

Privacy Games for Syntactic Privacy Notions

Abstract. It is well understood that the huge volumes of data captured in recent years have the potential to underpin significant research developments in many fields. But, to realise these benefits, all relevant parties must be comfortable with how this data is shared. At the heart of this is the notion of privacy — which is recognised as being somewhat difficult to define. Previous authors have shown how privacy notions such as anonymity, unlinkability and pseudonymity might be combined into a single formal framework. In this paper we use and extend this work by defining privacy games for individual and group privacy within distributed environments. More precisely, for each privacy notion, we formulate a game that an adversary has to win in order to break the notion. Via these games, we aim to clarify understanding of, and relationships between, different privacy notions; we also aim to give an unambiguous understanding of adversarial actions. Additionally, we extend previous work via the notion of *unobservability*.

1 Introduction

The continuous digitalisation of our daily lives has led to the collection of huge volumes of data. The sharing and use of that data can be beneficial in many ways. However, challenges arise from the processing of such data, one of which is guaranteeing privacy to individuals and groups.

One of the complexities in this respect is defining privacy. By its very nature, privacy is multi-dimensional and multi-faceted, and can mean different things to different people — resulting in a variety of notions capturing different perspectives on and properties of privacy.

Early efforts at establishing a consistent terminology in this regard include those of Pfitzmann and colleagues [30,29], who developed a precise terminology for privacy notions. Despite such efforts, an open problem with the formal treatment of privacy notions remained: discrepancies between formal models led to ill-defined relationships and incomparable notions. In addition, various attacks [24,21,15,32,16] influenced and challenged the formulation of new privacy notions. Bohli and Pashalidis addressed those disconnected notions in [7], as a first step towards uniting multiple privacy notions into a single formal framework.

We build upon the work of Bohli and Pashalidis [7]. As such, we give consideration to the challenges faced in distributed, heterogeneous contexts. In particular, for each privacy notion, we formulate a game an adversary has to win in order to break the notion. Via these games, we aim to clarify understanding of, and relationships between, different privacy notions; we also aim to give an *unambiguous understanding* of adversarial actions. In addition, we extend the

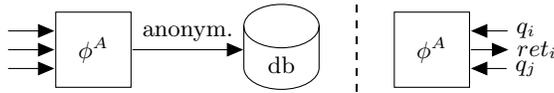


Fig. 1: Illustration of the two privacy-preserving data release modes: non-interactive (left) and interactive (right).

privacy framework of Bohli and Pashalidis [7] by introducing the notion of *unobservability*. Concretely, this notion guarantees to prevent an adversary from distinguishing if a system invocation has taken place or not (by observing a real or ‘generated’ outcome element). As such, this notion gives similar guarantees as the prominent notion of differential privacy [10]. However, instead of dealing with probabilistic privacy notions, we focus on syntactical notions for privacy. These notions are associated to the privacy properties that a system can inherit. More precisely, they are defined on properties of the system release.

The presented notions and games may be used in privacy-preserving systems to reason about their privacy properties. As such, we consider systems that may be stateful, stateless, or online. Moreover, we focus on individual and group privacy for interactive, non-interactive and statistical release scenarios — thereby reflecting a wide scope of applications.

The kind of system of interest to us is manifested by the deployment of Trustworthy Remote Entities (TREs) [4,20,28] — which we use later to motivate our contribution. A TRE provides input privacy by leveraging novel hardware security extensions, such as Intel’s SGX¹ [3,25,8] extensions and, further, may deploy differential privacy mechanisms to achieve release privacy.

To summarise, we claim the following contributions. We present games illustrating notions for individual and group privacy in distributed settings. Moreover, we extend [7] by introducing the notion of *unobservability*.

2 Background

2.1 Privacy-Preserving Data Release

In general, we consider two natural models in the privacy-preserving release of data: *non-interactive* and *interactive*, as illustrated in Figure 1. In the non-interactive setting, we assume data is to be shared by a trusted data curator, who publishes a ‘safe’ or ‘sanitised’ version of the database. In the literature this process is also known as *de-identification* or *anonymisation* of data. Traditionally, this can be achieved by the deployment of generalisation, suppression or perturbation mechanisms. Non-interactive data release gained in popularity in the mid 2000s with developments such as k -anonymity [31], ℓ -diversity [24] and t -closeness [21]. In the interactive setting, we assume users submit queries to the data curator. The data curator applies the query function on the database to obtain the so-called *true* answer. To protect privacy, for example, random noise is added to the true answer, with the subsequent result being returned to the user. Although there exist many other mechanisms to preserve privacy,

¹ Intel Software Guard Extensions (SGX), <https://software.intel.com/en-us/sgx>

the utilisation of data perturbation seems both natural and appealing (and predominantly used for explicit notions [12]). This can come in two flavours: input perturbation and output perturbation. While the former adds noise to the data before submission to the data curator, the latter adds noise only after the query mapping is applied (and added to the true answer). The first approach is stricter (in the sense of preserving privacy, although it limits utility) and also works if the data curator is untrusted. The latter requires a trustworthy data curator, but typically achieves more utility. Statistical data release is a variant — or, in fact, a specialisation — of the interactive approach: the query function computes and returns a single value containing statistical information about the database entries. The notion of differential privacy [10] emerged as a means of ensuring privacy in such contexts.

2.2 Syntactic Privacy vs. Probabilistic Privacy

By nature, privacy is multi-dimensional and multi-faceted, and can mean different things to different people. Therefore, defining privacy and notions representing specific properties related to privacy must be undertaken with care. Some privacy definitions ignore the knowledge of an adversary, and only give a requirement of the structure of the resulting data. These notions pertain to the class of *syntactic* methods. As such, they are mainly concerned with the syntactic form of the released outcome. In our context of concern, ignorance of a priori knowledge of an adversary, often called *background knowledge* or *auxiliary information*, may easily lead to privacy breaches if the adversary is able to obtain and combine information from other sources.

Many privacy attacks² focus on the model of non-interactive privacy, due to its release-and-forget approach (*i.e.* ‘re-sanitisation’ of a released (sanitised) version of a database in case of a privacy breach is not possible). Mechanisms to provide privacy in this setting have evolved, but this has subsequently led to the problem that those algorithms became increasingly complex. Examples for non-interactive release privacy include [31,24,21].

Another class of privacy notions pertains to *probabilistic* methods. Contrary to syntactic methods, these notions abstract the probabilities that govern the relationship between inputs to outputs. For example, to provide privacy in the interactive setting, the notion of *differential privacy* (DP) [10] emerged. Soon, it became established as a *de facto* standard due to its mathematically rigorous definition and provable privacy guarantees. This powerful notion (or, more precisely, property of an algorithm) allows for privacy-preserving views over statistical databases, while providing provable guarantees that the distribution of noisy query answers changes only negligibly with addition or deletion of any tuple in the database. However, while the mathematics of differential privacy are straightforward, an intuitive understanding of the notion can be elusive, as

² Typical attacks targeting non-interactive mechanisms include linking attacks (for example, homogeneity attacks [24], skewness attacks [21] or similarity attacks [21]), composition attacks [15], minimality attacks [32] or the attack by deFinetti [16].

can characterisations of its relationships with, for example, anonymity, participation hiding or unlinkability. Therefore, many variations have been proposed to provide different assurances³.

