

Many-to-One Trapdoor Functions and Their Relation to Public-Key Cryptosystems

Mihir Bellare¹, Shai Halevi², Amit Sahai³, and Salil Vadhan³

¹ Dept. of Computer Science & Engineering, University of California at San Diego
9500 Gilman Drive, La Jolla, CA 92093, USA

mihir@cs.ucsd.edu

<http://www-cse.ucsd.edu/users/mihir>.

² IBM T. J. Watson Research Center
P.O. Box 704, Yorktown Heights, NY 10598, USA

shaih@watson.ibm.com.

³ MIT Laboratory for Computer Science
545 Technology Square, Cambridge, MA 02139, USA

amits@theory.lcs.mit.edu

salil@math.mit.edu.

URL:<http://www-math.mit.edu/~salil>.

Abstract. The heart of the task of building public key cryptosystems is viewed as that of “making trapdoors;” in fact, public key cryptosystems and trapdoor functions are often discussed as synonymous. How accurate is this view? In this paper we endeavor to get a better understanding of the nature of “trapdooriness” and its relation to public key cryptosystems, by broadening the scope of the investigation: we look at general trapdoor functions; that is, functions that are not necessarily injective (ie., one-to-one). Our first result is somewhat surprising: we show that non-injective trapdoor functions (with super-polynomial pre-image size) can be constructed from any one-way function (and hence it is unlikely that they suffice for public key encryption). On the other hand, we show that trapdoor functions with polynomial pre-image size are sufficient for public key encryption. Together, these two results indicate that the pre-image size is a fundamental parameter of trapdoor functions. We then turn our attention to the converse, asking what kinds of trapdoor functions can be constructed from public key cryptosystems. We take a first step by showing that in the random-oracle model one can construct injective trapdoor functions from any public key cryptosystem.

1 Introduction

A major dividing line in the realm of cryptographic primitives is that between “one-way” and “trapdoor” primitives. The former effectively means the primitives of private key cryptography, while the latter are typically viewed as tied to public key cryptosystems. Indeed, the understanding is that the problem of building public key cryptosystems is the problem of “making trapdoors.”

Is it really? It is well known that injective (ie. one-to-one) trapdoor functions suffice for public key cryptography [Ya,GoMi]. We ask: is the converse true as well, or can public key cryptosystems exist under a weaker assumption? We take a closer look at the notion of a trapdoor, in particular from the point of view of how it relates to semantically secure encryption schemes, and discover some curious things. Amongst these are that “trapdoor one-way functions” are not necessarily hard to build, and their relation to public key encryption is more subtle than it might seem.

1.1 Background

The main notions discussed and related in this paper are one-way functions [DiHe], trapdoor (one-way) functions [DiHe], semantically secure encryption schemes [GoMi], and unapproximable trapdoor predicates [GoMi].

Roughly, a “one-way function” means a family of functions where each particular function is easy to compute, but most are hard to invert; trapdoor functions are the same with the additional feature that associated to each particular function is some “trapdoor” information, possession of which permits easy inversion. (See Section 2 for formal definitions.)

In the study of one-way functions, it is well appreciated that the functions need not be injective: careful distinctions are made between “(general) one-way functions”, “injective one-way functions,” or “one-way permutations.” In principle, the distinction applies equally well to trapdoor one-way functions. (In the non-injective case, knowledge of the trapdoor permits recovery of *some* pre-image of any given range point [DiHe].) However, all attention in the literature has focused on injective trapdoor functions, perhaps out of the sense that this is what is necessary for constructing encryption schemes: the injectivity of the trapdoor function guarantees the unique decryptability of the encryption scheme.

This paper investigates general (ie. not necessarily injective) trapdoor one-way functions and how they relate to other primitives. Our goal is to understand exactly what kinds of trapdoor one-way functions are necessary and sufficient for building semantically secure public key encryption schemes; in particular, is injectivity actually necessary?

Among non-injective trapdoor functions, we make a further distinction based on “the amount of non-injectivity”, measured by pre-image size. A (trapdoor, one-way) function is said to have pre-image size $Q(k)$ (where k is the security parameter) if the number of pre-images of any range point is at most $Q(k)$. We show that pre-image size is a crucial parameter with regard to building public-key cryptosystems out of a trapdoor function.

