Entropic Security and the Encryption of High Entropy Messages*

Yevgeniy Dodis[1] and Adam Smith[2]

[1] New York University, dodis@cs.nyu.edu
[2] Weizmann Institute of Science, adam.smith@weizmann.ac.il

Abstract. We study *entropic security*, an information-theoretic notion of security introduced by Russell and Wang [24] in the context of encryption and by Canetti et al. [5, 6] in the context of hash functions. Informally, a probabilitic map $Y = \mathcal{E}(X)$ (e.g., an encryption sheme or a hash function) is entropically secure if knowledge of Y does not help predicting any predicate of X, whenever X has high min-entropy from the adversary's point of view. On one hand, we strengthen the formulation of [5, 6, 24] and show that entropic security in fact implies that Y does not help predicting any *function* of X (as opposed to a predicate), bringing this notion closer to the conventioonal notion of semantic security [10]. On the other hand, we also show that entropic security is equivalent to *indistinguishability* on pairs of input distributions of sufficiently high entropy, which is in turn related to *randomness extraction* from non-uniform distributions [21].

We then use the equivalence above, and the connection to randomness extraction, to prove several new results on entropically-secure encryption. First, we give two general frameworks for constructing entropically secure encryption schemes: one based on expander graphs and the other on XOR-universal hash functions. These schemes generalize the schemes of Russell and Wang, yielding simpler constructions and proofs, as well as improved parameters. To encrypt an n-bit message of min-entropy t while allowing at most ϵ-advantage to the adversary, our best schemes use a shared secret key of length $k = n - t + 2\log\left(\frac{1}{\epsilon}\right)$. Second, we obtain lower bounds on the key length k for entropic security and indistinguishability. In particular, we show near tightness of our constructions: $k > n - t$. For a large class of schemes — including all the schemes we study — the bound can be strengthened to $k \geq n - t + \log\left(\frac{1}{\epsilon}\right) - O(1)$.

1 Introduction

If X and Y are random variables, the statement "Y leaks no information about X" is normally formalized by requiring that X and Y be almost statistically independent. Equivalently, one can require that the Shannon mutual information $\mathbf{I}(X, Y)$ be very small. In this work we study situations where information leakage

* A more complete version of this paper may be found on IACR Cryptology ePrint Archive, report 2004/219, at http://eprint.iacr.org/2004/219/ [9].

is unavoidable — that is, $\mathbf{I}(X, Y)$ is large — yet we still want a guarantee that no *useful* information about X is leaked by Y, even to a computationally unbounded adversary.

Consider an alternative notion of security, inspired by semantic security of encryptions [10]. We say Y *hides all functions of* X if for every function f, it is nearly as hard to predict $f(X)$ given Y as it is without Y, regardless of the adversary's computing power. If $Y = \mathcal{E}(X)$ for some probabilistic map $\mathcal{E}()$ (for example, an encryption scheme), then we say the map $\mathcal{E}$ is *entropically secure* if $\mathcal{E}(X)$ hides all functions of X whenever X has sufficiently high entropy.

A seemingly weaker variant of this definition has produced surprising results in at least two contexts so far: Canetti, Micciancio and Reingold [5, 6] constructed hash functions whose outputs leak no partial information about the input. Russell and Wang [24] showed how one can construct entropically-secure symmetric encryption schemes with keys much shorter than the length of the input, thus circumventing Shannon's famous lower bound on key length.

Our contributions can be divided into two areas.

- We elucidate the notion of entropic security. Our results apply to all entropically-secure primitives, including encryption schemes and hash functions. We provide two new variants on entropic security, one closer in spirit to semantic security of encryptions [10], and the other along the lines of indistinguishability of encryptions. The proofs that these various notions are equivalent give us new tools for handling entropic security and highlight a relationship with "randomness extraction" from non-uniform distributions.
- We use the connection to randomness extraction to obtain new constructions and lower bounds for encryption of high-entropy messages with a short key. First, we give two general frameworks for constructing entropically secure encryption schemes: one based on expander graphs and the other on XOR-universal hash functions. These schemes generalize the schemes of Russell and Wang, yielding simpler constructions and proofs, as well as improved parameters. Second, we obtain nearly tight lower bounds on the key length k for entropic security and indistinguishability.

1.1 Background

Although our general results apply to all entropically-secure primitives, we present entropic security (and our results) in the context of symmetric-key one-time encryption. Alice and Bob share a secret key K and Alice wants to securely send some message X to Bob over a public channel. X is assumed to come from some a-priori distribution on $\{0, 1\}^n$ (e.g., uniform), and the goal is to compute a ciphertext Y which: (a) allows Bob to extract X from Y using K; (b) reveals "no information" about X to the adversary Eve beyond what she already knew. Below, we write $Y \leftarrow \mathcal{E}(X, K)$ and $X = \mathcal{D}(Y, K)$.

Perfect and Computational Security. The first formalization of this problem came in a fundamental work of Shannon [25], who defined "no information" by requiring that X and Y be independent as random variables: using information theoretic notation, $\mathbf{I}(X;Y) = 0$, where $\mathbf{I}$ is the mutual information. He showed a lower bound on key length for his definition: encrypting messages of length n requires at least n bits of shared key (more formally, the Shannon entropy of the key must be at least that of message distribution: $\mathbf{H}_{sh}(K) \geq \mathbf{H}_{sh}(X)$). This bound is tight when the message is chosen uniformly from all strings of a fixed length n, since one can use a one-time pad. The bound was extended to the interactive setting by Maurer [19].

Goldwasser and Micali [10] relaxed the notion of perfect security to the *computational* setting: namely, any *efficient* Eve can extract only negligible "information" about X from Y. They had to properly redefine the notion of "information", since mutual information or conditional probabilities do not make much sense in a computationally-bounded world. They suggested two now classical definitions. Consider the following, equivalent version of Shannon's definition: for any two messages x_0 and x_1, the two corresponding distributions on ciphertexts should be identical, that is $\mathcal{E}(x_0) = \mathcal{E}(x_1)$ (as distributions). The first definition of Goldwasser and Micali, called *computational indistinguishability of encryptions*, generalizes this version of perfect security: they require that no efficient (polynomial-time adversary) can distinguish the encryptions of x_0 and x_1 with advantage more than ϵ over random guessing, where ϵ is some negligible quantity. Their second notion is called *semantic security*: for *any* distribution on messages X and any function $f()$, the adversary can predict $f(X)$ given $\mathcal{E}(X)$ with probability only negligibly better than she could without seeing $\mathcal{E}(X)$. The first definition is easier to work with, but the second definition seems to capture a stronger, more intuitive notion of security: for example, indistinguishability is the special case of semantic security when the message distribution X is restricted to uniform distributions over two points $\{x_0, x_1\}$. In fact, Goldwasser and Micali showed that the two definitions are equivalent. Thus, distributions with entropy 1 are in some sense the hardest to deal with for semantic security.

Statistical Security? A natural intermediate notion of security between perfect and computational security would be some kind of *statistical security*: Eve is computationally unbounded, as in the perfect setting, but can potentially recover some negligible amount of "information", as in the computational setting. At first glance, it seems like there is no gain in this notion, no matter how we interpret "information". For example, following Shannon's approach we could require that $\mathbf{I}(X;Y) \leq \epsilon$ instead of being 0. Unfortunately, Shannon's proof still implies that $\mathbf{H}_{sh}(K) \geq \mathbf{H}_{sh}(X) - \epsilon$. Similarly for indistinguishability: since the distribution $\mathcal{E}(x)$ should look almost the same for *any* fixed x, one can argue that $\mathbf{I}(Y;X) = \mathbf{H}_{sh}(\mathcal{E}(X)) - \mathsf{Exp}_x\left[\mathbf{H}_{sh}(\mathcal{E}(x))\right]$ still has to be negligible, and so the key must again have entropy almost $\mathbf{H}_{sh}(X)$.

In his original work Shannon envisioned applications where Eve has a lot of uncertainty about the message. To get a pessimistic bound that $\mathbf{H}_{sh}(K) \geq n$, one only has to take X to be uniformly distributed in $\{0,1\}^n$. In fact, in the perfect setting, security against the uniform distribution implies security against *any* distribution on messages. On the other hand, the notions of indistinguishability and semantic security primarily deal with entropy 1 distributions, and the straightforward extension of Shannon's bound to the statistical versions of these notions crucially uses this fact. Thus, it is natural to ask if we can meaningfully define (statistical) semantic security and/or indistinguishability for high entropy distributions (say, uniform), similar in spirit to the original work of Shannon. And if yes,

1. How do these notions relate to Shannon's (statistical) notion, $\mathbf{I}(X;Y) \leq \epsilon$? Most importantly, does the pessimistic bound on the key length still extend to these notions?
2. How do these notions relate to each other? Are semantic security and indistinguishability still equivalent when the message is guaranteed to have high entropy?

