c_1 + c_2 + \dots + c_{\kappa-1} + 2, \dots, n\},
\end{aligned}$$

and the ‘global inverse permutation’ Π as the permutation that maps the serial number of all elements that are output during the experiment to the serial number of their corresponding invocation. That is, Π permutes $(1, 2, \dots, n)$ such that, for all $1 \leq i \leq n$,

$$\Pi(i) = \pi_j^{-1}\left(i - \sum_{j'=1}^{j-1} c_{j'}\right) + \sum_{j'=1}^{j-1} c_{j'},$$

where $j \in \{1, 2, \dots, \kappa\}$ is such that $i \in I_j$. Finally, the functions f_0, f_1 are defined such that, for all $i \in \{1, 2, \dots, n\}$, $f_0(i) = u_{0, \Pi(i)}$ and $f_1(i) = u_{1, \Pi(i)}$, where $(u_{0, \Pi(i)}, u_{1, \Pi(i)})$ is the user pair from \mathcal{A} ’s $\Pi(i)$ th input query.

3 Hierarchy of privacy notions

This section formalises the privacy notions introduced in Section 2.2 and shows how they relate to each other. We begin by defining the following nine notions of function distinguishability.

Definition 1. *Two functions $f_0, f_1 \in \mathfrak{F}$, $f_0 \neq f_1$, are said, with respect to a subset of invocations $I \subseteq \{1, 2, \dots, n\}$, to be*

SA-indistinguishable	<i>in any case,</i>
SUP-indistinguishable	<i>if and only if $U_{f_0, I} = U_{f_1, I}$,</i>
SUU-indistinguishable	<i>if and only if $U_{f_0, I} = U_{f_1, I}$,</i>
WUP-indistinguishable	<i>if and only if $C_{f_0, I} = C_{f_1, I}$,</i>
WUU-indistinguishable	<i>if and only if $U_{f_0, I} = U_{f_1, I}$ and $C_{f_0, I} = C_{f_1, I}$,</i>
WU-indistinguishable	<i>if and only if $Q_{f_0, I} = Q_{f_1, I}$,</i>
PS-indistinguishable	<i>if and only if $P_{f_0, I} = P_{f_1, I}$,</i>
AN-indistinguishable	<i>if and only if $U_{f_0, I} = U_{f_1, I}$ and $P_{f_0, I} = P_{f_1, I}$,</i>
WA-indistinguishable	<i>if and only if $Q_{f_0, I} = Q_{f_1, I}$ and $P_{f_0, I} = P_{f_1, I}$.</i>

We are now ready to present our main privacy definitions. The intuition behind these definitions is that a privacy notion $X \in \{\text{SA}, \text{SUP}, \text{SUU}, \text{WUP}, \text{WUU}, \text{WU}, \text{PS}, \text{AN}, \text{WA}\}$ is achieved, if no \mathcal{A} , when restricted to X -indistinguishable input functions, can win the experiment with any significant advantage,

Definition 2. A privacy protecting system Φ^A is said to unconditionally (resp. statistically) provide privacy notion $X \in \{\text{SA}, \text{SUP}, \text{SUU}, \text{WUP}, \text{WUU}, \text{WU}, \text{PS}, \text{AN}, \text{WA}\}$ if and only if the adversary \mathcal{A} is restricted to invocation sequences $(u_{0,i}, \alpha_{0,i})_i$ and $(u_{1,i}, \alpha_{1,i})_i$ such that f_0 and f_1 are X -indistinguishable with respect to all $I \in 2^{\{I_1, \dots, I_\kappa\}}$, and for all such adversaries \mathcal{A} , it holds that $\text{Adv}_{\Phi^A, \mathcal{A}}^{X^x}(k) = 0$ (resp. $\text{Adv}_{\Phi^A, \mathcal{A}}^{X^x}(k) \leq \epsilon(k)$) for a negligible function $\epsilon(k)$. Moreover, Φ^A is said to computationally provide privacy notion X if and only if, for all \mathcal{A} with a running time polynomial in k , it holds that $\text{Adv}_{\Phi^A, \mathcal{A}}^{X^x}(k) \leq \epsilon(k)$ for a negligible function $\epsilon(k)$.

The above privacy notions are very strong because they require that \mathcal{A} does not obtain any advantage neither by corrupting users, nor on the basis of the parameter values that it passes in its input queries. We therefore require weaker notions that take corrupted users into account and that limit \mathcal{A} 's ability to distinguish between the two worlds on the basis of parameter values. The following notion of function indistinguishability is therefore necessary.

Definition 3. Two functions $f_0, f_1 \in \mathfrak{F}$ are said to be indistinguishable with respect to a subset of (corrupted) users \hat{U} , denoted $f_0 \approx_{\hat{U}} f_1$, if and only if, for all $i \in \{1, 2, \dots, n\}$ such that $f_0(i) \in \hat{U}$ or $f_1(i) \in \hat{U}$, $f_0(i) = f_1(i)$.

Let $\hat{U} \subseteq \mathcal{U}$ denote the set of users that \mathcal{A} has corrupted up to the point in time it outputs g , and let $A_0 = (\alpha_{0,1}, \alpha_{0,2}, \dots, \alpha_{0,n})$ and $A_1 = (\alpha_{1,1}, \alpha_{1,2}, \dots, \alpha_{1,n})$ denote the parameter sequences in the two worlds. We now present our weaker, more realistic notions.


Definition 4. A privacy protecting system Φ^A is said to unconditionally (resp. statistically, computationally) provide privacy notions X^x for $X \in \{\text{SA}, \text{SUP}, \text{SUU}, \text{WUP}, \text{WUU}, \text{WU}, \text{PS}, \text{AN}, \text{WA}\}$ and $x \in \{\circ, +, *\}$, if and only if \mathcal{A} is restricted concerning X as in Definition 2 and, additionally, concerning x as shown in Table 1, and, for all such \mathcal{A} , it holds that $\text{Adv}_{\Phi^A, \mathcal{A}}^{X^x}(k) = 0$ (resp. $\text{Adv}_{\Phi^A, \mathcal{A}}^{X^x}(k) \leq \epsilon(k)$ for a negligible function $\epsilon(k)$). Moreover, Φ^A is said to computationally provide privacy notion X^x if and only if for all \mathcal{A} with a running time polynomial in k , it holds that $\text{Adv}_{\Phi^A, \mathcal{A}}^{X^x}(k) \leq \epsilon(k)$ for a negligible function $\epsilon(k)$.

Privacy notion	Restrictions
X°	$A_0 = A_1$
X^+	$f_0 \approx_{\hat{U}} f_1$
$X^* := X^{\circ+}$	$A_0 = A_1$ and $f_0 \approx_{\hat{U}} f_1$

Table 1: Adversary restrictions for \circ , $+$, and $*$ notions. Compare to Figure 1 and Definition 4.

3.1 Relations between notions

For all $X \in \{\text{SA}, \text{SUP}, \text{SUU}, \text{WUP}, \text{WUU}, \text{WU}, \text{PS}, \text{AN}, \text{WA}\}$ it trivially holds that if \mathcal{A} is more restricted in its choices, the corresponding notion gets weaker. Therefore, X describes the strongest and X^* the weakest notion. The resulting hierarchy is displayed in the following lattice.



The privacy notions X^* are, perhaps, the most typical ones as they are concerned with the amount and type of information the system leaks *exclusively* on the basis of the identities of honest users. The notions X^+ are stronger, in the sense that a system providing some notion X^+ must not enable \mathcal{A} to distinguish between system invocations on the basis of the parameters passed to the system; the system must ensure that the output corresponding to different users is indistinguishable, irrespective of the two users' potentially different input.

The privacy notions X° can be seen as a form of ‘forward/backward privacy’, analogous to notions of forward and backward security for encryption schemes. *Forward privacy* means that, even if a user is compromised via a **corrupt** query, the user’s system interactions that occurred prior to this corruption remain private. Similarly, *backward privacy* means that system interactions of a user remain private, even if the user was corrupted prior to these interactions. Section 4.1 shows that the established privacy notion for group signatures is a forward/backward privacy notion.

The privacy notions X are very strong, in the sense that a system providing X protects the privacy of *all* users, honest and corrupted alike, *and* does not enable \mathcal{A} to distinguish between system invocations on the basis of the parameters passed to the system. That is, a system provides the notion X only if it provides X^+ and X° at the same time.

