Does Privacy Require True Randomness?

Carl Bosley* Yevgeniy Dodis†

October 3, 2006

Abstract

Most cryptographic primitives require randomness (for example, to generate their secret keys). Usually, one assumes that perfect randomness is available, but, conceivably, such primitives might be built under weaker, more realistic assumptions. This is known to be true for many authentication applications, when entropy alone is typically sufficient. In contrast, all known techniques for achieving privacy seem to fundamentally require (nearly) perfect randomness. We ask the question whether this is just a coincidence, or, perhaps, privacy inherently requires true randomness?

We completely resolve this question for the case of (information-theoretic) private-key encryption, where parties wish to encrypt a b-bit value using a shared secret key sampled from some imperfect source of randomness $\mathscr{S}$. Our main result shows that if such n-bit source $\mathscr{S}$ allows for a secure encryption of b bits, then one can deterministically extract roughly $(b - \log n)$ nearly perfect random bits from $\mathscr{S}$. Further, this bound is nearly tight: there exist sources $\mathscr{S}$ allowing one to perfectly encrypt $(\log n - \operatorname{loglog} n)$ bits, but not to deterministically extract even a single slightly unbiased bit.

Hence, to a large extent, *true randomness is inherent for encryption*: either the key length must be exponential in the message length b, or one can deterministically extract nearly b almost unbiased random bits from the key. In particular, *the one-time pad scheme is essentially "universal"*.

*Dept. of Computer Science, New York University. bosley@cs.nyu.edu. Supported by NSF Graduate Research Fellowship.

†Dept. of Computer Science, New York University. dodis@cs.nyu.edu. Supported in part by NSF Grants #0515121, #0133806, #0311095.

1 Introduction

Randomness is important in many areas of computer science. It is especially indispensible in cryptography: secret keys must be random, any many cryptographic taks, such as public-key encryption, secret sharing or commitment, require randomness for every use. Typically, 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] are easily seen to be non-extractable. One can then ask a natural question whether perfect randomness is indeed inherent for the considered application, or perhaps one can do with weaker, more realistic assumptions. Clearly, the answer depends on the application.

Positive Results. For one such application domain, a series of celebrated results [VV85, SV86, CG88, Zuc96, ACRT99] showed that entropy sources are 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 later extended to interactive protocols by Dodis et al. [DOPS04].

Moving to cryptographic applications, entropy sources are typically sufficient for authentication applications, since entropy is enough to ensure unpredictability. For example, in the non-interactive (i.e., one-message) setting Maurer and Wolf [MW97] show that, for a sufficiently high entropy rate (specifically, more than 1/2), entropy sources are indeed sufficient for unconditional one-time authentication (while Dodis and Spencer [DS02] showed that smaller rate entropy sources are not sufficient to authenticate even a single bit). Moreover, in the interactive setting, Renner and Wolf [RW03] show information-theoretic authentication protocols capable of tolerating any constant-fraction entropy rate. Finally, Dodis et al. [DOPS04] consider the existence of computationally secure digital signature (and thus also message authentication) schemes, and, under (necessarily) strong, but plausible computational assumptions, once again show that entropy sources are enough to build such signature schemes. From a different angle, [DS02] also show that for all entropy levels (in particular, below 1/2) there exist "severely non-extractable" imperfect sources which are nevertheless sufficient for non-trivial non-interactive authentication. Thus, good sources for authentication certainly do not require perfect randomness.

Randomness for Privacy? The situation is much less clear for privacy applications, whose security definitions include some kind of indistinguishability. Of those, the most basic and fundamental is the question of (private-key) encryption, whose definition requires that the encryptions of any two messages are indistinguishable. (Indeed, this will be the subject of this work.)

With one exception (discussed shortly), all the known results indictate that true randomness might be inherent for privacy applications, such as encryption. First, starting with the Shannon's one-time scheme [Sha49], all the exisiting methods for building secure encryptions schemes, as well as other privacy primitives, crucially rely on perfect randomness somewhere in their design. And this is true even in the computational setting. Second, attempts to build secure encryption schemes (and other privacy primitives)

[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 independence between different symbols of the source. In this sense they are the most general sources one would wish to tolerate, since cryptography clearly requires entropy.

based on known "non-extractable" sources, such as various entropy sources, *provably failed*, indicating that such sources are indeed insufficient for privacy. For example, 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 subsequently 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", such as commitment, zero-knowledge and others).

The only reassuring result in the other direction is the work of Dodis and Spencer [DS02], who considered the setting of symmetric encryption, where the shared secret key comes from an imperfect random source, instead of being truly random. In this setting, they constructed a particular non-extractable imperfect source, nevertheless allowing one to perfectly encrypt *a single bit*. On a surface, this might seem to solve our question in the negative, suggesting that true randomness is not inherently required, at least for secure encryption. However, this conclusion is somewhat rushed (and will actually be disproved by our results). Indeed, we typically care about encrypting considerably more than a single bit. In such cases, it is certainly unreasonable to expect that, say, encryption of b bits will necessarily imply extraction of *exactly* b bits (which was indeed disproved by [DS02] for $b = 1$). One would actually *expect* that an implication, if true, would lose at least a few bits (perhaps depending on the statistical distance ε from the uniform distribution that we want our extraction to achieve).

For example, a source consisting of a single uniform distribution on $N = 2^b(1 + \varepsilon)$ values clearly allows one to perfectly encrypt b bits (via masking the message by adding to it the secret key modulo N). However, it is a simple exercise that any extractor attempting to extract more that $(b - \log\left(\frac{1}{\varepsilon}\right))$ bits would have a contant statistical distance from the uniform distribution.[2] On the other hand, this example certainly does not show that such random source has no "true randomness" in it, because it does: first, it *is* a uniform distribution (albeit not on an even power of 2 values), and, second, one can trivially extract *almost* b bits which are statistically close to uniform (concretely, $b - \log\left(\frac{1}{\varepsilon}\right)$ bits of statistical distance ε from uniform).

In particular, the results of [DS02] leave open the following extreme possibilities: (a) perhaps any source encrypting already two bits must be extractable; or (b) perhaps there exists an n-bit source allowing one to perfectly encrypt almost n bits, and yet not to extract even a single bit. Clearly, possibility (a) would show that true randomness *is* inherent for encryption, while possibility (b) that it is *not*. As we will see shortly, both (a) and (b) happen to be false, but our point is that the results of [DS02] regarding *one-bit* encryption and extraction do not answer what we feel is the "real" question for private-key encryption:

Assume some imperfect source allows for a secure encryption of b bits.
Does it neccesarily allow for extraction of at least one (and, hopefully, close to b) nearly perfect bits?