1.2 Entropic Security

Russell and Wang [24] introduced the idea of statistical security for encryption of high-entropy message spaces. They considered the first question above, though they focused on a weakened version of semantic security. Their definition, *entropic security of encryptions for predicates*, is natural: for any distribution X of min-entropy[1] at least t and any predicate $g : \{0,1\}^n \to \{0,1\}$, Eve can predict $g(X)$ using Y only negligibly better than she could without Y (here n is the message length and t is a parameter). Russell and Wang showed that Shannon's lower bound does *not* extend to this new notion: they presented two schemes beating Shannon's bound on key length, which we describe further below. Entropic security also arose earlier in work of Canetti [5] and Canetti, Micciancio and Reingold [6]. They constructed probabilistic hash functions whose output reveals no partial information about their input as long as it had sufficiently high entropy.

We discuss a stronger version of the definition of [5, 6, 24], which requires that the adversary gain no significant advantage at predicting any *function* of the input (as opposed to a predicate). One of our results is the equivalence of their notion of security to the one described here.

Definition 1 (Entropic Security). *The probabilistic map Y hides all functions of X with leakage ϵ if for every adversary $\mathcal{A}$, there exists some adversary $\mathcal{A}'$ such that for all functions f,*

$$\big|\Pr[\mathcal{A}(Y(X)) = f(X)] - \Pr[\mathcal{A}'() = f(X)]\big| \leq \epsilon.$$

The map $Y()$ is called (t,ϵ)-entropically secure if $Y()$ hides all functions of X, whenever the min-entropy of X is at least t.

[1] The *min-entropy* of a random variable A is a measure of the uncertainty of its outcome. It is the negative logarithm of the probability that one can predict A ahead of time: $\mathbf{H}_\infty(A) = -\log(\max_a \Pr(A = a))$.

One gets some insight about this definition by thinking of it as an information-theoretic reformulation of semantic security of encryptions [10], although restricted to high-entropy message spaces. Alternatively, it might be instructive to view this definition as saying that Y leaks no *a-priori* information about X. Here "a-priori" refers to the fact that the function f has to be specified *before* the pair (X, Y) is sampled. In other words, although f is arbitrary, it cannot depend on the outcome of Y. This should be contrasted with *a-posteriori* information, where first the pair (X, Y) is sampled, then the adversary is given the outcome y of Y, and can choose a function f_y which is supposedly easier to predict when given y. In this latter case it is not very hard to see that Y leaks almost no a-posteriori information about X if and only if X and Y are essentially independent, i.e. the quantity $\mathbf{I}(X; Y)$ is "low". Thus, the results of [24, 5, 6] could be interpreted by saying that leakage of no a-priori information — although for the moment restricted to predicates rather than general functions — can be achieved in situations where it is *impossible* to leak no a-posteriori information.

1.3 Contributions of This Paper

This paper carefully studies and eludicates the notion of entropic security, obtaining several new insights into this notions, as well as simplifying and improving previous results of [24, 5, 6].

A Strong Definition of Security. As we mentioned, the definition we propose (Definition 1) is seemingly stronger than previously studied formulations of entropic security [5, 6, 24], which only considered the adversary's ability to predict *predicates* instead of all possible functions of the secret input. This definition may not be quite satisfying from several points of view. First, it states only that no predicate of the input is leaked, and provides no explicit guarantees about other functions. In contrast, the original semantic security definition of Goldwasser and Micali held for all functions. Second, there is no guarantee that the new adversary $\mathcal{A}'()$ is polynomial time, even in the case where, say, $\mathcal{A}$ runs polynomial time and X is samplable in polynomial time. We show that (a) entropic security for predicates *does* imply that for arbitrary functions (see Lemma 2), and (b) in the definition of entropic security one can always set $\mathcal{A}'()$ to be $\mathcal{A}(U_n)$, where U_n is the uniform distribution on n bits.

The equivalence between predicates and functions is not trivial. Consider, for example, the special case where f is the identity function. One might hope that a good predictor for X must imply a good predictor for some bit X_i of X. However, this is false. As a counterexample, assume X is equal to U_n and Y is equal to X with probability $1/2$, and to the bit-wise complement of X otherwise. Clearly, Y reveals X with probability at least $1/2$ (which is much larger than 2^{-n}), although no physical bit of X can be predicted with probability better than its natural probability $1/2$. Of course, in this case one can predict any even parity of the bits of X with probability 1, but this shows that a more sophisticated approach is needed. As we show in Proposition 1, for this function we can choose a "Goldreich-Levin" predicate at random, that is we can use the predicate $g_r(x) = r \odot x$ where r is a random n-bit string and $\odot$ is the binary inner product $r \odot x = \sum_i r_i x_i \mod 2$. For general functions, a more complicated construction is required (Lemma 2). This general equivalence between predicting predicates and predicting arbitrary functions could be of independent interest, as it provides an information-theoretic converse to the Goldreich-Levin hardcore bit construction.

An Equivalence to Indistinguishability. We also define a new *indistinguishability* notion, analogous to indistinguishability of (computationally secure) encryptions [10]. Namely, we say that the map $Y()$ is t-indistinguishable, if for any distributions X_1 and X_2 of min-entropy at least t, the distribution $Y(X_1)$ is statistically close to $Y(X_2)$. A bit more formally, indistinguishability is stated in terms of the statistical difference $\mathbf{SD}(A, B)$ between a pair of random variables A, B. This is half the L_1 distance between the distributions, $\mathbf{SD}(A, B) = \frac{1}{2}\sum_z |\Pr[A = z] - \Pr[B = z]|$. It also has an operational meaning: given a sample from either A or B (at random), the optimal adversary's chance of correctly guessing which distribution the sample came from is exactly $\frac{1}{2} + \frac{1}{2}\mathbf{SD}(A, B)$. This distance measure satisfies the triangle inequality, and so all the distributions $Y(X)$ must actually be close to a single distribution $G = Y(U_n)$, where U_n is the uniform distribution. We arrive at the following:

Definition 2. *A randomized map $Y()$ is* (t, ϵ)-indistinguishable *if there is a random variable G such that for every distribution on messages X over $\{0,1\}^n$ with min-entropy at least t, we have*

$$\mathbf{SD}(Y(X),\ G) \leq \epsilon.$$

As we can see, the notion of entropic security seems to be well motivated, but hard to work with. On the other hand, indistinguishability seems to be a much easier definition to work with, but might be less intuitively meaningful. Our main result is that the definitions are in fact equivalent:

Theorem 1. *Let Y be a randomized map with inputs of length n. Then*

1. *(t, ϵ)-entropic security for predicates implies $(t-1, 4\epsilon)$-indistinguishability.*
2. *$(t-2, \epsilon)$-indistinguishability implies $(t, \epsilon/8)$-entropic security for **all functions** when $t \geq 2\log\left(\frac{1}{\epsilon}\right) + 1$.*

In particular, since entropic security with respect to predicates is trivially implied by entropic security for all functions, Theorem 1 states that all three notions of security discussed above are equivalent up to small changes in the parameters. Although this result is inspired by a similar looking result of Goldwasser and Micali [10] (for computational encryption), our proof is considerably different and does not seem to follow from the techniques in [10].

The equivalence not only reconciles two natural definitions, but has several nice implications. First, in Definition 1 we can always take $\mathcal{A}'()$ to be $\mathcal{A}(Y(U_n))$, where U_n is the uniform distribution on $\{0,1\}^n$. Thus, the "simulated" adversary is as efficient as the original.

Second, the equivalence provides a new application of *randomness extractors* [21] to cryptography. Recall that an extractor takes as input an arbitrary, high entropy random source and a tiny random seed, and outputs uniformly random bits. The output bits are guaranteed to be almost uniformly distributed as long as the min-entropy of the input is above some threshold t. In other words, an extractor Y is precisely a t-indistinguishable map — in the sense of Definition 2 — with G being the uniform distribution. Thus, Theorem 1 implies that an extractor for t-sources hides all a-priori information about sources of min-entropy at least $t+2$. From the constructive point of view, it also suggests that to design an appropriate entropically secure scheme for a given task, such as encryption, it is sufficient to design a "special purpose" randomness extractor. In the case of encryption the extractor should be invertible when given the seed, since the seed corresponds to the shared secret key.

Finally, and most importantly, our equivalence simplifies the design and analysis of entropically secure schemes, yielding improvements over known schemes, new lower bounds, simpler proofs, and a stronger security guarantee. We illustrate these points for the case of entropically secure encryption.