3.2 Relations among notions of one type

Figure 2 shows the relations between different privacy notions of one type. These relations follow from the facts that knowledge of Q_f implies knowledge of U_f and C_f , knowledge of P_f implies knowledge of C_f , and that knowledge of either U_f or C_f implies knowledge of $|U_f|$. The same hierarchy also applies to the privacy notions X^+ , X° , and X^* .

The following theorem states that, if a system provides two mutually distinct privacy notions, then it also provides the combined privacy notion as indicated in Figure 2. This shows that our privacy hierarchy is closed.

Theorem 3.1. *For any system Φ^A it holds that*


- (a) *if Φ^A provides AN and WU, then Φ^A provides WUU,*
- (b) *if Φ^A provides PS and WUU, then Φ^A provides WUP,*
- (c) *if Φ^A provides WUP and SUU, then Φ^A provides SUP.*

Proof. (a) Assume there exists an adversary \mathcal{A} that has advantage $\mathbf{Adv}_{\Phi^A, \mathcal{A}}^{\text{WUU}}(k) > 0$ (resp. $> \epsilon(k)$) in the experiment $\mathbf{Exp}_{\Phi^A, \mathcal{A}}^{\text{WUU}-b}(k)$. The input sequences $(u_{0,i}, \alpha_{0,i})_i$ and $(u_{1,i}, \alpha_{1,i})_i$, chosen by \mathcal{A} , are therefore such that the induced functions f_0 and f_1 are WUU-indistinguishable. That is, $C_{f_0} = C_{f_1}$ and $U_{f_0} = U_{f_1}$. We assume that $P_{f_0} \neq P_{f_1}$ and $Q_{f_0} \neq Q_{f_1}$, i.e. that \mathcal{A} cannot break AN or WU.

We construct an adversary $\bar{\mathcal{A}}$ on either WU or AN that uses \mathcal{A} as a black box. First, a random coin determines to which experiment, $\mathbf{Exp}_{\Phi^A, \mathcal{A}}^{\text{WUU}-b}(k)$ or $\mathbf{Exp}_{\Phi^A, \mathcal{A}}^{\text{AN}-b}(k)$, $\bar{\mathcal{A}}$ will be connected to. Then $\bar{\mathcal{A}}$ collects input queries from \mathcal{A} . Whenever \mathcal{A} issues a **nextBatch** query, $\bar{\mathcal{A}}$ constructs a sequence $(u_{2,i}, \alpha_{2,i})_i$ such that the resulting function f_2 is, simultaneously, AN-indistinguishable from f_0 and WU-indistinguishable from f_1 . That is, f_2 is constructed such that $U_{f_2} = U_{f_1} = U_{f_0}$, $P_{f_2} = P_{f_0}$ and $Q_{f_2} = Q_{f_1}$. Note that constructing f_2 in this way is always possible because, since f_0 and f_1 are WUU-indistinguishable, $C_{f_0} = C_{f_1}$; f_2 can be constructed, for example, by first adopting the partition P_{f_0} and then ‘filling in’ the users according to the frequencies Q_{f_1} .

Now, if playing $\mathbf{Exp}_{\Phi^A, \mathcal{A}}^{\text{WUU}-b}(k)$, then $\bar{\mathcal{A}}$ uses $(u_{0,i}, \alpha_{0,i})_i$ and $(u_{2,i}, \alpha_{2,i})_i$, in the input queries to the challenger. Otherwise, i.e. if playing $\mathbf{Exp}_{\Phi^A, \mathcal{A}}^{\text{AN}-b}(k)$, it uses the input sequences $(u_{2,i}, \alpha_{2,i})_i$ and $(u_{1,i}, \alpha_{1,i})_i$. The system’s response is forwarded to \mathcal{A} . The **corrupt** queries are simply passed through by $\bar{\mathcal{A}}$. Once \mathcal{A} outputs a guess bit g , $\bar{\mathcal{A}}$ uses g as its own guess.

Let R be the probability that \mathcal{A} outputs $g = 0$ when facing the system output on the sequence $(u_{2,i}, \alpha_{2,i})_i$. Then the advantage of $\bar{\mathcal{A}}$ in $\mathbf{Exp}_{\Phi^A, \mathcal{A}}^{\text{WUU}-b}(k)$ is $\mathbf{Adv}_{\Phi^A, \bar{\mathcal{A}}}^{\text{WU}}(k) = |\Pr[\mathbf{Exp}_{\Phi^A, \mathcal{A}}^{\text{WUU}-0}(k) = 0] - R|$ and in $\mathbf{Exp}_{\Phi^A, \mathcal{A}}^{\text{AN}-b}(k)$



Remark 4. The privacy notion ‘unobservability’, as e.g. considered in [35, 27], is not part of our model. Intuitively, ‘unobservability’ is a privacy notion that ensures that \mathcal{A} cannot determine whether or not a system invocation takes place. A system can only provide unobservability if it supports the notion of a ‘void’ invocation. That is, potentially unobservable systems must accept ‘normal’ invocations, i.e. invocations that are associated with some user/parameter pair from $\mathcal{U} \times \mathcal{A}$, and void invocations, i.e. invocations that are not associated with anything. A system can only be unobservable if it produces an element for void invocations that is indistinguishable from the elements it produces as a result of normal invocations. Since our system model described in Section 2.1 does not support systems that accept void invocations, our framework does not include an ‘unobservability’ privacy notion. However, this discussion demonstrates that extending the framework in this direction is straight-forward.

Figure 2: Relations between privacy notions. The arrow labels indicate the property about f that the system may reveal. From left to right, more information about the linking relation P_f is revealed; from top to bottom, more information about the user involvement Q_f is revealed.

$\mathbf{Adv}_{\Phi^A, \mathcal{A}}^{\text{AN}}(k) = |R - \Pr[\mathbf{Exp}_{\Phi^A, \mathcal{A}}^{\text{WUU}^{-1}}(k) = 0]|$. Thus we have $\mathbf{Adv}_{\Phi^A, \mathcal{A}}^{\text{WU}}(k) + \mathbf{Adv}_{\Phi^A, \mathcal{A}}^{\text{AN}}(k) \geq \mathbf{Adv}_{\Phi^A, \mathcal{A}}^{\text{WUU}}(k)$ and therefore $\max(\mathbf{Adv}_{\Phi^A, \mathcal{A}}^{\text{WU}}(k), \mathbf{Adv}_{\Phi^A, \mathcal{A}}^{\text{AN}}(k)) > \mathbf{Adv}_{\Phi^A, \mathcal{A}}^{\text{WUU}}(k)/2$. Thus, if $\mathbf{Adv}_{\Phi^A, \mathcal{A}}^{\text{WUU}}(k)$ is positive (resp. non-negligible), so is $\max(\mathbf{Adv}_{\Phi^A, \mathcal{A}}^{\text{WU}}(k), \mathbf{Adv}_{\Phi^A, \mathcal{A}}^{\text{AN}}(k))$ and therefore Φ^A does not provide AN or WU (or neither).

The proofs for (b) and (c) are analogous. In (b), given f_0 and f_1 that are WUP-indistinguishable, f_2 is constructed such that f_0 and f_2 are PS-indistinguishable and f_1 and f_2 are WUU-indistinguishable. In (c), given f_0 and f_1 that are SUP-indistinguishable, f_2 is constructed such that f_0 and f_2 are WUP-indistinguishable and f_1 and f_2 are SUU-indistinguishable. \square

3.3 Online systems

This section examines systems that process every input individually, i.e. systems that have a constant batch size equal to one. While such systems, which we call ‘online’ systems, enable \mathcal{A} to trivially keep track of the mapping of `input` queries and the elements produced by the system, our definition still requires \mathcal{A} to determine whether it is interacting in the left or the right world. Nevertheless, the mere fact that \mathcal{A} can unambiguously determine which output elements correspond to which invocation serial numbers, has implications to the introduced hierarchy of privacy notions.

Lemma 3.2. Consider two functions $f_0, f_1 \in \mathfrak{F}$. If $U_{f_0, I} = U_{f_1, I}$ for all $I \in \{\{1\}, \{2\}, \dots, \{n\}\}$, then $f_0 = f_1$.