Other frameworks that aim to unify privacy notions (mainly focused on probabilistic privacy) include Pufferfish [19] and the framework of membership privacy [22]. Importantly, definitions within these frameworks satisfy the fundamental axioms for modern privacy design guidelines: *transformation invariance* and *convexity* [17]. Within the context of syntactic privacy notions, to provide a clear understanding of privacy, its properties and implications to systems, it is important to accurately model adversarial actions. Game-based notions formulate games illustrating the necessary actions for an adversary to break said notion. Examples include [2,5]. Such game-based definitions have been widely used in the cryptography community to provide simpler and more easily verifiable proofs. Similarly, via our games, we aim to clarify understanding of, and relationships between, different privacy notions.

3 Analysis of the Framework of Bohli and Pashalidis

Many privacy models (such as those of [31], [24], [21] and [10]) tend to assume that data to be shared originates from a centralised database. In this context, users must always trust the data curator to keep their data secure (and, consequently, preserve their privacy). Such models, however, do not capture use cases associated with distributed environments — in which data remains split over several devices and is stored in databases *distributed* over a large network. Yet, constant improvements in technology allow increasingly more operations to be executed in distributed environments (which seek to preserve the privacy of those individuals involved in such processing). The following provides an overview of the privacy model and notions of Bohli and Pashalidis [7] and, importantly, elaborates and motivates our distributed privacy model.

3.1 Privacy Model

While their system model is clearly focused on a centralised system, Bohli and Pashalidis [7] do not explicitly state any assumptions pertaining to the underlying (centralised or distributed) structure of their privacy model. Indeed, the privacy definitions follow the syntactic approach: as such, certain privacy notions may be achieved dependent on the information an adversary gains access to (or is allowed to learn). They do, though, consider systems that follow the non-interactive release mode. Batches of input invocations trigger a sequence of outputs, where output values represent a permutation of the input values. In particular, this permutation obfuscates the association to a user identifier

³ Contributions that present criticisms and relaxations of the pure mode of DP are summarised in [9] and include (ϵ, δ) -DP [11], probabilistic (ϵ, δ) -DP [23], distributional privacy [6], computational DP [27], no-free lunch DP [18], personalised DP [13], and mutual-information DP [9].

— essentially, the goal, in order to release data in a privacy-preserving way. Furthermore, [7] gives considerations to stateful, stateless and online systems.

In their system model, Bohli and Pashalidis model a function that maps the serial number of the output values to the identifier of the user associated with it. By means of this function, they define privacy properties that abstract the relation between the output element and the user identifier. In particular, they define the notions introduced below and show how these notions are related to each other.

3.2 Privacy Notions

The privacy notions considered in this paper are formally defined in Section 6, but we introduce them first in terms of natural language to aid comprehension (additionally, Figure 15 in Appendix C further illustrates the privacy notions in terms of our system model). Informally, we say that an adversary (simplified) is allowed to interact with a system in one of two ways:

- (a) it inputs information, or⁴
- (b) it obtains an outcome — in our case, a vector of input values (permuted, to ensure privacy) and an aggregated value (computed over the input values).

Then, the objectives of the privacy notions of Bohli and Pashalidis [7] are as follows:

- *Strong anonymity* prevents an adversary from learning any interesting relationship, from the observation of the outcome, with regards to the input.
- *Participation hiding* prevents an adversary from learning the relation of an outcome value to a specific user.
- *Unlinkability* prevents an adversary from learning (*i.e.* identifying) any outcome elements that are related to each other (*i.e.* pertaining to the same user).
- *Pseudonymity* allows an adversary to associate groups of outcome elements with a pseudonym, but prevents her from learning the actual user identifier.
- *Anonymity* prevents an adversary from learning of the relationship between groups of elements and/or groups of user identifiers.

Additionally, the objective of our extended notion of unobservability is as follows:

- *Unobservability* prevents an adversary to distinguish if a system invocation has taken place or not (by observing a real or ‘generated’ outcome element).

3.3 An Extension to a Semi-Distributed Model

A fully-distributed system represents a collection of single systems. As such, data is to be stored at each of the distributed parties and the privacy properties are manifested in each system’s individual release and in the combined release.

⁴ ‘or’ stands here in an inclusive context. Furthermore, we want to clarify that this is not a one-shot interaction, but a user (or adversary) may access the system oracles multiple times. A more detailed interpretation is given in Section 5 when we introduce the privacy games.

A user interacts via separate **input** and **view** interfaces of each individual system. Consequently, the benefits of such an approach are: (a) data remains at the data provider; (b) a reduced communication overhead; and (c) there is no single point of failure. Examples for such (privacy-preserving) systems include deep learning [1], crowdsourcing statistics [14], federated learning⁵ [26], or usage patterns collections (*i.e.* the approach by Apple⁶ in iOS 10).

Yet, reasoning about privacy properties in such systems may be complex and challenging. One needs to take into consideration that an adversary may obtain, corrupt or control a certain threshold of system outcomes. Furthermore, privacy notions need to be abstracted to include such ‘collaborative’ outcomes and relations between outcomes of individual systems. Thus, as a first step we consider a semi-distributed model. Data is still split over multiple parties, however, users gain access to the system via a global **input** and **view** interface. As such, privacy properties may be considered only over the whole ‘combined’ system release. In other words, for the privacy modelling, we consider the underlying structure as a ‘black box’ and argue (with our notions and games) about the release of this black box. Yet, we consider data to be distributed among multiple systems (as such, that the games to abstract the privacy notions take such a distributed modelling into consideration).

Our modelling of the privacy notions and games was undertaken with a focus on the requirements of a distributed system. In particular, the kind of system of interest to us are Trustworthy Remote Entities (TREs) [28], which operate in a distributed context and enable privacy-preserving release of data, while simultaneously ensuring verifiability (and therefore trustworthiness) of its algorithms. A brief overview of TREs is given in Appendix B. Cloud computing, the smart grid, collaborative learning (as well as others — see [4]) represent use cases that operate predominantly in distributed settings and benefit from the deployment of privacy-preserving systems (such as a TRE). Consequently, our modelling was dedicated to such kinds of systems, yet, we do not exclude other systems that fit within our system model. In the following, we review considerations of individual and group privacy, as well as the underlying structure (*i.e.* stateful, stateless or online) of the kind of systems of interest to us.

Individual and Group Privacy An adversary may be able to infer knowledge about an individual by directly or by indirectly targeting her (*e.g.* by grouping users according to similar properties such as their behaviour, location or user type; and gaining knowledge from the observation of the group). The former case we further denote as breaching an *individual’s privacy* and the latter as breaching *group privacy*. To illustrate, consider the following examples: a breach of individuals’ privacy would mean that an adversary is able, for any single element e_x a system releases, to infer its relation to the associated user u_x . Drawn from a genomics use case [4], e_x may represent a disease that user u_x has.

⁵ <https://research.googleblog.com/2017/04/federated-learning-collaborative.html>

⁶ <https://wired.com/2016/06/apples-differential-privacy-collecting-data>

A breach of group privacy would mean that an adversary is able to relate two or more elements e_x, e_y , where e_x pertains to a user u_x , e_y pertains to a user u_y , and both users are member of a group $u_x, u_y \in \mathbf{u}_g$. In this case (assuming that e_x and e_y are types of diseases), the adversary may infer that e_x and e_y pertains to a group of similar diseases. Another example would be: e_x is a specific disease, and if the adversary infers that $u_x, u_y \in \mathbf{u}_g$, then she may learn that u_y is somehow associated to e_x . More concretely, e_x is a specific inheritable disease and \mathbf{u}_g is a family, with u_x and u_y members of that family — u_y , therefore, may also be susceptible to the disease. Hence, we model our notions and games such that they give consideration to both individual and group privacy.