Rather than directly working with public-key cryptosystems, it will be more convenient to work with a more basic primitive called an unapproximable trapdoor predicate. Unapproximable trapdoor predicates are equivalent to semantically secure public key schemes for encrypting a single bit, and these in turn are equivalent to general semantically secure cryptosystems [GoMi].




Fig. 1. Illustrating our results: *Solid lines are standard implications; the dotted line is an implication in the random oracle model.*

1.2 Results

We have three main results. They are displayed in Fig. 1 together with known relations. We now discuss them.

One-way functions imply trapdoor functions. Our first result, given in Theorem 1, may seem surprising at first glance: we show that one-way functions imply trapdoor functions. We present a general construction which, given an arbitrary one-way function, yields a trapdoor (non-injective) one-way function.

Put in other words, we show that trapdoor functions are not necessarily hard to build; it is the combination of trapdooriness with “structural” properties like injectivity that may be hard to achieve. Thus the “curtain” between one-way and trapdoor primitives is not quite as opaque as it may seem.

What does this mean for public key cryptography? Impagliazzo and Rudich [ImRu] show that it would be very hard, or unlikely, to get a proof that one-way functions (even if injective) imply public key cryptosystems. Hence, our result shows that it is unlikely that any known technique can be used to construct public key encryption schemes from generic, non-injective, trapdoor functions. As one might guess given [ImRu], our construction does not preserve injectivity, so even if the starting one-way function is injective, the resulting trapdoor one-way function is not.

Trapdoor functions with poly pre-image size yield cryptosystems. In light of the above, one might still imagine that injectivity of the trapdoor functions is required to obtain public key encryption. Still, we ask whether the injectivity condition can be relaxed somewhat. Specifically, the trapdoor one-way functions which we construct from one-way functions have super-polynomial pre-image size. This leads us to ask about trapdoor functions with polynomially bounded pre-image size.

Our second result, Theorem 2, shows that trapdoor functions with polynomially bounded pre-image size suffice to construct unapproximable trapdoor predicates, and hence yield public key cryptosystems. This belies the impression that injectivity of the trapdoor function is a necessary feature to directly build a public key cryptosystem from it, and also suggests that the super-polynomial pre-image size in the construction of Theorem 1 is necessary.

From trapdoor predicates to trapdoor functions. We then turn to the other side of the coin and ask what kinds of trapdoor functions must necessarily exist to have a public key cryptosystem. Since unapproximable trapdoor predicates and semantically secure public key cryptosystems are equivalent [GoMi] we consider the question of whether unapproximable trapdoor predicates imply injective trapdoor functions.

In fact whether or not semantically secure public key cryptosystems imply injective trapdoor functions is not only an open question, but seems a hard one. (In particular, a positive answer would imply injective trapdoor functions based on the Diffie-Hellman assumption, a long standing open problem.) In order to get some insight and possible approaches to it, we consider it in a random oracle model (cf. [ImRu,BeRo]). Theorem 3 says that here the answer is affirmative: given an arbitrary secure public key cryptosystem, we present a function that has access to an oracle H , and prove the function is injective, trapdoor, and one-way when H is random.

The construction of Theorem 3 is quite simple, and the natural next question is whether the random oracle H can be replaced by some constructible cryptographic primitive. In the full version of the paper [BHSV], we show that this may be difficult, by showing that a cryptographically strong pseudorandom bit generator [BiMi,Ya], which seems like a natural choice for this construction, does not suffice. The next step may be to follow the approach initiated by Canetti [Ca]: find an appropriate cryptographic notion which, if satisfied by H , would suffice for the correctness of the construction, and then try to implement H via a small family of functions. However, one should keep in mind that replacement of a random oracle by a suitable constructible function is not always possible [CGH]. Thus, our last result should be interpreted with care.

1.3 Discussion and Implications

Theorems 1 and 2 indicate that pre-image size is a crucial parameter when considering the power of trapdoor functions, particularly with respect to constructing public-key cryptosystems. The significance and interpretation of Theorem 3, however, requires a bit more discussion.