Proof. Assume that $f_0 \neq f_1$, i.e. that there exists at least one $i \in \{1, 2, \dots, n\}$ such that $f_0(i) \neq f_1(i)$. Then $U_{f_0, \{i\}} \neq U_{f_1, \{i\}}$, contradicting the assumption. \square \square

The implication of Lemma 3.2 is that, for online systems, there exist no functions f_0 and f_1 , $f_0 \neq f_1$, that are X -indistinguishable for any $X \in \{\text{SUU}, \text{WUU}, \text{AN}, \text{WA}\}$; moreover, the notions SUP, WUP, and PS coincide. This can be seen easily: if, for all $I \in 2^{\{1, 2, \dots, n\}}$, $|U_{f_0, I}| = |U_{f_1, I}| = 1$, then the equivalence classes of the partition of elements must be identical in f_0 and f_1 . Hence, $P_{f_0} = P_{f_1}$ and $C_{f_0} = C_{f_1}$. The resulting collapsed hierarchy of privacy notions is sketched in Figure 3.

$$\text{SA} \xrightarrow{P_f} \{\text{SUP}, \text{WUP}, \text{PS}\} \xrightarrow{f} \perp$$

Figure 3: Privacy notions for online systems.

3.4 Shuffling vs. non-shuffling systems

This section analyses when and how the act of shuffling is beneficial for the privacy of the system. We informally argue that, (a) depending on the extent to which the output elements (e and β) leak information about their corresponding invocations, shuffling may or may not be beneficial, and (b) for certain systems, secret shuffling is sufficient to provide WU.

Consider a system where each of its outputs e does not depend on any invocation other than the one that caused it. We call such systems *stateless* because, except possibly for the computation of β , they do not have to keep any state about the history of events within a batch. A system is furthermore said to be *shuffling* if, before outputting elements, it applies a random permutation to them. Shuffling certainly makes only sense if neither β , nor the elements themselves, reveal the order in which the system produces its outputs. We note that a stateless *non-shuffling* system Φ^A can be transformed into a shuffling one by means of a shuffler, as follows.

Definition 5. A *shuffler* $\mathcal{S}(\Phi^A)$ is an algorithm that, on input a batch of invocations $(u_1, \alpha_1), (u_2, \alpha_2), \dots, (u_c, \alpha_c) \in (\mathcal{U} \times A)^c$, invokes a system Φ^A on this batch, and outputs $((e_{\pi'(1)}, e_{\pi'(2)}, \dots, e_{\pi'(c)}), \beta)$, where $((e_1, e_2, \dots, e_c), \beta)$ is Φ^A 's output and π' is a uniformly at random chosen permutation.

Consider a system whose output elements e enable \mathcal{A} to extract information about the parameters α , but no information about the user identifiers, or their corresponding invocations. Also assume that β leaks no further information. The strongest privacy notion that this system may provide is SA° , because X° notions are the strongest privacy notions that require the parameter sequences A_0 and A_1 to be identical (hence, by definition, the parameters do not reveal the secret bit b). Observe that this holds irrespective of whether or not the system shuffles²; since \mathcal{A} can map output elements to their corresponding invocations (by parameter values) anyway, shuffling does not buy any privacy advantage for such systems.

Consider, on the other hand, a system whose output elements e do enable \mathcal{A} to extract information about the identifier of the user of the corresponding invocation, but where neither the elements e , nor β reveal any information about the corresponding input parameters. In this case, \mathcal{A} learns Q_{f_b} . That is, \mathcal{A} can recognise which user invoked the system how often. Nevertheless, such a system may still provide some privacy. In fact, the strongest privacy notion that a *non-shuffling* system of this type may achieve is WA because, since the system does not shuffle, it does not hide the partition P_{f_b} —its outputs are partitioned in the same way as its inputs. (Online systems also provide WA as long as neither the elements e , nor β leak any information about the input parameters α .) A *shuffling* system of this kind may, on the other hand, maximally achieve WU, since the shuffle may hide P_f (but not Q_f). If \mathcal{A} is also able to obtain information

²As long as, if it shuffles, the chosen permutations do not depend on user identifiers.

about parameter values from β (which is not linked to individual elements), then only WA° is possible (for non-shuffling systems), and WU° (for shuffling systems).³

Theorem 3.3. *If a stateless system Φ^A provides WA^x , then $\mathcal{S}(\Phi^A)$ provides WU^x , unless Φ^A 's output contains extractable information about the per-user partitioning of the invocations.*

Proof. Assume that $\mathcal{S}(\Phi^A)$ does not provide WU^x . Then \mathcal{A} is able to extract some information about b based on either P_{f_b} , or some other property of f_b . However, since Φ^A provides WA^x , \mathcal{A} extracts some information about P_{f_b} . Since π' ensures that the *order* in which $\mathcal{S}(\Phi^A)$ outputs elements carries no information, Φ^A 's output must enable such information extraction. The result follows. \square

Consider a stateless system that, for every batch I , encodes the variables (Q_{f_I}, P_{f_I}) in β .⁴ If the system's output elements e do not carry extractable information about their corresponding invocations, i.e. user identifiers and parameter values (this could be achieved, for example, by probabilistic encryption), then this system provides WA^x , where x depends on the details of the system. This, in conjunction with Theorem 3.3, shows that, for any stateless system that does not leak information about parameter values (e.g. by means of probabilistic encryption) *and* that does not leak information about f_I in β beyond (Q_{f_I}, P_{f_I}) , superimposing a shuffler \mathcal{S} on the system is *sufficient* to provide WU .

4 Applications

This section places the privacy definitions concerning group signature, anonymous communication, and secret voting systems into the hierarchy introduced in the previous section.

4.1 Group signatures

Group signatures represent an important class of privacy protecting system. A group signature system consists of four algorithms (GKg , GSig , GVf , Open), as follows [5].

- The randomised *group key generation algorithm* GKg takes as input a security parameter $k \in \mathbb{N}$, and returns a tuple $(gpk, gmsk, gsk)$, where gpk is the *group public key*, $gmsk$ is the *group manager's secret key*, and gsk is an vector of keys where $gsk[u]$ is the *secret signing key* of user identified by $u \in \mathcal{U}$, and where the length of the vector is polynomially bounded in k .
- The randomised *group signing algorithm* GSig takes as input a secret signing key $gsk[u]$ and a message $m \in \mathcal{M}$, where \mathcal{M} is the system's message space, and returns a signature of m under $gsk[u]$ ($u \in \mathcal{U}$).
- The deterministic *group signature verification algorithm* GVf takes as input the group public key gpk , a message m , and a candidate signature σ for m , and returns either 1 or 0.
- The deterministic *opening algorithm* Open takes as input the group manager secret key $gmsk$, a message m , and a signature σ of m , and returns an identifier $u \in \mathcal{U}$ or the symbol \perp to indicate failure.

Full anonymity: We briefly revisit the definition of ‘full anonymity’ as defined in [5]. Full anonymity is defined by means of an FA -experiment between an adversary \mathcal{A}_{FA} and a challenger, which proceeds as follows. Initially, \mathcal{A}_{FA} is given gsk and gpk , and access to an opening oracle $\text{Open}(\cdot, \cdot)$ that, on input a message/signature pair (m, σ) , outputs $\text{Open}(gmsk, m, \sigma)$. At some point in time, \mathcal{A}_{FA} outputs a triple (u_0, u_1, m') and the challenger returns $\sigma' = \text{GSig}(gsk[u_b], m')$, where $b \in \{0, 1\}$ is chosen uniformly at random. \mathcal{A}_{FA} is then required to output a guess for b ; before doing this, it may again query the $\text{Open}(\cdot, \cdot)$

³If \mathcal{A} learns the input parameter α that corresponds to an individual element, then no privacy is possible. This is because \mathcal{A} would then know a (u, α) pair that caused a specific element, but cannot be forced to have the same pair present in both worlds. In fact, in reality, systems that reveal both users and their input should not be considered privacy-preserving.

⁴In practice, an adversary may infer the partition and/or the invocation frequency from context, background, and/or network layer information.

oracle, albeit not on (\cdot, σ') . \mathcal{A}_{FA} wins if its guess is correct, and the system is said to provide ‘full anonymity’, denoted FA, if no adversary can win the game with non-negligible advantage over random guessing.

We now examine how FA relates to the privacy notions as defined in the previous section. Translated to our model, the parameter space of a group signature scheme is its message space. That is, $A_{\text{gs}} = \mathcal{M}$. Since users compute and independently release signatures by themselves, the adversary is able to observe isolated system invocations. Thus, group signature schemes are online systems, and, hence, the only applicable privacy notions are SA and PS.⁵ The specification of the `corrupt`(\cdot) query for group signatures systems is as follows.

- `corrupt`(\cdot): on input $u \in \mathcal{U}$, the challenger outputs the secret key of the user identified by u , i.e. $gsk[u]$.