Encryption of High-Entropy Messages. As we mentioned, Russell and Wang [24] provided two constructions of entropically-secure encryption schemes which use short keys. Let ϵ denote the leakage — that is, the advantage which we allow the adversary. First, [24] give a deterministic scheme of the form $\mathcal{E}(X,K) = X \oplus p(K)$, which is secure only when X is uniformly distributed on $\{0,1\}^n$, where K has length only $k = 2\log n + 3\log\left(\frac{1}{\epsilon}\right) + O(1)$ and $p(K)$ is a random point in a δ-biased spaces [20] (where [24] used $\delta = \epsilon^{3/2}$). Thus, $p(K)$ could be viewed as a very sparse one-time pad which nevertheless hides any a-priori specified function $f(X)$. Second, for general min-entropy t, Russell and Wang gave a *randomized* scheme of the form $(\psi, \psi(X)+K) \leftarrow \mathcal{E}(X,K)$, where ψ is chosen at random from a family of 3-wise independent permutations (and the addition is defined over some appropriate space). The analysis in [24] shows that this second scheme needs key length $n-t+3\log\left(\frac{1}{\epsilon}\right)+O(1)$. While less than n for nontrivial settings of t, this key length again becomes $\Omega(n)$ when $n-t=\Omega(n)$. [24] left it open whether such dependence on $n-t$ is necessary.
We obtain the following new results:

1. Lower bounds on the key length k for entropic security and indistinguishability. In particular, we show near tightness of Russell-Wang constructions: $k > n-t$. (In fact, for a large class of schemes $k \geq n-t+\log\left(\frac{1}{\epsilon}\right)$.)
2. Two general frameworks for designing entropically secure encryption schemes: one based on expander graphs and the other on XOR-universal hash functions. These schemes generalize the schemes of Russell and Wang, yielding simpler constructions and proofs, as well as improved parameters. Namely, both constructions can yield keys of size $k = n-t+2\log\left(\frac{1}{\epsilon}\right)$.

Our Techniques. All our results critically use the equivalence between entropic security and indistinguishability.

On one hand, we use it to show that the general construction of Russell and Wang is nearly optimal: *any* entropically secure scheme must have $k > n - t$. In fact, for a special case of *public-coin* schemes, where the ciphertext contains the randomness used for encryption, we get an even stronger bound: $k \geq n - t + \log\left(\frac{1}{\epsilon}\right) - O(1)$. The latter result is proven by relating the notion of indistinguishability to that of *randomness extractors* [21]: namely, any indistinguishable public-coin scheme almost immediately yields a corresponding extractor. Using the optimal lower bounds on extractors [23], we get our stronger bound as well. In fact, if the ciphertext is statistically close to uniform (i.e., $G = U_n$ meaning that the encryption is actually a randomness extractor), we get a lower bound which *exactly* matches our upper bounds: $k \geq n-t+2\log\left(\frac{1}{\epsilon}\right)-O(1)$. The schemes in [24] and this work are all public-coin and have random ciphertexts.

On the other hand, the indistinguishability view allows us to give a general framework for constructing entropically secure encryption schemes. Specifically, assume we have a d-regular expander G on 2^n vertices V with the property that for any subset T of 2^t vertices, picking a random vertex v of T and taking a random neighbor w, we obtain an almost uniform distribution on V. Then, we almost immediately get an encryption scheme with key length $k = \log d$ which is indistinguishable for message spaces of min-entropy t. Looking at this from another perspective, the above encryption scheme corresponds to a randomness extractor which takes a source X of length n and min-entropy t, invests $\log d$ extra random bits K, and extracts n almost random bits Y (with the additional property that the source X is recoverable from Y and K). From this description, it is clear that the key length of this paradigm must be at least $n - t$ (which we show is required in any entropically secure encryption scheme). However, using optimal expanders we can (essentially) *achieve* this bound, and in several ways. First, using Ramanujan expanders [17], we get the best known construction with key length $k = n - t + 2\log\left(\frac{1}{\epsilon}\right)$. Second, using δ-biased spaces [20] (for appropriate $\delta = \delta(\epsilon, n, t)$ explained later), we get a general construction with slightly larger but still nearly optimal key length $k = n - t + 2\log n + 2\log\left(\frac{1}{\epsilon}\right)$. This last result generalizes (and slightly improves) to any value of t the special case of the uniform message distribution ($n - t = 0$) obtained by Russell and Wang [24]. Our approach also gives clearer insight as to why small-biased spaces are actually useful for entropic security.

While the deterministic constructions above are nearly optimal and quite efficient, we also observe that one can get simpler constructions by allowing the encryption scheme to be *probabilistic*. In our approach, this corresponds to having a *family* of "average case" expanders $\{G_i\}$ with the property that for any set T of size at least 2^t, picking a random graph G_i, a random v in T and taking a random neigbor w of v in G_i, we get that w is nearly uniform, *even given the graph index i*. By using any family of pairwise independent hash functions h_i (resp. permutations ψ_i) and a new variant of the leftover hash lemma [15], we get a probabilistic scheme of the form $\langle i,\ X \oplus h_i(K)\rangle$ (resp. $\langle i,\ \psi_i(X) \oplus K\rangle$) with a nearly optimal key length $k = n - t + 2\log\left(\frac{1}{\epsilon}\right)$. As a concrete example of this approach, we get the following simple construction: $\mathcal{E}(X, K; i) = (i, X + i \cdot K)$, where the local randomness i is a random element in $GF(2^n)$, $K \in \{0,1\}^k$ is interpreted as belonging to $GF(2^k) \subseteq GF(2^n)$, and addition and multiplication are done in $GF(2^n)$.

Once again, the above result (with permutations ψ_i) improves and simplifies the intuition behind the second scheme of Russell and Wang [24]. Indeed, the latter work had to assume that the ψ_i's come from a family of 3-wise independent permutations — which are more compicated and less efficient than 2-wise independent permutations (or functions) — and presented a significantly more involved analysis of their scheme.

1.4 A Caveat: Composing Entropically-Secure Constructions

A desirable property of definitions of security of cryptographic primitives is *composability*: once some protocol or algorithm has been proven secure, you would like to be able to use it as a building block in other protocols with your eyes closed—without having to worry about effects that violate the intuitive notion of security, but which are not covered by the original definition.

Composability is difficult to guarantee, since it is not clear how to translate it into a mathematical property. There are various formalizations of composability, most notably "Reactive Systems" [22], "Universal Composability" [7] and several frameworks based on logic algebras for automated reasoning (see [14] and the references therein). Finding protocols that are provably secure in these general frameworks is difficult, and sometimes provably impossible. A more common approach is to prove that a particular definition remains intact under a few straightforward types of composition, say by proving that it is still secure to encrypt the same message many times over.

The main weakness of entropic security, as defined above, is that it does not ensure composability, even in this straightforward sense. If $Y()$ and $Y'()$ are independent versions of the same entropically-secure mapping, then the map which outputs the pair $Y(X), Y'(X)$ may be insecure to the point of revealing X completely. In the case of encryption, this means that encrypting the same message twice may be problematic. (Given the first value $Y(X)$, the entropy of X may be too low for the security guarantee of $Y'()$ to hold).

For example, suppose that $Y(x)$ consists of the pair $\langle M, Mx\rangle$, where M is a random $\frac{3n}{4} \times n$ binary matrix M and $x \in \{0,1\}^n$. We will see later that $Y()$ is entropically secure whenever the entropy of X is close to n. However, the pair $Y(x), Y'(x)$ provides a set of $\frac{3n}{2}$ randomly chosen linear constraints on x. With high probability, these determine x completely, and so the pair $Y(), Y'()$ is insecure under any reasonable definition.

Given these issues, entropically-secure primitives must be used with care: one must ensure that the inputs truly have enough entropy for the security guarantee to hold. Requiring entropy is natural in many situations (e.g. when the input is a password), but the issue of composability nonetheless raises a number of interesting open questions for future research.

The generality and intuitive appeal of entropic security, as well as the variety of contexts in which it has arisen, make it an important concept to understand. We hope that the present work provides a major step in this direction.

2 Entropic Security and Indistinguishability

In this section we sketch the proof of Theorem 1, that is of the equivalence between entropic security for functions/predicates and indistinguishability.

First, some notation. Fix a distribution X on $\{0,1\}^n$. For a function $f : \{0,1\}^n \to \{0,1\}^*$, let $\mathsf{pred}_{f,X}$ be the maximum probability of any particular outcome, that is the maximum probability of predicting $f(X)$ without having any information about X: $\mathsf{pred}_{f,X} \stackrel{def}{=} \max_z \Pr[f(X) = z]$.
(When X is clear from the context, we may simply write pred_f.) We may rephrase entropic security as follows: for every function f and adversary $\mathcal{A}$, the probability of $\mathcal{A}$ predicting $f(X)$ given $Y(X)$ is at most $\mathsf{pred}_f + \epsilon$:

$$\Pr[\mathcal{A}(Y(X)) = f(X)] \leq \mathsf{pred}_{f,X} + \epsilon.$$

2.1 From Entropic Security to Indistinguishability

The first statement of Theorem 1 is the easier of the two to prove, and we give the intuition here: given two distributions X_0, X_1, we can define a predicate $g(x)$ which captures the question "is x more likely to have come from X_0 or X_1?" If X is a equal mixture of X_0 and X_1, then the adversary which makes the maximum likelihood guess at $g(X)$ given $Y(X)$ will have success probability $\frac{1}{2} + \frac{1}{2}\mathbf{SD}\left(Y(X_0), Y(X_1)\right)$. On the other hand, with no access to $Y(X)$, the adversary can succeed with probability at most $\mathsf{pred}_P = \frac{1}{2}$. Entropic security implies that the advantage over random guessing, and hence the statistical distance, must be small. The formal proof is more involved, and is given below.