Our Result. We completely resolve the above question. Our main result shows that if an n-bit source $\mathscr{S}$ allows for a secure encryption (even slightly biased), then one can deterministically extract roughly $(b - \log n)$ nearly perfect random bits from $\mathscr{S}$.[3] Moreover, this bound is essentially tight: there exists imperfect sources allowing one to perfectly encrypt $b \approx \log n - \operatorname{loglog} n$ bits, from which one cannot

[2] In the extreme case $b = 1$, this says that a uniform distribution on $\{0, 1, 2\}$ allows one to encrypt a bit, but not to extract a bit, which is obvious. Indeed, the actual contribution of [DS02] was not to show that the separation between one bit encryption and one bit extraction *exists* — as we just saw, this is trivial — but to show (a) a very strong level of separation and, more difficultly, (b) that the separation holds even if one additionally requires all the distributions in the imperfect source to have high entropy.

[3] A bit more precisely, $b - \log n - 2\log\left(\frac{1}{\varepsilon}\right)$ bits within statistical distance $(\varepsilon + \delta)$ from uniform, where δ is the bias of the encryption scheme; see Theorem 1(a).

deterministically extract even a single slightly unbiased bit (see Theorem 1(b)).[4] Hence, to a large extent, *true randomness is inherent for (information-theoretic) private-key encryption*:

Either the key length n must be exponential in the message length b, or
One can deterministically extract nearly b almost unbiased random bits from the key.

In particular, in the case when b is large enough, so that it is infeasible to sample more than 2^b (imperfect) bits for one's secret key, our result implies the following. In order to build a secure b-bit encryption scheme, one must come up with a source of randomness from which can anyway deterministically extract almost b nearly random bits! Notice, since such extracted bits can then be used as a one-time pad, we get that any b-bit encryption scheme can in principle be converted to a "one-time-pad-like" scheme capable of encrypting nearly b bits! In this sense, our results show that, for the purpose of *information-theoretically* encrypting a "non-trivial" number of bits, the one-time pad scheme is essentially *"universal"*.[5]

Organization. We define the needed notation in Section 2, which also allows us to formally state our main result (Theorem 1). In Section 3 we prove that encryption of b bits using an n-bit key implies extraction of roughly $b - \log n$ random bits. In Section 4, which is the main technical section (and is further split into subsections), we show that encryption of up to $(\log n - \operatorname{loglog} n)$ bits does not necessarily imply extraction of even a single bit. Finally, in Section 5 we conclude and state some open problems.

2 Notation and Definitions

We use calligraphic letters, like $\mathcal{X}$, to denote finite sets. The corresponding large letter X is then used to denote a random variable over $\mathcal{X}$, while the lowercase letter x — a particular element from $\mathcal{X}$. $U_{\mathcal{X}}$ denotes the uniform distribution over $\mathcal{X}$. A source $\mathscr{S}$ over $\mathcal{X}$ is a set of distributions over $\mathcal{X}$. We write $X \in \mathscr{S}$ to state that $\mathscr{S}$ contains a distribution X.

The statistical distance $\mathsf{SD}(X_1, X_2)$ between two random variables X_1, X_2 is

$$\mathsf{SD}(X_1, X_2) = \frac{1}{2} \sum_{x \in \mathcal{X}} \big| \Pr[X_1 = x] - \Pr[X_2 = x] \big| \tag{1}$$

If $\mathsf{SD}(X_1, X_2) \leq \varepsilon$, this means that no (even computationally unbounded) distinguisher D can tell apart a sample from X_1 from a sample from X_2 with an advantage greater than ε.

Definition 1 A random variable R over $\mathcal{R}$ is *ε-fair* if $\mathsf{SD}(R, U_{\mathcal{R}}) \leq \varepsilon$. Given a source $\mathscr{S}$ over some set $\mathcal{K}$, a function $\mathsf{Ext} : \mathcal{K} \to \mathcal{R}$ is an *$(\mathscr{S}, \varepsilon)$-extractor* if for all $K \in \mathscr{S}$, $\mathsf{Ext}(K)$ is ε-fair:

$$\mathsf{SD}(\mathsf{Ext}(K), U_{\mathcal{R}}) \leq \varepsilon \tag{2}$$

If such Ext exists for $\mathscr{S}$, we say that $\mathscr{S}$ is *$(\mathcal{R}, \varepsilon)$-extractable*. ♢

Definition 2 An *encryption scheme* $\mathcal{E}$ over message space $\mathcal{M}$, key space $\mathcal{K}$ and ciphertext space $\mathcal{C}$ 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 $\mathsf{Dec}(k, \mathsf{Enc}(k, m)) = m$.

[4] This result can also be seen as a non-trivial extension of the separation result of [DS02] from 1-bit to (roughly) $(\log n)$-bit encryption. Indeed, without the entropy constraints, our proof is considerably more involved than that of [DS02].

[5] Of course, it is completely possible that the deterministic extractor that we show exists is much less efficient than the original encryption scheme. Moreover, we do lose $\log n + 2 \log\left(\frac{1}{\varepsilon}\right)$ bits by such turning of our scheme into a "one-time pad".

Given a source $\mathscr{S}$ over $\mathcal{K}$, we say that the encryption scheme $\mathcal{E}$ is $(\mathscr{S}, \delta)$-*secure* if for all messages $m_1, m_2 \in \mathcal{M}$ and all distributions $K \in \mathscr{S}$ we have

$$\mathsf{SD}(\mathsf{Enc}(K, m_1), \mathsf{Enc}(K, m_2)) \leq \delta \tag{3}$$

If $\mathscr{S}$ admits some $(\mathscr{S}, \delta)$-secure encryption $\mathcal{E}$ over $\mathcal{M}$, we say that $\mathscr{S}$ is $(\mathcal{M}, \delta)$-*encryptable*.
When $\delta = 0$, we say that $\mathcal{E}$ is *perfect* on $\mathscr{S}$, and $\mathscr{S}$ is perfectly encryptable (on $\mathcal{M}$). $\diamondsuit$

Throughout we will use the following capital letters to denote the cardinalities of various sets: key set cardinality $|K| = N$, message set cardinality $|\mathcal{M}| = B$, ciphertext set cardinality $|\mathcal{C}| = S$, and extraction space cardinality $|\mathcal{R}| = L$. Although our results are general, for historical reasons it is customary to translate the results into "bit-notation". To accomodate these conventions, we let $b = \log B$, $\ell = \log L$, $n = \log N$ (here and elsewhere, all the logarithms are base 2), and will use the terms "b-bit encryption", "ℓ-bit extraction" or "n-bit key" with the obvious meanings attached. Moreover, we will slightly abuse the terminology and say that a source $\mathscr{S}$ is (1) *n-bit* if it is over a set $\mathcal{K}$ and $|\mathcal{K}| = N$; (2) (ℓ, ε)-*extractable* if it is $(\mathcal{R}, \varepsilon)$-extractable and $|\mathcal{R}| = L$, and (2) (b, δ)-*encryptable* if it is $(\mathcal{M}, \delta)$-encryptable and $|\mathcal{M}| = B$. Clearly, when b, ℓ or n are integers, this terminology is consistent with our intuitive understanding.