Note that our model lacks the `Open` oracle. This oracle, which has the flavour of a ‘reveal’ oracle, does not take official invocations of the system as input but rather private computations of the adversary. That is, according to the model in [5], a cryptographic group signature scheme can be invoked locally by the adversary and is not a central system as we assume in our model. Our model is nevertheless useful in analysing group signatures. To this end, we simply add the `Open` oracle to our model without modifications. The anonymity notion without the `Open` oracle, which would be the straightforward notion in our model, is called CPA anonymity [9].

- `Open`(\cdot, \cdot): on input (m, σ) where σ is a valid group signature of m this oracle applies the `Open`($gmsk, m, \sigma$) algorithm and returns an identifier $u \in \mathcal{U}$ or the symbol \perp to indicate failure.

We now show why certain privacy notions do not apply to group signature systems, while others are equivalent.

Lemma 4.1. *No group signature system provides SA^* , PS^* , SA^+ , or PS^+ . Moreover, for group signature systems, SA° and PS° are equivalent, and SA and PS are distinct, privacy notions.*

Proof. No group signature system can provide SA^* , PS^* , SA^+ , PS^+ because the signed message is published along with its group signature; \mathcal{A} can trivially win an $\text{Exp}_{\mathbb{F}, \mathcal{A}}^{X^x - b}(k)$ experiment of providing different messages in the left and the right world. SA° and PS° are equivalent by Lemma 4.2, and SA and PS are distinct by Remark 5. □

Lemma 4.2. *FA, computational SA° , and computational PS° , are equivalent.*

Proof. We first show that PS° implies FA. An $\mathcal{A}_{\text{PS}^\circ}$ adversary with access to an \mathcal{A}_{FA} adversary starts by corrupting all users, obtaining their secret keys, which it feeds into \mathcal{A}_{FA} . \mathcal{A}_{FA} ’s `Open` queries are passed on to the `Open` oracle. \mathcal{A}_{FA} makes only one sign query which $\mathcal{A}_{\text{PS}^\circ}$ passes on as `input` query to the challenger and gives the response back to \mathcal{A}_{FA} . The restriction $P_{f_0} = P_{f_1}$ is satisfied, since any two functions with a singleton domain induce the same partition on their domain. $\mathcal{A}_{\text{PS}^\circ}$ answers in the same way as \mathcal{A}_{FA} .

We show that FA also implies SA° by constructing an adversary \mathcal{A}_{FA} that has a non-negligible advantage in the FA-experiment, given black-box access to an adversary $\mathcal{A}_{\text{SA}^\circ}$ with non-negligible advantage in a (computational) SA° -experiment. \mathcal{A}_{FA} proceeds as follows. It uniformly at random selects a value $i \in \mathbb{N}$ such that $1 \leq q(k)$, where $q(k)$ is the upper bound on the number of queries that $\mathcal{A}_{\text{SA}^\circ}$ may issue. Using its knowledge of gsk , it answers $\mathcal{A}_{\text{SA}^\circ}$ ’s first $i - 1$ `input`((u_0, m), (u_1, m)) queries with `GSig`($gsk[u_0], m$). Before answering $\mathcal{A}_{\text{SA}^\circ}$ ’s i th `input`((u_0, m'), (u_1, m')) query, it queries the challenger with the triple (u_0, u_1, m') with values taken from $\mathcal{A}_{\text{SA}^\circ}$ ’s query. \mathcal{A}_{FA} returns the challenger’s answer σ' to $\mathcal{A}_{\text{SA}^\circ}$. Using its knowledge of gsk , it answers $\mathcal{A}_{\text{SA}^\circ}$ ’s remaining `input`((u_0, m), (u_1, m)) queries with `GSig`($gsk[u_1], m$), and finally outputs the same value as $\mathcal{A}_{\text{SA}^\circ}$. Using a standard hybrid argument [25], it can be shown that \mathcal{A}_{FA} ’s success probability is $(1/2) + \delta/q$, where q and δ are the number of queries issued by $\mathcal{A}_{\text{SA}^\circ}$ and $\mathcal{A}_{\text{SA}^\circ}$ ’s advantage, respectively. Since, as shown in Section 3.3, SA° implies PS° , a group signature system that provides FA also provides PS° . □

⁵We choose PS to represent $\{\text{SUP}, \text{WUP}, \text{PS}\}$.

The fact that there exists only a single (computational) forward/backward privacy notion for group signatures, explains, perhaps, why [5] claims that ‘anonymity and unlinkability are technically the same property’.

Remark 5. *From our framework it is now obvious that weaker privacy notions for group signatures exist; it is possible to refrain from forward/backward privacy, and optionally in addition tolerate the group signatures of the same signer being linkable. Traceable group signature schemes were to our knowledge first considered in [30]. We modify a traceable scheme from [10] to construct an instance of a group signature scheme that provides PS but not SA. The required modification is minor, as it merely consists in setting a particular parameter of the scheme to one. We now briefly review the modified scheme; for a complete description see [10]. The scheme uses a bilinear group pair (G_1, G_2) consisting of cyclic groups G_1 and G_2 of prime order p with an efficiently computable isomorphism from G_2 to G_1 , and an efficiently computable non-degenerate bilinear map $e : G_1 \times G_2 \rightarrow G_t$. For (G_1, G_2) the strong Diffie-Hellman assumption (see. [10]) has to hold. A public group key gpk is given by a triple of group elements (g_1, g_2, g_2^γ) , where $g_1 \in G_1$, $g_2 \in G_2$ are randomly chosen from the respective groups and act as generators, and γ is secretly and uniformly at random chosen from \mathbb{Z}_p . Then a private signing key for a user U_i is given by $(A_i = g_1^{1/(\gamma+x_i)}, x_i)$ for a uniformly at random chosen element $x_i \in \mathbb{Z}_p$. Furthermore, a hash function $H : \{0, 1\}^* \rightarrow G_1 \times G_2$ is given. The signing procedure is as follows:*

1. $(u, v) \leftarrow H(gpk)$
2. Choose $\alpha \leftarrow \mathbb{Z}_p$ uniformly at random and compute $T_1 \leftarrow u^\alpha$, $T_2 \leftarrow A_i v^\alpha$
3. Compute $c, s_\alpha, s_x, s_\delta$ as a witness indistinguishable proof of knowledge for correct computation of T_1, T_2 with respect to the private key A_i . This is done with Fiat-Shamir heuristic [20] and involves the message being signed.

The signature of a message is then given by $(T_1, T_2, c, s_\alpha, s_x, s_\delta)$. Now any two group signatures of the same signer can be linked by computing $e(A_i, u) = e(T_2, u)/e(T_1, v)$ which is a value that, since the parameters u, v are common to the entire group, depends only on the signer’s private key. Since this value cannot be traced back to any particular public key, some privacy remains.

Remark 6. *We are not aware of any group signature scheme that provides SA but not FA at the same time, i.e. a scheme without forward-/backward privacy. The following example, however, demonstrates the existence of such a scheme. Consider a group signature scheme that provides FA. We modify this scheme as follows. Every signer is given a pseudorandom number generator whose seed is part of the user’s secret key and, in order to sign a message, the user replaces all random choices by pseudo-randomness. As a result, every user behaves deterministically and, as long as the adversary does not know a user’s seed, the produced signatures are computationally indistinguishable from those based on true randomness. Moreover, once the adversary calls **corrupt** on a user and learns his seed, all past and future signatures of this user become linkable; the modified scheme no longer provides forward-/backward privacy. In fact, it provides computational SA. Note that this scheme might well apply to smart-card group signature implementations where replacing randomness by pseudo-randomness is a common option.*

4.2 Anonymous communication

Anonymous communication systems are modelled as protocols that transmit messages from senders to receivers. The input to an anonymous communication system is a sequence of triples of the form $(\sigma, \rho, m) \in \mathcal{U} \times \mathcal{U} \times \mathcal{M}$, where $\sigma, \rho \in \mathcal{U}$ are identifiers of the sender and the intended recipient, respectively, \mathcal{M} is the system’s message space, and $m \in \mathcal{M}$ is the message that is to be transmitted from σ to ρ . The output that is associated to an input triple of this form, is the bitstring that the system produces as a result of this input, and that the adversary can observe.

For anonymous communication systems we define two variants of the base experiment $\mathbf{Exp}_{\Phi^A, \mathcal{A}}^{X^x - b}(k)$, depending on whether the experiment is intended to capture the privacy of senders or the privacy of recipients.

