On Cryptographic Anonymity and Unpredictability in Secret Sharing

Anat Paskin-Cherniavsky[a], Ruxandra F. Olimid[b]

[a] *Department of Computer Science and Mathematics. Ariel University, Ariel, Israel*
anatpc@ariel.ac.il
[b] *Department of Computer Science University of Bucharest, Romania and Applied Cryptography Group, Orange*
ruxandra.olimid@fmi.unibuc.ro

Abstract

We revisit the notions of cryptographic anonymity and share unpredictability in secret sharing, introducing more systematic and fine grained definitions. We derive tight negative and positive results characterizing access structures with respect to the generalized definitions.

Keywords: secret sharing, anonymity, unpredictability

1. Introduction

Secret sharing is a cryptographic primitive that allows a secret to be shared among a set $\mathcal{P}$ of parties such that qualified subsets succeed, while unqualified subsets fail secret reconstruction. Sometimes reconstruction is questionable, so the parties that agree reconstruction want their identities to remain private; such a scheme is called *cryptographically anonymous*[1] [2].

1.1. Previous work

Guillermo et al. laid out the definitional framework for cryptographic anonymity [2]. They study the case where the adversary is given the share vector $\mathbf{S}^T$ held by a set T of parties (coupled with their identities) and needs to distinguish between sets of pairs of parties T_1, T_2 disjoint from T given the unordered set of T_i's shares. They also consider the enhanced setting of *strong cryptographic anonymity* where the adversary is also given the shared

[1] *Cryptographic anonymity* strengthens *anonymity*, which only states that identities are not *needed* at reconstruction, but not necessarily *hidden* [1].

secret. Given a sharing scheme, every set T induces a partition of $2^{\mathcal{P}\setminus T}$ into equivalence classes where all T_i's in every equivalence class are indistinguishable. Clearly, every such class contains sets of equal size; in particular, *participant anonymity* only considers T_i's of size 1. For this framework, they give several positive results with respect to various T's and access structures f. Another interesting setting considered in [2] extends cryptographic anonymity such that the size of T_i is unknown to the adversary. Non-trivial schemes here exist in a relaxed setting where parties in T_i may locally modify their shares (the interaction pattern relaxed so that upon reconstruction each participant is aware of the other reconstructing parties, but is not allowed to communicate with them). They provide strong positive results in this setting.

1.2. Our contribution

First, we naturally generalize the definition in [2] to include statistical and computational cryptographic anonymity (share-indistinguishability based). Additionally, we refine it to account for various a-priori distributions of secrets rather than require anonymity to hold for all initial distributions. For simplicity, our definitions and results are restricted to participant anonymity, but generalization to larger sets T_i is not hard.

Second, we make a simple observation (first made in [3]) that extends applicability by preserving cryptographically anonymity (at least computationally) against a coalition that may include the dealer: the dealer does not have to know the shares, but it could rather contribute the secret to a MPC (Multi-Party Protocol) that evaluates the sharing functionality. To achieve statistical anonymity (even against semi-honest parties), one should account both for bounds on MPC with statistical security, and the existence of a sharing *functionality* satisfying the anonymity requirement.

Finally, we revise unpredictability of shares, a notion introduced in [4] along with some positive results, but only used as a property of a subprotocol and not elucidated as an abstract concept. Although this is incomparable to anonymity, the notions are somehow related to the same use-case scenario (so we sometimes call both anonymity notions).

Negative and positive results. We derive tight negative results on statistical share-indistinguishable and share-unpredictability secret sharing. Establishing impossibility results (even for statistical, rather than perfect schemes, as considered before) is one of our main contribution over previous work. Another main contribution in terms of feasibility results is that indistinguishability of sets T_1, T_2 may hold relatively to some T even if they are not

isomorphic (in the sense that there exists a permutation of $\mathcal{P}$ mapping T_1 onto T_2 such that the access structure remains unchanged), as in all positive results of [2]. As example, consider f a weighted leveled threshold function over a set of 10 employees and 5 managers, where an employee share weights 1 and a manager share weights 3. For threshold equals 10, every pair of external sets of size 1 (either containing an employee or a manager) is indistinguishable to T that includes 3 managers. Such indistinguishability is satisfied in our settings, while it fails in the settings of [2] under the strong isomorphism requirement (which is independent of the choice of T). The positive results (derived by construction), along with the negative results characterize the access structures with respect to both our generalized definitions.