At first glance, it may seem that public key cryptosystems “obviously imply” injective trapdoor functions. After all, a public key cryptosystem permits unique decryptability; doesn’t this mean the encryption algorithm is injective? No, because, as per [GoMi], it is a *probabilistic* algorithm, and thus not a function. To make it a function, you must consider it a function of two arguments, the message and the coins, and then it may no longer be injective, because two

coin sequences could give rise to the same ciphertext for a given message. Moreover, it may no longer have a (full) trapdoor, since it may not be possible to recover the randomness from the ciphertext. (Public key cryptosystems in the Diffie and Hellman sense [DiHe] imply injective trapdoor one-way functions as the authors remark, but that's because encryption there is deterministic. It is now understood that secure encryption must be probabilistic [GoMi].)

Theorem 3 has several corollaries. (Caveat: All in the random oracle model). First, by applying a transformation of [BeRo], it follows that we can construct non-malleable and chosen-ciphertext secure encryption schemes based on the Ajtai-Dwork cryptosystem [AjDw]. Second, combining Theorems 3 and 2, the existence of trapdoor functions with polynomially bounded pre-image size implies the existence of injective trapdoor functions. (With high probability over the choice of oracle. See Remark 5.) Third, if the Decisional Diffie-Hellman problem is hard (this means the El Gamal [ElG] cryptosystem is semantically secure) then there exists an injective trapdoor function.

Note that in the random oracle model, it is trivial to construct (almost) injective one-way functions: a random oracle mapping, say, n bits to $3n$ bits, is itself an injective one-way function except with probability 2^{-n} over the choice of the oracle. However, random oracles do not directly or naturally give rise to trapdoors [ImRu]. Thus, it is interesting to note that our construction in Theorem 3 uses the oracle to “amplify” a trapdoor property: we convert the weak trapdoor property of a cryptosystem (in which one can only recover the message) to a strong one (in which one can recover both the message and the randomness used).

Another interpretation of Theorem 3 is as a demonstration that there exists a model in which semantically secure encryption implies injective trapdoor functions, and hence it may be hard to prove a separation result, in the style of [ImRu], between injective trapdoor functions and probabilistic encryption schemes.

2 Definitions

We present definitions for one-way functions, trapdoor functions, and unapproximable trapdoor predicates.

PRELIMINARIES. If S is any probability distribution then $x \leftarrow S$ denotes the operation of selecting an element uniformly at random according to S , and $[S]$ is the support of S , namely the set of all points having non-zero probability under S . If S is a set we view it as imbued with the uniform distribution and write $x \leftarrow S$. If A is a probabilistic algorithm or function then $A(x, y, \dots; R)$ denotes the output of A on inputs x, y, \dots and coins R , while $A(x, y, \dots)$ is the probability distribution assigning to each string the probability, over R , that it is output. For deterministic algorithms or functions A , we write $z := A(x, y, \dots)$ to mean that the output of $A(x, y, \dots)$ is assigned to z . The notation $\Pr[E : R_1; R_2; \dots; R_k]$ refers to the probability of event E after the random processes R_1, \dots, R_k are performed in order. If x and y are strings we write their concatenation as $x||y$

or just xy . “Polynomial time” means time polynomial in the security parameter k , PPT stands for “probabilistic, polynomial time”, and “efficient” means computable in polynomial time or PPT.

2.1 One-Way and Trapdoor Function Families

We first define families of functions, then say what it means for them to be one-way or trapdoor.

FAMILIES OF FUNCTIONS. A *family of functions* is a collection $F = \{F_k\}_{k \in \mathbb{N}}$ where each F_k is probability distribution over a set of functions. Each $f \in [F_k]$ has an associated domain $\text{Dom}(f)$ and range $\text{Range}(f)$. We require three properties of the family:

- *Can generate*: The operation $f \leftarrow F_k$ can be efficiently implemented, meaning there is a PPT *generation* algorithm $F\text{-Gen}$ that on input 1^k outputs a “description” of a function f distributed according to F_k . This algorithm might also output some auxiliary information aux associated to this function (this is in order to later model trapdoors).
- *Can sample*: $\text{Dom}(f)$ is efficiently samplable, meaning there is a PPT algorithm $F\text{-Smp}$ that given $f \in [F_k]$ returns a uniformly distributed element of $\text{Dom}(f)$.
- *Can evaluate*: f is efficiently computable, meaning there is a polynomial time evaluation algorithm $F\text{-Eval}$ that given $f \in F_k$ and $x \in \text{Dom}(f)$ returns $f(x)$.

