Separating Sources for Encryption and Secret Sharing

Yevgeniy Dodis[1⋆], Krzysztof Pietrzak[2⋆⋆], and Bartosz Przydatek[2]

[1] Department of Computer Science, New York University
New York, NY, USA
dodis@cs.nyu.edu
[2] Department of Computer Science, ETH Zurich
8092 Zurich, Switzerland
{pietrzak,przydatek}@inf.ethz.ch

Abstract. Most cryptographic primitives such as encryption, authentication or secret sharing require randomness. Usually one assumes that perfect randomness is available, but those primitives might also be realized under weaker assumptions. In this work we continue the study of building secure cryptographic primitives from imperfect random sources initiated by Dodis and Spencer (FOCS'02). Their main result shows that there exists a (high-entropy) source of randomness allowing for perfect encryption of a bit, and yet from which one cannot extract even a single weakly random bit, separating encryption from extraction. Our main result separates encryption from 2-out-2 secret sharing (both in the information-theoretic and in the computational settings): any source which can be used to achieve one-bit encryption also can be used for 2-out-2 secret sharing of one bit, but the converse is false, even for high-entropy sources. Therefore, possibility of extraction strictly implies encryption, which in turn strictly implies 2-out-2 secret sharing.

1 Introduction

For many important tasks, such as cryptography, randomness is indispensable. Usually one assumes that all parties have access to a perfect random source, but this assumption is at least debatable, and the question what kind of imperfect random sources can be used in various applications has attracted a lot of attention.

Extraction. The easiest such class of sources consists of *extractable* sources for which one can deterministically extract nearly perfect randomness, and then use it in any application. Although examples of such non-trivial sources are known [vN51,Eli72,Blu86,LLS89,CGH$^+$85,BBR88,AL93,CDH$^+$00,DSS01,KZ03,TV00], most natural sources such as the so called entropy sources[1] [SV86,CG88,Zuc96]

⋆ Supported in part by NSF career award CCR-0133806 and NSF grant CCR-0311095.
⋆⋆ Supported by the Swiss National Science Foundation, project No. 200020-103847/1.

[1] Informally, entropy sources guarantees that every distribution in the family has a non-trivial amount of entropy (and possibly more restrictions), but do not assume

are easily seen to be non-extractable. One can then ask a natural question whether perfect randomness is indeed needed for the considered application. Clearly, the answer depends on the application. In particular, the natural fundamental question is to understand the extent to which a given application can be based on imperfect randomness, and also to compare the randomness requirements for different applications.

Probabilistic Algorithms and Interactive Protocols. For example, a series of celebrated results [VV85,SV86,CG88,Zuc96,ACRT99] showed that entropy sources are necessary and sufficient for simulating probabilistic polynomial-time algorithms — namely, problems which do not *inherently* need randomness, but which could potentially be sped up using randomization. Thus, extremely weak imperfect sources can still be tolerated for this application domain. This result was recently extended to interactive protocols by Dodis et al. [DOPS04].

Encryption. On the other hand, McInnes and Pinkas [MP90] showed that unconditionally secure symmetric encryption cannot be based on entropy sources, even if one is restricted to encrypting a single bit. This result was recently strengthened by Dodis et al. [DOPS04], who showed that entropy sources are not sufficient even for *computationally* secure encryption (as well as essentially any other task involving "privacy"). On the opposite side, Dodis and Spencer [DS02] showed that randomness extraction is not necessary for the existence of secure encryption (at least when restricted to a single bit). Specifically, they show that there are sources which can be used to perfectly encrypt a bit but cannot be used to extract a single bit. This even holds if one additionally requires all the distributions in the imperfect source to have high min-entropy. Thus, good sources for encryption lie strictly in between extractable and entropy sources.

Authentication. In the usual non-interactive (i.e., one-message) setting, Maurer and Wolf [MW97] show that for sufficiently high entropy rate (specifically, more than 1/2), even general entropy sources are sufficient for unconditional one-time authentication, while Dodis and Spencer [DS02] showed that smaller rate entropy sources are indeed insufficient to authenticate even a single bit. On the other hand, [DS02] also show that for all entropy levels (in particular, below 1/2) there exist "severely non-extractable" imperfect sources which are sufficient for non-trivial authentication. Thus good sources for authentication once again lie strictly in between extractable and entropy sources. The relation to encryption sources is currently open (see Section 5). On a related note, [DOPS04] considered the existence of computationally secure digital signature (and thus also message authentication) schemes, and show that the latter seem to be possible even with general entropy sources, at least under very strong but seemingly reasonable computational assumptions. In the interactive setting, Renner and Wolf [RW03] show (indeed, highly interactive) information-theoretic authentication protocols capable of tolerating any constant-fraction entropy rate.

---

independence between different symbols of the source. In this sense they are the most general sources one would wish to tolerate.

Secret Sharing? In this work we consider for the first time another cryptographic primitive which inherently requires randomness: secret sharing. In particular, we concentrate on the simplest case of 2-out-2 (denoted simply 2-2) secret sharing: one wants to split a message m into shares S_1 and S_2 so that neither share leaks any information about m, and yet m can be reconstructed from both shares. We first observe that (either information-theoretic or computational) encryption implies the existence of a corresponding 2-2 secret sharing: one simply sets S_1 to be the decryption key, and S_2 to be the encryption of the message M under this key. Our main technical result is to show that the converse of this statement is false, at least when restricting to one-bit message. Namely, there exist imperfect sources sufficient for perfect secret sharing of a bit, but for which any bit encryption scheme can be insecure with constant distinguishing probability (on a positive note, we show that one cannot push this probability too close to 1). Additionally, just like in the case of separation between encryption and extraction [DS02], our separation can be extended to hold even if one additionally requires all the distributions in the imperfect source to have high min-entropy.[2] Moreover, our information-theoretic separation above can be extended even to the computational setting. This means that there exist high-entropy sources for which one can build *efficient* 2-out-2 secret sharing, but any (efficient) encryption scheme can be broken by an *efficient* distinguisher on an *efficiently-samplable* distribution from our source.[3]

To summarize (see Figure 1), extraction strictly implies encryption [DS02] which in turn, as we show, strictly implies 2-2 secret sharing.

Comparing Cryptographic Primitives. As we see, our work continues the approach initiated by Dodis and Spencer [DS02] to compare different cryptographic tasks according to how they utilize randomness. Namely, given a block length n of our randomness source, and the (min-)entropy threshold $m \leq n$, we say that primitive A implies primitive B if whenever an imperfect source $\mathscr{S}$ of length n and (min-)entropy m is sufficient to implement A, then one could also implement B with $\mathscr{S}$. When $m = n$, we get back to the case of perfect randomness, where primitive A implies B if and only if the smallest number of truly random bits needed to implement A is at least as large as the smallest number of truly random bits to implement B. As was shown by [DS02,DOPS04] and continued here, many implications true in the perfect case simply stop being true the moment we allow for slightly imperfect random sources (i.e., allow $m < n$). In other words, these implications inherently rely on perfect randomness. On the other hand, some implications continue to hold (at least to some extent) even with imperfect randomness, implying they have more to do with the cryptographic aspect of the problem rather than the availability of true randomness. We believe that such comparison between cryptographic primitives sheds more

[2] In particular, we construct sources with min-entropy only a constant away from the maximal entropy (cf. Lemma 3 and Theorem 2).