In particular, the variant that captures sender privacy is denoted by $\mathbf{Exp}_{\Phi^A, \mathcal{A}}^{\text{S-priv-b}}(k)$, and the variant that captures recipient privacy by $\mathbf{Exp}_{\Phi^A, \mathcal{A}}^{\text{R-priv-b}}(k)$. In both variants, the parameter space is $A_{\text{ac}} = \mathcal{M} \times \mathcal{U}$. The difference between the two variants is the way in which the challenger assigns the sender and receiver roles to the users indicated in an $\text{input}((\cdot, \cdot), (\cdot, \cdot))$ query; in all other respects the two variants are identical to the base experiment.

Definition 6. *On reception of an $\text{input}((u_0, (m_0, u'_0)), (u_1, (m_1, u'_1)))$ query in the context of an $\mathbf{Exp}_{\Phi^A, \mathcal{A}}^{\text{S-priv-b}}(k)$ (resp. $\mathbf{Exp}_{\Phi^A, \mathcal{A}}^{\text{R-priv-b}}(k)$) experiment, the challenger first increases the input counter i_c by one, and then re-members $(u_b, (m_b, u'_b))$ (resp. $(u'_b, (m_b, u_b))$) as (u_{i_c}, α_c) .*

In other words, the parameter $\alpha = (m, u) \in A_{\text{ac}}$ either specifies a message together with (the identifier of) its intended recipient (in the context of an $\mathbf{Exp}_{\Phi^A, \mathcal{A}}^{\text{S-priv-b}}(k)$ experiment), or a message together with (the identifier of) its sender (in the context of an $\mathbf{Exp}_{\Phi^A, \mathcal{A}}^{\text{R-priv-b}}(k)$ experiment). We now extend our generic definition for the context of anonymous communication.

Definition 7. *An anonymous communication system $\Phi^{A_{\text{ac}}}$ is said to unconditionally (resp. statistically, computationally) provide ‘sender- X ’, denoted S/X (resp. ‘recipient- X ’, denoted R/X) for some privacy notion $X \in \{Y^*, Y^\circ, Y^+, Y\}$ where $Y \in \{\text{SA}, \text{SUP}, \text{SUU}, \text{WUP}, \text{WUU}, \text{WU}, \text{PS}, \text{AN}, \text{WA}\}$, if and only if it unconditionally (resp. statistically, computationally) provides X with respect to an $\mathbf{Exp}_{\Phi^A, \mathcal{A}}^{\text{S-priv-b}}(k)$ (resp. $\mathbf{Exp}_{\Phi^A, \mathcal{A}}^{\text{R-priv-b}}(k)$) experiment.*

It trivially follows from the definition that S/SA^+ and R/SA^+ , as well as S/SA^* and R/SA^* , are equivalent notions. We define one more privacy notion, namely unlinkability, denoted UL . UL is specific to anonymous communication systems, and, like SA^* , its sender and recipient versions are equivalent. Unlinkability is the notion that ensures that \mathcal{A} cannot learn anything about f beyond what follows from knowledge of how many messages each sender sent, and how many messages each receiver received. Let $A_0 = ((\cdot, u'_{0,1}), (\cdot, u'_{0,2}), \dots, (\cdot, u'_{0,n}))$ and $A_1 = ((\cdot, u'_{1,1}), (\cdot, u'_{1,2}), \dots, (\cdot, u'_{1,n}))$ denote parameter sequences issued by the adversary during an $\mathbf{Exp}_{\Phi^A, \mathcal{A}}^{X^x-b}(k)$ experiment.

Definition 8. *An anonymous communication system $\Phi^{A_{\text{ac}}}$ is said to unconditionally (resp. statistically, computationally) provide privacy notion UL^* (resp. $\text{UL}^\circ, \text{UL}^+, \text{UL}$), called unlinkability, if and only if it unconditionally (resp. statistically, computationally) provides WU^* (resp. $\text{WU}^\circ, \text{WU}^+, \text{WU}$) with respect to an $\mathbf{Exp}_{\Phi^A, \mathcal{A}}^{\text{S-priv-b}}(k)$ and an $\mathbf{Exp}_{\Phi^A, \mathcal{A}}^{\text{R-priv-b}}(k)$ experiment where, for all $i \in \{1, 2, \dots, n\}$, $u'_{0,i} = u'_{1,i}$.*

4.2.1 Existing notions

We briefly revisit the privacy notions defined in [27] in order to examine how they relate to the ones defined above. [27] defines privacy by means of an experiment between an adversary and a challenger. The adversary specifies in advance two collections C^0 and C^1 of triples of the form $(\sigma, \rho, m) \in \mathcal{U}^2 \times \mathcal{M}$.⁶ The two collections are then given to the challenger, which selects a bit $b \in \{0, 1\}$ uniformly at random, and simulates $\Phi^{A_{\text{ac}}}$ on input the triples in C^b . The adversary, given $\Phi^{A_{\text{ac}}}$'s output, then produces a guess g for b and wins if and only if $g = b$; its advantage is defined in the usual way.

Let $S^b = \{\sigma \in \mathcal{U} : (\sigma, \cdot, \cdot) \in C^b\}$ and $R^b = \{\rho \in \mathcal{U} : (\cdot, \rho, \cdot) \in C^b\}$ denote the set of senders and receivers according to C^b . For all $\sigma \in S^b$ (resp. $\rho \in R^b$), we denote by $\text{sent}_\sigma^b = (\uplus m \in \mathcal{M} : (\sigma, \cdot, m) \in C^b)$ (resp. $\text{rcvd}_\rho^b = (\uplus m \in \mathcal{M} : (\cdot, \rho, m) \in C^b)$) the multiset of messages sent by σ (resp. received by ρ) according to C^b . The different privacy notions defined in [27] arise due to restrictions imposed on the adversary in the construction of C^0 and C^1 . In particular, an anonymous communication system is said to provide privacy notion $N \in \{\text{SUL}, \text{RUL}, \text{UL}, \text{SA}, \text{RA}, \text{SA}^*, \text{RA}^*, \text{SRA}, \text{UO}\}$ if no adversary, when restricted to choose C^0 and C^1 such that the conditions shown in Table 2 are satisfied, has a non-negligible advantage in the above experiment.

⁶In [27] these collections are called ‘message matrices’, and are encoded as matrices.

Privacy notion	Label	Conditions
<i>Sender Unlinkability</i>	$\overline{\text{SUL}}$	$S = S^0 = S^1, R = R^0 = R^1,$ $\forall \sigma \in S, \text{sent}_\sigma^0 = \text{sent}_\sigma^1 ,$ and $\forall \rho \in R, \text{rcvd}_\rho^0 = \text{rcvd}_\rho^1$
<i>Receiver Unlinkability</i>	$\overline{\text{RUL}}$	$S = S^0 = S^1, R = R^0 = R^1,$ $\forall \sigma \in S, \text{sent}_\sigma^0 = \text{sent}_\sigma^1,$ $\forall \rho \in R, \text{rcvd}_\rho^0 = \text{rcvd}_\rho^1 ,$ and
<i>Unlinkability</i>	$\overline{\text{UL}}$	$S = S^0 = S^1, R = R^0 = R^1,$ $\forall \sigma \in S, \text{sent}_\sigma^0 = \text{sent}_\sigma^1 ,$ and $\forall \rho \in R, \text{rcvd}_\rho^0 = \text{rcvd}_\rho^1 $
<i>Sender Anonymity</i>	$\overline{\text{SA}}$	$R = R^0 = R^1$ and, $\forall \rho \in R \text{rcvd}_\rho^0 = \text{rcvd}_\rho^1$
<i>Receiver Anonymity</i>	$\overline{\text{RA}}$	$S = S^0 = S^1$ and, $\forall \sigma \in S, \text{sent}_\sigma^0 = \text{sent}_\sigma^1$
<i>Strong Sender Anonymity</i>	$\overline{\text{SA}^*}$	$R = R^0 = R^1$ and, $\forall \rho \in R \text{rcvd}_\rho^0 = \text{rcvd}_\rho^1 $
<i>Strong Receiver Anonymity</i>	$\overline{\text{RA}^*}$	$S = S^0 = S^1$ and, $\forall \sigma \in S, \text{sent}_\sigma^0 = \text{sent}_\sigma^1 $
<i>Sender-Receiver Anonymity</i>	$\overline{\text{RA}^*}$	$ C^0 = C^1 $
<i>Unobservability</i>	$\overline{\text{UO}}$	none

Table 2: Conditions according to privacy definitions by Hevia et al.