For an element $y \in \text{Range}(f)$ we denote the set of *pre-images* of y under f by

$$f^{-1}(y) = \{x \in \text{Dom}(f) : f(x) = y\}.$$

We say that F is *injective* if f is injective (ie. one-to-one) for every $f \in [F_k]$. If in addition $\text{Dom}(f) = \text{Range}(f)$ then we say that F is a family of *permutations*. We measure the amount of “non-injectivity” by looking at the maximum pre-image size. Specifically we say that F has *pre-image size* bounded by $Q(k)$ if $|f^{-1}(y)| \leq Q(k)$ for all $f \in [F_k]$, all $y \in \text{Range}(f)$ and all $k \in \mathbb{N}$. We say that F has *polynomially bounded pre-image size* if there is a polynomial $Q(k)$ which bounds the pre-image size of F .

ONE-WAYNESS. Let F be a family of functions as above. The *inverting probability* of an algorithm $I(\cdot, \cdot)$ with respect to F is a function of the security parameter k , defined as $\text{InvProb}_F(I, k) \stackrel{\text{def}}{=} \Pr[x' \in f^{-1}(y) : f \leftarrow F_k ; x \leftarrow \text{Dom}(f) ; y \leftarrow f(x) ; x' \leftarrow I(f, y)]$.

F is *one-way* if $\text{InvProb}_F(I, k)$ is negligible for any PPT algorithm I .

TRAPDOORNESS. A family of functions is said to be trapdoor if it is possible, while generating an instance f , to simultaneously generate as auxiliary output “trapdoor information” tp , knowledge of which permits inversion of f . Formally, a family of functions F is *trapdoor* if $F\text{-Gen}$ outputs pairs (f, tp) where f is the “description” of a function as in any family of functions and tp is auxiliary

trapdoor information. We require that there exists a probabilistic polynomial time algorithm $F\text{-Inv}$ such that for all k , all $(f, tp) \in [F\text{-Gen}(1^k)]$, and all points $y \in \text{Range}(f)$, the algorithm $F\text{-Inv}(f, tp, y)$ outputs an element of $f^{-1}(y)$ with probability 1. A family of trapdoor functions is said to be one-way if it is also a family of one-way functions.

A good (candidate) example of a trapdoor, one-way function family which is non-injective is the *Rabin family* [Rab]: here each function in F_k is four to one. (Traditionally, this function is used as the basis of a public key cryptosystem by first modifying it to be injective.)

Remark 1. It is well known that one can define one-way functions either in terms of function families (as above), or in terms of a single function, and the two are equivalent. However, for trapdoor functions, one must talk of families. To maintain consistency, we use the family view of one-way functions as well.

2.2 Trapdoor Predicate Families

We define unapproximable trapdoor predicate families [GoMi]. Recall that such a family is equivalent to a semantically secure public-key encryption scheme for a single bit [GoMi].

A *predicate* in our context means a *probabilistic* function with domain $\{0, 1\}$, meaning a predicate p takes a bit b and flips coins r to generate some output $y = p(b; r)$. In a *trapdoor predicate family* $P = \{P_k\}_{k \in \mathbb{N}}$, each P_k is a probability distribution over a set of predicates, meaning each $p \in [P_k]$ is a predicate as above. We require:

- *Can generate:* There is a generation algorithm $P\text{-Gen}$ which on input 1^k outputs (p, tp) where p is distributed randomly according to P_k and tp is trapdoor information associated to p . In particular the operation $p \leftarrow P_k$ can be efficiently implemented.
- *Can evaluate:* There is a PPT algorithm $P\text{-Eval}$ that given p and $b \in \{0, 1\}$ flips coins to output y distributed according to $p(b)$.

