

Efficient Encryption for Rich Message Spaces under General Assumptions

Alexander Russell

acr@cse.uconn.edu

Hong Wang

hongmuw@cse.uconn.edu

Department of Computer Science and Engineering
University of Connecticut

April 3, 2001

Abstract

We present a new family of public-key encryption schemes which combine modest computational demands with provable security guarantees under only general assumptions. The schemes may be realized with any one-way trapdoor permutation, and provide a notion of security corresponding to semantic security under the condition that the message space has sufficient entropy. Furthermore, these schemes can be implemented with very few applications of the underlying one-way permutation: schemes which provide security for message spaces in $\{0, 1\}^n$ with minimum entropy $n - \ell$ can be realized with $\ell + w(k) \log k$ applications of the underlying one-way trapdoor permutation. Here k is the security parameter and $w(k)$ is any function which tends to infinity. In comparison, extant systems offering full semantic security require roughly n applications of the underlying one-way trapdoor permutation. Finally, we give a simplified proof of a fundamental “elision lemma” of Goldwasser and Micali.

1 Introduction

Given the current state of affairs in complexity theory, the study of encryption has adopted a somewhat axiomatic approach. A primary goal of the study is understand the basic relationship between the (complexity-theoretic) *assumptions* upon which encryption schemes can be based, and the *efficiency* and *privacy* guarantees offered by such schemes. Naturally, the most desirable encryption scheme is one which makes the most modest assumptions and offers efficient encryption with strong privacy guarantees.

A variety of complexity-theoretic assumptions have been studied, which range from general assumptions, like the existence of a one-way function, to strong assumptions about specific (often number-theoretic) functions. In this article, we will focus on the development of asymmetric encryption schemes under general (i.e., weak) assumptions. In particular, we will assume the existence of a one-way trapdoor permutation. (The constructions work under weaker assumptions, for example the existence of a one-way function, though in this case the target schemes must be private-key.)

A traditionally accepted notion of security for encryption schemes is that of *semantic security* [13], though a number of stronger (and important) notions exist (see, e.g., [4, 9, 22, 24]). A system with semantic security guarantees that observation of $E(m)$, the encryption of a message m , offers essentially no advantage to a bounded adversary in predicting *any piece of partial information about the message m* . A piece of partial information may be some specific bit of m , or, perhaps, a complicated function capturing some global property of m . Furthermore, this guarantee can be offered *regardless of the a priori distribution of*

the message m . Given a one-way trapdoor permutation f and a hard core predicate¹ b for f (see e.g., [7]), a semantically secure encryption for a message m with n bits can be realized with n applications of f . In general, if it is possible to extract $s_f(k)$ simultaneously secure bits from a single application of f to strings with security parameter k , then $n/s_f(k)$ evaluations of f suffice. If, for example, f is taken to be RSA, then it is known that if RSA is difficult to invert then $s_{RSA}(k) = \Omega(\log \log k)$ [1, 15], so that this scheme can be realized with $O(n/\log \log k)$ applications of RSA. The encryption schemes described below offer an efficient analogue of semantic security for the case when the adversary’s a priori knowledge of the message is limited.

We say that an encryption scheme offers *entropically bounded security* if for all message distributions with sufficient entropy, and all pieces of partial information $h : \{0, 1\}^* \rightarrow \{0, 1\}$, observation of $E(m)$ offers no bounded adversary any nonnegligible advantage in prediction of $h(m)$. If the definition is strengthened so that it applies for all message spaces, then we exactly recover the definition of semantic security. (See the next section for precise definitions.) We show that for message spaces with minimum entropy $n - \ell$, an encryption scheme offering entropically bounded security can be realized with very few applications of f ; in particular, $(\ell + w(k) \log k)/s_f(k)$ applications (inversions) suffice for encryption (decryption), where k is the security parameter, $w(k)$ is any function that tends to infinity, and $s_f(k)$ is the number of simultaneously secure bits which can be extracted from a single application of f (a one-way trapdoor permutation). When the message space is uniform, then, this results in a system which requires only $O(w(k) \log k/s_f(k))$ applications of the one-way permutation, for any function w which tends to infinity. The systems also involve a certain amount of “overhead,” which in each case does not exceed $O(n \text{ poly}(\log n))$ time.

The above results express the complexity of encryption as a function of both n , the message length, and k , the security parameter. This is somewhat unusual for asymmetric schemes, which are typically used to encrypt a key for a private-key scheme (typically of length k), as private-key schemes are generally (much) more efficient than a public-key schemes. Under the assumptions we consider, however, there is no (known) benefit to be had by applying a private-key system after key exchange, so we keep everything “under one hood”. (Alternatively, the results which follow can be cast in a private-key setting, as mentioned above.)

It is interesting to compare these results with known results adopting stronger assumptions. If factoring is difficult, then a scheme of Blum and Goldwasser [6] based on the Rabin functions ($x \mapsto x^2 \bmod pq$) encrypt (in a semantically secure fashion) an n -bit message in time $O(nk \text{ poly}(\log k))$. In comparison, the above scheme offers a weaker guarantee, analogous to semantic security when the message space has $n - \ell$ min entropy, in time $O([\ell + w(k) \log k]k \text{ poly}(\log k) + n \text{ poly}(\log n))$, where w is any function which tends to infinity. Under assumptions of a somewhat stronger flavor, Cramer and Shoup [8] show that a constant number of exponentiations over a group suffice to encrypt a message of length k , in such a way that the resulting system is secure against even (adaptive) chosen ciphertext attack. In particular, they assume that the Diffie-Hellman decision problem is hard (i.e., that the El Gamal scheme [10] is semantically secure). Previous work has also constructed efficient, secure encryption schemes (with quite strong notions of security) under the strong assumption of availability of an ideal hash function [5].

The two main theorems in the article, Theorem 3 and Theorem 4, are both instantiations of common paradigms in cryptography. The first is an information-theoretic variant on the standard practice of encrypting a short seed which is then used for a pseudorandom generator (in our case, this will be an ϵ -biased space). The second is a variant of the “simple embedding schemes” often used in practice, where a message is encrypted by applying a one-way permutation after a suitable (bijective) hash function. The scheme of Bellare and Rogaway [5] is also theoretical evidence for the quality of such systems.

In Section 2 we give basic definitions, including a brief discussion of ϵ -biased spaces, universal hash functions, and the Fourier analysis of \mathbb{Z}_2^n , which will be used in the main results, presented in Sections 3

¹A *hard-core predicate* b for a one-way function f is a efficiently computable Boolean function so that $b(x)$ is difficult to predict from $f(x)$.

and 4.

2 Definitions

For basic definitions of one-way trapdoor permutations and hard-core predicates we refer the reader to, e.g., [26, 19]. For a one-way permutation f , we shall let $s_f(k)$ denote a lower bound on the number of simultaneously secure hard-core predicates for f (see, e.g., [11]).

Definition 1. A public key encryption scheme is a triple (G, E, D) , where

- G is an efficient probabilistic key generation algorithm, which, on input 1^k , produces a pair of keys, (P, S) ; here P denotes the "public key" and S the "secret key."
- E is an efficiently computable encryption algorithm which, given a message m and public key P , outputs c , an encryption of the message m using the key P . We will consider probabilistic encryption schemes, where E may also depend on a sequence of random bits, R . The encryption of m with public key P and random string R is denoted $E(m; P, R)$.
- D is an efficiently computable decryption algorithm which, given a ciphertext c and secret key S , produces a message m for which $E(m; P, R) = c$ for some R .

As mentioned in introduction, *semantic security* is a standard notion of privacy for encryption schemes.