4.2.2 Comparison to existing notions

The adversarial model in [27] does not consider corrupted users, and does not consider adaptive adversaries. Translated to our system model, this amounts to the setting where \mathcal{A} issues only a single `nextBatch` query, and no `corrupt`(\cdot) queries. Due to this discrepancy of the adversarial models, the privacy notions defined in this paper are not directly comparable to the ones defined in [27]. If, however, \mathcal{A} is restricted to observe only a single batch and is allowed no corruptions, then the following notions are equivalent.

Lemma 4.3. *If, during an $\text{Exp}_{\Phi^{\mathcal{A}}, \mathcal{A}}^{\text{S-priv-b}}(k)$ or $\text{Exp}_{\Phi^{\mathcal{A}}, \mathcal{A}}^{\text{R-priv-b}}(k)$ experiment, \mathcal{A} does not issue any `corrupt`(\cdot) queries and at most a single `nextBatch` query, then $\overline{\text{SUL}}$ and S/WU , $\overline{\text{RUL}}$ and R/WU , $\overline{\text{UL}}$ and UL , $\overline{\text{SA}}$ and S/SA , $\overline{\text{RA}}$ and R/SA , $\overline{\text{SA}^*}$ and S/WU^+ , $\overline{\text{RA}^*}$ and R/WU^+ , and $\overline{\text{RA}^*}$ and $(\text{S/R})\text{SA}^+$, are equivalent privacy notions.*

Proof. Consider an adversary $\mathcal{A}_{\overline{\text{SUL}}}$. We construct an adversary \mathcal{A}_{WU} that wins an $\text{Exp}_{\Phi^{\mathcal{A}}, \mathcal{A}}^{\text{S-priv-b}}(k)$ experiment if and only if $\mathcal{A}_{\overline{\text{SUL}}}$ wins. Let C^0 and C^1 denote the collections output by $\mathcal{A}_{\overline{\text{SUL}}}$. Due to the applicable restrictions $S = S^0 = S^1, R = R^0 = R^1, |\text{sent}_\sigma^0| = |\text{sent}_\sigma^1|$ for all $\sigma \in S$, and $\text{rcvd}_\rho^0 = \text{rcvd}_\rho^1$ for all $\rho \in R$ (see Table 2), for each triple $(\sigma_0, \rho_0, m_0) \in C^0$ there exists exactly one ‘corresponding’ triple in C^1 , i.e. a triple (σ_1, ρ_1, m_1) such that $\rho_1 = \rho_0$ and $m_1 = m_0$. For each triple in $(\sigma_0, \rho_0, m_0) \in C_0$, \mathcal{A}_{WU} issues the query `input`(($\sigma_0, (\rho_0, m_0)$), ($\sigma_1, (\rho_1, m_1)$)) where σ_1, ρ_1 and m_1 are the values from the corresponding triple in C^1 . \mathcal{A}_{WU} then issues a `nextBatch` query, forwards the challenger’s output to $\mathcal{A}_{\overline{\text{SUL}}}$, and, finally outputs a guess that is identical to $\mathcal{A}_{\overline{\text{SUL}}}$ ’s guess. Clearly, \mathcal{A}_{WU} wins if and only if $\mathcal{A}_{\overline{\text{SUL}}}$ wins.

Consider adversary \mathcal{A}_{WA} of an $\text{Exp}_{\Phi^{\mathcal{A}}, \mathcal{A}}^{\text{S-priv-b}}(k)$ experiment. We construct an adversary $\mathcal{A}_{\overline{\text{SUL}}}$ who wins if and only if \mathcal{A}_{WA} wins. For every `input`(($\sigma_0, (\rho_0, m_0)$), ($\sigma_1, (\rho_1, m_1)$)) query issued by \mathcal{A}_{WA} , $\mathcal{A}_{\overline{\text{SUL}}}$ adds the triple $(\sigma_0, (\rho_0, m_0))$ to C^0 and the triple $(\sigma_1, (\rho_1, m_1))$ to C^1 . When \mathcal{A}_{WA} issues the `nextBatch` query, $\mathcal{A}_{\overline{\text{SUL}}}$ starts its experiment with C^0 and C^1 . Note that, due to the restrictions that apply in the experiment of \mathcal{A}_{WA} , the collections C^0 and C^1 , too, satisfy the required restrictions. $\mathcal{A}_{\overline{\text{SUL}}}$ then forwards the challenger’s output to \mathcal{A}_{WA} , and, finally outputs a guess that is identical to \mathcal{A}_{WA} ’s guess. Clearly, $\mathcal{A}_{\overline{\text{SUL}}}$ wins if and only if \mathcal{A}_{WA} wins. Thus, $\overline{\text{SUL}}$ and S/WU are equivalent privacy notions. Showing the validity of the other equivalences is analogous. \square \square

Remark 7. *The above privacy notions form a hierarchy, described in [27], that is separate from the one*

described in Section 3. Moreover, [27] demonstrates that one can construct anonymous communication systems that offer a particular privacy notion by appropriately augmenting a system that provides a weaker notion, with encryption techniques and/or dummy traffic. Since, according to the model in [27], the adversary may observe only a single communication batch, these transformations do not necessarily suffice in the face an adversary that may adaptively influence the system over multiple communication batches, i.e. in the model considered in this paper.

Since in our model, \mathcal{A} may issue multiple `nextBatch` queries, the notions $\overline{\text{S/WU}}$, $\overline{\text{R/WU}}$, $\overline{\text{UL}}$, $\overline{\text{S/SA}}$, $\overline{\text{R/SA}}$, $\overline{\text{S/WU}^+}$, $\overline{\text{R/WU}^+}$, and $\overline{(\text{S/R})\text{SA}^+}$, are all strictly stronger than $\overline{\text{SUL}}$, $\overline{\text{RUL}}$, $\overline{\text{UL}}$, $\overline{\text{SA}}$, $\overline{\text{RA}}$, $\overline{\text{SA}^*}$, $\overline{\text{RA}^*}$, and $\overline{\text{RA}^*}$, respectively. Consider, for example, an anonymous communication system that provides notion $\overline{\text{RA}}$, i.e. a system where, for an adversarially chosen batch of communications (where certain conditions hold), the adversary may be able to determine which messages were received by which receivers, but no information beyond this. In contrast to this, the system would only provide notion $\overline{\text{S/SA}}$ if it does not leak any such information even for multiple, adversarially and adaptively chosen batches of communication (where certain conditions hold). This suggests that an anonymous communication system provides a privacy notion in $\{\overline{\text{S/WU}}$, $\overline{\text{R/WU}}$, $\overline{\text{UL}}$, $\overline{\text{S/SA}}$, $\overline{\text{R/SA}}$, $\overline{\text{S/WU}^+}$, $\overline{\text{R/WU}^+}$, $\overline{(\text{S/R})\text{SA}^+}\}$ only if it is, effectively immune to ‘intersection’ (also known as ‘disclosure’ or ‘hitting set’) attacks [2, 29], while privacy notions in $\{\overline{\text{SUL}}$, $\overline{\text{RUL}}$, $\overline{\text{UL}}$, $\overline{\text{SA}}$, $\overline{\text{RA}}$, $\overline{\text{SA}^*}$, $\overline{\text{RA}^*}$, $\overline{\text{SRA}}\}$ can be achieved without such immunity.

4.3 Voting systems

A typical voting process has, among other things, a phase where users cast their votes, called the ‘voting phase’, and a phase where the results of the elections are computed and published, called the ‘tallying phase’. Although secret voting systems have to fulfill a multitude of requirements, in this paper we focus only on *ballot secrecy*, since this is a privacy notion. Moreover, we assume that each eligible voter is allowed to cast at most one vote.⁷

Voting systems are modelled as follows. During the voting phase, the system accepts pairs $(u, w) \in \mathcal{U} \times \mathcal{C}$, where \mathcal{U} is the set of (the identifiers of) the eligible voters, and \mathcal{C} is the set of candidates in the election. The adversary can cause a ballot to be cast by issuing an `input` query, and indicate the end of the voting phase by issuing a `nextBatch` query. The bitstring β that the system produces as part of its output contains the tally and additional information that may be needed for verification. Since no further voting is possible after the tally is computed, voting systems produce only a single batch.

Ballot secrecy can be modelled in two ways within our framework. The first, perhaps more intuitive one, is as follows. First, the parameter space of voting systems is defined as $A_{\text{vs}}^1 = \mathcal{C}$. The challenger, on reception of an `input` $((u_0, w_0), (u_1, w_1))$ query, then invokes the system with input (u_b, w_b) . The `corrupt` (u) queries are not needed. Any (cryptographic) secrets known to a user should be given to the adversary straight away, as for coercion resistance, the privacy may not depend on this secret information. Particularly those voting schemes that are based on blind signatures (e.g. [23]) fit in this model.