2. Preliminaries

Let $[n] = \{1, \ldots, n\}$. We consider distributions D over finite sets and denote sampling a value x according to D by $x \leftarrow D$. We use PPT (Probabilistic Polynomial Time) and negligible functions in the usual sense.

2.1. Secret Sharing

Let $\mathcal{P} = \{P_0, P_1, \ldots, P_n\}$ be a set of n players, where P_0 is a distinguished party called the *dealer* and $T \subseteq \mathcal{P}$ a subset of players. Let f be a finite function defining an access structure on $\mathcal{P}$ such that $f(T) = 1$ if T is qualified and $f(T) = 0$ if T is unqualified. We assume f is monotone and denote by a minterm of f a minimal qualified set.[2] For $T \subset \mathcal{P}$ and $\mathbf{S} = (S_0, \ldots, S_n)$, $\mathbf{S}^T$ denotes the set $\{S_i\}_{i \in T}$. A secret sharing functionality for f Sh_f takes an input $s \in \mathbb{F}$ (for a finite domain $\mathbb{F}$) from P_0 and distributes shares to the players such that it satisfies the standard requirements of secret sharing:

1. *Privacy.* For all $T \subseteq \mathcal{P}$ such that $f(T) = 0$ (T is unqualified), all secrets $s_1, s_2 \in \mathbb{F}$ such that $s_1 \neq s_2$ and all distinguishers D_T, the following holds for a negligible function ϵ:

$$|\Pr[D_T(\mathbf{S}_1^T) = 1] - \Pr[D_T(\mathbf{S}_2^T) = 1]| \leq \epsilon(k)$$

   Here and elsewhere, $\mathbf{S}_i$ denotes the random variable resulting from applying $\mathsf{Sh}_f(s_i)$. Sh_f is *perfectly private* if $\epsilon(k) = 0$, and *computationally private* if D_T is PPT.

[2] Unlike traditionally in secret sharing, we decouple the share distribution from the dealer, as this would limit achievable anonymity levels. We view P_0 as one of the participants, which constitutes a qualified set.

2. *(Statistical) Correctness.* There exists a function Rec_f and a negligible function $\epsilon(k)$, such that the following holds for all $T \subseteq \mathcal{P}$ such that $f(T) = 1$ (T is qualified) and all $s \in \mathbb{F}$:

$$\Pr[\mathsf{Rec}_f(\mathbf{S}^T) = s] \geq 1 - \epsilon(k)$$

Correctness is *perfect* if $\epsilon(k) = 0$.[3]

We refer to $(\mathsf{Sh}_f, \mathsf{Rec}_f)$ where Sh_f and Rec_f as above as a secret sharing scheme for an access structure f.

3. Definitional framework and settings

In addition to the standard requirements in the previous section, we are interested in various notions of hiding shares of players from a coalition T of other parties (that might be a qualified subset and possibly include P_0).

Definition 1 (Share-Indistinguishability)**.** *Consider a set $T \subseteq \mathcal{P}$ and a distribution D supported on $\mathbb{F}$ and R_T some equivalence relation on $\mathcal{P} \setminus T$. We say Sh_f is (T, D, R_T)-share-indistinguishable if for all $(P_i, P_j) \in R_T$ and all distinguishers D_T there exists a negligible function $\epsilon(k)$ such that:*

$$|\Pr_{s \leftarrow D}[D_T(\mathbf{S}^T, S[i]) = 1] - \Pr_{s \leftarrow D}[D_T(\mathbf{S}^T, S[j]) = 1]| \leq \epsilon(k).$$

. Here D_T receives S^T as an ordered sequence (in the sense that unlike Rec_f it also knows the index i of each S_i). Share-indistinguishability is perfect if $\epsilon(k) = 0$ and computational if D_T is PPT. Sh_f is $\mathcal{T}$-share-indistinguishable if it is $(T, D, \{(i,j) | i \neq j \in \mathcal{P} \setminus T\})$-share-indistinguishable for all $T \in \mathcal{T}$.