Definition 2. We say that an encryption scheme (G, E, D) possesses semantic security if for every message generator M and every probabilistic polynomial-time Turing machine A , there is a probabilistic polynomial-time Turing machine B , such that for every polynomial Q , there exists an integer k_0 such that $\forall k > k_0$ and $\forall h : \{0, 1\}^* \rightarrow \{0, 1\}^*$, $\Pr[A(1^k, P, E(m; P, R)) = h(m)] \leq \Pr[B(1^k) = h(m)] + \frac{1}{Q(k)}$ where the first probability is taken over $m \leftarrow M(1^k)$, $(P, S) \leftarrow G(1^k)$, R (the coin tosses of E), and coin tosses of A . The second probability is taken over all choices of $m \leftarrow M(1^k)$ and coin tosses of B .

We borrow the $\square \leftarrow \square$ notation from [14]: when x is a variable and S a random variable, $x \leftarrow S$ denotes the assignment of x according to S . If S is simply a set, we abuse the notation by allowing S to represent the random variable uniform on S . In the sequel, we will use the term "algorithm" to refer to a probabilistic polynomial time Turing machine. Furthermore, "message generators," as in the above definition, are algorithms which, for each $k \in \mathbb{N}$, produce a output in the set $\{0, 1\}^n$ (determined by the random coins of M), where n is polynomially bounded in k . Whenever a probability is expressed, as in the above definitions, it is understood that the random coins of any algorithm appearing inside the brackets are to be included in the probability space. When the underlying probability space of a variable x is clear from context, we may simply write $\Pr_x[P(x)]$, or elide x altogether.

Definition 3. We say that an encryption scheme possesses indistinguishability of encryptions if for every message generator M , every algorithm A , and for every polynomial Q , there exists an integer k_0 such that $\forall k > k_0$, $\Pr[A(1^k, P, m_0, m_1, E(m_i; P, R)) = i] < \frac{1}{2} + \frac{1}{Q(k)}$, this probability being taken over $m_0 \leftarrow M(1^k)$, $m_1 \leftarrow M(1^k)$, $(P, S) \leftarrow G(1^k)$, $i \leftarrow \{0, 1\}$, and selection of R .

Theorem 1. An encryption scheme is semantically secure if and only if it offers indistinguishability of encryptions.

The reverse implication was proven in [13]. The forward implication appears in [20, 12]. We shall require a strengthened version of the reverse implication, which we refer to as an "elision" lemma. This

strengthened version, discussed in Section 2.4, was originally proved in [12]. We give a streamlined proof of this result, which avoids the sampling present in existing proofs. One consequence of this equivalence is that if the piece of partial information h in the definition of semantic security is restricted to be a Boolean function, the notion of security is unchanged.

A random variable m taking values in $\{0, 1\}^n$ has *minimum entropy* $n - \ell$ when $\forall m_o \in \{0, 1\}^n$, $\Pr[m = m_o] \leq 2^{-n+\ell}$. A message generator M , which produces messages of length $n = n(k)$ given 1^k , is said to have minimum entropy $n - \ell$ when the random variable $M(1^k)$ possesses this property.

Definition 4. We say that an encryption system possesses $\ell(k)$ -entropic security if for every message generator M with minimum entropy $n - \ell(k)$, and every algorithm A , there is an algorithm B , such that for every polynomial Q , there exists an integer k_0 such that $\forall k > k_0$ and $\forall h : \{0, 1\}^* \rightarrow \{0, 1\}$,

$$\Pr[A(1^k, P, E(m; P, R)) = h(m)] \leq \Pr[B(1^k) = h(m)] + \frac{1}{Q(k)}$$

where the first probability is taken over $m \leftarrow M(1^k)$, $(P, S) \leftarrow G(1^k)$, and R . The second probability is taken over $m \leftarrow M(1^k)$.

Observe that a semantically secure encryption scheme possesses $p(k)$ -entropic security for every polynomial p . We will construct two encryption schemes, (G_u, E_u, D_u) and (G_b, E_b, D_b) , based on any one-way trapdoor permutation, so that

- E_u possesses 0-entropic security (i.e., provides security when the message space is uniform) and requires

$$O\left(\frac{w(k) \log(k)}{s_f(k)} t_f(k) + w(k) n \log^{1+\varepsilon} k\right)$$

time to encrypt a message, where $t_f(k)$ is the time required to compute a single application of the one-way permutation f to a string with security parameter k , $s_f(k)$ is the number of simultaneously secure bits which can be extracted from an application of f to strings with security parameter k , $w(k)$ is any function which tends to infinity in k , and $\varepsilon > 0$. In particular, the time cannot exceed $O((n + t_f(k)) \log^2 k)$. (The $w(k) n \log^{1+\varepsilon} k$ term may in fact be replaced by $n \log k \log \log k \log \log \log k$.)

- E_b possesses ℓ -entropic security (i.e., provides security when the message space has minimum entropy $n - \ell$) and requires

$$O\left(\frac{w(k) \log k + \ell}{s_f(k)} t_f(k) + n \log^2 n \log \log n\right)$$

time to encrypt a message, where $t_f(k)$, s_f , and w are as above.

For simplicity we will focus on the time taken to *encrypt* in these schemes, often simply focusing our attention on the number of applications of the underlying one-way permutation required. In these cases, decryption involves inverting the one-way permutation on a like number of elements (and the same ‘‘overhead’’ terms: $O(n \log n \log \log n)$ in the above case).

Our constructions make use of ε -biased sample spaces and universal hash functions, defined below.

2.1 ε -biased Sample Spaces

Definition 5. A sample space $S \subseteq \{0, 1\}^n$ is called ε -biased if for all nonempty $\alpha \subset [n] = \{1, \dots, n\}$,

$$\left| \text{Exp}_{s \in S} \left[\prod_{a \in \alpha} (-1)^{s_a} \right] \right| \leq \varepsilon.$$

Small probability spaces with these properties were initially constructed by Naor and Naor [21] and Peralta [23]. We will use a construction, due to Alon, Goldreich, Håstad and Peralta [2], which gives an ε -biased sample space in $\{0, 1\}^n$ of size about $(\frac{n}{\varepsilon})^2$. The sample space is given as the image of a certain function $\sigma_{n,m} : \mathbb{F}_{2^m} \times \mathbb{F}_{2^m} \rightarrow \{0, 1\}^n$. (Here \mathbb{F}_{2^m} denotes the finite field with 2^m elements.) To define σ , let $\text{bin} : \mathbb{F}_{2^m} \rightarrow \{0, 1\}^m$ be a bijection satisfying $\text{bin}(0) = 0^m$ and $\text{bin}(x + y) = \text{bin}(x) \oplus \text{bin}(y)$, where $\alpha \oplus \beta$ denotes the componentwise exclusive or of α and β . Then $\sigma(x, y) = r = (r_0, \dots, r_{n-1})$, where $r_i = \langle \text{bin}(x^i), \text{bin}(y) \rangle_2$, the inner product, modulo two, of x^i and y . The size of the sample space is 2^{2m} . Let $S_{n,m} \subset \{0, 1\}^n$ be the collection of points so defined. They show that

Theorem 2 ([2]). $S_{n,m} = \mathbf{im} \sigma_{m,n}$ is $\frac{n-1}{2^m}$ -biased.

Observe that when $m = \lceil \log n \varepsilon^{-1} \rceil$, $\frac{n-1}{2^m} \leq \varepsilon$. As we will be constructing elements of $S_{m,n}$ during the encryption (and decryption) phase of our encryption scheme, we analyze the complexity of computing the function above. First, we need to find an irreducible polynomial p of degree m over the finite field \mathbb{F}_2 . As the degree of the polynomial will correspond to the block length of the encryption scheme, we can be somewhat flexible concerning the degree of the irreducible polynomial and use an explicit construction (rather than rely on an algorithm²):

Fact 1. For each $c \in \mathbb{N}$, the polynomial $p_c(x) = x^{2^m} + x^m + 1$, where $m = 3^c$, is irreducible over \mathbb{F}_2 .