Lemma 4.4. *Assuming that each eligible voter may vote at most once, the strongest achievable privacy notion for voting schemes is, in the above modelling, unconditional SUP.*

Proof. The tally reveals the number of participants, as this matches the total number of ballots. Thus, SA-indistinguishable functions can be distinguished from the system output β . \square \square

Unfortunately, not all voting systems can be modelled in the above way. This is because the elements produced by certain systems are homomorphically encrypted ballots whose correspondence to voters is not at all hidden; we call the systems that produce such elements ‘homomorphic voting systems’ [1].⁸ Homomorphic voting systems provide privacy by ensuring that the identity of the candidate that is encoded in each (encrypted) ballot remains hidden.

⁷Surprisingly, there seem to exist no formal treatments of ballot secrecy in the literature; our formal model below follows the intuition provided by the notion of ‘perfect ballot secrecy’ that is informally described in [31].

⁸In order to provide ‘individual verifiability’, a property that enables each voter to verify that his own vote has been accounted for in the tally, these systems do not hide the correspondence between outputs and invocations.

Remark 8. *A basic homomorphic voting scheme operates as follows. Each element it produces is an encryption of the identity of the candidate that was selected in the corresponding invocation, where encryption is performed using an asymmetric homomorphic encryption scheme. In the tallying phase, it first computes the election result based on the homomorphism of the encrypted elements, and then outputs all elements (i.e. encrypted ballots), together with the election result and a proof of correct decryption.*

Homomorphic voting systems are modelled in our framework through a reversal of the roles of voters and candidates. That is, instead of treating ‘ballot secrecy’ as a privacy property concerning voters, it is treated as a privacy property concerning candidates. A system is then said to provide ballot secrecy if it hides the identities of the users that have voted for any given candidate; it is easy to see that this is equivalent to the case where it hides the identity of the candidate any given user has voted for. This reversal, however, enables us to alternatively define \mathcal{A} ’s interaction with the challenger based on candidates rather than voters. In particular, the parameter space for voting system as $A_{\text{vs}}^2 = \mathcal{U}$, where \mathcal{U} is the identifier set of eligible voters, and \sqsubseteq is used as the set of users. In particular, on reception of an `input`((w_0, u_0), (w_1, u_1)) query, the challenger invokes the system $\Phi^{A_{\text{vs}}^2}$ on input (u_b, w_b). As previously explained, a `corrupt`() query is not defined. In order to break privacy, \mathcal{A} would have to identify two (encrypted) ballots for the same candidate, or reveal the candidate of one ballot. We now show why the strongest possible notion for ballot secrecy is weak unlinkability.

Lemma 4.5. *In the above modelling, computational WU is the strongest achievable ballot secrecy notion for homomorphic voting systems.*

Proof. Homomorphic encryption schemes hide the plaintext only from computationally bounded adversaries, and the tally reveals Q_{f_b} . □ □

Lemma 4.5 and the results on stateless and shuffling systems discussed in Section 3.4 suggest, that voting schemes that are stateless, based on homomorphic encryption with an underlying encryption scheme is IND-CPA [25] and uniformly permute the encrypted ballots before publishing them will provide WU and therefore achieve the highest possible privacy for such a system.

5 Conclusions and open questions

We presented an application-agnostic hierarchy of privacy notions that describe potentially different degrees to which the correspondence between digital elements and the users that cause their appearance remains hidden from an adversary. Compared to an earlier version new privacy notions have been introduced to complete the structure. A theorem showed the orthogonality and completeness of the new hierarchy.

Previously isolated privacy notions pertaining to group signature, anonymous communication, and secret voting systems have been placed into this hierarchy, and thereby effectively made comparable. It is possible that privacy definitions pertaining to other system types, such as anonymous credentials, data anonymisation systems, and sensor information systems, can also be placed into our framework. Examining this possibility is subject of future research.

Our framework provides valuable insights into the relations and structure of different privacy notions, and highlights a largely unexplored space of such notions. Exemplarily, we identified two new notions for group signatures and pointed out how group signatures that match these definitions look like. Identifying useful schemes providing other ‘new’ notions, perhaps by trading off privacy against other features, is subject of future research. Of particular interest are techniques that transform systems achieving a given privacy notion into systems that provide another, perhaps stronger one in the adaptive adversarial model considered in this paper. We expect that the framework will also be useful in the construction and analysis of ‘multi-layer’ privacy protecting systems, i.e. systems that combine, for example, anonymous communication with group signing.

Finally, constructing ‘soft’, probabilistic privacy metrics for each of the notions in our framework is subject of current research. Such metrics will enable us to compare privacy protecting systems with considerably

higher granularity than is possible with definitions that are based on asymptotic polynomial indistinguishability.

References

- [1] Ben Adida. Advances in Cryptographic Voting Systems. PhD thesis, Massachusetts Institute of Technology, 2006.
- [2] Dakshi Agrawal and Dogan Kesdogan. Measuring Anonymity: The Disclosure Attack. *IEEE Security & Privacy*, 1(6):27–34, 2003.
- [3] Christer Andersson and Reine Lundin. On the Fundamentals of Anonymity Metrics. In *The Future of Identity in the Information Society*, IFIP International Federation for Information Processing. Springer Science & Business Media, 2008.
- [4] Mihir Bellare, Anand Desai, Eron Jorjani, and Phillip Rogaway. A Concrete Security Treatment of Symmetric Encryption. In *Proceedings of the 38th Annual Symposium on Foundations of Computer Science (FOCS'97)*, pages 394–403, 1997.
- [5] Mihir Bellare, Daniele Micciancio, and Bogdan Warinschi. Foundations of Group Signatures: Formal Definitions, Simplified Requirements, and a Construction Based on General Assumptions. In Eli Biham, editor, *Advances in Cryptology – EUROCRYPT 2003, International Conference on the Theory and Applications of Cryptographic Techniques, Warsaw, Poland, May 4–8, 2003, Proceedings*, volume 2656 of *Lecture Notes in Computer Science*, pages 614–629. Springer, 2003.
- [6] Mihir Bellare and Phillip Rogaway. Entity Authentication and Key Distribution. In Douglas R. Stinson, editor, *Advances in Cryptology – CRYPTO '93*, volume 773 of *Lecture Notes in Computer Science*, pages 232–249. Springer, 1993.
- [7] Ron Berman, Amos Fiat, and Amnon Ta-Shma. Provable Unlinkability Against Traffic Analysis. In Ari Juels, editor, *Financial Cryptography, 8th International Conference, FC 2004, Key West, FL, USA, February 9–12, 2004. Revised Papers*, volume 3110 of *Lecture Notes in Computer Science*, pages 266–280. Springer Verlag, Berlin, February 2004.
- [8] Jens-Matthias Bohli and Andreas Pashalidis. Relations Among Privacy Notions. In *Financial Cryptography and Data Security, FC'09*, volume 5628 of *Lecture Notes in Computer Science*, pages 362–380. Springer, 2009.
- [9] Dan Boneh, Xavier Boyen, and Hovav Shacham. Short Group Signatures. In Matthew K. Franklin, editor, *Advances in Cryptology – CRYPTO 2004*, volume 3152 of *Lecture Notes in Computer Science*, pages 41–55. Springer, 2004.
- [10] Dan Boneh and Hovav Shacham. Group signatures with verifier-local revocation. In *ACM Conference on Computer and Communications Security, CCS 2004*, pages 168–177. ACM, 2004.
- [11] J. Camenisch and A. Lysyanskaya. An Efficient System for Non-transferable Anonymous Credentials with Optional Anonymity Revocation. In Birgit Pfitzmann, editor, *Advances in Cryptology – EUROCRYPT 2001, International Conference on the Theory and Application of Cryptographic Techniques, Innsbruck, Austria, May 6–10, 2001, Proceedings*, volume 2045 of *Lecture Notes in Computer Science*, pages 93–118. Springer Verlag, Berlin, 2001.
- [12] David Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–90, 1981.