We say P has *decryption error* $\delta(k)$ if there is a PPT algorithm $P\text{-Inv}$ who, with knowledge of the trapdoor, fails to decrypt only with this probability, namely

$$\text{DecErr}_P(P\text{-Inv}, k) \stackrel{\text{def}}{=} \Pr [b' \neq b : p \leftarrow P_k ; b \leftarrow \{0, 1\} ; y \leftarrow p(b) ; b' \leftarrow P\text{-Inv}(p, tp, y)] \quad (1)$$

is at most $\delta(k)$. If we say nothing it is to be assumed that the decryption error is zero, but sometimes we want to discuss families with non-zero (and even large) decryption error.

UNAPPROXIMABILITY. Let P be a family of trapdoor predicates as above. The *predicting advantage* of an algorithm $I(\cdot, \cdot)$ with respect to P is a function of the security parameter k , defined as $\text{PredAdv}_P(I, k) \stackrel{\text{def}}{=} \Pr [b' = b : p \leftarrow P_k ; b \leftarrow \{0, 1\} ; y \leftarrow p(b) ; b' \leftarrow I(p, y)] - \frac{1}{2}$.

We say that P is *unapproximable* if $\text{PredAdv}_P(I, k)$ is negligible for any PPT algorithm I .

3 From One-Way Functions to Trapdoor Functions

In this section we establish the following result:

Theorem 1. *Suppose there exists a family of one-way functions. Then there exists a family of trapdoor, one-way functions.*

This is proved by taking an arbitrary family F of one-way functions and “embedding” a trapdoor to get a family G of trapdoor functions. The rest of this section is devoted to the proof.

3.1 Proof Sketch of Theorem 1

Given a family $F = \{F_k\}_{k \in \mathbb{N}}$ of one-way functions we show how to construct a family $G = \{G_k\}_{k \in \mathbb{N}}$ of trapdoor one-way functions.

Let us first sketch the idea. Given $f \in F_k$ we want to construct g which “mimics” f but somehow embeds a trapdoor. The idea is that the trapdoor is a particular point α in the domain of f . Function g will usually just evaluate f , except if it detects that its input contains the trapdoor; in that case it will do something trivial, making g easy to invert given knowledge of the trapdoor. (This will not happen often in normal execution because it is unlikely that a randomly chosen input contains the trapdoor.) But how exactly can g “detect” the trapdoor? The first idea would be to include α in the description of g so that it can check whether its input contains the trapdoor, but then g would no longer be one-way. So instead the description of g will include $\beta = f(\alpha)$, an image of the trapdoor under the original function f , and g will run f on a candidate trapdoor to see whether the result matches β . (Note that we do not in fact necessarily detect the real trapdoor α ; the trivial action is taken whenever some pre-image of β under f is detected. But that turns out to be OK.)

In the actual construction, g has three inputs, y, x, v , where v plays the role of the “normal” input to f ; x plays the role of the candidate trapdoor; and y is the “trivial” answer returned in case the trapdoor is detected. We now formally specify the construction and sketch a proof that it is correct.

A particular function $g \in [G_k]$ will be described by a pair (f, β) where $f \in [F_k]$ and $\beta \in \text{Range}(f)$. It is defined on inputs y, x, v by

$$g(y, x, v) = \begin{cases} y & \text{if } f(x) = \beta \\ f(v) & \text{otherwise.} \end{cases} \quad (2)$$

Here $x, v \in \text{Dom}(f)$, and we draw y from some samplable superset S_f of $\text{Range}(f)$. (To be specific, we set S_f to the set of all strings of length at most $p(k)$ where $p(k)$ is a polynomial that bounds the lengths of all strings in $\text{Range}(f)$.) So the domain of g is $\text{Dom}(g) = S_f \times \text{Dom}(f) \times \text{Dom}(f)$.

We now give an intuitive explanation of why G is one-way and trapdoor. First note that for any z it is the case that (z, α, α) is a preimage of z under g ,

so knowing α enables one to invert in a trivial manner, hence G is trapdoor. For one-wayness, notice that if $g(y, x, v) = z$ then either $f(v) = z$ or $f(x) = \beta$. Thus, producing an element of $g^{-1}(z)$ requires inverting f at either z or β , both of which are hard by the one-wayness of F . A formal proof that G satisfies the definition of a family of one-way trapdoor functions can be found in the full version of this paper [BHSV].