[3] In fact, even the process of finding such an efficiently-samplable distribution can be done efficiently (with exponentially high probability), given only the oracle access to the encryption oracle.

light on how they utilize randomness, and also serves as a stepping stone toward classifying imperfect sources sufficient for different cryptographic tasks.

ORGANIZATION. We give the preliminary definitions of our primitives in Section 2. Our main technical result comparing sources for (information-theoretic) encryption and 2-2 secret sharing is given in Section 3. In Section 4 we extend our results to the computational setting. Finally, in Section 5 we take a brief look at authentication and discuss some open problems considering imperfect sources sources sufficient for various cryptographic applications.

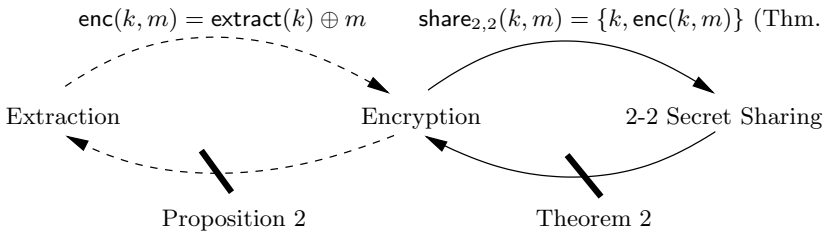

Fig. 1. The solid arrows indicate the implication and the separation we will prove.

2 Notation and Definitions

We use calligraphic letters like $\mathcal{X}$ to denote sets. The corresponding large letter X usually denotes a random variable over $\mathcal{X}$ and the small letter x an element from $\mathcal{X}$. We use $H(X)$ to denote the Shannon entropy of random variable X.

$X \in_\Omega \mathcal{X}$ means that X is a random variable whose distribution is Ω and $x \in_\Omega \mathcal{X}$ means that x is a value sampled from $\mathcal{X}$ with distribution Ω. $U_\mathcal{X}$ denotes the uniform distribution over $\mathcal{X}$. We write U_n to denote $U_{\{0,1\}^n}$, the uniform distribution over n-bit strings. A source $\mathscr{S}$ over $\mathcal{X}$ is a set of distributions over $\mathcal{X}$.

Definition 1 *A distribution Ω over $\mathcal{K}$ has* **min-entropy** *d if no element has probability more than 2^{-d}, i.e. $\max_{k \in \mathcal{K}} \mathsf{Pr}(k = k' | k' \in_\Omega \mathcal{K}) \leq 2^{-d}$. The largest such d is denoted $H_\infty(\Omega)$.* **A source $\mathscr{S}$ over $\mathcal{K}$ has min-entropy** *d if it only contains distributions with min-entropy at least d. The d-***weak source** *over $\mathcal{K}$ is the source which contains all distributions over $\mathcal{K}$ with min-entropy at least d.*

Definition 2 *A random variable B over $\{0,1\}$ is ϵ-***fair** *if*

$$\min\{\mathsf{Pr}(B = 0), \mathsf{Pr}(B = 1)\} \geq \epsilon$$

(so a uniform random bit is $1/2$-fair and a constant bit is 0-fair). A source $\mathscr{S}$ over $\mathcal{K}$ is ϵ-fair if there exists a one-bit extractor (which is simply a function $\mathsf{extract} : \mathcal{K} \to \{0,1\}$) such that $\mathsf{extract}(K)$, where $K \in_\Omega \mathcal{K}$, is ϵ-fair for all $\Omega \in \mathscr{S}$.

Definition 3 *An* **encryption scheme** *is a pair of algorithms* $\mathsf{enc} : \mathcal{K} \times \mathcal{M} \to \mathcal{C}$ *and* $\mathsf{dec} : \mathcal{K} \times \mathcal{C} \to \mathcal{M}$ *which for all keys* $k \in \mathcal{K}$ *and messages* $m \in \mathcal{M}$ *satisfies*

$$\mathsf{dec}(k, \mathsf{enc}(k, m)) = m \tag{1}$$

A source $\mathscr{S}$ *over* $\mathcal{K}$ *allows for* **perfect encryption** *of* $\mathcal{M}$ *if there is an encryption scheme such that for all distributions* $\Omega \in \mathscr{S}$ *the ciphertexts leak no information about the encrypted message* M*, i.e. for any random variable* M

$$\forall\, \Omega \in \mathscr{S} : H(M \mid \mathsf{enc}(K, M)) = H(M) \text{ where } K \in_\Omega \mathcal{K} \tag{2}$$

A source $\mathscr{S}$ *over* $\mathcal{K}$ *allows for* δ**-encryption** *if there is an encryption scheme such that for all distributions* $\Omega \in \mathscr{S}$ *the statistical distance of the encryption of any two distinct messages* m_1 *and* m_2 *is at most* δ*, i.e.*

$$\max_{\Omega \in \mathscr{S}, m_1 \neq m_2} \frac{1}{2} \sum_{c \in \mathcal{C}} |\mathsf{Pr}_{k \in_\Omega \mathcal{K}}(\mathsf{enc}(k, m_1) = c) - \mathsf{Pr}_{k \in_\Omega \mathcal{K}}(\mathsf{enc}(k, m_2) = c)| \leq \delta \tag{3}$$

Note that perfect encryption is 0-encryption and sending the plaintext is 1-encryption.

Definition 4 *For* $t, n \in \mathbb{Z}, t \leq n$ *a* t*-*n **secret sharing** *is a pair of algorithms* $\mathsf{share}_{t,n} : \mathcal{K} \times \mathcal{M} \to \mathcal{X}^n$ *and* $\mathsf{reconstruct}_{t,n} : \mathcal{X}^t \to \mathcal{M}$ *which for all keys* $k \in \mathcal{K}$ *and all* $m \in \mathcal{M}$ *satisfies*

$$\forall \mathcal{T} \subseteq \mathsf{share}_{t,n}(k, m) \text{ where } |\mathcal{T}| = t \text{ we have } \mathsf{reconstruct}_{t,n}(\mathcal{T}) = m \tag{4}$$

A source $\mathscr{S}$ *over* $\mathcal{K}$ *allows for* **perfect t-n secret sharing** *of* $\mathcal{M}$ *if any set of less than* t *shares does not reveal any information about the shared* M*, i.e. for all* $\Omega \in \mathscr{S}$ *and all* $1 \leq i_1 < i_2 < \ldots < i_{t-1} \leq n$ *we have for distributions* M

$$H(M \,|\, S_{i_1}, S_{i_2}, ..., S_{i_{t-1}}) = H(M) \text{ where } K \in_\Omega \mathcal{K}, \{S_1, ..., S_n\} \leftarrow \mathsf{share}_{t,n}(K, M) \tag{5}$$