Stateful, Stateless and Online Systems The design and application of privacy notions depends heavily on the underlying structure of a system, be it *stateful*, *stateless*, or *online*. We discuss examples following a permutation as privacy-preserving mechanism — yet any other mechanisms may apply similarly.

In a stateful system, released data may be shuffled according to a state- (or data-) dependent, potentially secret, permutation. Additionally, statistical queries may result in state-dependent answers (for example, identical queries may not be answered, result in the same answer, or be dependent on the privacy mechanism to include an increased noise value). In a stateless system, release data may be shuffled according to a permutation chosen uniformly at random. Because stateless systems may not rely on any state (from the perspective of the system) queries submitted to it are independent and any input has no influence on the behaviour of the system. In the design of privacy notions and privacy-preserving mechanisms we must cope with such restrictions (for example, submission of two identical queries might disclose information). Finally, in an online system, data is processed and output individually (*i.e.* before an invocation of another input), due to a limited the buffer size of the system. Therefore, a permutation may not be applicable and other sources (such as noise addition) may need to be applied to preserve privacy in such a setting.

4 System Model

We consider systems, denoted by ϕ , that are sequentially invoked a finite number of times, and require that they return values e_i and an aggregated value β_i computed over a finite number of inputs α_i . Further, we require that each input α_i must be associated to a unique user identifier u_i and a computational party p_i . Inputs α_i associated to the same user identifier u_i may collectively be denoted in vector notation $\boldsymbol{\alpha}_i$. The same applies for the output values e_i , with \mathbf{e}_i for outputs associated to the same u_i . Tuples of the form (u_i, α_i) are stored within datasets at the specified party p_i . The inputs to, as well as the outputs from, ϕ are drawn from a parameter space $\alpha_i, e_i \in A$ that is specific to the system. Similarly defined is the parameter space of the aggregated output $\beta_i \in A$. Each user identifier u_i and party identifier p_i is assumed to be unique, and drawn from an identifier space.

Invocations of ϕ , in the form of a query q_i , produce output batches (e_i, β_i) of potentially varying size. Thus, a query q_i is essentially a mapping of input batches to (e_i, β_i) , denoted by π^ϕ for values e_i and π_{q_i} for the aggregated value β_i . In particular, ϕ accepts input batches of the form

$$(u_1, \alpha_1, p_1), (u_2, \alpha_2, p_2), \dots, (u_c, \alpha_c, p_c)$$

for c invocations of the input oracle $\mathbf{input}(\cdot, \cdot, \cdot)$. Furthermore, for each invocation of the output oracle, $\mathbf{view}^\pi(\cdot)$, ϕ applies a potentially secret permutation π^ϕ on the datasets held by each party p_i and outputs a sequence

$$(e_1, \dots, e_c), \beta$$

where $e \leftarrow \pi^\phi(\alpha)$ and $\beta \leftarrow \pi_{q_i}(\alpha)$.

In the notation above, we assumed that each input (u_i, α_i) is distributed to a different party p_i . However, we do not limit our model in such contexts, but allow a data owner⁷ to freely decide to which party p_i she wants to send her inputs. Our model is assumed to handle any data distribution along parties p_i .

Building on the privacy model of Bohli and Pashalidis [7], our model includes the possibility of a ‘void’ invocation of the system, *i.e.* an invocation of ϕ , where the outcome (e_i, β_i) is unrelated to any u_i, α_i or p_i and drawn uniformly at random from the parameter space A . However, we require that the outcome of a void invocation must remain indistinguishable to the outcome of a ‘normal’ invocation of ϕ . This property is essential in the context of system unobservability.

Similarly to Bohli and Pashalidis, we consider systems that are stateful, stateless or online. As such, we take these properties into account when designing privacy games in Section 6. Moreover, our stated notions and games must remain independent of the data (*i.e.* the inputs and output of ϕ), and abstract to the mappings $\pi^\phi : \alpha_i \rightarrow e_i$ and $\pi_{q_i} : \alpha_i \rightarrow \beta_i$.

Our model allows single users or a collection of users to reason about their privacy assurances. Precisely, a notion or game can be applied to capture privacy properties about a mapping between:

1. a single user and an attribute: $u_x \rightarrow e_x$
2. a single user and multiple attributes: $u_x \rightarrow (e_{x,1}, \dots, e_{x,c})$
3. multiple users and multiple attributes: $u_x \rightarrow e_x, u_y \rightarrow (e_{y,1}, \dots, e_{y,c})$

where the last case denotes group privacy.

We denote the function f as the mapping between output values and user identifiers, and the function f^* as the mapping between dataset entries and user identifiers. More formally, $f : e_i \rightarrow u_i$ takes as an input an element (or vector) e_i (or e_i) of ϕ and maps it to the corresponding user u_i . Similarly, this holds

⁷ Data submitted to the system may include (or, pertain to) several data subjects, that are independent of the actual data owner. In our modelling, so far, we assume that each input value α_i is associated to a user identifier u_i , which identifies the data subject (which may or may not be the data owner). We further assume that data subjects have given consent to data owners to perform certain data-mining tasks on the committed data, while preserving the privacy of the data subjects. In any case, the data owner decides with which party p_i to share the data.

for the internal processing of ϕ , $f^* : \alpha_i \rightarrow u_i$, with α_i (or α_i) (stored in the dataset of any party p_i). Through the processing of a query q_i (*i.e.* essentially through π_{q_i}), ϕ obfuscates the function f . An adversary \mathcal{A} must not be able to directly learn f by observing the inputs (u_i, α_i, p_i) to ϕ or the released outcome (e_i, β_i) . Nevertheless, the adversary’s goal remains to determine f , or any other interesting property about f . Yet, the privacy notions and games reveal, to varying degrees, information about f .

We denote $p_i \in P$, with P the set of all computational parties. Inputs of ϕ are stored in datasets $\alpha_i \rightarrow p_i.db$, given any possible distribution. Moreover, we say $\alpha = \cup_{i=1}^P p_i.db$, that is, α is the set of all inputs $\alpha_i \in \alpha$ for all parties p_i . Similarly, it holds that $e = \cup_{i=1}^n e_i$ represents the union of all outcome values. Following the definitions of Bohli and Pashalidis, we denote the following properties of the outcome in the form of sets (the relationship between those sets is further illustrated on the right-hand side of Figure 15 in Appendix C):

Definition 1. (*Participant Set*): $U_f = \{f(e_i) : e_i \in e_u\}$ denotes the set of participants in ϕ . Each element $e_i \in e$ is uniquely associated to a user identifier $f(e_i) = u_i$.

Definition 2. (*Usage Frequency Set*): $Q_f = \{(u_i, \#u_i) : u_i \in U_f\}$, with $\#u_i = |\{e_i \in e : f(e_i) = u_i\}|$, denotes the relation of the number of elements $e_i \in e$ each user u_i is associated with. In other words, $\#u_i$ is the number of e_i user u_i has contributed to ϕ .