- [13] Sebastian Clauß. A Framework for Quantification of Linkability Within a Privacy-Enhancing Identity Management System. In Günter Müller, editor, *Emerging Trends in Information and Communication Security, International Conference, ETRICS 2006, Freiburg, Germany, June 6-9, 2006, Proceedings*, volume 3995 of *Lecture Notes in Computer Science*, pages 191–205. Springer Verlag, 2006.
- [14] Sebastian Clauß and Stefan Schiffner. Structuring anonymity metrics. In *DIM '06: Proceedings of the second ACM workshop on Digital identity management*, pages 55–62, New York, NY, USA, 2006. ACM Press.
- [15] C. Díaz, S. Seys, J. Claessens, and B. Preneel. Towards measuring anonymity. In R. Dingledine and P. F. Syverson, editors, *Privacy Enhancing Technologies, Second International Workshop, PET 2002, San Francisco, CA, USA, April 14–15, 2002, Revised Papers*, number 2482 in *Lecture Notes in Computer Science*, pages 54–68. Springer Verlag, Berlin, 2002.
- [16] Claudia Díaz. Anonymity Metrics Revisited. In *Anonymous Communication and its Applications, number 05411 in Dagstuhl Seminar Proceedings*, 2005.
- [17] Claudia Díaz, Joris Claessens, Stefaan Seys, and Bart Preneel. Information Theory and Anonymity. In B. Macq and J. Quisquater, editors, *Proceedings of the 23rd Symposium on Information Theory in the Benelux*, pages 179–186, 2002.
- [18] Matthew Edman, Fikret Sivrikaya, and Bülent Yener. A Combinatorial Approach to Measuring Anonymity. In *Proceedings of the 2007 IEEE International Conference on Intelligence and Security Informatics*. IEEE, 2007.
- [19] Joan Feigenbaum, Aaron Johnson, and Paul Syverson. Probabilistic analysis of onion routing in a black-box model. In *Proceedings of the 2007 ACM workshop on Privacy in electronic society*, pages 1–10. ACM Press, 2007.
- [20] Amos Fiat and Adi Shamir. How to Prove Yourself: Practical Solutions to Identification and Signature Problems. In Andrew M. Odlyzko, editor, *Advances in cryptology – CRYPTO '86*, volume 263 of *Lecture Notes in Computer Science*, pages 186–194. Springer, 1988.
- [21] Lars Fischer, Stefan Katzenbeisser, and Claudia Eckert. Measuring Unlinkability Revisited. In Marianne Winslett and Ragib Hasan to be published, editors, *Proceedings of the 2008 ACM Workshop on Privacy in the Electronic Society, WPES 2008, Alexandria, VA, USA, October 27, 2008*, pages 00–00. ACM, 2008.
- [22] Matthias Franz, Bernd Meyer, and Andreas Pashalidis. Attacking Unlinkability: The Importance of Context. In Nikita Borisov and Philippe Golle, editors, *Privacy Enhancing Technologies, 7th International Symposium, PET 2007 Ottawa, Canada, June 20-22, 2007, Revised Selected Papers*, volume 4776 of *Lecture Notes in Computer Science*, pages 1–16. Springer Verlag, Berlin, 2007.
- [23] Atsushi Fujioka, Tatsuaki Okamoto, and Kazuo Ohta. A Practical Secret Voting Scheme for Large Scale Elections. In *ASIACRYPT '92*, number 718 in *Lecture Notes in Computer Science*, pages 244–251. Springer, 1993.
- [24] Benedikt Gierlichs, Carmela Troncoso, Claudia Díaz, Bart Preneel, and Ingrid Verbauwhede. Revisiting A Combinatorial Approach Toward Measuring Anonymity. In *Proceedings of the 2008 ACM Workshop on Privacy in the Electronic Society, WPES 2008, Alexandria, VA, USA, October 27, 2008*, pages 111–116. ACM, 2008.
- [25] Shafi Goldwasser and Silvio Micali. Probabilistic Encryption. *J. Comput. Syst. Sci.*, 28(2):270–299, 1984.
- [26] Joseph Y. Halpern and Kevin R. O’Neill. Anonymity and Information Hiding in Multiagent Systems. *Journal of Computer Security*, 13(3):483–514, 2005.

- [27] Alejandro Hevia and Daniele Micciancio. An Indistinguishability-based Characterization of Anonymous Channels. In Nikita Borisov and Ian Goldberg, editors, *Privacy Enhancing Technologies Symposium, Eighth International Workshop, PET 2008, Leuven, Belgium, July 23–25, 2008, Revised Selected Papers*, volume 5134 of *Lecture Notes in Computer Science*, pages 24–43. Springer Verlag, Berlin, 2008.
- [28] D. Hughes and V. Shmatikov. Information Hiding, Anonymity and Privacy: a Modular Approach. *Journal of Computer Security*, 12(1):3–36, 2004.
- [29] Dogan Kesdogan, Dakshi Agrawal, and Stefan Penz. Limits of Anonymity in Open Environments. In F.A.P. Petitcolas, editor, *Information Hiding, 5th International Workshop, IH 2002, Noordwijkerhout, The Netherlands, October 7–9, 2002, Revised Papers*, volume 2578 of *Lecture Notes in Computer Science*, pages 53–69. Springer Verlag, Berlin, 2003.
- [30] Aggelos Kiayias, Yiannis Tsiounis, and Moti Yung. Traceable Signatures. In *Advances in Cryptology – EUROCRYPT 2004*, volume 3027 of *Lecture Notes in Computer Science*, pages 571–589. Springer, 2004.
- [31] Aggelos Kiayias and Moti Yung. Self-tallying Elections and Perfect Ballot Secrecy. In David Naccache and Pascal Paillier, editors, *5th International Workshop on Practice and Theory in Public Key Cryptosystems, PKC 2002 Paris, France, February 12–14, 2002 Proceedings*, number 2274 in *Lecture Notes in Computer Science*, pages 141–158. Springer Verlag, Berlin, 2002.
- [32] Greg Maitland, Jason Reid, Ernest Foo, Colin Boyd, and Ed Dawson. Linkability in Practical Electronic Cash Design. In Josef Pieprzyk, Eiji Okamoto, and Jennifer Seberry, editors, *Information Security, Third International Workshop, ISW 2000, Wollongong, NSW, Australia, December 20–21, 2000, Proceedings*, volume 1975 of *Lecture Notes in Computer Science*, pages 149–163. Springer Verlag, 2000.
- [33] Richard E. Newman, Ira S. Moskowitz, Paul Syverson, and Andrei Serjantov. Metrics for Traffic Analysis Prevention. In Roger Dingledine, editor, *Privacy Enhancing Technologies, Third International Workshop, PET 2003, Dresden, Germany, March 26–28, 2003, Revised Papers*, volume 2760 of *Lecture Notes in Computer Science*, pages 48–65. Springer Verlag, Berlin, 2003.
- [34] Andreas Pashalidis. Measuring the Effectiveness and the Fairness of Relation Hiding Systems. In *Proceedings of the First International Workshop on Multimedia, Information Privacy and Intelligent Computing Systems*. IEEE Press, 2008.
- [35] Andreas Pfitzmann and Marit Köhntopp. Anonymity, Unobservability, and Pseudonymity: A Proposal for Terminology. In Hannes Federrath, editor, *Designing Privacy Enhancing Technologies, International Workshop on Design Issues in Anonymity and Unobservability, Berkeley, CA, USA, July 25–26, 2000. Proceedings*, volume 2009 of *Lecture Notes in Computer Science*, pages 1–9. Springer, 2000.
- [36] A. Serjantov. On the Anonymity of Anonymity Systems. Phd thesis, 2004.
- [37] A. Serjantov and G. Danezis. Towards an Information Theoretic Metric for Anonymity. In R. Dingledine and P. F. Syverson, editors, *Privacy Enhancing Technologies, Second International Workshop, PET 2002, San Francisco, CA, USA, April 14–15, 2002, Revised Papers*, volume 2482 of *Lecture Notes in Computer Science*, pages 41–53. Springer Verlag, Berlin, 2002.
- [38] Vitaly Shmatikov and Ming-Hsiu Wang. Measuring relationship anonymity in mix networks. In *WPES '06: Proceedings of the 5th ACM workshop on Privacy in electronic society*, pages 59–62, New York, NY, USA, 2006. ACM Press.
- [39] S. Steinbrecher and S. Köpsell. Modelling Unlinkability. In R. Dingledine, editor, *Privacy Enhancing Technologies, Third International Workshop, PET 2003, Dresden, Germany, March 26–28, 2003, Revised Papers*, volume 2760 of *Lecture Notes in Computer Science*, pages 32–47. Springer Verlag, Berlin, 2003.

- [40] Gergely Tóth, Zoltán Hornák, and Ferenc Vajda. Measuring Anonymity Revisited. In Sanna Liimatainen and Teemupekka Virtanen, editors, *Proceedings of the Ninth Nordic Workshop on Secure IT Systems*, pages 85–90, Espoo, Finland, November 2004.