Definition 1 models cryptographic anonymity from [2], with D generalizing the (binary) concept of strong cryptographic anonymity.

Definition 2 (Decisional Share-Unpredictability)**.** *Consider a set $T \subseteq \mathcal{P}, P_i \notin T$ and a distribution D supported on $\mathbb{F}$. We say Sh_f is decisional (T, D, i)-share-unpredictable if for all distinguishers D_T there exists a negligible function $\epsilon(k)$ such that:*

$$|\Pr_{s \leftarrow D}[D_T(\mathbf{S}^T, i, S[i]) = 1) - \Pr_{s \leftarrow D, s' \leftarrow D}[D_T(\mathbf{S}^T, i, S'[i]) = 1)| \leq \epsilon(k)$$

Secrets s and s' are independently sampled. Share-unpredictability is perfect if $\epsilon(k) = 0$ and computational if D_T is PPT.

[3] Unlike in standard secret sharing schemes, the set T is not an input to Rec_f (that is, Sh_f is anonymous [2]).

Definition 2 extends unpredictability from [4] in several ways: it is decisional rather than computational (which makes it stronger) and considers any T rather than unqualified sets only.

3.1. The full secure computation setting and a concrete application

Besides of the appealing theoretical study of schemes that fulfill the above properties, they should be useful in real-life scenarios (where no trusted parties exist). A possible use case is as follows. A company holds a secret algorithm developed by P_0, which is stored in a shared manner on the computers of employees $P_1, \ldots, P_n$. Whenever the employees need to use the algorithm, they perform recovery of the secret. There is disagreement on whether it should be kept private or released as open source and hence the employees that wish to make it public want to remain anonymous. It is decided to resolve the problem by running the following 2-step "voting" functionality: (1) The sharing phase (where P_0 supplies s) is implemented via MPC (implementation may be interactive); (2) Each voter for release puts its share in a public directory (non-interactive). If a qualified subset votes for release, the secret can be recovered by anyone, otherwise it remain private. We assume that the parties are honest (but curious) and do not disrupt the sharing protocol in any way (step 1 is implemented by a program running on each computer and nobody tinkers with their software). In the aftermath, we want to ensure that even if a subset T of employees who put their shares in the public repository choose to identify their shares, (regardless of whether the secret is revealed or not), the identities of the other participants can not be deduced by inspection of the submitted set (we also assume those deciding to vote for release do not coordinate it among them to preserve anonymity, so, wlog. everyone just submits their share as is).

Claim 1. *Let Sh_f be (statistically) $\mathcal{T}$-share-indistinguishable. Then there exists an implementation of a voting protocol as above, such that for all $T \in \mathcal{T}$ and submitting sets $T \cup \{P_i\}$ where $P_i \notin T$, the distributions $(View(T), \mathbf{S}^T, S[i])$ and $(View(T), \mathbf{S}^T, S[j])$ ($View(T)$ are the views in the MPC protocol evaluating Sh_f) are computationally indistinguishable. Furthermore, if $\mathcal{T}$ is a Q_2 set, the indistinguishability is statistical [5].*

The proof of the claim is rather straightforward, and follows from general MPC protocols against semi-honest adversaries (e.g [5, 6]), along with the anonymity properties of Sh_f.

Similarly, the property of (decisional) share unpredictability of Sh_f implies the following application in the above scenario. Assume the distribution of a single $S[i]$ appears like inconspicuous random noise. Then, even if

some curious employee wants to check how P_i voted by going through his outgoing communication, and even if all other submitted shares are known (including identities), the submitted share $S[i]$ still appears like a fresh $S[i]$ (thus, still inconspicuous, and does not indicate P_i submitted his share).