Note that (4) means that from any t shares one can reconstruct m. In terms of *perfect randomness*, the uniform distribution over $\{0,1\}^n$ is necessary and sufficient to perfectly encrypt $\mathcal{M} = \{0,1\}^n$ (i.e. n-bit strings) for example by using the key k as a one time pad:

$$\mathsf{enc}(k, m) = k \oplus m \qquad \mathsf{dec}(k, c) = c \oplus m$$

where $\oplus$ denotes the bitwise XOR. U_n is also necessary and sufficient (as the dealer's randomness) to construct a perfect 2-2 secret sharing of $\{0,1\}^n$, for example as:

$$\mathsf{share}_{2,2}(k, m) = \{k, k \oplus m\} \qquad \mathsf{reconstruct}_{2,2}(s_1, s_2) = s_1 \oplus s_2$$

In the next section we will show that in terms of non-perfect randomness these two tasks are no longer equivalent. The sources which allow for perfect encryption also allow for 2-2 secret sharing (of the same message space) but not vice-versa. More precisely, we show that every source which allows for perfect 2-2 secret sharing of one bit allows for 1/2-encryption of one bit, but in general not for δ-encryption of one bit for $\delta < 1/3$. This even holds if we require the source to have high min-entropy.

3 Separating Encryption from Secret Sharing

We can now formally state the results of [MP90] and [DS02].

Proposition 1 ([MP90]) *The $(n-2)$-weak source over $\{0,1\}^n$ does not allow for δ-encryption of even 1 bit for any $\delta \neq 0$.*

So for every one-bit encryption scheme with key-space $\{0,1\}^n$ there exists a distribution for the keys with min-entropy $n-2$ such that the ciphertext always completely reveals the message.

Proposition 2 ([DS02]) *There is a source over $\{0,1\}^n$ which allows for perfect encryption of one bit, but which is not $2^{-n/2}$-fair.*

This separation holds even if we require the source to have high min-entropy: for any $\epsilon > 2^{-n/2+1}$ there is a source $\mathscr{S}$ over $\{0,1\}^n$ with min-entropy $n - \log(1/\epsilon) - O(1)$ which allows for perfect encryption of one bit but which is not ϵ-fair.

3.1 Encryption → 2-2 Secret Sharing

Theorem 1 *Any source $\mathscr{S}$ over $\mathcal{K}$ which allows for perfect encryption of $\mathcal{M}$ allows for perfect 2-2 secret sharing of $\mathcal{M}$.*

Proof: For $\mathsf{enc}, \mathsf{dec}$ which satisfy properties (1) and (2) we define for all $k \in \mathcal{K}, m \in \mathcal{M}$

$$\mathsf{share}_{2,2}(k,m) = (k, \mathsf{enc}(k,m)) \quad \text{and} \quad \mathsf{reconstruct}_{2,2}(s_1,s_2) = \mathsf{dec}(s_1,s_2).$$

Property (1) implies immediately that this scheme satisfies (4). It also satisfies property (5) as for any random variables M and $\Omega \in \mathscr{S}, K \in_\Omega \mathcal{K}$ we have that $H(M \mid K) = H(M)$ as K is independent of M and $H(M \mid \mathsf{enc}(K,M)) = H(M)$ follows from (2). ■

In the following section we show that for $\mathcal{M} = \{0,1\}$, the converse is not true.

3.2 2-2 Secret Sharing ↛ Encryption

In this section we will prove our main technical result (Theorem 2 below), namely that sources which allow for 2-2 secret sharing do not allow for encryption in general. We split the proof of the theorem into the following three lemmas.

Lemma 1 *There is a source which allows for perfect 2-2 secret sharing of a bit but does not allow for δ-encryption of a bit for any $\delta < 1/3$.*

This separation is in some sense not so strong as the separation for encryption from extraction where a source was shown which allows perfect encryption but not even a weak form of extraction. The question arises if we can get something as $\delta \leq 1 - o(1)$ (and not just $\delta < 1/3$) here too. The answer is no, since already $\delta \leq 1/2$ is not achievable as shown in the next lemma.

Lemma 2 *Any source which allows for perfect 2-2 secret sharing of a bit allows for* $1/2$*-encryption of a bit.*

We prove Lemma 1 by showing a concrete source which contains only four distributions over a domain of size six. Here the question arises whether this separation only works for such toy examples and possibly breaks down when we require the source to have high min-entropy. This is not the case: we show how one can turn such a toy-example into a high min-entropy source with the same parameters.

Lemma 3 *For any* $t \in \mathbb{N}$ *there is a source as in Lemma 1, where the distributions in the source have range of size* $6t$ *and the min-entropy of each distribution is at least* $\log(6t) - \log(192)$.

Combining Lemma 2 and Lemma 3 we get the following theorem

Theorem 2 *There are sources over any* $\mathcal{K}$ *with min-entropy* $\log|\mathcal{K}| - 11$ *which allow for perfect 2-2 secret sharing but do not allow for* δ*-encryption of one bit for any* $\delta < 1/3$.

From the positive side, any source which allows for perfect 2-2 secret sharing of a bit allows for $1/2$*-encryption of one bit.*

Theorem 2 is stated for sources over any $\mathcal{K}$ and not just for sets of size $6t$ as in Lemma 3. This is compensated for by an additional factor of $\log(6)$ in the min-entropy gap (i.e. we have a gap of $11 > \log(192) + \log(6)$).

Proof of Lemma 1: Let $\mathscr{S}$ be a source over $\mathcal{K} = \{k_1, \ldots, k_6\}$ which contains 4 distributions $\Omega_1, \ldots, \Omega_4$ where each Ω_i is the uniform distribution over $\mathcal{S}_i \subset \mathcal{K}$ with $\mathcal{S}_1 = \{k_1, k_2\}, \mathcal{S}_2 = \{k_3, k_4\}, \mathcal{S}_3 = \{k_1, k_3, k_5\}$ and $\mathcal{S}_4 = \{k_1, k_4, k_6\}$ respectively. Lemma 1 follows from the two claims below.

Claim 1 $\mathscr{S}$ *allows for perfect 2-2 secret sharing of one bit.*

Proof: We define the sharing $\mathsf{share}_{2,2} : \mathcal{K} \times \{0,1\} \to \mathcal{A} \times \mathcal{B}$, where $\mathcal{A} = \{a_1, a_2, a_3, a_4\}$ and $\mathcal{B} = \{b_1, b_2, b_3, b_4\}$ as shown in Figure 2. A key k_i is represented by a pair of directed edges, where the edge from $\mathcal{A}$ to $\mathcal{B}$ corresponds to the shares of 0, and the edge from $\mathcal{B}$ to $\mathcal{A}$ to the shares of 1. For example $\mathsf{share}_{2,2}(k_1, 0) = (a_3, b_2)$ and $\mathsf{share}_{2,2}(k_1, 1) = (a_1, b_2)$.