With this in mind, our main result can be restated as follows:

Theorem 1 *Secure encryption of b bits with an n-bit key requires nearly perfect randomness (in fact, almost b random bits!) if and only if b is greater than* $\log n$. *More precisely,*

(a) $\forall \varepsilon > 0$, if $\mathscr{S}$ is (b, δ)-encryptable, then $\mathscr{S}$ is $(b - \log n - 2\log\left(\frac{1}{\varepsilon}\right), \varepsilon + \delta)$-extractable. Thus, encryption of $b > \log n$ bits implies extraction of roughly $(b - \log n)$ nearly perfect bits.

(b) For any $b \leq \log n - \operatorname{loglog} n - 2$,[6] there exists $\mathscr{S}$ which is $(b, 0)$-encryptable, but not $(1, \varepsilon)$-extractable, where $\varepsilon = \frac{1}{2} - 2^{(2b - \frac{n}{2^b})} \geq \frac{1}{2} - \frac{1}{16n^2}$. Thus, even perfect encryption of nearly $\log n$ bits does not imply extraction of even a single slightly unbiased bit.

3 Encryption $\Rightarrow$ Extraction if $b > \log n$

In this section we prove the implication given in Theorem 1(a), which shows that encryption of b bits implies extraction of nearly b bits. Assume $\mathcal{E} = (\mathsf{Enc}, \mathsf{Dec})$ is $(\mathscr{S}, \delta)$-secure over message space $\mathcal{M}$, ciphertext space $\mathcal{C}$ and key space $\mathcal{K}$. Let $\ell = b - \log n - 2\log\left(\frac{1}{\varepsilon}\right)$, $L = 2^\ell = B\varepsilon^2 / \log N$, and $\mathcal{R}$ be an arbitrary set of cardinality L. Also, for convenience let us identify the message space $\mathcal{M}$ with $\{0, \ldots, B-1\}$.

We construct our extractor $\mathsf{Ext} : \mathcal{K} \to \mathcal{R}$ in two steps. First we construct a special "ciphertext packing" function $f : \mathcal{C} \to \mathcal{R}$. One can view such f as throwing ciphertexts (or "balls") c into "bins" $f(c)$. Pretend we have the needed packing f. Now let us fix any key $k \in \mathcal{K}$ and see where (necessarily distinct)[7] ciphertexts $\mathsf{Enc}(k, 0), \ldots, \mathsf{Enc}(k, B-1)$ end up under our packing f. To see this, given any bin $r \in \mathcal{R}$, let $X_{k,r}(f)$ denote the number of messages $m \in \mathcal{M}$ whose "balls" $\mathsf{Enc}(k, m)$ were thrown into bin r; i.e., m's such that $f(\mathsf{Enc}(k, m)) = r$. Since there are B balls and L bins, on average every bin should contain roughly B/L ciphertext balls $\mathsf{Enc}(k, m)$. Of course, depending on f, in reality some bins will have slightly more, and some slightly less than B/L ciphertexts under key k. What will determine a suitable f for us is exactly the fact that *every* bin should have very close to B/L ciphertexts.

[6] The formula also holds for $b = \log n - \operatorname{loglog} n - 1$, but yields a slightly smaller $\varepsilon = \frac{1}{2} - \frac{1}{4\log n}$.

[7] This follows from the fact that all decryptions $\mathsf{Dec}(k, \mathsf{Enc}(k, m)) = m$ are distinct.

Lemma 1 *For our setting of $|\mathcal{R}| = L$, there exists a packing $f : \mathcal{C} \to \mathcal{R}$ such that for* all *$k \in \mathcal{K}$ and $r \in \mathcal{R}$ we have*

$$\left| X_{k,r}(f) - \frac{B}{L} \right| \leq 2\varepsilon \cdot \frac{B}{L} \tag{4}$$

Proof: We prove the lemma by showing that the probability that a *random* f does not satisfy the condition above is less than 1. Therefore, there exists an f which satisfies the condition.

Let us fix a particular key $k \in \mathcal{K}$ and a particular output bin $r \in \mathcal{R}$. Then the number $X_{k,r}(f)$ is a random variable (over the choice of f) which counts the number of messages m such that $f(\mathsf{Enc}(k, m)) = r$. We can rewrite $X_{k,r}(f) = X_0 + \ldots + X_{B-1}$, where the indicator variable X_m is 1 if $f(\mathsf{Enc}(k, m)) = r$, and 0 otherwise. Since f is random and all balls $\mathsf{Enc}(k, 0), \ldots, \mathsf{Enc}(k, B-1)$ are distinct, all X_m's are independent and $\Pr[X_m = 1] = 1/L$. Thus, $\mathbb{E}[X_{k,r}(f)] = B/L$, and using the standard Chernoff's bound (e.g., see [Sho05], theorem 6.13.iii), we get

$$\Pr\left[\left| X_{k,r}(f) - \frac{B}{L} \right| \geq 2\varepsilon \cdot \frac{B}{L} \right] \leq 2 \cdot e^{-2\varepsilon^2 B/L} \tag{5}$$

Recalling that we set $L = B\varepsilon^2 / \log N$, the above probability is at most $2/N^2$. Taking now the union bound over all $k \in \mathcal{K}$ and $r \in \mathcal{R}$, we get that the probability that random f fails for any k and r is at most $2L/N < 1$, as needed. □

Corollary 2 *For our setting of $|\mathcal{R}| = L$, there exists a packing $f : \mathcal{C} \to \mathcal{R}$ such that for* all *distributions K on $\mathcal{K}$, we have*

$$\mathsf{SD}(f(\mathsf{Enc}(K, U_{\mathcal{M}})), U_{\mathcal{R}}) \leq \varepsilon \tag{6}$$