Remark 2. One can verify that the trapdoor functions g produced in the above construction are *regular* (ie. the size of $g^{-1}(y)$ is the same for all $y \in \text{Range}(g)$) if the original one-way functions f are regular. Thus, adding regularity as a requirement is *not* likely to suffice for making public-key cryptosystems.

4 From Trapdoor Functions to Cryptosystems

Theorem 1 coupled with [ImRu] says that it is unlikely that general trapdoor functions will yield semantically secure public-key cryptosystems. However, in our construction of Section 3.1 the resulting trapdoor function was “very non-injective” in the sense that the pre-image size was exponential in the security parameter. So, we next ask, what is the power of trapdoor function families with polynomially bounded pre-image size? We show a positive result:

Theorem 2. *If there exist trapdoor one-way function families with polynomially bounded pre-image size, then there exists a family of unapproximable trapdoor predicates with exponentially small decryption error.*

Theorem 2 extends the well-known result of [Ya,GoMi] that injective trapdoor functions yield semantically secure public-key cryptosystems, by showing that the injectivity requirement can be relaxed. Coupled with [ImRu] this also implies that it is unlikely that the analogue of Theorem 1 can be shown for trapdoor functions with polynomially bounded pre-image sizes.

4.1 Proof of Theorem 2

Let $F = \{F_k\}_{k \in \mathbb{N}}$ be a family of trapdoor one-way functions with pre-image size bounded by a polynomial Q . The construction is in two steps. We first build an unapproximable family of trapdoor predicates P with decryption error $1/2 - 1/\text{poly}(k)$, and then reduce the decryption error by repetition to get the family claimed in the theorem.

The first step uses the Goldreich-Levin inner-product construction [GoLe]. This construction says that if f is a one-way function, one can securely encrypt a bit b via $f(x), r, \sigma$ where $\sigma = b \oplus (x \odot r)$ with r a random string, $x \in \text{Dom}(f)$, and \odot denoting the inner-product mod 2. Now, if f is an *injective* trapdoor function, then with the trapdoor information, one can recover b from $f(x), r$, and σ by finding x and computing $b = \sigma \oplus (x \odot r)$. If instead f has polynomial-size pre-images, the “correct” x will only be recovered with an inverse polynomial probability. However, we will show that the rest of the time, the success probability is exactly 50%. This gives a noticeable $(\frac{1}{2} + \frac{1}{\text{poly}(k)})$ bias towards the right

value of b . Now, this slight bias needs to be amplified, which is done by repeating the construction many times in parallel and having the decryptor take the majority of its guesses to the bit in the different coordinates. A full description and proof follow.

We may assume wlog that there is a polynomial $l(k)$ such that $\text{Range}(f) \subseteq \{0, 1\}^{l(k)}$ for all $f \in [F_k]$ and all $k \in \mathbb{N}$. We now describe how to use the Goldreich-Levin inner-product construction [GoLe] to build $P = \{P_k\}_{k \in \mathbb{N}}$. We associate to any $f \in [F_k]$ a predicate p defined as follows:

Predicate $p(b)$ // Takes input a bit b
 $x \leftarrow \text{Dom}(f)$ // Choose x at random from the domain of f
 $r \leftarrow \{0, 1\}^{l(k)}$ // Choose a random $l(k)$ -bit string
 $\sigma := b \oplus (x \odot r)$ // XOR b with the GL bit
Output $(f(x), r, \sigma)$

Here \oplus denotes XOR (ie. addition mod 2) and \odot denotes the inner-product mod 2. The generator algorithm for P will choose $(f, tp) \leftarrow F\text{-Gen}(1^k)$ and then output (p, tp) with p defined as above. Notice that p is computable in PPT if f is.

The inversion algorithm $P\text{-Inv}$ is given p , the trapdoor tp , and a triple (y, r, σ) . It first runs the inversion algorithm $F\text{-Inv}$ of F on inputs f, tp, y to obtain x' , and then outputs the bit $b' = \sigma \oplus (x' \odot r)$. It is clear that the inversion algorithm is not always successful, but in the next claim we prove that it is successful appreciably more often than random guessing.