For any (a_i, b_j) there is at most one possible $m \in \{0,1\}$ such that $(a_i, b_j) = \mathsf{share}_{2,2}(k, m)$ for some $k \in \mathcal{K}$. Thus for any random variable M it always holds that $H(M \mid \mathsf{share}_{2,2}(k, M)) = 0$.

Note that for any $i, 1 \leq i \leq 4$, Ω_i is the uniform distribution over some subset of $\mathcal{K}$ whose corresponding directed edges (as shown in Figure 2) form a directed cycle, where the edges alternate between $\mathcal{A}$ and $\mathcal{B}$ (e.g. for Ω_1 we have the cycle $a_3 \to b_2 \to a_1 \to b_4 \to a_3$). So the distribution on $\mathcal{A}$ is the same no matter if we choose a random edge from $\mathcal{A}$ to $\mathcal{B}$ (a sharing of the secret 0) or from $\mathcal{B}$ to $\mathcal{A}$ (a sharing of the secret 1) on this cycle. This proves that the random variable A defined as $(A, B) = \mathsf{share}_{2,2}(k \in_{\Omega_i} \mathcal{K}, M)$ is independent of M and $H(M \mid A) = H(M)$ (and similarly for $H(M \mid B) = H(M)$). $\square$

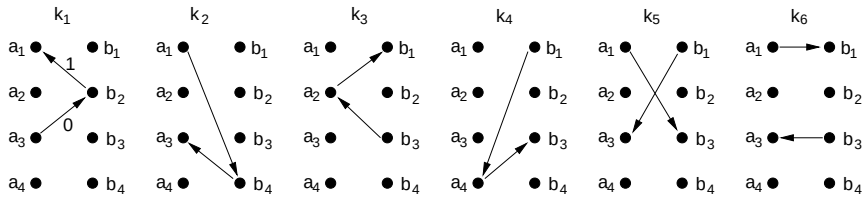

Fig. 2. The mapping $\mathsf{share}_{2,2}$ from the proof of Lemma 1.

Claim 2 *$\mathscr{S}$ does not allow for δ-encryption of a bit for any $\delta < 1/3$.*

Proof: Consider any mapping $\mathsf{enc} : \mathcal{K} \times \{0,1\} \to \mathcal{C}$. We will prove that for our source $\mathscr{S}$ this enc cannot satisfy Definition 3 with $\delta < 1/3$. Recall that $\mathscr{S}$ contains the distributions $\Omega_i, 1 \leq i \leq 4$, where Ω_i is uniform over $\mathcal{S}_i$. Consider the graphs G_i, with $V(G_i) = \mathcal{C}$, $E(G_i) = \{(\mathsf{enc}(k,0), \mathsf{enc}(k,1)) | k \in \mathcal{S}_i\}$, where each edge is labeled with the corresponding $k \in \mathcal{K}$. We will show that $E(G_i)$ does not form a directed cycle for at least one $i, 1 \leq i \leq 4$.

Suppose this is not the case, then $E(G_1)$ forms a cycle of length 2, say $k_1 = (c_1, c_2), k_2 = (c_2, c_1)$. And similarly for $E(G_2)$, say $k_3 = (c_3, c_4), k_4 = (c_4, c_3)$. If $E(G_3)$ forms a cycle then (because of the above) either $c_1 = c_4$ or $c_2 = c_3$ but not both must hold (e.g. if $c_1 = c_4$ then we can set $k_5 = (c_2, c_3)$). Similarly if $E(G_4)$ is a cycle then either $c_1 = c_3$ or $c_2 = c_4$ but not both must hold. We can write this two conditions as $(c_1 = c_4 \oplus c_2 = c_3) \wedge (c_1 = c_3 \oplus c_2 = c_4) = true$, which cannot be satisfied and we have a contradiction.

For an i where $E(G_i)$ does not form a directed cycle (such an i exists as we just proved) there are vertices $c', c'' \in V(G_i)$ such that $indegree(c') > outdegree(c')$ and $outdegree(c'') > indegree(c'')$, which using $|E(G_i)| = |\mathcal{S}_i| \leq 3$ gives

$$\frac{1}{2} \sum_{c \in \mathcal{C}} |Pr_{k \in_{\Omega_i} \mathcal{K}}(\mathsf{enc}(k,0) = c) - Pr_{k \in_{\Omega_i} \mathcal{K}}(\mathsf{enc}(k,1) = c)| \geq 1/3.$$

So because of this Ω_i property (3) cannot be satisfied with $\delta < 1/3$. □

■

Proof of Lemma 3: We will now show how to make a high min-entropy source out of a toy example like the one from Lemma 1. Let $\mathcal{K}$ and $\mathcal{S}_1, \ldots, \mathcal{S}_4 \subset \mathcal{K}$ be as in the proof of Lemma 1. For $t \geq 2$ and for each $i, 1 \leq i \leq t$ let $\mathcal{K}_i = \{k_{i,1}, \ldots, k_{i,6}\}$ be a copy of $\mathcal{K}$, and let $\mathcal{S}_{i,j} = \{k_{i,x} : k_x \in \mathcal{S}_j\}$ denote the corresponding subsets. The key-space $\widetilde{\mathcal{K}}$ of our source is

$$\widetilde{\mathcal{K}} = \mathcal{K}_1 \cup \mathcal{K}_2 \cup \ldots \cup \mathcal{K}_t \,. \tag{6}$$

For a set $I \subseteq \{1, \ldots, t\}$ and a mapping $\sigma : \{1, \ldots, t\} \to \{1, \ldots, 4\}$ we define

$$\mathcal{T}_{I,\sigma} = \bigcup_{i \in I} \mathcal{S}_{i,\sigma(i)}.$$

Our source $\mathscr{S}$ contains all uniform distributions over the sets $\mathcal{T}_{I,\sigma}$ with $|I| \geq \lceil t/64 \rceil$. That is, our source contains all the uniform distributions over sets which are constructed by taking the union of subsets $\mathcal{S}_{i,j} \subset \mathcal{K}_i$ from at least a 1/64'th fraction of the $\mathcal{K}_i$'s. Since each distribution is uniform over a set of size at least $2t/64$, $\mathscr{S}$ has min-entropy at least $\log(t/32) = \log(6t) - \log(192)$.

The source $\mathscr{S}$ allows for perfect 2-2 secret sharing of one bit: On input a key $k_{i,j} \in \mathcal{K}_i$ the dealer can compute the shares as he would for the key $k_j \in \mathcal{K}$ in Claim 1 in the proof of Lemma 1. One can assume (but this is not necessary) that the dealer also publishes i, then we have the same situation as in Claim 1 and it follows from (the proof of) Claim 1 that this is indeed a perfect 2-2 secret sharing.