Proof: Fix f satisfying Equation (4) in Lemma 1, and let $X_{k,r} = X_{k,r}(f)$. Then for any k and r, it is easy to see that Equation (4) can be rewritten as follows:

$$\Pr_{U_{\mathcal{M}}}[f(\mathsf{Enc}(k, U_{\mathcal{M}})) = r] = \frac{X_{k,r}}{B} \in \left[(1 - 2\varepsilon) \cdot \frac{1}{L},\ (1 + 2\varepsilon) \cdot \frac{1}{L} \right] \tag{7}$$

Now, take any distribution K and let $p_k = \Pr[K = k]$. Then we get [8]

$$\begin{aligned}
\mathsf{SD}(f(\mathsf{Enc}(K, U_{\mathcal{M}})), U_{\mathcal{R}}) &= \frac{1}{2} \sum_r \left| \sum_k p_k \cdot \Pr_{U_{\mathcal{M}}}[f(\mathsf{Enc}(k, U_{\mathcal{M}})) = r] - \frac{1}{L} \right| \\
&= \frac{1}{2} \sum_r \left| \sum_k p_k \cdot \left(\Pr_{U_{\mathcal{M}}}[f(\mathsf{Enc}(k, U_{\mathcal{M}})) = r] - \frac{1}{L} \right) \right| \\
&\leq \frac{1}{2} \sum_r \sum_k p_k \cdot \left| \Pr_{U_{\mathcal{M}}}[f(\mathsf{Enc}(k, U_{\mathcal{M}})) = r] - \frac{1}{L} \right| \\
&\overset{\text{Eq.(7)}}{\leq} \frac{1}{2} \sum_r \sum_k p_k \cdot \frac{2\varepsilon}{L} = \varepsilon
\end{aligned}$$

□

[8] The derivation below is simply the general relation between the ℓ_1- and ℓ_∞-norms. See also Remark 1.

So far we only used the fact the $\mathcal{E}$ is correctly decryptable, but did not use the security of $\mathcal{E}$ under the source $\mathscr{S}$. Recall, the $(\mathscr{S}, \delta)$-security of $\mathcal{E}$ implies that for any $K \in \mathscr{S}$ and $m \in M$ we have $\mathsf{SD}(\mathsf{Enc}(K, 0), \mathsf{Enc}(K, m)) \leq \delta$. Since this holds for any m, it also holds for any distribution M on m, including the uniform distribution $U_{\mathcal{M}}$:

$$\mathsf{SD}(\mathsf{Enc}(K, 0), \mathsf{Enc}(K, U_{\mathcal{M}})) \leq \delta \tag{8}$$

We can finally collect all the pieces together. We fix f given by Corollary 2 and extend the ciphertext packing f to an extractor $\mathsf{Ext} : \mathcal{K} \to \mathcal{R}$ by simply setting $\mathsf{Ext}(k) = f(\mathsf{Enc}(k, 0))$. We now claim that Ext is an $(\mathscr{S}, \varepsilon + \delta)$-extractor. Take any $K \in \mathscr{S}$. Then, by the triangle inequality,

$$\begin{aligned} \mathsf{SD}(\mathsf{Ext}(K), U_{\mathcal{R}}) &= \mathsf{SD}(f(\mathsf{Enc}(K, 0)), U_{\mathcal{R}}) \\ &\leq \mathsf{SD}(f(\mathsf{Enc}(K, 0)), f(\mathsf{Enc}(K, U_{\mathcal{M}}))) + \mathsf{SD}(f(\mathsf{Enc}(K, U_{\mathcal{M}})), U_{\mathcal{R}}) \end{aligned}$$

Finally, the first term is at most δ by Equation (8), since the application of f can only reduce the statistical distance, while the second term is at most ε directly from Equation (6). This means that for all $K \in \mathscr{S}$, $\mathsf{SD}(\mathsf{Ext}(K), U_{\mathcal{R}}) \leq \delta + \varepsilon$, which completes the proof.

Remark 1 *Assume the encryption scheme $\mathcal{E}$ satisfies a slightly stronger property (which is always satisfied by perfect encryption with $\delta = 0$). Namely, for every $K \in \mathscr{S}$, every messages $m_1, m_2 \in \mathcal{M}$ and* every *ciphertext $c \in \mathcal{C}$, we have*

$$|\Pr[\mathsf{Enc}(K, m_1) = c] - \Pr[\mathsf{Enc}(K, m_2) = c]| \leq 2\delta \cdot \Pr[\mathsf{Enc}(K, m_2) = c]$$

then we would prove the existence of extractor Ext *such that for every $K \in \mathscr{S}$ and* every *$r \in \mathcal{R}$,*

$$\Pr[\mathsf{Ext}(K) = r] \in \left[\frac{1 - (\varepsilon + \delta)}{L}, \frac{1 + (\varepsilon + \delta)}{L}\right]$$

4 Encryption $\not\Rightarrow$ Extraction if $b < \log n - \operatorname{loglog} n$

In this section we prove the non-implication given in Theorem 1(b), which shows that even perfect encryption of nearly $\log n$ bits does not necessarily imply extraction of even a single bit. For that we need to define a specific b-bit encryption scheme $\mathcal{E} = (\mathsf{Enc}, \mathsf{Dec})$ and a source $\mathscr{S}$, such that $\mathscr{S}$ is perfect on $\mathcal{E}$, but "non-extractable". The proof will proceed in several stages.

4.1 Defining Good Encryption $\mathcal{E}$

As the first observation, we claim that we only need to define the encryption scheme $\mathcal{E}$, and then let the source $\mathscr{S} = \mathscr{S}(\mathcal{E})$ be the set of all key distributions K making $\mathcal{E}$ perfect:

$$\mathscr{S}(\mathcal{E}) = \{K \mid \forall\ m_1, m_2 \in \mathcal{M}, c \in \mathcal{C} \Rightarrow \Pr[\mathsf{Enc}(K, m_1) = c] = \Pr[\mathsf{Enc}(K, m_2) = c]\}$$

Indeed, $\mathscr{S}(\mathcal{E})$ is the largest source which is $(b, 0)$-encryptable by means of $\mathcal{E}$, so it is the hardest one to extract even a single bit from. We call distributions in $\mathscr{S}(\mathcal{E})$ *perfect* (for $\mathcal{E}$).