Claim. P is an unapproximable trapdoor predicate family, with decryption error at most $(1/2) - 1/[2Q(k)]$.

Proof. We know that F is one-way. Thus, the inner product is a hardcore bit for F [GoLe]. This implies that P is unapproximable. It is left to show that the decryption error of P is as claimed, namely that $\text{DecErr}_P(P\text{-Inv}, k)$ (as defined in Equation (1)) is at most $(1/2) - 1/[2Q(k)]$.

Fix f, tp, b , let x, r be chosen at random as by $p(b)$, let $y = f(x)$, let $\sigma = b \oplus (x \odot r)$, let $x' \leftarrow F\text{-Inv}(f, tp, y)$, and let $b' = \sigma \oplus (x' \odot r)$. Notice that if $x' = x$ then $b' = b$, but if $x' \neq x$ then the random choice of r guarantees that $b' = b$ with probability at most $1/2$. (Because $F\text{-Inv}$, who generates x' , gets no information about r .) The chance that $x = x'$ is at least $1/Q(k)$ (because $F\text{-Inv}$ gets no information about x other than that $f(x) = y$) so

$$\text{DecErr}_P(P\text{-Inv}, k) \leq \left(1 - \frac{1}{Q(k)}\right) \cdot \frac{1}{2}$$

as desired. □

Now, we can iterate the construction $q(k) \stackrel{\text{def}}{=} \Theta(kQ(k)^2)$ times independently and decrypt via a majority vote to reduce the decryption error to e^{-k} . In more detail, our final predicate family $P^q = \{P_k^q\}_{k \in \mathbb{N}}$ is like this. An instance $p^q \in [P_k^q]$ is still described by a function $f \in [F_k]$ and defined as $p^q(b) = p(b) \parallel \dots \parallel p(b)$, meaning it consists of $q(k)$ repetitions of the original algorithm p on independent

coins. The inversion algorithm $P^q\text{-Inv}$ is given the trapdoor tp and a sequence of triples

$$(y_1, r_1, \sigma_1) \parallel \cdots \parallel (y_{q(k)}, r_{q(k)}, \sigma_{q(k)}) .$$

For $i = 1, \dots, q(k)$ it lets $b'_i = P\text{-Inv}(p, tp, (y_i, r_i, \sigma_i))$. It outputs b' which is 1 if the majority of the values $b'_1, \dots, b'_{q(k)}$ are 1, and 0 otherwise. Chernoff bounds show that $\text{DecErr}_{P^q}(P^q\text{-Inv}, k) \leq e^{-k}$. Furthermore standard “hybrid” arguments [GoMi,Ya] show that P^q inherits the unapproximability of P . This concludes the proof of Theorem 2.

Remark 3. Notice that Theorem 2 holds even if the family F only satisfies a very weak trapdoor property — namely, that $F\text{-Inv}$ produces an element of $f^{-1}(y)$ with probability at least $1/p(k)$ for some polynomial p . Essentially the same proof will show that $P\text{-Inv}$ can guess b correctly with probability at least $1/2 + 1/[2Q(k)p(k)]$.

5 From Cryptosystems to Trapdoor Functions

In this section we investigate the relation between semantically secure public key cryptosystems and injective trapdoor functions. It is known that the existence of unapproximable trapdoor predicates is equivalent to the existence of semantically secure public-key encryption [GoMi]. It is also known that injective trapdoor one-way functions can be used to construct unapproximable trapdoor predicates [Ya] (see also [GoLe]). In this section, we ask whether the converse is true:

Question 1. Can unapproximable trapdoor predicates be used to construct injective trapdoor one-way functions?

Note the importance of the injectiveness condition in Question 1. We already know that non-injective trapdoor functions can be constructed from trapdoor predicates (whether the latter are injective or not) because trapdoor predicates imply one-way functions [ImLu] which in turn imply trapdoor functions by Theorem 1.

We suggest a construction which requires an additional “random looking” function G and prove that the scheme is secure when G is implemented as a random oracle (to which the adversary also has access). Hence, IF it is possible to implement using one-way functions a function G with “sufficiently strong randomness properties” to maintain the security of this scheme, then Question 1 would have a positive answer (as one-way functions can be constructed from unapproximable trapdoor predicates [ImLu]).