It only remains to show that $\mathscr{S}$ does not allow for δ-encryption of one bit with $\delta < 1/3$. For this consider any mapping $\mathsf{enc} : \widetilde{\mathcal{K}} \times \{0,1\} \to \mathcal{C}$. As shown in the proof of Lemma 1, for each $i, 1 \leq i \leq t$, there is a distribution Ω_i which is uniform over the set $\mathcal{S}_{i,j}$ for some $j, 1 \leq j \leq 4$, satisfying

$$\frac{1}{2}\sum_{c\in\mathcal{C}} |\mathsf{Pr}_{k\in_{\Omega_i}\widetilde{\mathcal{K}}}(\mathsf{enc}(k,0)=c) - \mathsf{Pr}_{k\in_{\Omega_i}\widetilde{\mathcal{K}}}(\mathsf{enc}(k,1)=c)| \geq 1/3. \tag{7}$$

Let $\mathcal{X}_i$ denote such a $\mathcal{S}_{i,j}$. Consider a random mapping $\phi : \mathcal{C} \to \{-1,1\}$. We say that $\mathcal{X}_i$ is *good* if

$$\forall c \in \mathcal{C}: \quad 0 \leq \phi(c)\cdot(\,|\{k \in \mathcal{X}_i : \mathsf{enc}(k,0)=c\}| - |\{k \in \mathcal{X}_i : \mathsf{enc}(k,1)=c\}|\,). \tag{8}$$

As $|\mathcal{X}_i| \leq 3$ the rhs of eq. (8) can be nonzero for at most 6 different $c \in \mathcal{C}$ and thus $\mathcal{X}_i$ is good with probability at least 2^{-6}. This shows (as a simple application of the probabilistic method) that there is a ϕ for which at least $\lceil t2^{-6} \rceil$ of the $\mathcal{X}_i$'s are good. Fix such a ϕ and let $\mathcal{X}$ be the union of good sets. The uniform distribution over $\mathcal{X}$, $\Omega_{\mathcal{X}}$, is in $\mathscr{S}$, but it does not allow for δ-encryption with $\delta < 1/3$: Let γ_i be the event that $k \in \mathcal{X}_i$ (below all not explicitly labeled probabilities are over $k \in_{\Omega_{\mathcal{X}}} \widetilde{\mathcal{K}}$).

$$\begin{aligned}
&\frac{1}{2}\sum_{c\in\mathcal{C}} |\mathsf{Pr}(\mathsf{enc}(k,0)=c) - \mathsf{Pr}(\mathsf{enc}(k,1)=c)| \\
&\overset{(8)}{=} \frac{1}{2}\sum_{c\in\mathcal{C}} \phi(c)(\mathsf{Pr}(\mathsf{enc}(k,0)=c) - \mathsf{Pr}(\mathsf{enc}(k,1)=c)) \\
&= \frac{1}{2}\sum_{c\in\mathcal{C}}\sum_{i=1}^{t} \phi(c)(\mathsf{Pr}(\gamma_i)\mathsf{Pr}(\mathsf{enc}(k,0)=c|\gamma_i) - \mathsf{Pr}(\gamma_i)\mathsf{Pr}(\mathsf{enc}(k,1)=c|\gamma_i)) \\
&\overset{(8)}{=} \sum_{i=1}^{t} \mathsf{Pr}(\gamma_i)\frac{1}{2}\sum_{c\in\mathcal{C}} |(\mathsf{Pr}(\mathsf{enc}(k,0)=c|\gamma_i) - \mathsf{Pr}(\mathsf{enc}(k,1)=c|\gamma_i)| \\
&= \sum_{i=1}^{t} \mathsf{Pr}(\gamma_i)\frac{1}{2}\sum_{c\in\mathcal{C}} |\mathsf{Pr}_{k\in_{\Omega_i}\widetilde{\mathcal{K}}}(\mathsf{enc}(k,0)=c) - \mathsf{Pr}_{k\in_{\Omega_i}\widetilde{\mathcal{K}}}(\mathsf{enc}(k,1)=c)| \\
&\overset{(7)}{\geq} \sum_{i=1}^{t} \mathsf{Pr}(\gamma_i)\frac{1}{3} \;=\; \frac{1}{3}
\end{aligned}$$

■

Proof of Lemma 2: Let $\mathscr{S}$ be a source with distributions over $\mathcal{K}$ which allows for perfect 2-2 secret sharing and $\mathsf{share}_{2,2} : \mathcal{K} \times \{0,1\} \to \mathcal{A} \times \mathcal{B}$ be an appropriate sharing (we can wlog. assume that $\mathcal{A} \cap \mathcal{B} = \emptyset$). To prove the lemma we first define a mapping $\mathsf{enc} : \mathcal{K} \times \{0,1\} \to \mathcal{C}$ (where $\mathcal{C} = \mathcal{A} \cup \mathcal{B}$), and then prove that it is a 1/2-encryption (i.e. satisfies eq. (3) with $\delta = 1/2$).

For $k \in \mathcal{K}$ and $m \in \{0,1\}$ let $(a_{m,k}, b_{m,k}) = \mathsf{share}_{2,2}(k, m)$, we set

$$\mathsf{enc}(k, m) = a_{m,k} \text{ if } a_{0,k} \neq a_{1,k} \text{ and } b_{m,k} \text{ otherwise}$$

We cannot have $a_{0,k} = a_{1,k}$ and $b_{0,k} = b_{1,k}$ simultaneously as otherwise the share $(a_{0,k}, b_{0,k})$ could be a share of either 0 or 1 which is impossible when the secret sharing is perfect. So we always have $\mathsf{enc}(k,0) \neq \mathsf{enc}(k,1)$ and decryption is always possible. We will now prove that this enc satisfies eq.(3) with $\delta \leq 1/2$ (as our plaintext-domain is only one bit we can set $m_1 = 0$ and $m_2 = 1$ in (3) wlog.). For any $\Omega \in \mathscr{S}$ we have (all probabilities are over $k \in_\Omega \mathcal{K}$) using $\mathcal{C} = \mathcal{A} \cup \mathcal{B}, \mathcal{A} \cap \mathcal{B} = \emptyset$

$$\frac{1}{2} \sum_{c \in \mathcal{C}} |\mathsf{Pr}(\mathsf{enc}(k,0) = c) - \mathsf{Pr}(\mathsf{enc}(k,1) = c)| \tag{9}$$

$$= \frac{1}{2} \sum_{a \in \mathcal{A}} |\mathsf{Pr}(\mathsf{enc}(k,0) = a) - \mathsf{Pr}(\mathsf{enc}(k,1) = a)| \tag{10}$$

$$+ \quad \frac{1}{2} \sum_{b \in \mathcal{B}} |\mathsf{Pr}(\mathsf{enc}(k,0) = b) - \mathsf{Pr}(\mathsf{enc}(k,1) = b)| \tag{11}$$