Although we are not required to do so, let us intuitively motivate our choice of $\mathcal{E}$ before actually defining it. For that it is very helpful to view our key space $\mathcal{K}$ in terms of the encryption scheme $\mathcal{E}$ as follows. Given any $\mathcal{E} = (\mathsf{Enc}, \mathsf{Dec})$, we identify each key $k \in \mathcal{K}$ with an ordered B-tuple of ciphertexts $(c_1, \ldots, c_B)$, where

$\mathsf{Enc}(k, m) = c_m$. Technically, some B-tuples might repeat for several keys, but it is easy to see that such "repeated" keys will only complicate our job.[9] More interestingly, some B-tuples might not correspond to valid keys. For example, this is the case when $c_i = c_j$ for some $i \neq j$, since then encryptions of i and j are the same under this key. Intuitively, however, the larger is the set of valid B-tuples of ciphertexts, the more variety we have in the set of perfect distributions $\mathscr{S}(\mathcal{E})$, and the harder it would be to extract from $\mathscr{S}(\mathcal{E})$. This suggests that every B-tuple $(c_1, \ldots, c_B)$ of ciphertexts should correspond to a potential key, except for the necessary constraint that all the c_m's must be distinct to allow unique decryption.

A bit more formally, we assume that N can be written as $N = S(S-1)\ldots(S-B+1)$ for some integer S.[10] Then we define the set $\mathcal{C} = \{1, \ldots S\}$ to be the set of ciphertexts, $\mathcal{M} = \{1, \ldots, B\}$ be the set of plaintexts, and view the key set $\mathcal{K}$ as the set of distinct B-tuples over $\mathcal{C}$:

$$\mathcal{K} = \{k = (c_1, \ldots c_B) \mid \forall\ i \neq j \Rightarrow c_i \neq c_j\}$$

We then define $\mathsf{Enc}((c_1 \ldots c_B), m) = c_m$, while $\mathsf{Dec}((c_1, \ldots, c_B), c)$ to be the (necessarily unique) m such that $c_m = c$, and arbitrarily if no such m exists. Notice, $N < S^B$, so that $S > N^{1/B}$, which is strictly greater than B when $b < \log n - \operatorname{loglog} n$. Thus, S contains enough ciphertexts to allow for B distinct encryptions.

Excluding 0-monochromatic Distributions. Let us now take an arbitrary bit extractor $\mathsf{Ext} : \mathcal{K} \to \{0,1\}$ and argue that it is not very good on the set of perfect distributions $\mathscr{S}(\mathcal{E})$. We say that a distribution K is 0-*monochromatic* if $\Pr[\mathsf{Ext}(K) = 0] = 1$. Clearly, if the set of perfect distributions $\mathscr{S}(\mathcal{E})$ contains a 0-monochromatic distribution K, then $\mathsf{SD}(\mathsf{Ext}(K), U_1) = \frac{1}{2}$ (here and below, U_1 is the uniform distribution of $\{0,1\}$), and we would be done. Thus, for the remainder of the proof we assume that $\mathscr{S}(\mathcal{E})$ *does not contain a* 0-*monochromatic distribution.* The heart of the proof then will consist of designing a perfect encryption distribution K such that

$$\Pr[\mathsf{Ext}(K) = 0] \leq \frac{B^2}{S} \tag{9}$$

Once this is done, recalling that $S > N^{1/B} = 2^{n/2^b}$ we immediately get

$$\mathsf{SD}(\mathsf{Ext}(K), U_1) = \left|\frac{1}{2} - \Pr[\mathsf{Ext}(K) = 0]\right| \geq \frac{1}{2} - 2^{(2b - \frac{n}{2^b})}$$

as claimed by Theorem 1(b). Thus, we concentrate on building a perfect distribution K satisfying Equation (9). For that, we need to (1) characterize perfect distributions using linear algebra; (2) use this characterization to understand the implication of the lack of 0-monochromatic perfect distributions; and (3) use this implication to construct the required perfect distribution K. We do so in the next three subsections.

4.2 Characterizing Perfect Distributions

Let K be any distribution on $\mathcal{K}$. Given a key $k = (c_1 \ldots c_B)$, let $p_k = x_{(c_1 \ldots c_B)} = \Pr[K = (c_1 \ldots c_B)]$ and p be the N-dimensional column vector whose k-th component is equal to p_k. Notice, being a probability vector, we know that $\sum p_k = 1$ and $p \geq 0$ (which is a shorthand for $p_k \geq 0$ for all k). Conversely, any such x defines a unique distribution K.

[9] We omit the argument, since it is not very illuminating. Essentially, such keys force us to consider more extractors when arguing lack of extraction, without expanding the "geometry" of perfect key distributions.

[10] If not, take largest S such that $N \geq S(S-1)\ldots(S-B+1)$, and work on the subset of $N' = S(S-1)\ldots(S-B+1)$ keys, but this will not change our bounds.

Assume now that K is a perfect encryption distribution for $\mathcal{E}$. This adds several more constraints on p. Specifically, a necessary and sufficient condition for a perfect encryption distribution is to require that for all $c \in \mathcal{C}$ and all $m > 1$, we have

$$\Pr[c_1 = c \mid (c_1 \dots c_B) \leftarrow K] = \Pr[c_m = c \mid (c_1 \dots c_B) \leftarrow K] \tag{10}$$

We can translate this into a linear equation by noticing that the left probability is equal to $\sum_{\{(c_1 \dots c_B):c_1=c\}} p_{(c_1 \dots c_B)}$, while the second — to $\sum_{\{(c_1 \dots c_B):c_m=c\}} p_{(c_1 \dots c_B)}$. Thus, Equation (10) can be rewritten as

$$\sum_{\{(c_1 \dots c_B):c_1=c\}} p_{(c_1 \dots c_B)} - \sum_{\{(c_1 \dots c_B):c_m=c\}} p_{(c_1 \dots c_B)} = 0 \tag{11}$$

We can then rewrite all these constraints on p into a more compact notation by defining a *constraint matrix* $V = \{v_{i,j}\}$, which has $(1+(B-1)S)$ rows (corresponding to the constraints) and N columns (corresponding to keys). The first row of V will consist of all 1's: $v_{1,k} = 1$ for all $k \in \mathcal{K}$. This would later correspond to the fact that $\sum p_k = 1$. To define the rest of V, which would correspond to $(B-1)S$ constraints from Equation (11), we first make our notation more suggestive. We index the N columns of V by tuples $(c_1, \dots c_B)$, and the remaining $(B-1)S$ rows of V by tuples (m, c), where $m \in \{2, \dots B\}$ and $c \in \{1 \dots S\}$. Then, we define

$$v_{(m,c),(c_1,\dots,c_B)} = \begin{cases} 1, & c = c_1, \\ -1, & c = c_m, \\ 0, & \text{otherwise.} \end{cases} \tag{12}$$