Proof. It is sufficient to prove indistinguishability for all distributions which are uniform on some set of 2^{t-1} points. To see why, recall that any distribution of min-entropy at least $t-1$ can be written as a convex combination of such *flat* distributions. If $X_0 = \sum \lambda_{0,i} X_{0,i}$ and $X_1 = \sum_j \lambda_{1,j} X_{1,j}$, where the $X_{0,i}$ and $X_{1,j}$ are all flat distributions, then the statistical distance $\mathbf{SD}\left(Y(X_0), Y(X_1)\right)$ is bounded above by $\sum_{i,j} \lambda_{0,i}\lambda_{1,j}\mathbf{SD}\left(Y(X_{0,i}), Y(X_{1,j})\right)$ (by the triangle inequality). If each of the pairs $Y(X_{0,i}), Y(X_{1,j})$ has distance at most ϵ, then the entire sum will be bounded by ϵ.

Now let X_0, X_1 be any two flat distributions over *disjoint* sets of 2^{t-1} points each (we will deal with non-disjoint sets below), and let X be an equal mixture of the two. That is, to sample from X, flip a fair coin B, and sample from X_B. Take g to be any predicate which is 0 for any sample from X_0 and 1 for any sample from X_1. A good predictor for g will be the adversary $\mathcal{A}$ who, given a string y as input, guesses as follows:

$$\mathcal{A}(y) = \begin{cases} 0 \text{ if } y \text{ is more likely under the distribution } Y(X_0) \text{ than under } Y(X_1) \\ 1 \text{ otherwise} \end{cases}$$

By the definition of statistical difference, this adversary guesses the predicate with probability exactly:

$$\Pr\left[\mathcal{A}(Y(X)) = B = g(X)\right] = \tfrac{1}{2} + \tfrac{1}{2}\mathbf{SD}\left(Y(X_0), Y(X_1)\right). \tag{1}$$

We can now apply the assumption that $Y()$ is (t,ϵ)-entropically secure to bound $\mathbf{SD}\left(Y(X_0), Y(X_1)\right)$. First, for any random variable G over $\{0,1\}$ which is independent of X, the probability that $G = g(X)$ is exactly $\frac{1}{2}$. The distribution X has min-entropy t by construction, and so by entropic security the probability that $\mathcal{A}(y)$ can guess $g(X)$ is bounded:

$$\Pr[\mathcal{A}(Y(X)) = g(X)] \leq \max_G \left\{\Pr[G = g(X)]\right\} + \epsilon = \tfrac{1}{2} + \epsilon. \tag{2}$$

Combining the last two equations, the statistical difference $\mathbf{SD}\left(Y(X_0), Y(X_1)\right)$ is at most 2ϵ. This takes care of the case where X_0 and X_1 have disjoint supports.

To get the general indistinguishability condition, fix any $\tilde{X}_0$ as above (flat on 2^{t-1} points). For any other flat distribution $\tilde{X}_1$, there is some third flat distribution X' which is disjoint from both $\tilde{X}_0$ and $\tilde{X}_1$. By the previous reasoning, both $\mathbf{SD}\left(Y(\tilde{X}_0), Y(X')\right)$ and $\mathbf{SD}\left(Y(X'), Y(\tilde{X}_1)\right)$ are less than 2ϵ. By the triangle inequality $\mathbf{SD}\left(Y(X_0), Y(X_1)\right) \leq 4\epsilon$. A more careful proof avoids the triangle inequality and gives distance 2ϵ even when the supports of X_0, X_1 overlap. $\square$

2.2 From Indistinguishability to Entropic Security

Proving that indistinguishability implies entropic security is considerably more delicate. We begin with an overview of the main ideas and notation.

The Case of Balanced Predicates. We say a function f is *balanced* (w.r.t. X) if it takes on all its possible values with equal probability, i.e. there are $\frac{1}{\mathsf{pred}_f}$ possible values and each occurs with probability pred_f. The reductions we consider are much easier for balanced functions. In fact, we start with balanced *predicates.*

Namely, suppose that $g()$ is a balanced predicate for distribution X, that is $\Pr[g(X) = 0] = \Pr[g(X) = 1] = \frac{1}{2}$, and that that $\mathcal{A}$ is an adversary contradicting entropic security for min-entropy $t = \mathbf{H}_\infty(X)$, that is $\Pr[\mathcal{A}(Y(X)) = g(X)] = \frac{1}{2} + \epsilon$. For $b \in \{0,1\}$, let X_b be the distribution of X conditioned on $g(X) = b$. The adversary's advantage over random guessing in distinguishing $Y(X_0)$ from $Y(X_1)$ is ϵ. However, that same advantage is also a lower bound for the statistical difference. We get:

$$\begin{aligned}\tfrac{1}{2} + \epsilon &= \Pr[\mathcal{A}(Y(X)) = g(X)] \\ &= \Pr[b \leftarrow \{0,1\} :\ \mathcal{A}(Y(X_b)) = b] \leq \tfrac{1}{2} + \tfrac{1}{2}\mathbf{SD}\left(Y(X_0), Y(X_1)\right),\end{aligned}$$

and so the distance between $Y(X_0)$ and $Y(X_1)$ is at least $\epsilon/2$. To see that this contradicts indistinguishability, note that since $g(X)$ is balanced, we obtain X_0 and X_1 by conditioning on events of probability at least $\frac{1}{2}$. Probabilities are at most doubled, and so the min-entropies of both X_0 and X_1 are at most $\mathbf{H}_\infty(X) - 1$.

Balancing Predicates. If the predicate $g()$ is not balanced on X, then the previous strategy yields a poor reduction. For example, $\Pr[g(X)=0]$ may be very small (potentially as small as ϵ). The probabilities in the distribution X_0 would then be a factor of $1/\epsilon$ bigger than their original values, leading to a loss of min-entropy of $\log(1/\epsilon)$. This argument therefore proves a weak version of Theorem 1: (t,ϵ) indistinguishability implies $(t+\log\left(\frac{1}{\epsilon}\right), 2\epsilon)$ entropic security for *predicates*.

This entropy loss is not necessary. We give a better reduction in Lemma 1 below. The idea is that to change the predicate $g()$ into a balanced predicate by flipping the value of the predicate on points on which the original adversary $\mathcal{A}$ performed poorly. By greedily choosing a set of points in $g^{-1}(0)$ of the right size, we show that there exists a balanced predicate $g'()$ on which the same adversary as before has advantage at least $\epsilon/2$, if the adversary had advantage ϵ for the original predicate.

Lemma 1. *$(t-2, 2\epsilon)$-indistinguishability implies (t,ϵ)-entropic security for **predicates** for $t \geq 2$.*

Proof. Suppose that the scheme is not (t,ϵ)-entropically secure. That is, there is a message distribution X with min-entropy at least t, a predicate g and an adversary $\mathcal{A}$ such that

$$\Pr[\mathcal{A}(Y(X)) = g(X)] > \epsilon + \max_{i=0,1}\{\Pr[g(X)=i]\} \tag{3}$$

We wish to choose two distributions of min-entropy $t-2$ and use the adversary to distinguish them, thus contradicting indistinguishability. It's tempting to choose the sets $g^{-1}(0)$ and $g^{-1}(1)$, since we know the adversary can predict g reasonably well. That attempt fails because one of the pre-images $g^{-1}(0), g^{-1}(1)$ might be quite small, leading to distributions of low min-entropy. Instead, we partition the support of X into sets of (almost) equal measure, making sure that the smaller of $g^{-1}(0)$ and $g^{-1}(1)$ is entirely contained in one partition.

Now let:

$$\begin{aligned} p &= \Pr[h(X)=1] \\ q_0 &= \Pr[\mathcal{A}(Y(X))=1 | g(X)=0] \\ q_1 &= \Pr[\mathcal{A}(Y(X))=1 | g(X)=1] \end{aligned}$$

Suppose without loss of generality that $p \geq 1/2$, i.e. that $g(X)=1$ is more likely than, or as likely as, $g(X)=0$ (if $p<1/2$, we can just reverse the roles of 0 and 1). The violation of entropic security (Eq. 3) can be re-written:

$$pq_1 + (1-p)(1-q_0) > p + \epsilon$$

In particular, $p - pq_1 > 0$ so we get:

$$(1-p)(q_1 - q_0) > \epsilon \tag{4}$$

Now we wish to choose two distributions A, B, each of min-entropy $t-2$. For now, fix any set $\mathcal{S} \subseteq g^{-1}(1)$, where $g^{-1}(1) = \{x \in \{0,1\}^n | g(x)=1\}$. We make the choice of § more specific below. Let $A_{\mathcal{S}}$ be the conditional distribution of X conditioned on $X \in \mathcal{S}$, and let $B_{\mathcal{S}}$ be distributed as X conditioned on $X \in \{0,1\}^n \setminus \mathcal{S}$. That is, $A_{\mathcal{S}}$ and $B_{\mathcal{S}}$ have disjoint supports and the support of $B_{\mathcal{S}}$ covers $g^{-1}(0)$ entirely.