We will show that the term (9) is $\leq 1/2$ (which is exactly the statement of the Lemma) by showing that (10) is equal to 0 and (11) is $\leq 1/2$. Let the random variables A_m, B_m be defined as $(A_m, B_m) = \mathsf{share}_{2,2}(k \in_\Omega \mathcal{K}, m)$, and let $\mathcal{K}^{\neq} = \{k \in \mathcal{K} | a_{0,k} \neq a_{1,k}\}$. From (5) we see that in a perfect 2-2 secret sharing the distribution of the share of each player is independent of the shared secret. This is used in the first step below (again all probabilities are over $k \in_\Omega \mathcal{K}$)

$$0 = \frac{1}{2} \sum_{a \in \mathcal{A}} |\mathsf{Pr}(A_0 = a) - \mathsf{Pr}(A_1 = a)| \tag{12}$$

$$= \frac{1}{2} \sum_{a \in \mathcal{A}} |\mathsf{Pr}(k \in \mathcal{K}^{\neq}) \left(\mathsf{Pr}(A_0 = a | k \in \mathcal{K}^{\neq}) - \mathsf{Pr}(A_1 = a | k \in \mathcal{K}^{\neq})\right) \tag{13}$$

$$+ \mathsf{Pr}(k \notin \mathcal{K}^{\neq}) \underbrace{\left(\mathsf{Pr}(A_0 = a | k \notin \mathcal{K}^{\neq}) - \mathsf{Pr}(A_1 = a | k \notin \mathcal{K}^{\neq})\right)}_{= 0 \text{ by the definition of } \mathcal{K}^{\neq}} | \tag{14}$$

$$= \frac{1}{2} \sum_{a \in \mathcal{A}} |\mathsf{Pr}(k \in \mathcal{K}^{\neq}) \left(\mathsf{Pr}(A_0 = a | k \in \mathcal{K}^{\neq}) - \mathsf{Pr}(A_1 = a | k \in \mathcal{K}^{\neq})\right) | \tag{15}$$

Here (15) is exactly (10) so (10) is 0. We now show that (11) is $\leq 1/2$:

$$\mathsf{Pr}(\mathsf{enc}(k,0) = b) = \mathsf{Pr}(B_0 = b \wedge k \notin \mathcal{K}^{\neq}) \leq \mathsf{Pr}(B_0 = b)$$

$$\mathsf{Pr}(\mathsf{enc}(k,1)=b)=\mathsf{Pr}(B_1=b\wedge k\notin\mathcal{K}^{\neq})\leq\mathsf{Pr}(B_1=b)$$

Perfect secret sharing implies $\mathsf{Pr}(B_0=b)=\mathsf{Pr}(B_1=b)$ and as the difference of two positive values cannot be larger than any those values we get

$$\frac{1}{2}\sum_{b\in\mathcal{B}}|\mathsf{Pr}(\mathsf{enc}(k,0)=b)-\mathsf{Pr}(\mathsf{enc}(k,1)=b)|\leq\frac{1}{2}\sum_{b\in\mathcal{B}}\mathsf{Pr}(B_0=b)=\frac{1}{2}.$$

■

4 Some Computational Aspects

Until now we have only considered an information theoretic setting. In particular, we did not care about whether the primitives, the attacks or the sampling considered can be efficiently realized. In this section, which we keep rather informal, we examine some computational aspects of the results from the previous section.

4.1 Computational Version of Theorem 1

The proof of Theorem 1, which states that any source which can be used for encryption can also be used for secret sharing, easily translates in the computational setting.

Proposition 3 (Computational version of Theorem 1) *(informal) Any source $\mathscr{S}$ over $\mathcal{K}$ which allows "computationally secure" encryption of $\mathcal{M}$ allows for "computationally secure" 2-2 secret sharing of $\mathcal{M}$.*

In the above proposition we left open what "computationally secure" exactly means. A direct translation from the information theoretic setting would advise the following security notion for encryption: the adversary can choose two messages m_0 and m_1 and then, given the encryption of m_b for a random b, should not be able to guess b (much better than with prob. $1/2$).[4] Then the security achieved for secret sharing is the following: First the adversary can choose two messages m_0 and m_1. Then, given one share of m_b for random b he cannot guess b (i.e. which message was shared). A stronger notion for encryption (e.g. semantic security) will result in a stronger security guarantee for secret sharing.

4.2 Computational Version of Theorem 2

We now take a look at Theorem 2 which follows from the Lemma 2 and Lemma 3. The proof of Lemma 2 translates into the computational setting.

[4] This notion is weaker than the notion of semantic security, where the adversary can additionally ask for encryptions of his choice except for m_0 and m_1.

Proposition 4 (Computational version of Lemma 2) *(informal)*
Any source which allows for "computationally secure" 2-2 secret sharing of a bit allows for "computationally secure" $1/2$*-encryption of a bit.*

As before, "computationally secure" can have several meanings (and a stronger notion for secret sharing implies a stronger notion for 1/2-encryption). Also the concept of δ-encryption has a natural meaning in the computational setting, where it means that the distinguishing advantage of any efficient adversary for the ciphertexts of two messages m_0 and m_1 is at most negligibly larger than $1/2 + \delta/2$.

Lemma 3 states that there is a source with high min-entropy which allows for 2-2 secret sharing but not for δ-encryption of one bit with $\delta < 1/3$. We can strengthen this lemma in several ways by considering computational aspects. In particular, we can require the following properties:

i. The secret sharing is efficient.
ii. For every encryption scheme the source contains an *efficiently samplable* distribution, for which the encryption-scheme is not 1/3-secure.
iii. There exists an efficient algorithm which breaks the 1/3-security of the encryption scheme under the efficiently samplable distribution from (ii).
iv. One can efficiently find the distribution from (ii).

We can achieve all four points simultaneously. However, to satisfy properties (ii) and (iv) we must be able to efficiently compute encryptions (either by getting a polynomial-size circuit or access to an oracle which computes encryptions given a key and a message). We now describe how one can adapt the proof of Lemma 3 (which we assume the reader is familiar with) to achieve these additional properties.

We can encode the key-space (see eq. (6)) as pairs of integers, i.e. $\widetilde{\mathcal{K}} \equiv [1, \ldots, t] \times [1, \ldots, 6]$. With this encoding property (i) (efficient secret sharing) is achieved: recall that the shares of $m \in \{0, 1\}$ under key (i, j) are $\mathsf{share}_{2,2}(k_j, m)$ with $\mathsf{share}_{2,2}(\cdot, \cdot)$ as defined in the proof of Lemma 1, which can be computed in constant time.