Definition 3. (*Linking Relation*): $P_f = \{e_{u_1}, e_{u_2}, \dots, e_{u_{|U_f|}}\} \vdash e$, denotes the linking relation of elements $e_i \in e$ which pertain to a certain user u_i . In other words, e is partitioned in non-overlapping subsets e_{u_i} , where e_{u_i} consists of all e_i for user u_i . It holds, that $\forall e_i, e_j \in e_{u_i}, f(e_i) = f(e_j) = u_i$. Precisely, $\forall e_i \in e_{u_i} : e_i \sim u_i$.

5 Adversarial Model and Principles of Privacy Games

In the following, we briefly state our adversarial model and then elaborate on the fundamental definitions and interactions of our privacy games.

5.1 Adversarial Model

An adversary, denoted by \mathcal{A} , is represented by an algorithm that runs in polynomial time. \mathcal{A} is given access to a limited number of oracles, as defined by the privacy game. An oracle, denoted by \mathcal{O} , gives \mathcal{A} access to certain knowledge about ϕ and relations between the outcome values (e, β) . The interaction of \mathcal{A} with an oracle \mathcal{O} is further denoted by $\mathcal{A}^{\mathcal{O}}$.

Overall, we say that \mathcal{A} is interested in learning the mapping function f , or any interesting properties about f , represented by the defined notions in Section 6. Each of these notions is represented as a game. We define a game G , and say that

\mathcal{A} is able to break the privacy notion associated with that game if G evaluates to **true** with a non-negligible probability. The advantage of \mathcal{A} is denoted by

$$\mathbf{Adv}(\mathcal{A}) = Pr[G_\star^{\mathcal{A}} \rightarrow \mathbf{true}]$$

where $\star \in \{\mathbf{SA}, \mathbf{PH}, \mathbf{SU}, \mathbf{WU}, \mathbf{PS}, \mathbf{AN}, \mathbf{WA}, \mathbf{UO}\}$ ⁸ denotes any of the games.

We assume \mathcal{A} to be semi-honest: when \mathcal{A} ‘plays’ a game, we assume that it behaves truthfully and follows the game’s rules. In other words, \mathcal{A} does not deviate from the game. In particular, this means that \mathcal{A} does not submit elements that it has not previously obtained from the system ϕ (e.g. it does submit a random value $e_x \notin \mathbf{e}$, where $\mathbf{e} \leftarrow \mathcal{A}^{\text{view}^{\pi_a}(\cdot)}$), or any of the other oracles $\mathcal{O}_{U_f}, \mathcal{O}_{Q_f}, \mathcal{O}_{P_f}$ or $\mathcal{O}_{|U_f|}$, revealing specific information about certain subsets of f .

Precisely, \mathcal{A} is allowed to interact adaptively with the oracles of ϕ , and, at a given point in time, must commit an input value pair (u_x, α_x) to which it receives the output tuple (e_x, β_x) from ϕ , which we further denote as *general observation*. Given these parameters (and access to some system-specific oracles, as defined by the notions below) it will be evaluated if \mathcal{A} is able to successfully break the given notion. Further, for individual privacy, we assume u_x to be fixed; for group privacy, we assume $u_x \in \mathbf{u}_g$, where \mathbf{u}_g is a fixed group of user identifiers.

These restrictions allow us to abstract (and in some cases omit) some of ϕ ’s specific operations and tasks in order to keep the following games concise and clearly readable. To this end, we do not include checks of validity (for example, $u_x \in U_f$, for input or validation procedures) for any of the values submitted to the oracles. Moreover, we omit the definition of the oracles $\mathcal{O}_{U_f}, \mathcal{O}_{Q_f}, \mathcal{O}_{P_f}$ or $\mathcal{O}_{|U_f|}$ here, as we see such definitions as a technicality and, as such, present them in Appendix A.

5.2 Principles of Privacy Games

Any game G_\star consists of a limited number of procedures. A procedure, denoted by **proc**, defines the computational steps an algorithm has to follow in order to terminate. Procedures are executed top-down and in order. Further, some procedures may only be executed once during an attack session, and others may be executed an arbitrary number of times, indicated by **proc**_○. This allows \mathcal{A} to ‘play’ around a bit and get used to the system, before, for a fixed set of inputs, it tries to break the privacy notion, by executing **proc** \star , where $\star \in \{\mathbf{SA}, \mathbf{PH}, \mathbf{SU}, \mathbf{WU}, \mathbf{PS}, \mathbf{AN}, \mathbf{WA}, \mathbf{UO}\}$.

In the following, we define a basic game G , which provides procedures used by all games G_\star . We define *inheritance* as follows: if game G_x inherits all procedures of game G_y , denoted $G_x \vDash G_y$, then G_x can, additionally to its own procedures, call any of the procedures of G_y . In our case: $G_\star \vDash G$. In Figure 2, we present game G . \mathcal{A} may play game G for an *arbitrary* number of times, before it plays a specific game G_\star *once*, where she aims to break the expressed notion.

In general, \mathcal{A} interacts with ϕ (‘plays’ the games) in the following way:

⁸ Here, the notation follows the notions of strong anonymity (**SA**), participation hiding (**PH**), strong unlinkability (**SU**), weak unlinkability (**WU**), pseudonymity (**PS**), anonymity (**AN**), weak anonymity (**WA**), and unobservability (**UO**).

```

proc initialise
  forall  $i$  in range(0,  $n$ ):  $p_i.db \leftarrow \emptyset$ 
proc◊ input( $u_i, \alpha_i, p_i$ )
   $p_i.db \leftarrow p_i.db \cup (u_i, \alpha_i)$ 
proc◊ view $\pi_q$ 
  forall  $i$  in range(0,  $n$ ):
     $e_i \leftarrow \pi^\phi(p_i.db); \beta_i \leftarrow \pi_q(p_i.db)$ 
   $e \leftarrow e_1 \circ e_2 \circ \dots \circ e_n$ 
   $\beta \leftarrow \beta_1 \circ \beta_2 \circ \dots \circ \beta_n$ 
  return ( $e, \beta$ )
proc◊ corrupt( $p_i$ )
  return  $p_i.db$ 

```

Fig. 2: Game G.

1. \mathcal{A} inputs data (α_i) into the system ϕ .
2. \mathcal{A} obtains⁹ output (e, β) from the system ϕ .
3. \mathcal{A} may gain additional information by corrupting a threshold of parties.
4. \mathcal{A} selects any value from the output, to which it believes it knows the relation (specified by the privacy notions) and submits that to the validation oracle (*e.g.* any $e_x \in e_x$, where \mathcal{A} believes to know u_x).
5. Game G_\star returns **true** or **false** according to the output of the validation oracle.

As already mentioned, our privacy model operates in a distributed setting, but it also supports a centralised setting. For the games, we define an integer n , which represents the number of parties $p \in P$, with P the set of all parties. Each party p_i holds a database, denoted as $p_i.db$, containing the input values submitted to it. Then, for a centralised model, we say $n = 1$ (meaning we have exactly one party), and for a distributed model $n > 1$.

6 Privacy Games

We now state the notions for privacy that we consider in our privacy model. For each notion, we formulate a game that an adversary has to win in order to break the associated notion. The privacy notions follow and extend the work of Bohli and Pashalidis [7]; the game-based definitions are novel. In formulating said notions and games, we take into account that they must abstract the behaviour of a system (*i.e.* it may be stateful, stateless or online), must abstract the locality of the computational parties (*i.e.* centralised or distributed), and must give consideration to individuals as well as groups.