(See [18, Exercise 3.96].) Computation of $\sigma = \sigma_{m,n}$ for a pair (x, y) is performed on a component by component basis: given x^i , computation of x^{i+1} requires a single multiplication in $\mathbb{F}_{2^m} \cong \mathbb{F}_2[x]/(p_c)$. Using fast polynomial multiplication, computing this product takes $O(m \log m \log \log m)$ time (see [28], or the discussion in [3, p. 232]). As p_c is sparse (it has only 3 nonzero terms), reducing this result modulo p_c requires $O(m)$ time. Hence computation of $\sigma(x, y)$ requires $O(nm \log m \log \log m)$ time. In order for $S = \mathbf{im} \sigma$ to be ε -biased, we may take $m = \lceil \log(n/\varepsilon) \rceil$, in which case the above running time is $O(n \log(n/\varepsilon) \log \log(n/\varepsilon) \log \log \log(n/\varepsilon))$. To simplify notation, we let $\sigma_{n,\varepsilon}$ denote $\sigma_{n,m}$ in the sequel, for $m = \lceil \log(n/\varepsilon) \rceil$.

2.2 k -wise Independent Permutations

A family of permutations $\mathcal{P} \subset \{f : X \rightarrow X\}$ is a family of k -wise independent permutations [33] if for all distinct $s_1, \dots, s_k \in X$ and all distinct $t_1, \dots, t_k \in X$,

$$\Pr_{\phi \in \mathcal{P}} [\forall i, \phi(s_i) = t_i] = \prod_{i=0}^{k-1} \frac{1}{|X| - i}.$$

We will use a family of 3-wise independent permutations, described below. See Rees [25] for a more detailed description.

Let V be a two-dimensional vector space over \mathbb{F} , a finite field. For two non-zero vectors \vec{v} and \vec{w} in this space, we write $\vec{v} \sim \vec{w}$ when $\vec{v} = c\vec{w}$ for some $c \in \mathbb{F}$ (so that the two vectors span the same one-dimensional subspace). This is an equivalence relation; we write $[\vec{v}]$ for the equivalence class containing \vec{v} . Projective 2-space over \mathbb{F} is then $P_2(\mathbb{F}) = \{[\vec{v}] \mid \vec{v} \neq \vec{0}\}$. We let $\text{GL}_2(\mathbb{F})$ denote the set of non-singular 2×2 matrices over \mathbb{F} , and $\text{PGL}_2(\mathbb{F}) = \text{GL}_2(\mathbb{F}) / \{cI \mid c \in \mathbb{F}\}$, where I is the identity matrix. An element ϕ of $\text{PGL}_2(\mathbb{F})$ acts on $P_2(\mathbb{F})$ in a natural (and well-defined) way, mapping $[\vec{v}]$ to $[\phi(\vec{v})]$. It is not difficult to show that for any distinct $[\vec{u}_1], [\vec{u}_2], [\vec{u}_3] \in P_2(\mathbb{F})$ and any distinct $[\vec{v}_1], [\vec{v}_2], [\vec{v}_3] \in P_2(\mathbb{F})$, there is in fact a unique $\phi \in \text{PGL}_2(\mathbb{F})$ so that $\phi([\vec{u}_i]) = [\vec{v}_i]$ for each i . In particular, $\text{PGL}_2(\mathbb{F})$ is a 3-wise independent family

²Irreducible polynomials over \mathbb{F}_2 of degree m can be found deterministically in $m^{4+\varepsilon}$ time for any $\varepsilon > 0$ [30]; a randomized algorithm is known [31] which finds such a polynomial time in expected $m^2 \log^{O(1)} m$ time.

of permutations. As multiplication and inversion in a finite field \mathbb{F}_p , for a prime p , may be accomplished in time $O(\log p(\log \log p)^2 \log \log \log p)$ time [29, 27, 16], evaluation of an element $\phi \in \text{PGL}_2(\mathbb{F}_p)$ also has this complexity.

2.3 Fourier Analysis of Boolean Functions

Let $L(\mathbb{Z}_2^n) = \{f : \mathbb{Z}_2^n \rightarrow \mathbb{R}\}$ denote the set of real valued functions on $\mathbb{Z}_2^n = \{0, 1\}^n$. Though our interest shall be in Boolean functions, it will be temporarily convenient to consider this richer space. $L(\mathbb{Z}_2^n)$ is a vector space over \mathbb{R} of dimension 2^n , and has a natural inner product: for $f, g \in L(\mathbb{Z}_2^n)$, define $\langle f, g \rangle = 2^{-n} \sum_{x \in \{0, 1\}^n} f(x)g(x)$. For a subset $\alpha \subset \{1, \dots, n\}$, define the function $\chi_\alpha : \{0, 1\}^n \rightarrow \mathbb{R}$ so that $\chi_\alpha(x) = \prod_{a \in \alpha} (-1)^{x_a}$. These functions χ_α are the *characters* of $\mathbb{Z}_2^n = \{0, 1\}^n$. Among their many wonderful properties is the fact that *the characters form an orthonormal basis for $L(\mathbb{Z}_2^n)$* . To see this, observe that $\forall \alpha \subset [n]$, $\sum_{x \in \{0, 1\}^n} \chi_\alpha(x) = 2^n$ when $\alpha = \emptyset$, and 0 otherwise. Furthermore, for $\alpha, \beta \subset [n]$, $\chi_\alpha(x)\chi_\beta(x) = \chi_{\alpha \oplus \beta}(x)$, where $\alpha \oplus \beta$ denotes the symmetric difference of α and β , so that $\langle \chi_\alpha, \chi_\beta \rangle = 1$ when $\alpha = \beta$, and 0 otherwise. Considering that there are 2^n characters, pairwise orthogonal, they span $L(\mathbb{Z}_2^n)$, as promised. Any function $f : \{0, 1\}^n \rightarrow \mathbb{R}$ may then be written in terms of this basis:

$$f = \sum_{\alpha \subset [n]} \hat{f}_\alpha \chi_\alpha$$

where $\hat{f}_\alpha = \langle f, \chi_\alpha \rangle$ is the projection of f onto χ_α . These coefficients \hat{f}_α , $\alpha \subset [n]$, are the *Fourier coefficients* of f , and, as we have above observed, uniquely determine the function f .

Given the above, it is easy to establish the *Plancherel* equality:

Proposition 1. *Let $f \in L(\mathbb{Z}_2^n)$. Then $\|f\|_2^2 = \sum_\alpha \hat{f}_\alpha^2$, where $\|f\|_2^2 = \langle f, f \rangle = \frac{1}{2^n} \sum_{x \in \{0, 1\}^n} f(x)^2$.*

As always, $\hat{f}_\emptyset = \text{Exp}[f]$ and, when the range of f is $\{\pm 1\}$, $\sum_\alpha \hat{f}_\alpha^2 = \|f\|_2^2 = 1$. See [32] for an excellent discussions of discrete Fourier analysis.

2.4 An Elision Lemma.

We will use an “elision” lemma for semantically secure encryption schemes, applied in the proofs of Sections 3 and 4. We use the term *elision lemma* to refer to an assertion that a cryptosystem offering indistinguishability of encryptions possesses the property that any efficient computation performed with observation of $E(m)$, an encryption, (and, perhaps, some related information) may as well have been performed without it.

The following lemma, which generalizes the original elision lemma of [13], is due to [12]. We give a streamlined proof which improves upon previous proofs in the sense that it *requires no sampling* on the part of the constructed algorithm (F , in the proof below). It gives an error bound which depends only on a natural 2-norm of the message distribution.