The first property we will need from $\mathcal{S}$ is that it split the mass of X somewhat evenly. If the probability mass p' of $\mathcal{S}$ under X was exactly $1/2$, then the min-entropies of $A_{\mathcal{S}}$ and $B_{\mathcal{S}}$ would both be exactly $t-1$. Depending on the distribution X, it may not be possible to have such an even split. Nonetheless, we can certainly get $\frac{1}{2} \leq p' < \frac{1}{2} + 2^{-t}$, simply by adding points one at a time to § until it gets just below $1/2$. The order in which we add the points is not important. For $t > 2$ (which is a hypothesis of this proof), we get $\frac{1}{2} \geq p' \geq \frac{3}{4}$. Hence, we can choose $\mathcal{S}$ so that the min-entropies of $A_{\mathcal{S}}$ and $B_{\mathcal{S}}$ are both at least $t-2$.

We will also need that $\mathcal{S}$ have other properties. For every point x in the support of X, we define $q_x = \Pr[\mathcal{A}(Y(x)) = 1]$. The average over $x \leftarrow X$, restricted to $g^{-1}(1)$, of q_x is exactly q_1, that is

$$\mathsf{Exp}_{x \leftarrow X}\left[q_x\right] = q_1$$

If we now the choose the set $\mathcal{S}$ greedily, always adding points which maximize q_x, we are guaranteed that the average over X, conditioned on $X \in \mathcal{S}$, is at least q_1. That is, there exists a choice of S with mass $p' \in [\frac{1}{2}, \frac{3}{4}]$ such that

$$\Pr[\mathcal{A}(Y(A_{\mathcal{S}})) = 1] = \mathsf{Exp}_{x \leftarrow A_{\mathcal{S}}}\left[q_x\right] \geq q_1.$$

We can also now compute the probability that $\mathcal{A}(Y(B_{\mathcal{S}}))$ is 1:

$$\Pr[\mathcal{A}(Y(B_{\mathcal{S}})) = 1] = \frac{1-p}{1-p'}q_0 + \frac{p-p'}{1-p'}\Pr[\mathcal{A}(Y(X)) = 1 | X \notin \mathcal{S} \text{ and } g(X) = 0]$$

Now $\Pr[\mathcal{A}(Y(X)) = 1 | X \notin \mathcal{S} \text{ and } g(X) = 0]$ is at most q_1 (since by the greedy construction of $\mathcal{S}$, this is the average over elements in $g^{-1}(1)$ with the lowest values of q_x). Using $\mathcal{A}$ as a distinguisher for the distributions $Y(A_{\mathcal{S}})$ and $Y(B_{\mathcal{S}})$, we get:

$$\begin{aligned}\Big|\Pr\big[\mathcal{A}(Y(A_{\mathcal{S}})) = 1\big] - \Pr\big[\mathcal{A}(Y(B_{\mathcal{S}})) = 1\big]\Big| &\geq q_1 - \frac{1-p}{1-p'} \cdot q_0 - \frac{p-p'}{1-p'} \cdot q_1 \\ &= \frac{1-p}{1-p'} \cdot (q_1 - q_0)\end{aligned}$$

Since entropic security is violated (Eq. 4), we have $(1-p)(q_1-q_0)/(1-p') > \epsilon/(1-p')$. By construction, we have $p' > \frac{1}{2}$ so the advantage of the predictor is at least 2ϵ, that is:

$$\mathbf{SD}\left(Y(A_{\mathcal{S}}), Y(B_{\mathcal{S}})\right) \geq \Big|\Pr\big[\mathcal{A}(Y(A_{\mathcal{S}})) = 1\big] - \Pr\big[\mathcal{A}(Y(B_{\mathcal{S}})) = 1\big]\Big| \geq 2\epsilon$$

Since A and B each have min-entropy at least $t-2$, this contradicts $(t-2, 2\epsilon)$-indistinguishability, completing the proof. $\square$

From Predicates to Arbitary Functions. In order to complete the proof of Theorem 1, we need to show that entropic security for predicates implies entropic security for all functions. The reduction is captured by the following lemma, which states that for every function with a good predictor (i.e. a predictor with advantage at least ϵ), there exists a predicate for which nearly the same predictor does equally well. This is the main technical result of this section.

The reduction uses the predictor $\mathcal{A}(Y(X))$ as a black box, and so we will simply use the random variable $A = \mathcal{A}(Y(X))$.

Lemma 2 (Main Lemma). *Let X be any distribution on $\{0,1\}^n$ of min-entropy $t \geq \frac{3}{2}\log\left(\frac{1}{\epsilon}\right)$, and let A be any random variable (possibly correlated to X). Suppose there exists a function $f : \{0,1\}^n \to \{0,1\}^*$ such that $\Pr[A = f(X)] \geq \mathsf{pred}_f + \epsilon$. Then there exists a* predicate *$g : \{0,1\}^n \to \{0,1\}$ and an algorithm $B(\cdot)$ such that*

$$\Pr[B(A) = g(X)] \geq \mathsf{pred}_g + \epsilon/4.$$

Due to space limitations, the proof is only given in the full version [9]. We mention only that there are two main steps to proving the lemma:

- If A is a good predictor for an (arbitrary) function $f(\cdot)$, then there is a (almost) *balanced* function $f'(\cdot)$ and a good predictor A' of the form $g(A)$.
- If $f(\cdot)$ is a balanced (or almost balanced) function and A is a good predictor for $f(X)$, then there is a predicate $g(\cdot)$ of the form $g'(f(\cdot))$ such that $g'(A)$ is a good predictor for $g(X)$.

A More Efficient Reduction. Lemma 2 completes the proof of Theorem 1. However, it says nothing about the running time of $B(\cdot)$—in general, the reduction may yield a large circuit. Nonetheless, we may indeed obtain a polynomial-time reduction for certain functions f. If no value of f occurs with probability more than ϵ^2, then inner product with a random vector provides a good predicate. The idea behind the following proof has appeared in other contexts, e.g. [11].

Proposition 1. *Let X be any random variable distributed in $\{0,1\}^n$. Let $f : \{0,1\}^n \to \{0,1\}^N$ be a function such that $\mathsf{pred}_{f,X} \leq \epsilon^2/4$, and let A be a random variable with advantage ϵ at guessing $f(X)$. For $r \in \{0,1\}^N$, let $g_r(x) = r \odot f(x)$. If r is drawn uniformly from $\{0,1\}^N$, then*

$$\mathsf{Exp}_r\left[\Pr[r \odot A = g_r(X)] - \mathsf{pred}_{g_r}\right] \geq \epsilon/4.$$

In particular, there exists a value r and a $O(N)$-time circuit B which satisfy $\Pr[B(A) = g_r(X)] \geq \mathsf{pred}_{g_r} + \epsilon/4$.

Proof. We can calculate the expected advantage almost directly. Note that conditioned on the event $A = f(X)$, the predictor $r \odot A$ always agrees with $g_r(X)$. When $A \neq f(X)$, they agree with probability exactly $\frac{1}{2}$. Hence, we have

$$\mathsf{Exp}_r\left[\Pr[r \odot A = g_r(X)]\right] = \frac{1}{2} + \frac{1}{2}\Pr[A = f(X)] \geq \frac{1}{2}(1 + \mathsf{pred}_f + \epsilon)$$

We must still bound the expected value of pred_{g_r}. Let $r_z = (-1)^{z \odot r}$. For any particular, r, we can compute pred_{g_r} as $\frac{1}{2} + \frac{1}{2} \left| \sum_z p_z r_z \right|$. Using the fact $\mathsf{Exp}\left[|Z|\right] \le \sqrt{\mathsf{Exp}\left[Z^2\right]}$ for any random variable Z, we get:

$$\mathsf{Exp}_r\left[\mathsf{pred}_{g_r}\right] = \frac{1}{2} + \frac{1}{2}\mathsf{Exp}_r\left[\left|\sum_z p_z r_z\right|\right] \le \frac{1}{2} + \frac{1}{2}\sqrt{\mathsf{Exp}_r\left[\left(\sum_z p_z r_z\right)^2\right]}$$

By pairwise independence of the variables r_z, we have $\mathsf{Exp}\left[r_z r_a\right]$ is 1 if $z = a$ and 0 otherwise.

$$\mathsf{Exp}_r\left[\mathsf{pred}_{g_r}\right] \le \frac{1}{2} + \frac{1}{2}\sqrt{\sum_z p_z^2} \le \frac{1}{2} + \frac{1}{2}\sqrt{\mathsf{pred}_f}.$$

The last inequality holds since pred_f is the maximum of the values p_z, and the expression $\sum_z p_z^2$ is maximized when $p_z = \mathsf{pred}_f$ for all z (note that this sum is the collision probability of $f(X)$). Combining the two calculations we have

$$\mathsf{Exp}_r\left[\Pr[r \odot A = g_r(X)] - \mathsf{pred}_{g_r}\right] \ge \frac{1}{2}\left(\mathsf{pred}_f + \epsilon - \sqrt{\mathsf{pred}_f}\right)$$

Using the hypothesis that $\mathsf{pred}_f \le \epsilon^2/4$, we see that the expected advantage is at least $\epsilon/4$. □

3 Encryption of High-Entropy Sources

In this section, we discuss the results on entropic security to the encryption of mesages which are guaranteed to come from a high-entropy distribution. Roughly: if the adversary has only a small chance of guessing the message ahead of time, then one can design information-theoretically secure encryption (in the sense of hiding all functions, Definition 1) using a much shorter key than is usually possible—making up for the small entropy of the key using the entropy inherent in the message.