In Figures 3-10, we illustrate the games G_\star , with $\star \in \{\text{SA, PH, SU, WU, PS, AN, WA, UO}\}$. Most of the games are quite similar, and differ only in the validation oracle and by the access to the oracles $\mathcal{O}_{U_f}, \mathcal{O}_{Q_f}, \mathcal{O}_{P_f}$ or $\mathcal{O}_{|U_f|}$. Therefore, we can

⁹ \mathcal{A} obtains (e, β) by querying the **view** oracle. For each party's database a permutation is applied. Moreover, $e_x \circ e_y$ denotes the combination of two sets e_x, e_y under this permutation. This means that $\pi^\phi(e_x) \circ \pi^\phi(e_y)$ is equivalent to $\pi^\phi(e_x \cup e_y)$. Similarly, for β , \circ denotes the π_q operation.

```

proc validateSA( $e_x, u_x$ )
  if  $f(e_x) = u_x$  then: return true else:
    return false
proc SA
   $\mathcal{A}^{\text{input}(u_x, \alpha_x, p_x)}$ 
   $e_x, \beta_x \leftarrow \mathcal{A}^{\text{view}^{\pi_q}()}$ 
  return  $\mathcal{A}^{\text{validate}_{SA}(e_x, u_x)}$ 

```

Fig. 3: Game G_{SA} (Strong Anonymity).

group certain games and notions, as they aim to achieve similar objectives. In particular, we group the notions as follows: (SA, PH), (SU, WU), (PS, AN, WA) and (UO).

6.1 Privacy Games for Single Elements

The objective of group (SA, PH) is, from the outcome e_x, β_x , to find any value $e_x \in \mathbf{e}_x$ for which \mathcal{A} is able to identify the user identifier u_x . The validation functions of this group return **true** if \mathcal{A} is able to do so.

Definition 4. (*Strong Anonymity*): A system is said to provide strong anonymity, denoted by SA, if it does not leak any information about the mapping function f .

A system ϕ that satisfies the property of SA prevents \mathcal{A} relating output elements to their associated user identifiers. More precisely, \mathcal{A} must not be able, for any element $e_x \in \mathbf{e}_x$, to learn the associated identifier u_x , such that $f(e_x) = u_x$. \mathcal{A} is not given anything beyond the general observation.

Definition 5. (*Participation Hiding*): A system is said to hide participation, denoted by PH, if it does not leak any information about f beyond the size of the participant set $|U_f|$.

A system ϕ that satisfies the property of PH prevents \mathcal{A} relating output elements to their associated user identifiers. More precisely, \mathcal{A} must not be able, for any element $e_x \in \mathbf{e}_x$, to learn the associated identifier u_x such that $f(e_x) = u_x$. In addition to the general observation, we assume that \mathcal{A} is able to learn $|U_f|$, which might or might not help to find said relation.

6.2 Privacy Games for Multiple Elements

The objective of group (SU, WU) is, from the outcome e_x, β_x , to find any two values $e_x, e_y \in \mathbf{e}_x$ for which \mathcal{A} believes they are related to the user identifier u_x . The validation functions of this group return **true** if \mathcal{A} is able to do so.

Definition 6. (*Strong Unlinkability*): A system is said to provide strong unlinkability, denoted by SU, if it does not leak any information about f beyond the participant set U_f .

A system ϕ that satisfies the property of SU prevents \mathcal{A} relating output elements pertaining to the same user identifier. More precisely, \mathcal{A} must not be able, for a fixed user identifier u_x to find any two elements $e_x, e_y \in \mathbf{e}_x$, such that $f(e_x) = f(e_y) = u_x$. In addition to the general observation, we assume that \mathcal{A} is able to learn U_f , which might or might not help to find said relation.

Definition 7. (*Weak Unlinkability*): A system is said to provide weak unlinkability, denoted by WU, if it does not leak any information about f beyond the usage frequency set Q_f ¹⁰.

¹⁰ Note that knowledge of Q_f implicitly implies knowledge of U_f .

<pre> proc \mathcal{O}_{U_f} <u>return U_f</u> proc validate$_{SU}(e_x, e_y, u_x)$ if $f(e_x) = f(e_y) = u_x$ then: return true else: return false proc SU $\mathcal{A}^{\text{input}(u_x, \alpha_x, p_x)}$ $\mathcal{A}^{\text{input}(u_x, \alpha_y, p_x)}$ $e_x, \beta_x \leftarrow \mathcal{A}^{\text{view}^{\pi q}()}$ return $\mathcal{A}^{\text{validate}_{SU}(e_x, e_y, u_x)}$ </pre>	<pre> proc \mathcal{O}_{P_f} <u>return P_f</u> proc validate$_{PS}(e_{u_x}, u_x)$ if $f(e_{u_x}) = u_x$ then: return true else: return false proc PS $\mathcal{A}^{\text{input}(u_x, \alpha_x, p_x)}$ $e_x, \beta_x \leftarrow \mathcal{A}^{\text{view}^{\pi q}()}$ return $\mathcal{A}^{\text{validate}_{PS}(e_{u_x}, u_x)}$ </pre>
---	--

Fig. 4: Game G_{SU} (Strong Unlinkability).Fig. 5: Game G_{PS} (Pseudonymity).

A system ϕ that satisfies the property of WU prevents \mathcal{A} relating output elements pertaining to the same user identifier. More precisely, \mathcal{A} must not be able, for a fixed user identifier u_x , to find any two elements $e_x, e_y \in e_x$ such that $f(e_x) = f(e_y) = u_x$. In addition to the general observation, we assume that \mathcal{A} is able to learn U_f and Q_f . This allows \mathcal{A} to learn, for each user, the number of elements of its outcome set, which might or might not help to find said relation.

6.3 Privacy Games for Element Groups

The objective of group (PS, AN, WA) is, from the outcome e_x, β_x , given the ability to group certain outcome elements together, for this very group e_{u_x} , to identify the user identifier u_x . The validation functions of this group return **true** if \mathcal{A} is able to do so.

Definition 8. (*Pseudonymity*): A system is said to provide pseudonymity, denoted by PS, if it does not leak any information about f beyond the linking relation P_f .

A system ϕ that satisfies the property of PS prevents \mathcal{A} relating a group of output elements (or an element within that group), pertaining to the same user identifier, to the user identifier. More precisely, \mathcal{A} must not be able, for a fixed user identifier u_x , to find the group of related outcome elements $e_{u_x} \subseteq e_x$ such that $\forall e_x \in e_{u_x} | f(e_x) = u_x$. In addition to the general observation, we assume that \mathcal{A} is able to learn P_f . This allows \mathcal{A} to assign, for each partition of the overall outcome set e_x , a pseudonym, which might or might not help to find said relation.

Definition 9. (*Anonymity*): A system is said to provide anonymity, denoted by AN, if it does not leak any information about f beyond the participation set U_f and the linking relation P_f .