Lemma 1. *Let (G, E, D) denote an encryption scheme possessing indistinguishability of encryptions. Then for every message space M and algorithm A , there is an algorithm B so that for all polynomials Q_1 , all efficiently computable $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$, and every polynomial Q_2 , $\exists k_0, \forall k > k_0$ and $\forall h : \{0, 1\}^* \rightarrow \{0, 1\}^*$,*

$$\Pr[A(1^k, P, f(s, m), E(m; P, R)) = h(s, m)] \leq \Pr[B(1^k, f(s, m)) = h(s, m)] + \frac{1}{Q_2(k)}.$$

The first probability is taken over $m \leftarrow M(1^k)$, $(P, S) \leftarrow G(1^k)$, $s \leftarrow \{0, 1\}^{Q_1(k)}$, and R . The second probability is taken over $m \leftarrow M(1^k)$ and $s \leftarrow \{0, 1\}^{P(k)}$.

Proof. The algorithm B uses A as a black box: given 1^k and $f(s, m)$, B proceeds as follows:

1. Select $m' \leftarrow M(1^k)$, $(P, S) \leftarrow G(1^k)$, and choose a random string R of appropriate length,
2. Return $A(1^k, P, f(s, m), E(m'; P, R)) = v$.

Observe that $\Pr[B(1^k, f(s, m)) = h(s, m)] = \Pr[A(1^k, P, f(s, m), E(m'; P, R)) = h(s, m)]$. In this case, the lemma is a consequence of the following claim:

Claim 1. For every message space M , efficient algorithm A , every polynomial Q_1 , efficiently computable $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$, and every polynomial Q_2 , $\exists k_0, \forall k > k_0$ and $\forall h : \{0, 1\}^* \rightarrow \{0, 1\}^*$,

$$\Pr[A(1^k, P, f(s, m), E(m; P, R)) = h(s, m)] \leq \Pr[A(1^k, P, f(s, m), E(m'; P, R)) = h(s, m)] + \frac{1}{Q_2(k)},$$

where each probability is taken over $m \leftarrow M(1^k)$, $m' \leftarrow M(1^k)$, $(P, S) \leftarrow G(1^k)$, $s \leftarrow \{0, 1\}^{Q_1(k)}$, and R .

Proof of Claim. Suppose not. Then there is a polynomial Q_2 , a message space M , and an algorithm A , a polynomial Q_1 and a function f so that $\forall k_0, \exists k > k_0$,

$$\Pr_{s, m, R, P} [A(1^k, P, f(s, m), E(m; P, R)) = h(s, m)] > \Pr_{s, m', R, P} [A(1^k, P, f(s, m), E(m'; P, R)) = h(s, m)] + \varepsilon$$

where $\varepsilon = \varepsilon(k) = \frac{1}{Q_2(k)}$.

For a pair of messages m, m' , define $P_{m, m'} = \Pr_{s, R, P} [A(1^k, P, f(s, m), E(m'; P, R)) = h(s, m)]$ and $P_{m, *} = \text{Exp}_{m'} [P_{m, m'}]$. Observe, then, that $\Pr_{s, R, P} [A(1^k, P, f(s, m), E(m; P, R)) = h(s, m)] = P_{m, m}$, so that

$$\begin{aligned} \Pr_{s, m, R, P} [A(1^k, P, f(s, m), E(m; P, R)) = h(s, m)] &= \text{Exp}_m [P_{m, m}], \text{ and} \\ \Pr_{s, m, m', R, P} [A(1^k, P, f(s, m), E(m'; P, R)) = h(s, m)] &= \text{Exp}_m [P_{m, *}]. \end{aligned}$$

In particular, $\text{Exp}_m [P_{m, m}] - \text{Exp}_m [P_{m, *}] > \varepsilon$.

Now, we build an algorithm F which, given random m_0 and m_1 , can distinguish an encryption of m_0 from one of m_1 . (See Definition 3.) The algorithm F proceeds as follows: given m_0, m_1 and $\alpha = E(m_i; P, R)$,

- j is chosen uniformly in $\{0, 1\}$, s is chosen uniformly in $\{0, 1\}^{P(n)}$, and R is chosen uniformly among strings of appropriate length. $E(m_j; P, R)$ and $f(s, m_0)$ are computed.
- $A(1^k, f(s, m_0), E(m_j; P, R))$ is simulated, resulting in the value v_j . $A(1^k, f(s, m_0), \alpha)$ is simulated, resulting in the value v .
- If $v = v_j$, output j ; otherwise output $1 - j$.

Let $I_n = \{A(1^k, P, f(s, m'), E(m; P, R)) \mid m, m' \in \{0, 1\}^n, s \in \{0, 1\}^{Q_1(n)}, R\}$ be values that algorithm A can take, when restricted to those inputs possible when $|m| = n$. Then, for $v \in I_n$, let

$$D_{m', m}^s(v) = \Pr_{R, P} [A(1^k, P, f(s, m'), E(m; P, R)) = v],$$

so that $P_{m',m} = \text{Exp}_s[D_{m',m}^s(h(s, m'))]$. Now, for a particular pair m_0, m_1 ,

$$\begin{aligned} \Pr[F(m_0, m_1, \alpha) = i] &= \sum_{i'=0}^1 \sum_{j'=0}^1 \Pr[i = i' \wedge j = j'] \cdot \Pr[F(m_0, m_1, \alpha) = j' \mid i = i', j = j'] \\ &= \text{Exp}_s \left[\frac{1}{4} \left(\sum_v D_{m_0, m_0}^s(v)^2 + 2 \left(1 - \sum_v D_{m_0, m_0}^s(v) \cdot D_{m_0, m_1}^s(v) \right) + \sum_v D_{m_0, m_1}^s(v)^2 \right) \right] \\ &\geq \frac{1}{2} + \frac{1}{4} \left(\text{Exp}_s [D_{m_0, m_0}^s(h(s, m_0)) - D_{m_0, m_1}^s(h(s, m_0))] \right)^2 = \frac{1}{2} + \frac{1}{4} (P_{m_0, m_0} - P_{m_0, m_1})^2. \end{aligned}$$

where inequality \geq^* follows because $\text{Exp}[X]^2$ never exceeds $\text{Exp}[X^2]$ for any random variable. Then

$$\begin{aligned} \Pr_{m_0, m_1} [F(m_0, m_1, \alpha) = i] &\geq \text{Exp}_{m_0, m_1} \left[\frac{1}{2} + \frac{1}{4} \cdot (P_{m_0, m_0} - P_{m_0, m_1})^2 \right] \\ &\geq \frac{1}{2} + \frac{1}{4} \cdot \left(\text{Exp}_{m_0, m_1} [P_{m_0, m_0} - P_{m_0, m_1}] \right)^2 = \frac{1}{2} + \frac{1}{4} \cdot \left(\text{Exp}_{m_0} [P_{m_0, m_0}] - \text{Exp}_{m_1} [P_{m_0, m_1}] \right)^2 \\ &= \frac{1}{2} + \frac{1}{4} \cdot \left(\text{Exp}_{m_0} [P_{m_0, m_0} - P_{m_0, *}] \right)^2 = \frac{1}{2} + \frac{1}{4} \cdot \left(\text{Exp}_m [P_{m, m}] - \text{Exp}_m [P_{m, *}] \right)^2 \geq \frac{1}{2} + \frac{\epsilon^2}{4}. \end{aligned}$$

Hence (E, D) does not offer indistinguishability of encryptions. \square

As mentioned above, the Lemma follows immediately from the Claim. \square

3 Security for Uniformly Distributed Message Spaces

We begin by constructing an encryption scheme offering security in the case when the adversary has no a priori knowledge concerning the message (i.e., the message space is uniform).

As mentioned in the introduction, under the assumption that there exists a one-way permutation f , there is a semantically secure public-key cryptosystem, $C_f = (G, E, D)$, which encrypts a message $m \in \{0, 1\}^n$ with $n/s_f(k)$ applications of the function f .