We now describe how to efficiently sample a distribution from our source which breaks the 1/3-security of enc. The distribution from the lemma is not efficiently samplable as the ϕ used to define it cannot be computed efficiently; We have only shown that a suitable ϕ exists — where suitable means that at least $\lceil t2^{-6} \rceil$ of the $\mathcal{X}_i$'s are good — using the probabilistic method. The argument used there was that a random ϕ satisfies (8) with probability at least 2^{-6}, as the rhs of (8) is nonzero for at most 6 different $c \in \mathcal{C}$. Fortunately for this argument we don't need a *random* ϕ — in fact, 6-wise independence is enough. Therefore if $\tau : \mathcal{W} \times \mathcal{C} \to \{-1, 1\}$ is an (efficiently computable) 6-wise independent function, then there is some key $w \in \mathcal{W}$ such that $\tau(w, \cdot)$ is good for $\lceil t2^{-6} \rceil$ of the $\mathcal{X}_i$'s (as $t2^{-6}$ is a lower bound for the expected number of good $\mathcal{X}_i$'s for a 6-wise independent function).

With this efficient $\phi(\cdot) = \tau(w, \cdot)$, we can now efficiently sample a key (using uniform randomness) according to the distribution of Lemma 3 (i.e. a random key from the union of all good sets) as follows:

1. Choose an integer $i, 1 \leq i \leq t$ uniformly at random.
2. Find a $j, 1 \leq j \leq 4$ (say the smallest) such that Ω_i, the uniform distribution over $\mathcal{X}_i = \mathcal{S}_{i,j}$, satisfies (7).
3. Check if this $\mathcal{X}_i$ is good, i.e. satisfies (8). If it does not, return to step 1.
4. If $|\mathcal{X}_i| = 2$ then return to step 1 with probability 1/3. (This is done to equalize the proportional weights of the $\mathcal{X}_i$'s of size 2 and 3.)
5. Output a key chosen uniformly at random from $\mathcal{X}_i$.

Note that this sampling will terminate in expected polynomial time if we can compute enc (in Step 2) and ϕ (in Step 3) efficiently.

We now describe an efficient breaking algorithm for enc, thus satisfying property (iii). Equation (8) tells us that the encryption of 0 and 1 have statistical distance at least 1/3, and from (7) we see that given a ciphertext c of a message m, $\phi(c)$ is an optimal guess on m. So if $\phi(\cdot)$ can be efficiently computed (which is the case if we set it to $\tau(w, \cdot)$, as described before), then we can efficiently break the 1/3 security of enc.

Finally we come to property (iv), which now can be stated as how to find a key w for our 6-wise independent function τ, such that $\phi(\cdot) = \tau(w, \cdot)$ is good for a 1/64 fraction of the $\mathcal{X}_i$'s. Unfortunately, for a given w one can't efficiently check if $\tau(w, \cdot)$ is good on a 1/64 fraction as for that we would have to go over all $\mathcal{X}_i$ for $i = 1, \ldots, t$, but t is exponential. But we can efficiently find a w such that $\tau(w, \cdot)$ will be good on a slightly smaller subset, say a 1/66 fraction, with probability exponentially close to 1 as follows. Choose a random w and approximate the fraction on which $\tau(w, \cdot)$ is good by randomly sampling $i \in [1, \ldots, t]$ and checking if it is good for $\mathcal{X}_i$. Accept this w if it was good on, say at least a 1/65 fraction, of the $\mathcal{X}_i$'s. By the Chernoff bound, the probability we will accept a w which is not good on at least a 1/66 fraction of all $\mathcal{X}_i$'s is exponentially small in the number of samples we have drawn for the approximation. Further, by the Markov bound we are guaranteed that we pick a w which is good on at least a 1/65 fraction after a constant number of tries.

5 Open Problems

There are many interesting open questions considering imperfect sources for various cryptographic applications. In our opinion the most dazzling one is whether the reductions from Proposition 2 and Theorem 2 generalize to larger domains. Already if we only extend the domain of the message space from two to three we cannot even show a sub-constant bound for the fairness.

Open Problem 1 *Is there an $\epsilon(n) \in o(1)$ such that there exist sources over $\{0,1\}^n$ which allow for the encryption (or 2-2 secret sharing) of a trit[5] but cannot be used to extract an $\epsilon(n)$-fair bit (recall that for bits one can show $\epsilon(n) = 2^{-n/2}$).*

[5] A trit is like a bit but can take three and not just two values.

AUTHENTICATION. Another interesting primitive we did not consider so far is authentication. Here, we will only consider the one-bit case, which already leaves several interesting open questions.

Definition 5 *We say that a source $\mathscr{S}$ with distributions over some set $\mathcal{K}$ allows for τ-***authentication of one bit** *if there is a mapping $\mathsf{auth} : \mathcal{K} \times \{0,1\} \to \mathcal{A}$ such that for all distributions $\Omega \in \mathscr{S}$ and $k \in_\Omega \mathcal{K}$ we have $\min_k H_\infty(\mathsf{auth}(k,0) \mid \mathsf{auth}(k,1)) \geq -\log \tau$ and $\min_k H_\infty(\mathsf{auth}(k,1) \mid \mathsf{auth}(k,0)) \geq -\log \tau$.*

Note that τ-authentication of a bit simply means that given the authenticator $\mathsf{auth}(k,b)$ of a bit $b \in \{0,1\}$ the probability that one can guess $\mathsf{auth}(k,1-b)$ (the authenticator of the other bit) correctly is at most $2^{-\tau}$. Authentication is very undemanding in its randomness requirements, any source whose min-entropy is large enough will do.

Proposition 5 ([MW97]) *The $(n+\tau)$-weak source over $\{0,1\}^{2n}$ allows for τ-authentication of one bit.*

The authentication which achieves the above bound is extremely simple: use the first and the last n bits to authenticate 0 and 1 respectively. Note that any half has min-entropy at least τ even when given the other half as an n-bit string has min-entropy of at most n. As such weak-sources are not enough for encryption (see Proposition 1), this already shows that sources for authentication do not allow for encryption, and one can easily show that they do not allow for secret sharing and any other cryptographic primitive requiring privacy we could think of.

But how about the other direction? Can sources which allow for encryption or secret sharing always be used for authentication? Recall, in the case of perfect randomness the result of [DS02] implies that $2n$ uniform bits are both necessary and sufficient for achieving n-authentication of 1-bit (in particular, we need at least 2 bits to do anything non-trivial at all), which means we can only hope that encryption (or 2-2 secret sharing) of at least $2n$-bits might (or might not) imply n-authentication of even a single bit. More generally,

Open Problem 2 *Find a lower bound for $\tau(n)$ and an upper bound for $\gamma(n)$ in the following statement (bounds for $n = 1$ only are already interesting):*

> *A source (possibly with some guaranteed min-entropy) which allows for the encryption (or 2-2 secret sharing) of $2n$ bits must always allow for $\tau(n)$-authentication of one bit, but in general not for $\gamma(n)$-authentication.*

As we remarked, we know that $n \geq \gamma(n) \geq \tau(n)$. Interestingly, we observe below that 3-3 secret sharing *does* imply authentication. (As a sanity check, in case of perfect randomness both n-authentication of a bit and 3-3 secret sharing of n bits need $2n$ perfectly random bits.)