Now, Equation (11) simply becomes $\sum_k v_{(m,c),k} \cdot p_k = 0$. Finally, we define a $(1 + (B-1)S)$-column vector e by $e_1 = 1$ and $e_i = 0$ for $i > 1$. Combining all this notation, we finally get

Lemma 3 *An N-dimensional real vector p defines a perfect distribution K for $\mathcal{E}$ if and only if $Vp = e$ and $p \geq 0$.*

4.3 Using the Lack of 0-Monochromatic Distributions

Next, we use Lemma 3 to understand our assumption that no perfect distribution K is 0-monochromatic with respect to Ext. Before that, we remind a well known Farkas Lemma (e.g., see [Str80]):

Farkas Lemma. *For any matrix A and column vector e, the linear system $Ax = e$ has no solution $x \geq 0$ if and only if there exists a row vector y s.t. $yA \geq 0$ and $ye < 0$.*

Now, let $Z = \{k \mid \mathsf{Ext}(k) = 0\}$ be the set of "0-keys" under Ext, and let A denotes $(1 + (B-1)S) \times |Z|$-matrix equal to the constraint matrix V restricted its $|Z|$ columns in Z. Take any real vector p such that $p_k = 0$ for all $k \notin Z$. By Lemma 3, p corresponds to a (necessarily 0-monochromatic) perfect distribution K if and only if $Vp = e$ and $p \geq 0$. But since $p_k = 0$ for all $k \notin Z$, the above conditions are equivalent to saying that the $|Z|$-dimensional restriction $x = p|_Z$ of p to its coordinates in Z satisfies $Ax = e$ and $x \geq 0$. Conversely, any x satisfying the above constraints defines a 0-monochromatic perfect distribution p by letting $p|_Z = x$ and $p_k = 0$ for $k \notin Z$.

Thus, Ext defines no 0-monochromatic perfect distributions if and only if the constraints $Ax = e$ and $x \geq 0$ are unsatisfiable. But this is exactly the precondition to the Farkas' Lemma above! Using the Farkas Lemma on our A and e, we get the existence of the $(1 + (B-1)S)$-dimensional row vector y such that $yA \geq 0$ and $ye < 0$. Just like we did for the rows of V, we denote the first element of y by y_1, and use

the notation $y_{(m,c)}$ to denote the remaining elements of y. We now translate the constraints $yA \geq 0$ and $ye < 0$ using our specific choices of A and e.

Notice, since $e_1 = 1$ and $e_i = 0$ for $i > 1$, it means that $ye = y_1$, so the constraint that $ye < 0$ is equivalent to $y_1 < 0$. Next, recalling that A is just the restriction of V to its columns in Z, and that the first row of V is the all-1 vector, we get that $yA \geq 0$ is equivalent to saying that for all $(c_1, \ldots, c_B) \in Z$ we have

$$y_1 + \sum_{m>1} \sum_{c} y_{(m,c)} \cdot v_{(c_1,\ldots,c_B)} \geq 0 \tag{13}$$

Notice, since $y_1 < 0$, this equation implies that the double sum above is *strictly* greater than 0. Thus, recalling the definition of $v_{(c_1,\ldots,c_B)}$ given in Equation (12), we conclude that for all $k = (c_1, \ldots, c_B)$, such that $\mathsf{Ext}(k) = 0$, we have

$$\sum_{m>1} \left(y_{(m,c_1)} - y_{(m,c_m)} \right) > 0 \tag{14}$$

The last equation finally allows us to derive the implication we need:

Theorem 2 *Assume* Ext *defines no* 0*-monochromatic perfect distributions. Then there exist real numbers* $\left\{ y_{(m,c)} \mid m \in \{2 \ldots B\}, c \in \{1 \ldots S\} \right\}$ *such that the following holds. If a key* $k = (c_1, \ldots, c_B)$ *is such that*

$$y_{(m,c_1)} - y_{(m,c_m)} \leq 0 \qquad \text{for all } m > 1, \tag{15}$$

then $\mathsf{Ext}(k) = 1$.

Proof: Summing Equation (15) for all $m > 1$ we get a contradiction to Equation (14), which means that $\mathsf{Ext}(k) \neq 0$; i.e., $\mathsf{Ext}(k) = 1$. □

4.4 Building Non-Extractable yet Perfect K

We are ready to collect all the pieces together. We need to define a special perfect distribution K which contains many keys satisfying Equation (15), meaning that $\mathsf{Ext}(K)$ is very biased towards 1. We will construct such K having a very special form.

Definition 3 Assume $\pi_1, \ldots, \pi_d : \mathcal{C} \to \mathcal{C}$ are d permutations over the ciphertext space $\mathcal{C} = \{1 \ldots S\}$. We say that $\pi_1, \ldots, \pi_d$ are *d-valid* if for every $c \in C$, and distinct $i, j \in \{1 \ldots d\}$, we have $\pi_i(c) \neq \pi_j(c)$. ♢

The reason for this terminology is the following. Given any B-valid $\pi_1, \ldots, \pi_B$, where recall that $B = |\mathcal{M}|$, we can define S valid keys $k_1, \ldots, k_S \in \mathcal{K}$ by $k_c = (\pi_1(c), \ldots, \pi_B(c))$, where the B-validity constraint precisely ensures that all the B ciphertexts inside k_c are distinct, so that k_c is a legal key in $\mathcal{K}$. Now, we denote by $K_{(\pi_1,\ldots,\pi_B)}$ the uniform distribution over these S keys $k_1, \ldots, k_S$.

Lemma 4 *If* $\pi_1, \ldots, \pi_B$ *are* B*-valid permutations, then* $K_{(\pi_1,\ldots,\pi_B)}$ *is a perfect encryption distribution.*

Proof: For any message m, $\mathsf{Enc}(K_{(\pi_1,\ldots,\pi_B)}, m)$ is equivalent to outputting $\pi_m(U_{\mathcal{C}})$, where $U_{\mathcal{C}}$ is the uniform distribution over $\mathcal{C}$. Since each π_m is a permutation over $\mathcal{C}$, this is equivalent to $U_{\mathcal{C}}$. Thus, encryption of every message m yields a truly random ciphertext $c \in \mathcal{C}$, which means that $K_{(\pi_1,\ldots,\pi_B)}$ is perfect. □

Choosing Good Permutations. We will construct our perfect distribution $K = K_{(\pi_1,\ldots,\pi_B)}$ by carefully choosing a B-valid family $(\pi_1, \ldots, \pi_B)$ such that $\mathsf{Ext}(K)$ is very biased towards 1. We start by choosing π_1 to be the identity permutation $\pi_1(c) = c$ (for all c), and proceed by defining $\pi_2 \ldots \pi_B$ iteratively. After defining each π_d, we will maintain the following invariants which clearly hold for the base case $d = 1$:

(i) $\pi_1, \ldots, \pi_d$ are d-valid.