4. Results

4.1. Negative results

Theorem 1. *Fix a sharing functionality Sh_f and let $T = T' \cup T''$ where $1 = f(T' \cup \{P_i\}) \neq f(T' \cup \{P_j\}) = 0$ for some $P_i, P_j \in \mathcal{P} \setminus T$. Then, Sh_f is not (T, D, R_T)-share-indistinguishable for all R_T containing (P_i, P_j) and all but possibly a constant ($O(n^2)$) number of distributions D.*

Proof. We construct a distinguisher D_T that in the share-indistinguishability experiment is given $\mathbf{S}^{T'} \cup \{S\}$, where S is either $S[i]$ or $S[j]$ and wins with constant advantage for all but possibly a few distributions D.

$D_T(\mathbf{S}^T, S)$*:* Consider a distribution $R = \mathsf{Rec}_f({S_0}^{T'} \cup \{S_0[j]\})$ (induced by $s_0 \in \mathbb{F}$ selected arbitrarily). Assume that $\lim_{k\to\infty} \Delta(D, R) > 0$ (the assumption fails only for $D = lim_{k\to\infty} R$). Let s_1 denote a value such that $|\Pr_D[s_1] - \Pr_R[s_1]| = \Omega_k(1)$. Compute $o = \mathsf{Rec}_f(\mathbf{S}^{T'} \cup \{S\})$. Output 1 if $o = s_1$, and 0 otherwise.

Analysis: If $S = S[i]$, by correctness of Sh_f D_T recovers $o = s$ (where s is the sampled secret) with probability $\geq 1 - \epsilon(k) = 1 - o(1)$. Thus, D_T outputs 1 with probability $\Pr_D(s_1) \pm o_k(1)$. If $S = S[j]$, by privacy o satisfies $\Delta(o, R) = o_k(1)$. Thus, D_T outputs 1 with probability $\Pr_R[s_1] \pm o_k(1)$. Thus, D_T has $\Omega_k(1)$ distinguishing advantage as $\lim_{k\to\infty} \Delta(D, R) > 0$. □

Theorem 2. *Fix a sharing functionality Sh_f and let $T = T' \cup T''$ where $f(T') = 0$ and $f(T' \cup \{P_i\}) = 1$ for some $i \in \mathcal{P} \setminus T$. Assume further that at least one of the following conditions holds: (1) $f(\{P_i\}) = 0$. (2) $f(T) = 1$. Then, Sh_f is not (T, D, i)-share unpredicatble for all but a constant (poly(n)) number of distributions D.*

Proof. We construct a distinguisher D_T that in the share-unpredictability experiment is given $\mathbf{S}^T \cup \{S\}$, where S is either $S[i]$ or $S'[i]$ with advantage $\Omega(1)$ for all but a few distributions D. First assume condition (1) holds.

$D_T(\mathbf{S}^T, S)$*:* Let $R = \mathsf{Rec}_f(\mathbf{S}^{T'} \cup \{S'[i]\})$ denote a distribution induced by applying $\mathsf{Sh}_f(s_0)$ for some $s = s' = s_0 \in \mathbb{F}$ ($\mathbf{S}, \mathbf{S}'$ sampled independently). Let s_1 denote a value such that $|\Pr_D[s_1] - \Pr_R[s_1]| = \Omega_k(1)$ (exists for all $D \neq \lim_{k\to\infty} R$). Compute $o = \mathsf{Rec}_f(\mathbf{S}^{T'} \cup \{S\})$. Output 1 if $o = s_1$, and 0 otherwise.

Analysis. If $S = S[i]$, then $o = s$ holds with probability at least $1 - \epsilon_1(k)$ from correctness of Sh_f. Thus, it outputs 1 w.p $\Pr_D(s_1) \pm o_k(1)$.