Theorem 3. *Let f be a one-way trapdoor permutation, and $C_f = (G, E, D)$ the associated semantically secure encryption scheme. Define a new scheme (G_u, E_u, D_u) , where $G_u = G$, and $E_u(m; P, (R, s)) = (m \oplus \sigma_{n, \epsilon}(s), E(s; P, R))$ where $|m| = n$ and s is chosen randomly in the domain of $\sigma_{n, \epsilon}$. Decryption is immediate. Then for $\epsilon = k^{-\omega(1)}$, where k is the security parameter of the system, this encryption scheme offers 0-entropic security. Furthermore, the scheme requires $O(w(k) \log k / s_f(k))$ applications of f , where w is any function tending to infinity. The scheme has $O(n \log k \log \log k \log \log \log k)$ overhead.*

Proof. For simplicity, we treat h as a function with range $\{\pm 1\}$ rather than $\{0, 1\}$. From Lemma 1, we have for every message space M and algorithm A , there is an algorithm B' such that for every polynomial P , there exists an integer k_0 such that $\forall k > k_0$ and $\forall h : M \rightarrow \{-1, 1\}$

$$\Pr[A(1^k, P, m \oplus \sigma(s), E(s; P, R)) = h(m)] \leq \Pr[B'(1^k, m \oplus \sigma(s)) = h(m)] + \frac{1}{P(k)}$$

We use B' to construct an algorithm B , which can predict $h(m)$ nearly as well as can A , even without witnessing $E_u(m)$. The algorithm B , on input 1^k , proceeds as follows:

- Select $m' \in \{0, 1\}^n$ randomly.

- Return $B'(1^k, m') = v$.

Observe that $\Pr_m[B(1^k) = h(m)] = \Pr_{m,m'}[B'(1^k, m') = h(m)]$.

Claim 2. Let G_m be the random variable $\text{Exp}_s[h(m \oplus \sigma(s))] - \text{Exp}_{m'}[h(m')]$; then

$$\left| \Pr_{m,m'} [B'(1^k, m') = h(m)] - \Pr_{m,s} [B'(1^k, m) = h(m \oplus \sigma(s))] \right| \leq \frac{1}{2} \text{Exp}_m [|G_m|].$$

Proof. Let $c(m, m')$ be the random variable so that $c(m, m') = 1$ when $B'(1^k, m) = h(m')$ and 0 if $B'(1^k, m) \neq h(m')$. As $h(m')$ takes values in the set $\{\pm 1\}$, we can rewrite $c(m, m') = \frac{1}{2} + \frac{B'(1^k, m)h(m')}{2}$ and

$$\begin{aligned} & \left| \Pr_{m,m'} [B'(1^k, m') = h(m)] - \Pr_{m,s} [B'(1^k, m) = h(m \oplus \sigma(s))] \right| \\ &= \left| \text{Exp}_m \left[\frac{B'(1^k, m)}{2} \left(\text{Exp}_{m'} [h(m')] - \text{Exp}_s [h(m \oplus \sigma(s))] \right) \right] \right| \\ &\leq \frac{1}{2} \text{Exp}_m \left[\left| \text{Exp}_{m'} [h(m')] - \text{Exp}_s [h(m \oplus \sigma(s))] \right| \right] = \frac{1}{2} \text{Exp}_m [|G_m|]. \end{aligned}$$

□

We apply the second moment method to control $\text{Exp}_m[|G_m|]$. Observe that $\text{Exp}_m[h(m)] = \widehat{h}_0$, so

$$G_m = \text{Exp}_s \left[\sum_{\alpha \neq 0} \widehat{h}_\alpha \chi_\alpha(m \oplus \sigma(s)) \right] = \sum_{\alpha \neq 0} \widehat{h}_\alpha \text{Exp}_s [\chi_\alpha(m \oplus \sigma(s))] = \sum_{\alpha \neq 0} \widehat{h}_\alpha \chi_\alpha(m) \text{Exp}_s [\chi_\alpha(\sigma(s))]$$

Then $\text{Exp}_m[G_m] = \sum_{\alpha \neq 0} \widehat{h}_\alpha \text{Exp}_s [\chi_\alpha(\sigma(s))] \text{Exp}_m [\chi_\alpha(m)] = 0$. Now, the random variables $\widehat{h}_\alpha \chi_\alpha(m \oplus \sigma(s))$ and $\widehat{h}_\beta \chi_\beta(m \oplus \sigma(s))$ are pairwise independent so that

$$\begin{aligned} \text{Var}_m[G_m] &= \text{Var} \left[\sum_{\alpha \neq 0} \widehat{h}_\alpha \chi_\alpha(m) \text{Exp}_s [\chi_\alpha(\sigma(s))] \right] \\ &= \sum_{\alpha \neq 0} \widehat{h}_\alpha^2 \text{Exp}_s [\chi_\alpha(\sigma(s))]^2 \text{Var}[\chi_\alpha(m)] \leq \varepsilon^2 \sum_{\alpha \neq 0} \widehat{h}_\alpha^2 \leq \varepsilon^2 \end{aligned}$$

by the Plancherel equality (see Section 2.3) and the fact that $\text{Var}[\chi_\alpha(m)] = 1$. Now, applying Chebyshev's inequality, we have $\Pr_m[|G_m| > \lambda] < \varepsilon^2 \lambda^{-2}$.

Selecting $\lambda = \varepsilon^{\frac{2}{3}}$, we have

$$\text{Exp}_m [|G_m|] = \Pr_m [|G_m| > \lambda] \cdot \max_m |G_m| + \Pr_m [|G_m| \leq \lambda] \cdot \lambda \leq \frac{\varepsilon^2}{\lambda^2} \cdot 2 + \left(1 - \frac{\varepsilon^2}{\lambda^2}\right) \cdot \lambda \leq 3\varepsilon^{\frac{2}{3}}.$$

Hence

$$\left| \Pr_{m,s} [B'(1^k, m \oplus \sigma(s)) = h(m)] - \Pr_m [B(1^k) = h(m)] \right| < \frac{3}{2} \varepsilon^{\frac{2}{3}}$$

and $\Pr[A(1^k, m \oplus \sigma(s), E(s; P, R)) = h(m)] \leq \Pr[B(1^k) = h(m)] + \frac{1}{P(k)} + \frac{3}{2} \varepsilon^{\frac{2}{3}}$. As $\varepsilon = k^{-\omega(1)}$, this completes the proof. The bound on $|s|$ (and hence the number of applications of the underlying one-way permutation which are required and the running time) follows from Section 2.1. □

4 Security for Entropically Rich Message Spaces

For convenience, in this section we will assume that the message space is \mathbb{Z}_{p+1} for a (known) prime p . Now, we select an artificial bijection $L : \mathbb{Z}_{p+1} \rightarrow P_2(\mathbb{F}_p)$, so that

$$L(z) = \left[\begin{pmatrix} 1 \\ z \end{pmatrix} \right], \text{ for } 0 \leq z \leq p-1, \text{ and } \quad L(p) = \left[\begin{pmatrix} 0 \\ 1 \end{pmatrix} \right].$$

L can be computed in linear time; L^{-1} can be computed by single inversion modulo p . Having fixed this bijection, we will treat the functions $\text{PGL}_2(\mathbb{F}_p)$, described in Section 2.2, as if they act on \mathbb{Z}_{p+1} .