(ii) There exists a large set T_d of "good" ciphertexts (where, initially, $T_1 = \mathcal{C}$) of size $q_d > S - d^2$, which satisfies the following equation for all $c \in T_d$ and $1 < m \leq d$:[11]

$$y_{(m,c)} - y_{(m,\pi_m(c))} \leq 0 \tag{16}$$

Now, assuming inductively that we have defined $\pi_1 = id, \pi_2, \ldots, \pi_d$ which satisfy properties (i) and (ii) above, we will construct π_{d+1} still satisfying (i) and (ii).

This inductive step is somewhat technical, and we will come back to it in the next subsection. But first, assuming it is true, we show that we can easily finish our proof. Indeed, we apply the induction for $B-1$ iterations and get B permutations $\pi_1, \ldots, \pi_B$ satisfying properties (i) and (ii) above. Then, property (i) and Lemma 4 imply that $K_{(\pi_1,\ldots,\pi_B)}$ is a perfect encryption distribution. On the other hand, property (ii) and the definition of $k_c = \{c, \pi_2(c), \ldots, \pi_B(c)\}$ imply that any key $k_c \in T_B$ satisfies Equation (15). Thus, by Theorem 2 we get that $\mathsf{Ext}(k_c) = 1$ for every $c \in T_B$. Since, $|T_B| > S - B^2$, we get that at most B^2 out of S keys k_c extract to 0. Thus, since $K_{(\pi_1,\ldots,\pi_B)}$ is uniform over its S keys, we get

$$\Pr[\mathsf{Ext}(K_{(\pi_1,\ldots,\pi_B)}) = 0] \leq \frac{B^2}{S}$$

which shows Equation (9) and completes our proof (modulo the inductive step).

4.5 Completing the Inductive Step

We start by recalling some basic facts about bipartite graphs, which we will need soon. A (balanced) bipartite graph G is given by two vertex sets L and R of cardinality S and an edge set $E = E(G) \subseteq L \times R$. A *matching* P in G is a subset of node-disjoint edges of E. P is *perfect* if $|P| = S$. In this case every $i \in L$ is matched to a unique $j \in R$ and vice versa.

We say that a subset $L' \subseteq L$ is *matchable* (in G) if there exists a matching P containing L' as the set of its endpoints in L. In this case we also say that L' is *matchable with* R', where $R' \subseteq R$ is the set of P's endpoints in R. (Put differently, L' is matchable with R' precisely when the subgraph induced by L' and R' contains a perfect matching.) The famous Hall's marriage theorem gives a necessary and sufficient condition for L' to be matchable.

Hall's Marriage Theorem. *L' is matchable if and only if every subset A of L' contains at least $|A|$ neighbors in R. Notationally, if $\mathcal{N}(A)$ denotes the set of elements in R containing an edge to A, then L' is matchable iff $|\mathcal{N}(A)| \geq |A|$, for all $A \subseteq L'$.*

We will only use the following two special cases of Hall's theorem.

Corollary 5 *Assume every vertex $v \in L \cup R$ has degree at least $S - d$: $\mathsf{deg}_G(v) \geq S - d$. Then, for any $L' \subset L$ and $R' \subset R$ of cardinality $2d$, we have that L' is matchable with R'.*

Proof: Let us consider the $2d \times 2d$ bipartite subgraph G' of G induced by L' and R'. Clearly, that every vertex $v \in L' \cup R'$ has degree at least d in G', since each such v is not connected to at most d opposite

[11] To get some intuition, we will see shortly that "good" ciphertexts c will lead to keys k_c satisfying Equation (15), so that $\mathsf{Ext}(k_c) = 1$ by Theorem 2.

vertices in the entire G, let alone G'. We claim that L' meets the conditions of the Hall's theorem in G'. Consider any non-empty $A \subseteq L'$. If $|A| \leq d$, then any vertex v in A had $\mathsf{deg}_{G'}(v) \geq d \geq |A|$ neighbors, so $|\mathcal{N}(A) \geq |A|$. If $d < |A| \leq 2d$, let us assume for the sake of contradiction that $|\mathcal{N}(A)| < |A|$. Consider now any vertex $v \in R \backslash \mathcal{N}(A)$. Such v exists as $|\mathcal{N}(A)| < |A| \leq 2d = |R'|$. Then no element in A can be connected to v, since $v \notin \mathcal{N}(A)$. Thus, the degree of v can be at most $2d - |A| < d$, which is a contradiction. □

Corollary 6 *Assume L contains a subset $L' = \{c_1, \ldots, c_\ell\}$ such that $\mathsf{deg}_G(c_i) \geq i$, for $1 \leq i \leq \ell$. Then L' is matchable in G. In particular, G contains a matching of size at least ℓ.*

Proof: We show that L' satisfies the conditions of Hall's theorem. Assume $A = \{c_{i_1}, \ldots, c_{i_a}\}$, where $1 \leq i_1 < i_2 < \ldots < i_a \leq \ell$. Notice, this means $i_j \geq j$ for all j. Then the neighbors of A at least include the neighbors of i_a, so that $|\mathcal{N}(A)| \geq \mathsf{deg}_G(c_{i_a}) \geq i_a \geq a = |A|$. □

Mapping Induction into a Matching Problem. We now return to our induction. Recall, we are given permutations $\pi_1 = id, \pi_2, \ldots, \pi_d$ satisfying properties (i) and (ii), and need to construct π_{d+1} also satisfying (i) and (ii). We translate this task into some graph matching problem, starting with the property (i) first.