3.1 Using Expander Graphs for Encryption

Formally, a symmetric encryption scheme is a pair of randomized maps $(\mathcal{E}, \mathcal{D})$. The encryption takes three inputs, an n-bit message x, a k-bit key κ and r random bits i, and produces a N-bit ciphertext $y = \mathcal{E}(x, \kappa; i)$. Note that the key and the random bits are expected to be uniform random bits, and when it is not necessary to denote the random bits or key explicitly we use either $\mathcal{E}(x, \kappa)$ or $\mathcal{E}(x)$. The decryption takes a key κ and ciphertext $y \in \{0, 1\}^N$, and produces the plaintext $x' = \mathcal{D}(y, \kappa)$. The only condition we impose for $(\mathcal{E}, \mathcal{D})$ to be called an encryption scheme is completeness: for all keys κ, $\mathcal{D}(\mathcal{E}(x, \kappa), \kappa) = x$ with probability 1.

In this section, we discuss graph-based encryption schemes and show that graph expansion corresponds to entropically secure encryption schemes.

Graph-based Encryption Schemes. Let $G = (V, E)$ be a d-regular graph, and let $N(v, j)$ denote the j-th neighbor of vertex v under some particular labeling of the edges. We'll say the labeling is *invertible* if there exists a map N' such that $N(v, j) = w$ implies $N'(w, j) = v$.

By Hall's theorem, every d-regular graph has an invertible labeling.[2] However, there is a large class of graphs for which this invertibility is much easier to see. The Cayley graph $G = (V, E)$ associated with a group $\mathcal{G}$ and a set of generators $\{g_1, ..., g_d\}$ consists of vertices labeled by elements of $\mathcal{G}$ which are connected when they differ by a generator: $E = \{(u, u \cdot g_i)\}_{u \in V, i \in [d]}$. When the set of generators contains all its inverses, the graph is undirected. For such a graph, the natural labeling is indeed invertible, since $N(v, j) = v \cdot j$ and $N'(w, j) = w \cdot j^{-1}$. All the graphs we discuss in this paper are in fact Cayley graphs, and hence invertibly labeled.

Now suppose the vertex set is $V = \{0,1\}^n$ and the degree is $d = 2^k$, so that the neighbor function N takes inputs in $\{0,1\}^n \times \{0,1\}^k$. Consider the encryption scheme:

$$\mathcal{E}(x, \kappa) = N(x, \kappa). \tag{5}$$

Notice, $\mathcal{E}$ is a proper encryption scheme if and only if the labeling is invertible. In that case, $\mathcal{D}(y, \kappa) = N'(y, \kappa) = x$. For efficiency, we should be able to compute N and N' in polynomial time. We will show that this encryption scheme is secure when the graph G is a sufficiently good expander. The following definition is standard:

Definition 3. *A graph $G = (V, E)$ is a (t, ϵ)-extractor if, for every set S of 2^t vertices, taking a random step in the graph from a random vertex of S leads to a nearly uniform distribution on the whole graph. That is, let U_S be uniform on S, J be uniform on $\{1, ..., d\}$ and U_V be uniform on the entire vertex set V. Then for all sets S of size at least 2^t, we require that:*

$$\mathbf{SD}\left(\ N(U_S, J)\ ,\ U_V\ \right) \le \epsilon.$$

The usual way to obtain extractors as above is to use good expanders. This is captured by the following fact.

Fact 1 (Expander smoothing lemma [12]). *A graph G with second largest (normalized) eigenvalue $\lambda \le \epsilon 2^{-(n-t)/2}$ is a (t, ϵ)-extractor.*

The equivalence between entropic security and indistinguishability (Theorem 1) gives us the following result:

Proposition 2. *For a 2^k-regular, invertible graph G as above, the encryption scheme $(\mathcal{E}, \mathcal{D})$ given by N, N' is (t, ϵ)-entropically secure if G is a $(t-2, 2\epsilon)$-extractor (in particular, if G has second eigenvalue $\lambda \le \epsilon \cdot 2^{-(n-t-2)/2}$).*

---

[2] We thank Noga Alon for pointing out this fact. If $G = (V, E)$ is a d-regular undirected graph, consider the bipartite graph with $|V|$ vertices on each side and where each edge in E is replaced by the corresponding pair of edges in the bipartite graph. By Hall's theorem, there exist d disjoint matchings in the bipartite graph. These induce an invertible labeling on the original graph.

Proof. By Theorem 1, it suffices to show that $(t-2,\epsilon)$-indistinguishability. And this immediately follows from the lemma above and the fact that any min-entropy $(t-2)$ distribution is a mixture of flat distributions. □

We apply this in two ways. First, using optimal expanders (Ramanujan graphs) we obtain the best known construction of entropically-secure encryption schemes (Corollary 1). Second, we give a simpler and much stronger analysis of the original scheme of Russell and Wang (Corollary 2).

Corollary 1. *There exists an efficient deterministic (t,ϵ)-entropically secure scheme with $k = n - t + 2\log\left(\frac{1}{\epsilon}\right) + 2$.*

Proof. We apply Proposition 2 to *Ramanujan graphs.* These graphs are optimal for this particular construction: they achieve optimal eigenvalue $\lambda = 2\sqrt{d-1}$ for degree d [17]. The bound on k now follows. □

The main drawback of Ramanujan graphs is that explicit constructions are not known for all sizes of graphs and degrees. However, large families exist (e.g. graphs with $q+1$ vertices and degree $p+1$, where p and q are primes congruent to 1 mod 4). Below we show why the construction from Russell and Wang [24] using small-biased spaces is actually a special case of Proposition 2.

Using Small-biased Sets. A set S in $\{0,1\}^n$ is *δ-biased* if for all nonzero $\alpha \in \{0,1\}^n$, the binary inner product $\alpha \odot s$ is nearly balanced for s drawn uniformly in S:

$$\Pr_{s \leftarrow S}[\alpha \odot s = 0] \in \left[\frac{1-\delta}{2}, \frac{1+\delta}{2}\right] \text{ or, equivalently, } \left|\mathsf{Exp}_{s\leftarrow S}\left[(-1)^{\alpha\odot S}\right]\right| \leq \delta. \quad (6)$$

Alon et al. [1] gave explicit constructions of δ-biased sets in $\{0,1\}^n$ with size $O(n^2/\delta^2)$. Now suppose the δ-biased set is indexed $\left\{s_\kappa | \kappa \in \{0,1\}^k\right\}$. Consider the encryption scheme: $\mathcal{E}(x,\kappa) = x \oplus s_\kappa$. Russell and Wang introduced this scheme and showed that it is (n,ϵ)-entropically secure when $\delta = \epsilon^{3/2}$, yielding a key length of $k = 2\log n + 3\log\left(\frac{1}{\epsilon}\right)$. However, their analysis works only when the message is drawn uniformly from $\{0,1\}^n$.

We propose a different analysis: consider the Cayley graph for $\mathbb{Z}_2^n$ with generators S, where S is δ-biased. This graph has second eigenvalue $\lambda \leq \delta$ [20, 2]. Hence, by Proposition 2 the scheme above is (t,ϵ)-entropically secure as long as $\delta \leq \epsilon 2^{-(n-t-2)/2}$. This gives a version of the Vernam one-time pad for high-entropy message spaces, with key length $k = n - t + 2\log n + 2\log\left(\frac{1}{\epsilon}\right) + O(1)$. Unlike [24], this works for *all* settings of t, and also improves the parameters in [24] for $n = t$.