Claim 3 *Any source which allows for perfect 3-3 secret sharing of n bits allows for n-authentication of a bit.*

Proof: This reduction can be achieved as follows: first compute the sharing (S_1, S_2, S_3) for some constant message, say $m = 0^n$. Now use S_1 as the authentication of 0 and S_2 as the authentication of 1. We observe that the joint distribution of shares S_1 and S_2 when $m = 0^n$ is the same as when m is uniform over $\{0,1\}^n$ as otherwise the shares S_1 and S_2 would leak information on which is the case. So let M be uniform over $\{0,1\}^n$ and K be chosen according to a distribution from our source. With the above observation we now must only prove that

$$H_\infty(S_1|S_2) \geq n \quad \text{and} \quad H_\infty(S_2|S_1) \geq n \quad \text{where} \quad (S_1, S_2, S_3) = \mathsf{share}_{3,3}(K, M).$$

Here $H_\infty(S_1|S_2) \geq n$ means that $H_\infty(S_1|S_2 = s) \geq n$ for all s in the support of S_2 (and not as sometimes used that the expectation over S_2 is at least n, i.e. not $\sum_s \mathsf{Pr}(S_2 = s) H_\infty(S_1|S_2 = s) \geq n$). Now by the definition of perfect secret sharing we have

$$H_\infty(M|S_2S_3) = n \quad \text{and} \quad H_\infty(M|S_1S_2S_3) = 0, \tag{16}$$

which implies $H_\infty(S_1|S_2S_3) \geq n$. To see this assume that this was not true, i.e. we have for some s_1, s_2, s_3 that $\mathsf{Pr}(S_1 = s_1|S_2 = s_2, S_3 = s_3) > 2^{-n}$, but then for $m = \mathsf{reconstruct}_{3,3}(s_1, s_2, s_3)$ also $\mathsf{Pr}(m|S_2 = s_2, S_3 = s_3) > 2^{-n}$ which contradicts $H_\infty(M|S_2S_3) = n$. The desired $H_\infty(S_1|S_2) \geq n$ now easily follows from $H_\infty(S_1|S_2S_3) \geq n$, and $H_\infty(S_2|S_1) \geq n$ can be shown similarly. □

Now, as encryption implies 2-2 secret sharing and 3-3 secret sharing implies authentication, a proof that 2-2 secret sharing implies some non-trivial 3-3 secret sharing would immediately give a non-trivial bound for $\tau(n)$ from Open Problem 2. Moreover, we think that comparing 2-2 secret sharing of $2n$ bits with 3-3 secret sharing of n bits is interesting in its own right, since it would show that different t-m secret sharing schemes have (or have not) different requirements on the way they utilize randomness, even if the same amount of perfect randomness is required for them.

References

[ACRT99] Alexander Andreev, Andrea Clementi, Jose Rolim, and Luca Trevisan. Dispersers, deterministic amplification, and weak random sources. *SIAM J. on Computing*, 28(6):2103–2116, 1999.

[AL93] Miklós Ajtai and Nathal Linial. The influence of large coalitions. *Combinatorica*, 13(2):129–145, 1993.

[BBR88] Charles H. Bennett, Gilles Brassard, and Jean-Marc Robert. Privacy amplification by public discussion. *SIAM J. on Computing*, 17(2):210–229, 1988.

[Blu86] Manuel Blum. Independent unbiased coin flips from a correlated biased source — a finite state Markov chain. *Combinatorica*, 6(2):97–108, 1986.

[CDH+00] Ran Canetti, Yevgeniy Dodis, Shai Halevi, Eyal Kushilevitz, and Amit Sahai. Exposure-resilient functions and all-or-nothing transforms. In *Proc. EUROCRYPT'00*, pages 453–469, 2000.

[CG88] Benny Chor and Oded Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM J. on Computing*, 17(2):230–261, 1988.

[CGH+85] Benny Chor, Oded Goldreich, Johan Håstad, Joel Friedman, Steven Rudich, and Roman Smolensky. The bit extraction problem of t-resilient functions. In *Proc. 26th IEEE FOCS*, pages 396–407, 1985.

[DOPS04] Yevgeniy Dodis, Shien Jin Ong, Manoj Prabhakaran, and Amit Sahai. On the (im)possibility of cryptography with imperfect randomness. In *Proc. 45th IEEE FOCS*, pages 196–205, 2004.

[DS02] Yevgeniy Dodis and Joel Spencer. On the (non-)universality of the one-time pad. In *Proc. 43rd IEEE FOCS*, pages 376–388, 2002.

[DSS01] Yevgeniy Dodis, Amit Sahai, and Adam Smith. On perfect and adaptive security in exposure-resilient cryptography. In *Proc. EUROCRYPT'01*, pages 301–324, 2001.

[Eli72] Peter Elias. The efficient construction of an unbiased random sequence. *Ann. Math. Stat.*, 43(2):865–870, 1972.

[KZ03] Jess Kamp and David Zuckerman. Deterministic extractors for bit-fixing sources and exposure-resilient cryptography. In *Proc. 44th IEEE FOCS*, pages 92–101, 2003.

[LLS89] David Lichtenstein, Nathan Linial, and Michael Saks. Some extremal problems arising from discrete control processes. *Combinatorica*, 9(3):269–287, 1989.

[MP90] James L. McInnes and Benny Pinkas. On the impossibility of private key cryptography with weakly random keys. In *Proc. CRYPTO'90*, pages 421–436, 1990.

[MW97] Ueli Maurer and Stefan Wolf. Privacy amplification secure against active adversaries. In *Proc. CRYPTO'97*, pages 307–321, 1997.

[RW03] Renato Renner and Stefan Wolf. Unconditional authenticity and privacy from an arbitrary weak secret. In *Proc. CRYPTO'03*, pages 78–95, 2003.

[SV86] Miklos Santha and Umesh V. Vazirani. Generating quasi-random sequences from semi-random sources. *JCSS*, 33(1):75–87, 1986.

[TV00] Luca Trevisan and Salil Vadhan. Extracting randomness from samplable distributions. In *Proc. 41st IEEE FOCS*, pages 32–42, 2000.

[vN51] John von Neumann. Various techniques used in connection with random digits. *National Bureau of Standards, Applied Mathematics Series*, 12:36–38, 1951.

[VV85] Umesh V. Vazirani and Vijay V. Vazirani. Random polynomial time is equal to slightly-random polynomial time. In *Proc. 26th IEEE FOCS*, pages 417–428, 1985.

[Zuc96] David Zuckerman. Simulating BPP using a general weak random source. *Algorithmica*, 16(4/5):367–391, 1996.