Theorem 4. *Let f be a one-way trapdoor permutation, and $C_f = (G, E, D)$ the associated semantically secure encryption scheme. Define a new scheme (G_b, E_b, D_b) , where $G_b = G$, and $E_b(m; P, (R, z)) = (\phi(m) + z, \phi, E(z; P, R))$ where $m \in \mathbb{Z}_{p+1}$, z is chosen randomly in $S \subset \mathbb{Z}_{p+1}$, and ϕ is chosen randomly from $\text{PGL}_2(\mathbb{F}_p)$. (Here $+$ is modulo $p+1$ and S is an arbitrary, but fixed subset of size 2^s , for example, we may take $S = \{0, 1, \dots, 2^s - 1\}$.) Decryption is immediate. Then for $s = \ell + \omega(\log k)$, where k is the security parameter of the system, this encryption scheme offers ℓ -entropic security. Furthermore, the scheme can be realized with no more than $s/s_f(k)$ applications of f . Aside from these evaluations (inversions) of f , encryption (decryption) has $O(n \log^2 n \log \log n)$ overhead.*

Proof. From Lemma 1, we know that for every message space M and algorithm A , there is a algorithm B' such that for every polynomial P , there exists an integer k_0 such that $\forall k > k_0$ and $\forall h : M \rightarrow \{0, 1\}$

$$\Pr[A(1^k, P, \phi(m) + z, \phi, E(z; P, R)) = h(m)] \leq \Pr[B'(1^k, \phi(m) + z, \phi) = h(m)] + \frac{1}{P(k)}.$$

We use B' to construct an algorithm B , which can predict $h(m)$ (nearly) as well as can A , even without witnessing $E_b(m)$. The algorithm B , on input 1^k , proceeds as follows:

- Select m' according to M , select $\phi \in \text{PGL}_2(\mathbb{F}_p)$, and select z at random in S .
- Return $B'(1^k, \phi(m') + z, \phi) = v$.

Observe that $\Pr_m[B(1^k) = h(m)] = \Pr_{m, m', \phi, z}[B'(1^k, \phi(m') + z, \phi) = h(m)]$.

We begin by recording an analogue of Claim 2 for this cryptosystem which allows us to remove the dependence on the behavior of B' ; the proof is placed in an appendix.

Claim 3. *For any function h and all algorithms B' ,*

$$\left| \Pr_{m, m', \phi, z} \left[B'(1^k, \phi(m') + z, \phi) = h(m') \right] - \Pr_{m, \phi, z} \left[B'(1^k, \phi(m) + z, \phi) = h(m) \right] \right| \leq \frac{1}{2} \text{Exp}_{m, \phi, z} \left[\left| \text{Exp}_{m'}[h(m')] - \text{Exp}_{m', z'}[h(m') \mid \phi(m') + z' = \phi(m) + z] \right| \right]$$

Now, for an element $z_0 \in S$, let $S_{z_0} = \{z - z_0 \bmod (p+1) \mid z \in S\}$. Let

$$G_{m, \phi}^{z_0} = \text{Exp}_{m'}[h(m')] - \text{Exp}_{m', z' \in S_{z_0}}[h(m') \mid \phi(m') + z' = \phi(m)].$$

From above, it suffices to show that for any $z_0 \in S$, $\text{Exp}_{m, \phi}[|G_{m, \phi}^{z_0}|]$ is small. Now fix $z_0 \in S$. For a fixed message $m_0 \in \mathbb{Z}_{p+1}$ and an element $w \in \mathbb{Z}_{p+1}$, let $\mathcal{P}_w = \{\phi \in \text{PGL}_2(\mathbb{F}_p) \mid \phi(m_0) = w\}$. Let $p_{m_i} =$

$\Pr_m[m = m_i]$. For an element $w_0 \in \mathbb{Z}_{p+1}$ and a permutation $\phi \in \mathcal{P}_{w_0}$, let $A_{w_0}^\phi = \sum_{m \in \phi^{-1}(w_0 + S_{z_0})} p_m$ and $B_{w_0}^\phi = \sum_{m \in \phi^{-1}(w_0 + S_{z_0})} p_m \cdot h(m)$; then

$$\frac{B_{w_0}^\phi}{A_{w_0}^\phi} = \text{Exp}_m [h(m) \mid \phi(m) \in w_0 + S_{z_0}]. \quad (1)$$

(Here $w_0 + S_{z_0}$ denotes the set $\{w_0 + z \bmod (p+1) \mid z \in S_{z_0}\}$.)

Let X_m be the random variable taking the value p_m if $\phi(m) \in w_0 + S_{z_0}$, and 0 otherwise. Then $\sum_{m \in M} X_m = A_{w_0}^\phi$ and $\sum_{m \in h^{-1}(1)} X_m = B_{w_0}^\phi$ so that

$$\text{Exp}_{\phi \in \mathcal{P}_{w_0}} [A_{w_0}^\phi] = \text{Exp}_{\phi \in \mathcal{P}_{w_0}} [p_{m_0} + \sum_{m \neq m_0} X_m] = p_{m_0} + \sum_{m \neq m_0} (\Pr[\phi(m) \in w_0 + S_{z_0}] \cdot p_m) = p_{m_0} \left(1 - \frac{2^s - 1}{2^n}\right) + \frac{2^s - 1}{2^n}.$$

Hence $\frac{2^s - 1}{2^n} \leq \text{Exp}_{\phi \in \mathcal{P}_{w_0}} [A_{w_0}^\phi] \leq p_{m_0} \left(1 - \frac{2^s - 1}{2^n}\right) + \frac{2^s - 1}{2^n} \leq 2^{-n+s} + p_{m_0}$. Similarly, let $q = \Pr_m[h(m) = 1]$, then

$$q \cdot \frac{2^s - 1}{2^n} \leq \text{Exp}_{\phi \in \mathcal{P}_{w_0}} [B_{w_0}^\phi] \leq q \cdot \frac{2^s - 1}{2^n} + p_{m_0} \left(1 - \frac{2^s - 1}{2^n}\right).$$

Recalling that the distribution of m has minimum entropy $n - \ell$,

$$\frac{\text{Exp}_{\phi \in \mathcal{P}_{w_0}} [B_{w_0}^\phi]}{\text{Exp}_{\phi \in \mathcal{P}_{w_0}} [A_{w_0}^\phi]} - \text{Exp}_{m \in M} [h(m)] \leq \frac{2^\ell}{2^s - 1},$$

and similarly,

$$\frac{\text{Exp}_{\phi \in \mathcal{P}_{w_0}} [B_{w_0}^\phi]}{\text{Exp}_{\phi \in \mathcal{P}_{w_0}} [A_{w_0}^\phi]} - \text{Exp}_{m \in M} [h(m)] \geq \frac{q \cdot (2^s - 1)}{2^s + 2^n p_{m_0} - 1} - q \geq -2^{-s+\ell}.$$

Hence,

$$\left| \frac{\text{Exp}_{\phi \in \mathcal{P}_{w_0}} [B_{w_0}^\phi]}{\text{Exp}_{\phi \in \mathcal{P}_{w_0}} [A_{w_0}^\phi]} - \text{Exp}_{m \in M} [h(m)] \right| \leq 2 \cdot 2^{-s+\ell}.$$

We wish to insure that $A_{w_0}^\phi$ and $B_{w_0}^\phi$ are close to their expected values. In preparation for applying Chebyshev's inequality, we compute their variances. We have

$$\text{Var}_{\phi \in \mathcal{P}_{w_0}} [A_{w_0}^\phi] = \text{Var}[\sum_{m \in M} X_m] = \sum_{m \neq m_0} (\text{Var}[X_m]) + \sum_{m_1 \neq m_2 \neq m_0} \text{Cov}[X_{m_1}, X_{m_2}].$$

Now, $\sum_{m \neq m_0} \text{Var}[X_m] \leq \sum_{m \neq m_0} \text{Exp}[X_m^2] \leq 2^{-n+s} \cdot \sum_{m \in M} (p_m^2) \leq 2^{-2n+s+\ell}$, and these variables are pairwise negatively correlated (so that $\text{Cov}[X_{m_1}, X_{m_2}] < 0$ for $m_1 \neq m_2$, both distinct from m_0). Then, $\text{Var}_{\phi \in \mathcal{P}_{w_0}} [A_{w_0}^\phi] < 2^{-2n+s+\ell}$. Similarly, $\text{Var}_{\phi \in \mathcal{P}_{w_0}} [B_{w_0}^\phi] < q \cdot 2^{-2n+s+\ell}$. Observe that 3-wise independence is required here.