Corollary 2. *If $\left\{s_\kappa | \kappa \in \{0,1\}^k\right\}$ is a δ-biased set, then the encryption scheme $\mathcal{E}(x,\kappa) = x \oplus s_\kappa$ is (t,ϵ) indistinguishable when $\epsilon = \delta 2^{(n-t-2)/2}$. Using the costruction of [1], this yields a scheme with key length $k = n - t + 2\log\left(\frac{1}{\epsilon}\right) + 2\log(n) + O(1)$ (for any value of t).*

3.2 A Random Hashing Construction

This section presents a simpler construction of entropically secure encryption based on pairwise independent hashing. Our result generalizes the construction of Russell and Wang [24] for nonuniform sources, and introduces a new variant of the leftover-hash/privacy-amplification lemma [3, 15].

The idea behind the construction is that indistinguishability is the same as extraction from a weak source, except that the extractor must in some sense be invertible: given the key, one must be able to recover the message.

Let $\{h_i\}_{i \in I}$ be some family of functions $h_i : \{0,1\}^k \to \{0,1\}^n$, indexed over the set $I = \{0,1\}^r$. We consider encryption schemes of the form

$$\mathcal{E}(x, \kappa; i) = (\ i,\ x \oplus h_i(\kappa)\ ) \quad \text{(for general functions } h_i\text{), or} \tag{7}$$

$$\mathcal{E}'(x, \kappa; i) = (\ i,\ h_i(x) \oplus \kappa\ ) \quad \text{(when the functions } h_i \text{ are permutations)} \tag{8}$$

These schemes can be thought of as low-entropy, probabilistic one-time pads. Decryption is obviously possible, since the description of the function h_i is public. For the scheme to be (t, ϵ)-secure, we will see that it is enough to have $k = n - t + 2\log\left(\frac{1}{\epsilon}\right) + 2$, and for the function family to be pairwise independent. (This matches the result in Corollary 1.) In fact, a slightly weaker condition is sufficient: The following definition was introduced in the context of authentication [16]:

Definition 4 (XOR-universal function families). *A collection of functions $\{h_i\}_{i \in I}$ from n bits to n bits is* XOR-universal *if:* $\forall a, x, y \in \{0,1\}^n, x \neq y$: $\Pr_{i \leftarrow I}[h_i(x) \oplus h_i(y) = a] \leq \frac{1}{2^n - 1}$.

It is easy to construct XOR-universal families. Any (ordinary) pairwise independent hash family will do, or one can save some randomness by avoiding the "offset" part of constructions of the form $h(x) = ax + b$. Specifically, view $\{0,1\}^n$ as $\mathcal{F} = GF(2^n)$, and embed the key set $\{0,1\}^k$ as a subset of $\mathcal{F}$. For any $i \in \mathcal{F}$, let $h_i(\kappa) = i\kappa$, with multiplication in $\mathcal{F}$. This yields a family of linear maps $\{h_i\}$ with 2^n members. Now fix any $a \in \mathcal{F}$, and any $x, y \in \mathcal{F}$ with $x \neq y$. When i is chosen uniformly from $\{0,1\}^n$, we have $h_i(x) \oplus h_i(y) = i(x - y) = a$ with probability exactly 2^{-n}. If we restrict i to be nonzero, then we get a family of *permutations*, and we get $h_i(x) \oplus h_i(y) = a$ with probability at most $\frac{1}{2^n - 1}$.

Proposition 3. *If the family $\{h_i\}$ is XOR-universal, then the encryption schemes*

$$\mathcal{E}(x, \kappa; i) = (i, x \oplus h_i(\kappa)) \quad \textit{and} \quad \mathcal{E}'(x, \kappa; i) = (i, h_i(x) \oplus \kappa)$$

are (t, ϵ)-entropically secure, for $t = n - k + 2\log\left(\frac{1}{\epsilon}\right) + 2$. (However, $\mathcal{E}'$ is a proper encryption scheme only when $\{h_i\}$ is a family of permutations.)

This proposition proves, as a special case, the security of the Russell-Wang construction, with slightly better parameters (their argument gives a key length of $n - t + 3\log\left(\frac{1}{\epsilon}\right)$ since they used 3-wise independent permutations, which are also harder to construct). It also proves the security of the simple construction $\mathcal{E}(x, \kappa; i) = (i, x + i\kappa)$, with operations in $GF(2^n)$.

Proposition 3 follows from the following lemma of independent interest, which is closely related to the to the *leftover hash lemma* [13] (also called *privacy amplification*; see, e.g. [3, 4]), and which conveniently handles both the $\mathcal{E}$ and the $\mathcal{E}'$ variants.

Lemma 3. *If A, B are independent random variables such that $\mathbf{H}_\infty(A)+\mathbf{H}_\infty(B) \geq n + 2\log\left(\frac{1}{\epsilon}\right) + 1$, and $\{h_i\}$ is a XOR-universal family, then*

$$\mathbf{SD}\left(\langle i,\ h_i(A) \oplus B\rangle\ ,\ \langle\ i,\ U_n\rangle\right) \leq \epsilon,$$

where U_n and i are uniform on $\{0,1\}^n$ and $\mathcal{I}$.

Proof. Consider the collision probability of $(i,\ h_i(A) \oplus B)$. A collision only occurs if the same function h_i is chosen both times. Conditioned on that, one obtains a collision only if $h_i(A) \oplus h_i(A') = B \oplus B'$, for A', B' i.i.d. copies of A, B. We can use the XOR-universality to bound this last term:

$$\begin{aligned}\Pr[(i,\ h_i(A) \oplus B) = (i,\ h_i(A') \oplus B')]& \\ = \Pr[i = i']\Big(&\Pr[B = B'] \cdot \Pr[h_i(A) = h_i(A')] \\ &+ \sum_{a \neq 0} \Pr[B \oplus B' = a] \cdot \Pr[h_i(A) \oplus h_i(A') = a]\Big)\end{aligned} \tag{9}$$

Now let $t_a = \mathbf{H}_2(A)$, $t_b = \mathbf{H}_2(B)$. For $a \neq 0$, we have $\Pr[h_i(A) \oplus h_i(A') = a] \leq 1/(2^n - 1)$, by the conditions on $\{h_i\}$. On the other hand, by a union bound we have

$$\Pr[h_i(A) = h_i(A')] \leq \Pr[A = A'] + \frac{1}{2^n - 1} \leq 2^{-t_a} + \frac{1}{2^n - 1}$$

Hence, Eqn. 9 reduces to

$$\begin{aligned}&\frac{1}{|\mathcal{I}|}\left(2^{-t_b}\left(2^{-t_a} + \frac{1}{2^n - 1}\right) + \frac{1}{2^n - 1}\left(\sum_{a \neq 0} \Pr[B \oplus B' = a]\right)\right) \\ &\qquad\qquad \leq \frac{1}{|\mathcal{I}|2^n}\left(1 + 2^{n - t_a - t_b} + 2^{-t_b} + \frac{2}{2^n - 1}\right)\end{aligned}$$

Now $2^{n - t_a - t_b} \leq \epsilon^2/2$ by assumption, and we also have $2^{-n} \leq 2^{-t_b} \leq \epsilon^2/2$, since $t_a, t_b \leq n$ and $t_a + t_b \geq n + 2\log\left(\frac{1}{\epsilon}\right)$ (similarly, $n \geq 2\log\left(\frac{1}{\epsilon}\right)$). Hence Eqn. 9 reduces to $(1 + 2\epsilon^2)/|\mathcal{I}|2^n$. Any distribution on a finite set S with collision probability $(1 + 2\epsilon^2)/|S|$ is at statistical distance at most ϵ from the uniform distribution [15]. Thus, $(i, h_i(A) \oplus B)$ is ϵ-far from uniform. □

Note that the lemma gives a special "extractor by XOR" which works for product distributions $A \times B$ with at least n bits of min-entropy between them.

3.3 Lower Bounds on the Key Length

Proposition 4. *Any encryption scheme which is (t, ϵ)-entropically secure for inputs of length n requires a key of length at least $n - t$.*

Proof. We can reduce our entropic scheme to Shannon-secure encryption of strings of length $n - t + 1$. Specifically, for every $w \in \{0,1\}^{n-t+1}$, let X_w be the uniform over strings with w as a prefix, that is the set $\{w\} \times \{0,1\}^{t-1}$. Since X_w has min-entropy $t - 1$, any pair of distributions $\mathcal{E}(X_w), \mathcal{E}(X_{w'})$ are indistinguishable, and so we can use $\mathcal{E}()$ to encrypt strings of length $n - t + 1$. When $\epsilon < 1/2$, we must have key length at least $(n - t + 1) - 1 = n - t$ by the usual Shannon-style bound (the loss of 1 comes from a relaxation of Shannon's bounds to statistical security). $\square$

Bounds for Public-Coin Schemes via Extractors. In the constructions of Russell and Wang and that of Section 3.1 and Section 3.2, the randomness used by the encryption scheme (apart from the key) is sent *in the clear* as part of the ciphertext. That is, $\mathcal{E}(x, \kappa; i) = (i, \mathcal{E}'(x, \kappa; i))$. For these types of schemes, called *public-coin* schemes, the intuitive connection between entropic security and extraction from weak sources is pretty clear: encryption implies extraction. As a result, lower bounds on extractors [23] apply, and show that our construction is close to optimal.