A system ϕ that satisfies the property of AN prevents \mathcal{A} relating a group of output elements (or an element within that group), pertaining to the same user identifier, to the user identifier. More precisely, \mathcal{A} must not be able, for a fixed user identifier u_x , to find the group of related outcome elements $e_{u_x} \subseteq e_x$ such that $\forall e_x \in e_{u_x} | f(e_x) = u_x$. In addition to the general observation, we assume

14

$\frac{\mathbf{proc}_{\circlearrowleft} \mathcal{O}_{U_f}}{\text{return } U_f}$ $\frac{\mathbf{proc}_{\circlearrowleft} \mathcal{O}_{Q_f}}{\text{return } Q_f}$ $\frac{\mathbf{proc}_{\circlearrowleft} \mathcal{O}_{P_f}}{\text{return } P_f}$ $\frac{\mathbf{proc} \text{ validate}_{UO}(g)}{\text{return } (g == b)}$	$\mathbf{proc}_{\circlearrowleft} \text{ view}_{\text{refut.}}^{\pi_q}$ $b \stackrel{\$}{\leftarrow} \{0, 1\}$ $\text{forall } i \text{ in range}(0, n):$ $e_i \leftarrow \pi^\phi(p_i.db); \beta_i \leftarrow \pi_q(pi.db)$ $\text{if } b = 0 \text{ then: } e_{u_x} \stackrel{\$}{\leftarrow} A; \beta_{u_x} \stackrel{\$}{\leftarrow} A$ $e \leftarrow e_1 \circ e_2 \circ \dots \circ e_n$ $\beta \leftarrow \beta_1 \circ \beta_2 \circ \dots \circ \beta_n$ $\text{return } (e, \beta)$	$\mathbf{proc} \text{ UO}$ $\mathcal{A}^{\text{input}(u_x, \alpha_x, P_x)}$ $e_x, \beta_x \leftarrow \mathcal{A}^{\text{view}_{\text{refut.}}^{\pi_q}(\cdot)}$ $\text{return } \mathcal{A}^{\text{validate}_{UO}(g)}$
--	---	---

Fig. 6: Game G_{UO} (Unobservability).

that \mathcal{A} is able to learn U_f and P_f . This allows \mathcal{A} to learn the user identifiers associated with the outcome values, and the ability to learn partitions of the overall outcome set e_x , however not the relationship between those sets, which might or might not help to find said relation.

Definition 10. (*Weak Anonymity*): A system is said to provide weak anonymity, denoted by **WA**, if it does not leak any information about f beyond the participation set U_f , the usage frequency set Q_f and the linking relation P_f .

A system ϕ that satisfies the property of **WA** prevents \mathcal{A} relating a group of output elements (or an element within that group), pertaining to the same user identifier, to the user identifier. More precisely, \mathcal{A} must not be able, for a fixed user identifier u_x , to find the group of related outcome elements $e_{u_x} \subseteq e_x$ such that $\forall e_x \in e_{u_x} | f(e_x) = u_x$. In addition to the general observation, we assume that \mathcal{A} is able to learn U_f, Q_f and P_f . This allows \mathcal{A} to learn the user identifiers associated with the outcome values, partitions of the overall outcome set e_x and information about on the number of elements each user has contributed in the outcome set, which might or might not help to find said relation.

Remark 1. In the case that each user contributes a unique amount of elements to ϕ , and \mathcal{A} is able to infer this knowledge, then it is easy to see that \mathcal{A} is able to determine $f(e_x) = u_x$ by the mere combination of the sets U_f, Q_f and P_f . However, in other cases, where this condition is not satisfied (for example, the number of elements each user contributes is uniformly distributed), \mathcal{A} only partially gains information by the combination of the previously mentioned sets.

6.4 Unobservability or Privacy Games for System Behaviour

Finally, the objective group (**UO**) is, from the outcome e_x, β_x , to distinguish if the values e_{u_x}, β_{u_x} were generated as a result of a system invocation, or generated purely at random. The validation functions of this group return **true**, if \mathcal{A} is able to do so.

Definition 11. (*Unobservability*): A system is said to be unobservable, denoted by **UO**, if the output is indistinguishable from a random output in the range of the output parameter set, even with knowledge of the participation set U_f , the usage frequency set Q_f and the linking relation P_f .

A system ϕ that satisfies the property of $\mathbf{U0}$ prevents \mathcal{A} from distinguishing that, for a fixed tuple of user identifier and input value, a system invocation has taken place or not. More precisely, \mathcal{A} must not be able, for a fixed input tuple (u_x, α_x) , by observing the outcome (e_x, β_x) , be able to distinguish if $e_x \leftarrow \pi^\phi(\alpha_x), \beta_x \leftarrow \pi_q(\alpha_x)$ given $\alpha_x \in \alpha_x, f^*(\alpha_x) = u_x$, or, given $\alpha_x \notin \alpha_x, f^*(\alpha_x) \neq u_x$. Intuitively, it is easy to distinguish a system ϕ behaving in the above way, by simply determining $|e_x|$. If $\alpha_x \notin \alpha_x$, then $|e_x| < |e_y|$, where in the latter case $\alpha_x \in \alpha_x$. Therefore, even for a *void* invocation of ϕ , there must hold $\exists e_x \in e_x |e_x| \neq \pi^\phi(\alpha_x)$. e_x must be produced in a way that is indistinguishable of a system being invoked or not. In addition to the general observation, we assume that \mathcal{A} is able to learn U_f, Q_f and P_f . This allows \mathcal{A} to learn the user identifiers associated with the outcome values, partitions of the overall outcome set e_x and information about the number of elements each user has contributed in the outcome set, which might or might not help to find said relation.

Remark 2. The previous illustration of the privacy games is for individual privacy, where we assume u_x to be fixed. For group privacy, we assume $u_x \in \mathbf{u}_g$, where \mathbf{u}_g is a group of fixed user identifiers. Then, each objective may be extended to find elements $e_x \in e_x$ that relate to any u_x pertaining to the group of \mathbf{u}_g . For example, if \mathcal{A} was about to infer if any e_x is related with e_y , where $f(e_x) = u_x, f(e_y) = u_y$ and $u_x, u_y \in \mathbf{u}_g$, then this would violate group privacy.

7 Conclusions

We set out to provide a clear and concise understanding of various notions for privacy-preserving data release. Precisely, we build upon the syntactic notions of other authors and formulate those notions into privacy games, following the game-playing technique of provable security. Furthermore, we extend the framework of Bohli and Pashalidis [7] by the notion of unobservability — a notion that captures the capability of an adversary to distinguish if a system invocation has taken place or not. With the definition of these games, we aim to clarify understanding of, and relationships between, different privacy notions; we also aim to give an unambiguous understanding of adversarial actions (in order to win a game and, therefore, break a notion). With the definition of the privacy notions and games, we take into consideration the different requirements of privacy protections for individuals and groups. Overall, many of the current existing models for privacy focus on the preservation of privacy in centralised environments, though various use cases show the importance for concise privacy models and notions in distributed environments. Our notions, games, and system model are designed for distributed environments (yet are still compliant with centralised systems). For example, TRE systems (as outlined in Appendix B) may be a good fit of the contributions of this paper.

Future work will include an extension to externally distributed systems (reasoning about privacy aspects between the release of various systems), policies for the selection of privacy notions, exploration of primitives guaranteeing such notions and, most importantly, application to real-world case studies.