If $S = S'[i]$, since $f(\{i\}) = 0, f(\{T'\}) = 0$, by privacy of Sh_f, the distribution $\mathbf{S}^{T'} \cup \{S'[i]\}$, and thus $\mathsf{Rec}_f(\mathbf{S}^{T'} \cup \{S'[i]\})$ is statistically independent of s, s' selected during the experiment and is statistically close to R. Thus, in this case D_T recovers $o = s_1$ with probability $\Pr_R(s_1) \pm o_k(1)$. Overall, D_T's advantage is $|\Pr_D(s_1) - \Pr_R(s_1)| = \Omega_k(1)$. Finally, clearly this is the case iff. $D \neq \lim_{k\to\infty} R$.

Next, lets assume condition (2) holds, but condition (1) does not hold. Then D_T computes $o_1 = \mathsf{Rec}_f(\mathbf{S}^T)$ and $o_2 = \mathsf{Rec}_f(S)$, and outputs 1 iff. $o_1 = o_2$. As in this case $f(T) = f(\{P_i\}) = 1$, by correctness of Sh_f, we have that if $S = S[i]$, then D_T outputs 1 w.p $1 - o_k(1)$, and otherwise it outputs 1 w.p $\sum_{s\in\mathbb{F}}(\Pr_D[s])^2 \pm o_k(1)$. Thus, D_T it has distinguishing advantage $\Omega_k(1)$ for all but $|\mathbb{F}|$ distributions D (those with support of size 1). □

4.2. Positive results

We propose a general scheme that has tight parameters for both notions. The construction stems from similar ideas of the NIMPC construction in [7].

Construction 1. Sh_{f-gen}*: Let $T_1, \ldots, T_l$ be the minters of f in some order. To share a secret $s \in \mathbb{F}$, generate a sequence of vectors in a large enough vector space V (say over $\mathbb{F}_2$) as follows. There is a vector $S_{i,s',k}$ for every $s' \in \mathbb{F}$, $i \in [n]$ and $k \in [l]$. The vectors are random in the distribution under the following constraints: (1) $\forall j \in [l], \sum_{i\in T_j} S_{i,s,j} = 0$. (2) otherwise, all the vectors are independent (there are overall l linear dependencies). Pick a random permutation π on $[l]$. $S[i]$ is the matrix*

$$M_i = \begin{pmatrix} S_{i,s_1,\pi(1)} & \cdots & S_{i,s_1,\pi(l)} \\ \vdots & \ddots & \vdots \\ S_{i,s_{|\mathbb{F}|},\pi(1)} & \cdots & S_{i,s_{|\mathbb{F}|},\pi(l)} \end{pmatrix}$$

.

Theorem 3. *Fix a finite access structure $f : \mathcal{P} \to \{0, 1\}$. Then, Sh_{f-gen} in Construction 1 is a secret sharing functionality for f satisfying also:*

1. *It is (T, D, R_T)-share-indistinguishable for all T, R_T for which $(i, j) \in R_T$ implies $f(T' \cup \{i\}) = f(T' \cup \{j\})$ for all $T' \subseteq T$ and for all D (alternatively, it can achieve 0-error with expected polynomial time samplers).*
2. *It is (T, D, i)-share-unpredictable for all T, i for which the preconditions of Theorem 2 do not hold and all D.*

Both properties hold perfectly.[4]

Proof. Correctness: $\mathsf{Rec}_{f-gen}(\mathbf{S}^T)$ outputs s' if there exists $H \subseteq T, l'$ such that $\sum_{i \in H} S_{s',i,l'} = 0$. Otherwise outputs $\perp$. Clearly, Rec_{f-gen} is always correct if $f(T) = 1$ unless Sh_{f-gen} failed to sample the vectors so that all dependencies are as required, which results in a negligible error probability.

Privacy: Consider an unqualified $H \subseteq [n]$. The values $S_{s',i,j}$ held by H are random independent field elements (for all $s' \in \mathbb{F}$, including $s' = s$). Note that this analysis would hold even without applying π (indeed these are needed only for share-indistinguishability).