Proposition 5. *Any* public-coin*, (t, ϵ)-entropically secure encryption has key length $k \geq n - t + \log\left(\frac{1}{\epsilon}\right) - O(1)$ (as long as $t > 2\log\left(\frac{1}{\epsilon}\right)$).*

To prove the result, we first reduce to the existence of extractors:

Lemma 4. *Assume $(\mathcal{E}, \mathcal{D})$ is a* public-coin*, (t, ϵ)-entropically secure encryption scheme with message length n, key length k and r bits of extra randomness. Then there exists an extractor with seed length $k + r$, input length n and output length $n + r - \log\left(\frac{1}{\epsilon}\right)$, such that for any input distribution of min-entropy $t + 1$, the output is within distance 3ϵ of the uniform distribution.*

Proof. We combine three observations. First, when U is uniform over all messages in $\{0,1\}^n$, the entropy of the distribution $\mathcal{E}(U)$ must be high. Specifically: $\mathbf{H}_\infty(\mathcal{E}(U)) = n + r$. To see this, notice that there is a function $(\mathcal{D})$ which can produce R, K, U from $K, \mathcal{E}(U, K; R)$. Since the triple (R, K, U) is uniform on $\{0,1\}^{r+k+n}$, it must be that $(K, \mathcal{E}(U, K))$ also has min-entropy $r + k + n$, i.e. that any pair (κ, c) appears with probability at most $2^{-(n-k-r)}$. Summing over all 2^k values of κ, we see that any ciphertext value c appears with probability at most $\sum_\kappa 2^{-n-r-k} = 2^{-n-r}$, as desired.

The second observation is that there is a deterministic function ϕ which maps ciphertexts into $\{0,1\}^{n+r-\log\left(\frac{1}{\epsilon}\right)}$ such that $\phi(\mathcal{E}(U))$ is within distance ϵ of the uniform distribution. In general, any *fixed* distribution of min-entropy t can be mapped into $\{0,1\}^{t-\log(1/\epsilon)}$ so that the result is almost uniform (Simply assign elements of the original distribution one by one to strings in $\{0,1\}^{t-\log(1/\epsilon)}$, so that at no time do two strings have difference of probability more than 2^{-t}. The total variation from uniform will be at most $2^{t-\log(1/\epsilon)} \cdot 2^{-t} = \epsilon$.). Note that ϕ need not be efficiently computable, even if both $\mathcal{E}$ and $\mathcal{D}$ are straightforward. This doesn't matter, since we are after a combinatorial contradiction.

Finally, by Theorem 1, for all distributions of min-entropy $t-1$, we have $\mathbf{SD}(\mathcal{E}(U), \mathcal{E}(X)) \leq 2\epsilon$, and so $\mathbf{SD}(\phi(\mathcal{E}(U)), \phi(\mathcal{E}(X))) \leq 2\epsilon$. By the triangle inequality, $\phi(\mathcal{E}(X))$ is within 3ϵ of the uniform distribution on $n + r - \log\left(\frac{1}{\epsilon}\right)$ bits, proving the lemma. □

We can now apply the lower bound of Radhakrishnan and Ta-Shma [23], who showed that any extractor for distributions of min-entropy t with error parameter δ and d extra random bits can extract at most $t + d - 2\log(1/\delta) + O(1)$ nearly random bits. From Lemma 4, we get and extractor for min-entropy $t+1$, $\delta = 3\epsilon$, $k+r$ extra random bits, and output length $n+r-\log(1/\epsilon)$. Thus, $n+r-\log(1/\epsilon)$ is at most $t+1+k+r-2\log(1/\epsilon)+O(1)$, which immediately gives us Proposition 5.

Remark 1. We do not lose $\log(1/\epsilon)$ in the output length in Lemma 4 when the encryption scheme in indistinguishable from the uniform distribution (i.e., ciphertexts look truly random). For such public-coin schemes, we get $k \geq n-t+2\log\left(\frac{1}{\epsilon}\right) - O(1)$. Since all of our constructions are of this form, their parameters cannot be improved at all. In fact, we conjecture that $k \geq n-t+2\log\left(\frac{1}{\epsilon}\right) - O(1)$ for *all* entropically-secure schemes, public-coin or not.

Acknowledgements

We are grateful to many friends for helpful discussions on this work. We especially thank Noga Alon, Leonid Reyzin, Madhu Sudan, Salil Vadhan and Avi Wigderson for their insights.

References

1. Noga Alon, Oded Goldreich, Johan Håstad, René Peralta: Simple Constructions of Almost k-Wise Independent Random Variables. FOCS 1990: 544-553
2. Noga Alon and Yuval Roichman. Random Cayley graphs and expanders. *Random Structures & Algorithms* 5 (1994), 271–284.
3. C. Bennett, G. Brassard, and J. Robert. Privacy Amplification by Public Discussion. *SIAM J. on Computing*, **17**(2), pp. 210–229, 1988.
4. C. Bennett, G. Brassard, C. Crépeau, and U. Maurer. Generalized Privacy Amplification. *IEEE Transactions on Information Theory*, **41**(6), pp. 1915-1923, 1995.
5. R. Canetti. Towards realizing random oracles: Hash functions that hide all partial information. In *Crypto 1997*.
6. R. Canetti, D. Micciancio, O. Reingold. Perfectly One-Way Probabilistic Hash Functions. In *Proc. 30th ACM Symp. on Theory of Computing*, 1998, pp. 131–140.
7. Ran Canetti. Universally Composable Security: A New Paradigm for Cryptographic Protocols. *Proc. IEEE Symp. on Foundations of Computer Science*, 2001, pp. 136-145.

8. T. Cover, J. Thomas. *Elements of Information Theory.* Wiley series in telecommunication, 1991, 542 pp.
9. Y. Dodis, and A. Smith. Entropic Security and the Encryption of High Entropy Messages. Full version of this paper. Available at *IACR Cryptology ePrint Archive, report 2004/219*, at `http://eprint.iacr.org/2004/219/`.
10. S. Goldwasser and S. Micali. Probabilistic encryption. *JCSS*, **28**(2), pp. 270–299, April 1984.
11. Oded Goldreich, Salil Vadhan and Avi Wigderson. On Interactive Proofs with a Laconic Prover. *Computational Complexity*, 11(1-2): 1-53 (2002).
12. Oded Goldreich, Avi Wigderson: Tiny families of functions with random properties: A quality-size trade-off for hashing. Random Structures and Algorithms 11(4): 315-343 (1997)
13. J. Håstad, R. Impagliazzo, L. Levin, M. Luby. A Pseudorandom generator from any one-way function. In *Proc. 21st ACM Symp. on Theory of Computing*, 1989.
14. Jonathan Herzog. *Computational Soundness for Standard Assumptions of Formal Cryptography.* Ph.D. Thesis, Massachusetts Institute of Technology, May 2004.
15. R. Impagliazzo and D. Zuckerman. How to Recycle Random Bits. In *Proc. 30th IEEE Symp. on Foundations of Computer Science*, 1989.
16. H. Krawczyk. LFSR-Based Hashing and Authentication. In *Proc. CRYPTO '94*, p. 129–139, 1994.
17. A. Lubotzky, R. Phillips, P. Sarnak: Ramanujan graphs. Combinatorica 8(3): 261-277 (1988).
18. U. Maurer. Conditionally-Perfect Secrecy and a Provably-Secure Randomized Cipher. *J. Cryptology*, **5**(1), pp. 53–66, 1992.
19. U. Maurer. Secret Key Agreement by Public Discussion. *IEEE Trans. on Info. Theory*, 39(3):733–742, 1993.
20. J. Naor, M. Naor. Small-Bias Probability Spaces: Efficient Constructions and Applications. In *SIAM J. Comput.* 22(4): 838-856 (1993).
21. N. Nisan, D. Zuckerman. Randomness is Linear in Space. In *JCSS*, **52**(1), pp. 43–52, 1996.
22. B. Pfitzmann, M. Waidner. A Model for Asynchronous Reactive Systems and its Application to Secure Message Transmission. In *Proc. IEEE Symp. on Security and Privacy*, 2001, 184–200.
23. J. Radhakrishnan and A. Ta-Shma. Tight bounds for depth-two superconcentrators. In *Proc. 38th IEEE Symp. on Foundations of Computer Science*, 1997, pp. 585–594.
24. A. Russell and Wang. How to Fool an Unbounded Adversary with a Short Key. In *Advances in Cryptology — EUROCRYPT 2002.*
25. C. Shannon. Communication Theory of Secrecy systems. In *Bell Systems Technical J.*, 28:656–715, 1949. Note: The material in this paper appeared originally in a confidential report 'A Mathematical Theory of Cryptography', dated Sept. 1, 1945, which has now been declassified.