For every $c \in C$, we define the "forbidden" set $F_c = \{c, \pi_2(c), \ldots, \pi_d(c)\}$. Then, the $(d+1)$-validity constraint (i) is equivalent to requiring $\pi_{d+1}(c) \notin F_c$ for all $c \in \mathcal{C}$. Next we define a bipartite "constraint graph" G on two copies L and R of $\mathcal{C}$ containing all the non-forbidden edges: $(c, c') \in E(G)$ if and only if $c' \notin F_c$. We observe two facts about G. First,

Claim 1 *Every vertex $v \in L \cup R$ has degree at least $S - d$: $\mathsf{deg}_G(v) \geq S - d$. In particular, by Corollary 5 every two $2d$-element subsets of L and R are matchable with each other in G.*

Proof: This claim is obvious for $v \in L$ as $|F_v| = c$. It is also true for $v \in R$, since any value $v \in R$ is forbidden by exactly d (necessarily distinct) elements $v, \pi_2^{-1}(v), \ldots, \pi_d^{-1}(v)$. □

Second, any perfect matching P of G uniquely defines a permutation π on S elements such that $P = \{(c, \pi(c))\}_{c \in L}$. Since, by definition, $\pi(c) \notin F_c$, it is clear that this π will always satisfy constraint (i). Thus, we only need to find a perfect matching P for G which will define a permutation π_{d+1} satisfying condition (ii).

Notice, our inductive assumption implies the existence of a subset T_d of L (recall, L is just a copy of $\mathcal{C}$) of size $q_d > S - d^2$ such that Equation (16) is satisfied for all $c \in T_d$ and $1 < m \leq d$. Irrespective of the permutation π_{d+1} we will construct later, we will restrict T_{d+1} to be a *subset* of T_d. This means that Equation (16) will already hold for all $c \in T_{d+1}$ and $1 < m \leq d$. Thus, we will only need to ensure this equation for $m = d + 1$; i.e., that for all $c \in T_{d+1}$

$$y_{(d+1,c)} - y_{(d+1,\pi_{d+1}(c))} \leq 0 \tag{17}$$

This constraint motivates us to define a subgraph G' of our constraint graph G as follows. As edge $(c, c') \in E(G')$ if and only if $(c, c') \in E(G)$ (i.e., $c' \notin F_c$) and $y_{(d+1,c)} - y_{(d+1,c')} \leq 0$. In other words, we only leave edges (c, c') which would satisfy Equation (17) if we were to define $\pi_{d+1}(c) = c'$. The key property of G' turns out to be

Lemma 7 *G' contains a matching P' of size at least $S - d$.*

Proof: We will use Corollary 6. Let us sort the vertices $v_1 \dots v_S$ of L and R in the order of non-decreasing $y_{(d+1,\cdot)}$ values; i.e.

$$y_{(d+1,v_1)} \leq y_{(d+1,v_2)} \leq \dots \leq y_{(d+1,v_S)}$$

Then, the edge (v_i, v_j) satisfies $y_{(d+1,v_i)} - y_{(d+1,v_j)} \leq 0$ whenever $i \leq j$. Thus, such (v_i, v_j) belongs to G' if and only if it also belongs to the larger constraint graph G; i.e., $v_j \notin F_{v_i}$. But since each v_i has at most d forbidden edges in G, and $|\{j \mid j \geq i\}| = S - i + 1$, we have that $\mathsf{deg}_{G'}(v_i) \geq (S - i + 1) - d$. In particular, $\mathsf{deg}_{G'}(v_{S-d}) \geq 1, \dots, \mathsf{deg}_{G'}(v_1) \geq S - d$. By Corollary 6, $\{v_{S-d}, \dots, v_1\}$ is matchable in G', completing the proof. □

Finishing the Proof. Finally, we can collect all the pieces together and define a good matching P in G (corresponding to π_{d+1}). With an eye on satisfying property (ii), we start with a large (but not yet perfect) matching P' of G' of size at least $S - d$, guaranteed by Lemma 7. Ideally, we would like to extend P' to some perfect matching in the full graph G, by somehow matching the vertices currently unmatched by P'. Unfortunately, we do not know how to argue that such extension is possible, since there are at most d vertices unmatched, and we can only match arbitrary sets of size at least $2d$ by Claim 1. So we simply take an arbitrary sub-matching P'' of P' of size $S - 2d$, just throwing away any $|P'| - (S - 2d)$ edges of P'.

Notice, P'' is also a matching of G which has exactly $2d$ unmatched vertices on both sides. By Claim 1, we know that we can always match these missing vertices, and get a perfect matching P of the entire G. We finally claim that this perfect matching P defines a permutation π_{d+1} on $\mathcal{C}$ satisfying properties (i) and (ii).

Property (i) is immediate since P is a perfect matching of G. As for property (ii), let L' denote the $S - 2d$ endpoints of P'' in L. Now, every $c \in L'$ satisfies Equation (17), since this is how the graph G' was defined and $(c, \pi_{d+1}(c)) \in P'' \subseteq E(G')$. Thus, we can inductively define $T_{d+1} = T_d \cap L'$ and have T_{d+1} satisfy property (ii). We only need to argue that T_{d+1} is large enough, but this is easy. Since L' misses only $2d$ ciphertexts, we get by induction that

$$|T_{d+1}| \geq |T_d| - 2d > S - d^2 - 2d > S - (d+1)^2$$

completing the induction and the whole proof.

5 Conclusions and Open Problems

We study the question if true randomness is inherent for achieving privacy, and show a largely positive answer for the case of information-theoretic private-key encryption. Needless to say, many exciting questions are still open. First, we ignored the issue of efficiency. Could it be that a source allowing for *efficient* encryption always allows for *efficient* extraction? On another front, although private-key encryption is the simplest and most fundamental "privacy application", our work leaves open the question if other "privacy applications" inherently require true randomness as well. Perhaps 2-out-2 secret sharing (which is strictly implied by private-key encryption [DPP06]) is a good candidate where our information-theoretic techniques might be applicable. More challengingly, what about *computationally* secure primitives, such as public-key encryption, commitment and zero-knowledge, which will require a completely new technique?

Finally, we hope that our result and techniques will stimulate further interest in understanding the extent to which cryptographic primitives can be based on imperfect randomness.

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.

[DPP06] Yevgeniy Dodis, Krzysztof Pietrzak and Bartosz Przydatek. Separating Sources for Encryption and Secret-Sharing. In *Proc. Theory of Cryptography Conference (TCC)*, pages 601–616, 2006.

[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.

[Sha49] C. Shannon. Communication Theory of Secrecy systems. In *Bell Systems Technical J.*, 28:656–715, 1949.

[Sho05] Victor Shoup. A Computational Introduction to Number Theory and Algebra. *Cambridge University Press*, 2005.

[Str80] Gilbert Strang. Linear Algebra and Its Applications. *Academic Press*, London, 1980.

[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.