By Chebyshev's inequality we have

$$\Pr_{\phi \in \mathcal{P}_{w_0}} [|A_{w_0}^\phi - \text{Exp}[A_{w_0}^\phi]| \geq \delta_a] \leq \frac{\text{Var}[A_{w_0}^\phi]}{\delta_a^2} < \frac{2^s}{2^{2n-\ell} \cdot \delta_a^2}, \quad (2)$$

and

$$\Pr_{\phi \in \mathcal{P}_{w_0}} [|B_{w_0}^\phi - \text{Exp}[B_{w_0}^\phi]| \geq \delta_b] \leq \frac{\text{Var}[B_{w_0}^\phi]}{\delta_b^2} < \frac{q \cdot 2^s}{2^{2n-\ell} \cdot \delta_b^2}. \quad (3)$$

When both $\left|A_{w_0}^\phi - \text{Exp}_{\phi \in \mathcal{P}_{w_0}}[A_{w_0}^\phi]\right| \leq \delta_a$ and $\left|B_{w_0}^\phi - \text{Exp}_{\phi \in \mathcal{P}_{w_0}}[B_{w_0}^\phi]\right| \leq \delta_b$, we have, in particular, from equation (1),

$$\left| \frac{\text{Exp}_{\phi \in \mathcal{P}_{w_0}}[B_{w_0}^\phi]}{\text{Exp}_{\phi \in \mathcal{P}_{w_0}}[A_{w_0}^\phi]} - \text{Exp}_m[h(m) | \phi(m) \in w_0 + S_{z_0}] \right| \leq \frac{2(\delta_a + \delta_b)}{\text{Exp}_{\phi \in \mathcal{P}_{w_0}}[A_{w_0}^\phi]} \leq 2(\delta_a + \delta_b) \cdot 2^{n-s},$$

(assuming that $\delta_a, \delta_b < 1/2$). When $\delta_a = \delta_b = \frac{\varepsilon_1}{4 \cdot 2^{n-s}}$, this is the statement that

$$\left| \text{Exp}_m[h(m) | \phi(m) \in w_0 + S_{z_0}] - \text{Exp}_m[h(m)] \right| < \varepsilon_1 + 2 \cdot 2^{-s+\ell}$$

For this z_0 , we say that (ϕ, w) is an ε -concealed pair if $|\text{Exp}_m[h(m) | \phi(m) \in w + S_{z_0}] - \text{Exp}_m[h(m)]| < \varepsilon$.

From inequalities (2) and (3), for any fixed w_0 and $\varepsilon_1 > 0$ we see that

$$\Pr_{\phi \in \mathcal{P}_{w_0}} \left[(\phi, w_0) \text{ is not an } (\varepsilon_1 + 2 \cdot 2^{-s+\ell})\text{-concealed pair} \right] \leq \frac{2 \cdot 2^s}{2^{2n-\ell} \cdot \delta_a^2} < \frac{2^5}{\varepsilon_1^2 \cdot 2^{s-\ell}}.$$

Now, for any fixed m_0 and $\varepsilon > 2 \cdot 2^{-s+\ell}$,

$$\begin{aligned} \Pr_{\phi} [(\phi, \phi(m_0)) \text{ is an } \varepsilon\text{-concealed pair}] &= \sum_{w_i} \left(\Pr[\phi \in \mathcal{P}_{w_i}] \cdot \Pr_{\phi \in \mathcal{P}_{w_i}} [(\phi, w_i) \text{ is an } \varepsilon\text{-concealed pair}] \right) \\ &\geq 1 - \frac{2^5}{\varepsilon^2 \cdot 2^{s-\ell} - 4\varepsilon + 4 \cdot 2^{-s+\ell}}. \end{aligned}$$

For a random pair (m, ϕ) , we will (lower) bound the probability that $(\phi, \phi(m))$ is an ε -concealed pair for z_0 , since

$$\text{Exp}_{m, \phi} \left[\left| G_{m, \phi}^{z_0} \right| \right] \leq \varepsilon \Pr_{m, \phi} [(\phi, \phi(m)) \text{ is } \varepsilon\text{-concealed}] + (1 - \Pr_{m, \phi} [(\phi, \phi(m)) \text{ is } \varepsilon\text{-concealed}]).$$

For specific m and random choice of ϕ , we define Y_m to be the random variable taking the value 1 if $(\phi, \phi(m))$ is an ε -concealed pair, and 0 otherwise. Now,

$$\begin{aligned} \Pr_{m, \phi} [(\phi, \phi(m)) \text{ is an } \varepsilon\text{-concealed pair}] &= \text{Exp}_{m, \phi} [Y_m] = \sum_{m_i} p_{m_i} \cdot \text{Exp}_{\phi} [Y_{m_i}] \\ &= \Pr_{\phi} [(\phi, \phi(m)) \text{ is an } \varepsilon\text{-concealed pair}] \\ &\geq 1 - \frac{2^5}{\varepsilon^2 \cdot 2^{s-\ell} - 4\varepsilon + 4 \cdot 2^{-s+\ell}}, \end{aligned}$$

so that for all $\varepsilon > 2 \cdot 2^{-s+\ell}$, $\text{Exp}_{m, \phi} (|G_{m, \phi}^{z_0}|) \leq \varepsilon(1 - \delta) + \delta < \varepsilon + \delta$, where $\delta = \frac{2^5}{\varepsilon^2 \cdot 2^{s-\ell} - 4\varepsilon + 4 \cdot 2^{-s+\ell}}$. Select $\varepsilon = 4 \cdot 2^{\frac{-s+\ell}{3}}$. As $s = \ell + \omega(\log k)$, we can be guaranteed that $\varepsilon = k^{-\omega(1)}$ and so for all z , $\text{Exp}_{m, \phi} \left[\left| G_{m, \phi}^z \right| \right] = k^{-\omega(1)}$.

Hence, $\text{Exp}_{m, \phi, z} \left[\left| G_{m, \phi}^z \right| \right] = k^{-\omega(1)}$, which, considering the above claim, completes the proof. The bound on the number of applications of the underlying one-way permutation follows immediately from the definition of s . \square

5 Acknowledgements

We thank Oded Goldreich, Mats Näslund, and Réne Peralta for comments on an early draft of this paper. We thank David Zuckerman for referring us to the family of polynomials p_c of Fact 1 and the permutations $\text{PGL}_2(\mathbb{F})$.