Share-indistinguishability. Let $(P_i, P_j) \in R_T$. Thus $f(T' \cup \{P_i\}) = f(T' \cup \{P_j\})$ for all $T' \subseteq T$. We conclude that the distributions of the distinguishers' view in both cases of the experiment are identical. In more detail, for a given T', either $T' \cup \{P_i\}$ is not a minterm, in which case neither is $T' \cup \{P_j\}$, or both are (follows by easy case analysis). Thus, $T_a \subseteq T$ contributes a sequence $M_i[s, \pi(a)]$ summing up to 0 if T_a is a minterm, or $M_i[s, \pi(a)]$ and S sum up to 0 if $T_a \cup \{P_i\}$ ($T_a \cup \{P_j\}$) is a minterm.

Share-unpredicatability. There are two cases. In one case, $f(T' \cup \{i\}) = 0$ for any $T' \subseteq T$. In the second case, $f(T) = 0$ and $f(P_i) = 1$. In both cases, it is easy to see that the adversary's input distributions are the same for both $S = S'[i]$ and $S = S[i]$, for all D.

□

Note that in Theorem 3, players P_i and P_j should only be isomorphic relatively to T ($f(T' \cup \{i\}) = f(T' \cup \{j\})$ for all $T' \subseteq T$). In the positive result of [2], players P_i and P_j are required to be isomorphic in the stronger sense: $f(T' \cup \{i\}) = f(T' \cup \{j\})$ for all $T' \in \mathcal{P} \setminus \{P_i, P_j\}$.

5. Conclusions and open problems

We have obtained tight characterization of when an access structure f is (participant-) anonymous *relatively* to a set T of parties according to two types of anonymity we revisit (*share-indistinguisability* and *share-unpredictability*). It is not hard to generalize our results to sets T_1, T_2 rather than single parties.

An interesting generalization of this work is pushing (say general, ind.-based) anonymity to the limit. Consider the scenario where reconstructing parties are still not aware of each other, but the shares are further subdivided

[4] An implementation of Sh_f would need to run in expected polynomial time.

into sub-shares, and it is only recorded by Rec_f which sub-shares where submitted, but not which came from the same share. To what extent can we hide the identity of T_1 (say, starting with an adversary involving $T = \phi$), hiding all except for $f(T_1)$, including even the size of T_1? We have work under way providing general positive results, assuming we allow for large (but still $1 - \Omega_k(1)$) distinguishing advantage. The full version of this paper (to appear soon), includes our results in both directions.

A more practical direction is finding Sh_f which is linear (unlike our scheme in Theorem 3). Then, using general MPC protocols for semi-honestly computing linear functions against arbitrary adversaries would strengthen Claim 1 to ensure statistical anonymity without the Q_2 condition on $\mathcal{T}$.

Finally, it is interesting to understand when efficient indistinguishability-based anonymity is possible (even in the computational setting). Some positive results (including threshold access structures) are given in [2] for the perfect case.

References

[1] C. Blundo, D. R. Stinson, Anonymous secret sharing schemes, Discrete Applied Mathematics 77 (1) (1997) 13–28.

[2] M. Guillermo, K. M. Martin, C. M. O'Keefe, Providing anonymity in unconditionally secure secret sharing schemes, Des. Codes Cryptography 28 (3) (2003) 227–245.

[3] R. F. Olimid, How to split a secret into unknown shares, Proceedings of the Romanian Academy, Series A 16, Special Issue (2015) 321–328.

[4] P. Rogaway, M. Bellare, Robust computational secret sharing and a unified account of classical secret-sharing goals, in: Proceedings of the 2007 ACM Conference on Computer and Communications Security, CCS 2007, 2007, pp. 172–184.

[5] M. Hirt, U. M. Maurer, Player simulation and general adversary structures in perfect multiparty computation, J. Cryptology 13 (1) (2000) 31–60.

[6] O. Goldreich, S. Micali, A. Wigderson, How to play any mental game or A completeness theorem for protocols with honest majority, in: Proceedings of the 19th Annual ACM Symposium on Theory of Computing, 1987, pp. 218–229.

[7] A. Beimel, A. Gabizon, Y. Ishai, E. Kushilevitz, S. Meldgaard, A. Paskin-Cherniavsky, Non-interactive secure multiparty computation, in: Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Proceedings, Part II, 2014, pp. 387–404.