References

1. Abadi, M., Chu, A., Goodfellow, I., McMahan, H.B., Mironov, I., Talwar, K., Zhang, L.: Deep learning with differential privacy. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. pp. 308–318. CCS '16, ACM, New York, NY, USA (2016)
2. Achenbach, D., Huber, M., Müller-Quade, J., Rill, J.: Closing the Gap: A Universal Privacy Framework for Outsourced Data, pp. 134–151. Springer International Publishing, Cham (2016)
3. Anati, I., Gueron, S., Johnson, S., Scarlata, V.: Innovative Technology for CPU based Attestation and Sealing. In: Proceedings of the 2nd International Workshop on Hardware and Architectural Support for Security and Privacy. HASP 2013, vol. 13. ACM (2013)
4. Ankele, R., Küçük, A., Martin, A., Simpson, A., Paverd, A.: Applying the trustworthy remote entity to privacy-preserving multiparty computation: Requirements and criteria for large-scale applications. In: 2016 Intl IEEE Conferences on Ubiquitous Intelligence Computing, Advanced and Trusted Computing, Scalable Computing and Communications, Cloud and Big Data Computing, Internet of People, and Smart World Congress (UIC/ATC/ScalCom/CBDCCom/IoP/SmartWorld). pp. 414–422. IEEE (July 2016)
5. Bellare, M., Rogaway, P.: The game-playing technique. Cryptology ePrint Archive, Report 2004/331 (2004)
6. Blum, A., Ligett, K., Roth, A.: A learning theory approach to non-interactive database privacy. In: Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing, pp. 609–618. STOC '08, ACM, New York, NY, USA (2008)
7. Bohli, J.M., Pashalidis, A.: Relations among privacy notions. ACM Trans. Inf. Syst. Secur. 14(1), 4:1–4:24 (Jun 2011)
8. Corporation, I.: Intel Software Guard Extensions Programming Reference (2014), <https://software.intel.com/sites/default/files/managed/48/88/329298-002.pdf>
9. Cuff, P., Yu, L.: Differential privacy as a mutual information constraint. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. pp. 43–54. CCS '16, ACM, New York, NY, USA (2016)
10. Dwork, C.: Differential privacy. In: Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part II. pp. 1–12 (2006)
11. Dwork, C., Kenthapadi, K., McSherry, F., Mironov, I., Naor, M.: Our data, ourselves: Privacy via distributed noise generation. In: Proceedings of the 24th Annual International Conference on The Theory and Applications of Cryptographic Techniques. pp. 486–503. EUROCRYPT'06, Springer-Verlag, Berlin, Heidelberg (2006)
12. Dwork, C., Roth, A.: The algorithmic foundations of differential privacy. Found. Trends Theor. Comput. Sci. 9(3), 211–407 (Aug 2014), <http://dx.doi.org/10.1561/04000000042>
13. Ebadi, H., Sands, D., Schneider, G.: Differential privacy: Now it's getting personal. In: Proceedings of the 42Nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. pp. 69–81. POPL '15, ACM, New York, NY, USA (2015)
14. Erlingsson, U., Pihur, V., Korolova, A.: Rappor: Randomized aggregatable privacy-preserving ordinal response. In: Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security. pp. 1054–1067. CCS '14, ACM, New York, NY, USA (2014)

15. Ganta, S.R., Kasiviswanathan, S.P., Smith, A.: Composition attacks and auxiliary information in data privacy. In: Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 265–273. KDD '08, ACM, New York, NY, USA (2008)
16. Kifer, D.: Attacks on privacy and definetti's theorem. In: Proceedings of the 2009 ACM SIGMOD International Conference on Management of Data. pp. 127–138. SIGMOD '09, ACM, New York, NY, USA (2009)
17. Kifer, D., Lin, B.R.: Towards an axiomatization of statistical privacy and utility. In: Proceedings of the Twenty-ninth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems. pp. 147–158. PODS '10, ACM, New York, NY, USA (2010), <http://doi.acm.org/10.1145/1807085.1807106>
18. Kifer, D., Machanavajjhala, A.: No free lunch in data privacy. In: Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data. pp. 193–204. SIGMOD '11, ACM, New York, NY, USA (2011)
19. Kifer, D., Machanavajjhala, A.: Pufferfish: A framework for mathematical privacy definitions. *ACM Trans. Database Syst.* 39(1), 3:1–3:36 (Jan 2014), <http://doi.acm.org/10.1145/2514689>
20. Küçük, A., Paverd, A., Martin, A., Asokan, N., Simpson, A., Ankele, R.: Exploring the Use of Intel SGX for Secure Many-Party Applications. pp. 5:1–5:6. SysTEX 2016, ACM (2016)
21. Li, N., Li, T., Venkatasubramanian, S.: t-Closeness: Privacy Beyond k-Anonymity and l-Diversity. In: Proceedings of the 23rd IEEE International Conference on Data Engineering (2007)
22. Li, N., Qardaji, W., Su, D., Wu, Y., Yang, W.: Membership privacy: A unifying framework for privacy definitions. In: Proceedings of the 2013 ACM SIGSAC Conference on Computer and Communications Security. pp. 889–900. CCS '13, ACM, New York, NY, USA (2013), <http://doi.acm.org/10.1145/2508859.2516686>
23. Machanavajjhala, A., Kifer, D., Abowd, J., Gehrke, J., Vilhuber, L.: Privacy: Theory meets practice on the map. In: Proceedings of the 2008 IEEE 24th International Conference on Data Engineering. pp. 277–286. ICDE '08, IEEE Computer Society, Washington, DC, USA (2008)
24. Machanavajjhala, A., Kifer, D., Gehrke, J., Venkatasubramanian, M.: L-diversity: Privacy Beyond K-anonymity. *ACM Trans. Knowl. Discov. Data* (2007)
25. McKeen, F., Alexandrovich, I., Berenzon, A., Rozas, C.V., Shafi, H., Shanbhogue, V., Savagaonkar, U.R.: Innovative Instructions and Software Model for Isolated Execution. In: Proceedings of the 2nd International Workshop on Hardware and Architectural Support for Security and Privacy (2013)
26. McMahan, H.B., Moore, E., Ramage, D., Hampson, S., y Arcas, B.A.: Communication-efficient learning of deep networks from decentralized data. In: Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS) (2016), <http://arxiv.org/abs/1602.05629>
27. Mironov, I., Pandey, O., Reingold, O., Vadhan, S.: Computational differential privacy. In: Proceedings of the 29th Annual International Cryptology Conference on Advances in Cryptology. pp. 126–142. CRYPTO '09, Springer-Verlag, Berlin, Heidelberg (2009)
28. Paverd, A.: Enhancing Communication Privacy Using Trustworthy Remote Entities. D.Phil thesis, University of Oxford (2016)
29. Pfitzmann, A., Hansen, M.: A terminology for talking about privacy by data minimization: Anonymity, Unlinkability, Undetectability, Unobservability, Pseudonymity, and Identity Management (2009)

30. Pfitzmann, A., Köhntopp, M.: Anonymity, Unobservability, and Pseudonymity — A Proposal for Terminology, pp. 1–9. Springer Berlin Heidelberg, Berlin, Heidelberg (2001)
31. Sweeney, L.: K-anonymity: A Model for Protecting Privacy. Int. J. Uncertain. Fuzziness Knowl.-Based Syst. (2002)
32. Wong, R.C.W., Fu, A.W.C., Wang, K., Pei, J.: Minimality attack in privacy preserving data publishing. In: Proceedings of the 33rd International Conference on Very Large Data Bases. pp. 543–554. VLDB '07, VLDB Endowment (2007)