References

- [1] Werner Alexi, Benny Chor, Oded Goldreich, and Claus P. Schnorr. RSA/Rabin bits are $1/2 + 1/\text{poly}(\log N)$ secure. In *25th Annual Symposium on Foundations of Computer Science*, pages 449–457, Singer Island, Florida, 24–26 October 1984. IEEE.
- [2] Noga Alon, Oded Goldreich, Johan Håstad, and René Peralta. Simple constructions of almost k -wise independent random variables. In *31st Annual Symposium on Foundations of Computer Science*, volume II, pages 544–553, St. Louis, Missouri, 22–24 October 1990. IEEE.
- [3] Eric Bach and Jeffrey Shallit. *Algorithmic number theory. Vol. 1*. MIT Press, Cambridge, MA, 1996. Efficient algorithms.
- [4] M. Bellare, A. Desai, A. Pointcheval, and P. Rogaway. Relations among notions of public-key cryptosystems. In Krawczyk [17], page 540.
- [5] Mihir Bellare and Phillip Rogaway. Optimal asymmetric encryption. In Alfredo De Santis, editor, *Advances in Cryptology—EUROCRYPT 94*, volume 950 of *Lecture Notes in Computer Science*, pages 92–111. Springer-Verlag, 1995, 9–12 May 1994.
- [6] Manuel Blum and Shafi Goldwasser. An *efficient* probabilistic public-key encryption scheme which hides all partial information. In G. R. Blakley and David Chaum, editors, *Advances in Cryptology: Proceedings of CRYPTO 84*, volume 196 of *Lecture Notes in Computer Science*, pages 289–299. Springer-Verlag, 1985, 19–22 August 1984.
- [7] Manuel Blum and Silvio Micali. How to generate cryptographically strong sequences of pseudo-random bits. *SIAM Journal on Computing*, 13(4):850–864, November 1984.
- [8] Ronald Cramer and Victor Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In Krawczyk [17], pages 13–25.
- [9] Danny Dolev, Cynthia Dwork, and Moni Naor. Non-malleable cryptography (extended abstract). In *Proceedings of the Twenty Third Annual ACM Symposium on Theory of Computing*, pages 542–552, New Orleans, Louisiana, 6–8 May 1991.
- [10] T. El Gamal. A public key cryptosystem and signature scheme based on discrete logarithms. *IEEE Transactions on Information Theory*, 31:469–472, 1985.
- [11] Oded Goldreich. Three xor-lemmas – an exposition. This survey is available at <ftp://theory.lcs.mit.edu/pub/people/oded/xor.ps.>, 1991.
- [12] Oded Goldreich. A uniform-complexity treatment of encryption and zero-knowledge. *Journal of Cryptology*, 6(1):21–53, 1993.
- [13] Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270–299, April 1984.
- [14] Shafi Goldwasser, Silvio Micali, and Ronald L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal on Computing*, 17(2):281–308, April 1988.
- [15] Johan Håstad and Mats Näslund. The security of individual RSA bits. In *39th Annual Symposium on Foundations of Computer Science*, pages 510–519, Palo Alto, California, 8–11 November 1998. IEEE.

- [16] D. E. Knuth. The analysis of algorithms. In *Actes du Congrès International des Mathématiciens*, volume 3. Gauthier-Villars, Paris, 1971.
- [17] Hugo Krawczyk, editor. *Advances in Cryptology—CRYPTO '98*, volume 1462 of *Lecture Notes in Computer Science*. Springer-Verlag, 23–27 August 1998.
- [18] Rudolf Lidl and Harald Niederreiter. *Finite Fields*, volume 20 of *Encyclopedia of Mathematics and its Applications*. Addison-Wesley Publishing Company, Reading, Massachusetts, 1983.
- [19] Alfred J. Menezes, Paul C. van Oorschot, and Sctoo A. Vanstone. *Handbook of Applied Cryptography*. CRC Press, 1996.
- [20] Silvio Micali, Charles Rackoff, and Bob Sloan. The notion of security for probabilistic cryptosystems. *SIAM Journal on Computing*, 17(2):412–426, April 1988.
- [21] Joseph Naor and Moni Naor. Small-bias probability spaces: Efficient constructions and applications. *SIAM Journal on Computing*, 22(4):838–856, August 1993.
- [22] Moni Naor and Moti Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *Proceedings of the Twenty Second Annual ACM Symposium on Theory of Computing*, pages 427–437, Baltimore, Maryland, 14–16 May 1990.
- [23] Rene Peralta. On the distribution of quadratic residues and nonresidues modulo a prime number. *Mathematics of Computation*, 58(197):433–440, 1992.
- [24] Charles Rackoff and Daniel R. Simon. Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In J. Feigenbaum, editor, *Advances in Cryptology—CRYPTO '91*, volume 576 of *Lecture Notes in Computer Science*, pages 433–444. Springer-Verlag, 1992, 11–15 August 1991.
- [25] E. G. Rees. *Notes on Geometry*. Springer-Verlag, 1983.
- [26] Ronald L. Rivest. Cryptography. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume A. MIT Press, 1990.
- [27] A. Schönhage. Schnelle berechnung von kettenbruchentwicklungen. *Acta Informatica*, 1:139–144, 1971.
- [28] A. Schönhage. Schnelle Multiplikation von Polynomen über Körpern der Charakteristik 2. *Acta Informat.*, 7(4):395–398, 1976/77.
- [29] A. Schönhage and V. Strassen. Schnelle multiplikation großer zahlen. *Computing*, 7:281–292, 1971.
- [30] Victor Shoup. New algorithms for finding irreducible polynomials over finite fields. *Math. Comp.*, 54(189):435–447, 1990.
- [31] Victor Shoup. Fast construction of irreducible polynomials over finite fields. *J. Symbolic Comput.*, 17(5):371–391, 1994.
- [32] Audrey Terras. *Fourier Analysis on Finite Groups and Applications*. Number 43 in Student Texts. Cambridge University Press, Cambridge, UK, 1999.
- [33] Mark N. Wegman and J. Lawrence Carter. New classes and applications of hash functions. In *20th Annual Symposium on Foundations of Computer Science*, pages 175–182, San Juan, Puerto Rico, 29–31 October 1979. IEEE.

A Proof of Claim 3

Proof. Proof of Claim 3

$$\begin{aligned}
& \left| \Pr_{m,m',\phi,z} \left[B'(1^k, \phi(m) + z, \phi) = h(m') \right] - \Pr_{m,\phi,z} \left[B'(1^k, \phi(m) + z, \phi) = h(m) \right] \right| \\
&= \left| \text{Exp}_{m,m',\phi,z} \left[\frac{1 + B'(1^k, \phi(m) + z, \phi)h(m')}{2} - \frac{1 + B'(1^k, \phi(m) + z, \phi)h(m)}{2} \right] \right| \\
&= \frac{1}{2} \left| \text{Exp}_{m,m',\phi,z} \left[B'(1^k, \phi(m) + z, \phi)h(m') \right] - \text{Exp}_{m,\phi,z} \left[B'(1^k, \phi(m) + z, \phi)h(m) \right] \right| \\
&= \frac{1}{2} \left| \text{Exp}_{\phi,m,z} \left[B'(1^k, \phi(m) + z, \phi) \left(\text{Exp}_{m'}[h(m')] - h(m) \right) \right] \right| \\
&= \frac{1}{2} \left| \text{Exp}_{\phi} \left[\sum_{e \in \mathbb{Z}_{p+1}} \Pr_{m,z}[\phi(m) + z = e] \cdot B'(1^k, e, \phi) \text{Exp}_{m,z} \left[\left(\text{Exp}_{m'}[h(m')] - h(m) \right) \mathbb{1}_{\phi(m) + z = e} \right] \right] \right| \\
&\leq \frac{1}{2} \text{Exp}_{\phi} \left[\sum_{e \in \mathbb{Z}_{p+1}} \Pr[\phi(m) + z = e] \cdot \left| \text{Exp}_{m'}[h(m)] - \text{Exp}_{m,z}[h(m) \mid \phi(m) + z = e] \right| \right]
\end{aligned}$$

Observe now that for any functions $f : \mathbb{Z}_{p+1} \rightarrow \mathbb{R}$ and $g : \mathbb{Z}_{p+1} \rightarrow \mathbb{Z}_{p+1}$,

$$\sum_e \Pr_x[g(x) = e] \text{Exp}_x[f(g(x)) \mid g(x) = e] = \text{Exp}_x[f(g(x))].$$

Then the above is equal to

$$\frac{1}{2} \text{Exp}_{\phi,m,z} \left[\left| \text{Exp}_m[h(m)] - \text{Exp}_{m',z'}[h(m') \mid \phi(m') + z' = \phi(m) + z] \right| \right].$$

□