A Additional Privacy Games and Definitions of Oracles

A.1 Additional Privacy Games

```

proc⊙  $\mathcal{O}_{|U_f|}$ 
  return  $|U_f|$ 
proc validatePH( $e_x, u_x$ )
  if  $f(e_x) = u_x$  then: return true else:
    return false
proc PH
   $\mathcal{A}^{\text{input}(u_x, \alpha_x, p_x)}$ 
   $e_x, \beta_x \leftarrow \mathcal{A}^{\text{view}^{\pi q}()}$ 
  return  $\mathcal{A}^{\text{validate}_{PH}}(e_x, u_x)$ 

```

Fig. 7: Game G_{PH} (Participation Hiding).

```

proc⊙  $\mathcal{O}_{U_f}$ 
  return  $U_f$ 
proc⊙  $\mathcal{O}_{Q_f}$ 
  return  $Q_f$ 
proc validateWU( $e_x, e_y, u_x$ )
  if  $f(e_x) = f(e_y) = u_x$  then: return true
  else: return false
proc WU
   $\mathcal{A}^{\text{input}(u_x, \alpha_x, p_x)}$ 
   $\mathcal{A}^{\text{input}(u_x, \alpha_y, p_x)}$ 
   $e_x, \beta_x \leftarrow \mathcal{A}^{\text{view}^{\pi q}()}$ 
  return  $\mathcal{A}^{\text{validate}_{WU}}(e_x, e_y, u_x)$ 

```

Fig. 8: Game G_{WU} (Weak Unlinkability).

```

proc⊙  $\mathcal{O}_{U_f}$ 
  return  $U_f$ 
proc⊙  $\mathcal{O}_{P_f}$ 
  return  $P_f$ 
proc validateAN( $e_{u_x}, u_x$ )
  if  $f(e_{u_x}) = u_x$  then: return true else:
    return false
proc AN
   $\mathcal{A}^{\text{input}(u_x, \alpha_x, p_x)}$ 
   $e_x, \beta_x \leftarrow \mathcal{A}^{\text{view}^{\pi q}()}$ 
  return  $\mathcal{A}^{\text{validate}_{AN}}(e_{u_x}, u_x)$ 

```

Fig. 9: Game G_{AN} (Anonymity).

```

proc⊙  $\mathcal{O}_{U_f}$ 
  return  $U_f$ 
proc⊙  $\mathcal{O}_{Q_f}$ 
  return  $Q_f$ 
proc⊙  $\mathcal{O}_{P_f}$ 
  return  $P_f$ 
proc validateWA( $e_{u_x}, u_x$ )
  if  $f(e_{u_x}) = u_x$  then: return true else:
    return false
proc WA
   $\mathcal{A}^{\text{input}(u_x, \alpha_x, p_x)}$ 
   $e_x, \beta_x \leftarrow \mathcal{A}^{\text{view}^{\pi q}()}$ 
  return  $\mathcal{A}^{\text{validate}_{WA}}(e_{u_x}, u_x)$ 

```

Fig. 10: Game G_{WA} (Weak Anonymity).

A.2 Definition of Oracles

proc_⊙ \mathcal{O}_{U_f}
 forall i in range(0, $|e|$):
 $U_f \leftarrow U_f \cup f(e_i)$
 return U_f

Fig. 11: Oracle \mathcal{O}_{U_f} (Participant Set).

proc_⊙ \mathcal{O}_{Q_f}
 forall i in range(0, $|e|$):
 $Q_f \leftarrow Q_f \cup^a \{f(e_i)\}$
 return U_f

Fig. 12: Oracle \mathcal{O}_{Q_f} (Usage Frequency Set).

^a \cup is implemented such that if $f(e_i)$ already exists in $Q_f = \{(u_i, \#u_i)\}$, then it solely increases frequency counter $\#u_i$.

proc_⊙ \mathcal{O}_{P_f}
 $P_f \leftarrow \text{part}(e)$
 return P_f

Fig. 13: Oracle \mathcal{O}_{P_f} (Linking Relation).

proc_⊙ $\mathcal{O}_{|U_f|}$
 $|U_f| \leftarrow \text{sizeof}(\mathcal{A}^{\mathcal{O}_{U_f}})$
 return $|U_f|$

Fig. 14: Oracle $\mathcal{O}_{|U_f|}$ (Number of Participants).

B TRE systems

A Trustworthy Remote Entity (TRE) [28] is a system that enables privacy-preserving data release. The concept was introduced by Pavard [28] and extended by Ankele *et al.* [4] and Küçük *et al.* [20] to support the application of privacy-preserving data release in distributed settings. In its simplest form, a TRE operates as an intermediary between participating communication parties. Essentially a TRE portrays a trusted third party (TTP), but, contrary to a TTP that requires users to blindly trust all performed actions, a TRE can verify its trustworthy nature. This can be done, for example, by using Intel Software Guard Extensions (SGX) [3,25,8] as an underlying hardware security primitive, which features isolated execution of code and memory pages in so-called *enclaves*, encrypted storage outside of the immediate CPU package, and software attestation to promote its trustworthiness. A TRE is not bound to any architectural design choice and may be applied in a centralised or distributed fashion. Overall, a TRE provides input privacy to a system. To achieve release privacy, privacy-preserving mechanisms such as data permutation, data perturbation or data anatomisation must be applied.

We may use the syntactic privacy notions and games introduced in this paper to reason about privacy guarantees that such a TRE system is able to achieve. Of course, this is dependent on the underlying use case, the data, and the privacy mechanism leveraged by the TRE system to realise *release privacy*.

C Illustration of Privacy Properties abstracted by the Privacy Notions and Games

In Figure 15, we illustrate previously defined privacy notions. As depicted, the system ϕ is queried with a triple (u_x, α_x, p_i) and releases a tuple (e, β) on a `view` query. Dashed encircled are the sets that the adversary \mathcal{A} is given access to for each notion. Further, the dashed arrows show the objective of each notion. For simplicity, this representation includes only one database $p_i.db$. However, any finite number of databases are equally possible and supported by the model.

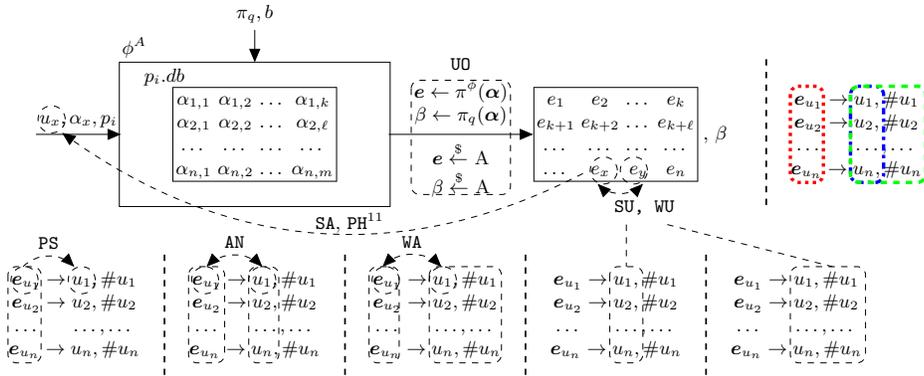


Fig. 15: Illustration of privacy notions, given access to the system ϕ , and oracle access to the sets \mathbf{U}_f (---), \mathbf{Q}_f (- - -) and \mathbf{P}_f (---).

¹¹ For the notion of PH, \mathcal{A} has additional access to an oracle giving the size of the participation set $|\mathbf{U}_f|$.