The key difference between trapdoor functions and trapdoor predicates is that predicates are *probabilistic*, in that their evaluation is a probabilistic process. Hence, our construction is essentially a de-randomization process.

Suppose we have a family P of unapproximable trapdoor predicates, and we want to construct a family F of injective one-way trapdoor functions from P . A

first approach would be to take an instance p of P and construct an instance f of F as

$$f(b_1 b_2 \cdots b_k \| r_1 \| \cdots \| r_k) = p(b_1; r_1) \| \cdots \| p(b_k; r_k),$$

where k is the security parameter. Standard direct product arguments [Ya] imply that F constructed in this manner is one-way. However, F may fail to be trapdoor; the trapdoor information associated with p only allows one to recover b_1, \dots, b_k , but not r_1, \dots, r_k .

Our approach to fixing this construction is to instead have r_1, \dots, r_k determined by applying some “random-looking” function G to b_1, \dots, b_k :

$$f(b_1 b_2 \cdots b_k) = p(b_1; r_1) \| \cdots \| p(b_k; r_k), \quad \text{where } r_1 \| \cdots \| r_k = G(b_1 \cdots b_k).$$

Since G must be length-increasing, an obvious choice for G is a pseudo-random generator. A somewhat circular intuitive argument can be made for the security of this construction: If one does not know b_1, \dots, b_k , then r_1, \dots, r_k “look random,” and if r_1, \dots, r_k “look random,” then it should be hard to recover b_1, \dots, b_k by the unapproximability of P . In the full version of the paper [BHSV], we show that this argument is in fact false, in that there is a choice of an unapproximable trapdoor predicate P and a pseudorandom generator G for which the resulting scheme is insecure.

However, it is still possible that there are choices of functions G that make the above secure. Below we show that the scheme is secure when G is implemented as a truly random function, ie. a random oracle (to which the adversary also has access). Intuitively, having access to the oracle does not help the adversary recover $b_1 \cdots b_k$ for the following reason: the values of the oracle are irrelevant except at $b_1 \cdots b_k$, as they are just random strings that have nothing to do with $b_1 \cdots b_k$ or $f(b_1 \cdots b_k)$. The adversary’s behavior is independent of the value of the oracle at $b_1 \cdots b_k$ unless the adversary queries the oracle at $b_1 \cdots b_k$. On the other hand, if the adversary queries the oracle at $b_1 \cdots b_k$, it must already “know” $b_1 \cdots b_k$. Specifically, if the adversary queries the oracle at $b_1 \cdots b_k$ with non-negligible probability then it can invert f with non-negligible probability without making the oracle call, by outputting the query. We now proceed with a more formal description of the random oracle model and our result.

THE RANDOM ORACLE MODEL. In any cryptographic scheme which operates in the random oracle model, all parties are given (in addition to their usual resources) the ability to make oracle queries [BeRo]. It is postulated that all oracle queries, independent of the party which makes them, are answered by a single function, denoted \mathcal{O} , which is uniformly selected among all possible functions (where the set of possible functions is determined by the security parameter).

The definitions of families of functions and predicates are adapted to the random oracle model in a straightforward manner: We associate some fixed polynomial Q with each family of functions or predicates, such that on security parameter k all the algorithms in the above definitions are given oracle access to a function $\mathcal{O} : \{0, 1\}^* \rightarrow \{0, 1\}^{Q(k)}$. The probabilities in these definitions are then taken over the randomness of these algorithms and also over the choice of \mathcal{O} uniformly at random among all such functions.

Theorem 3. *If there exists a family of unapproximable trapdoor predicates, then there exists a family of injective trapdoor one-way functions in the random oracle model.*

Remark 4. Theorem 3 still holds even if the hypothesis is weakened to only require the existence of a family of unapproximable trapdoor predicates *in the random oracle model*. To see that this hypothesis is weaker, note that a family of unapproximable trapdoor predicates (in the standard, non-oracle model) remains unapproximable in the random oracle model, as the oracle only provides randomness which the adversary can generate on its own.

See Sections 1.2 and 1.3 for a discussion of the interpretation of such a result. We now proceed to the proof.

5.1 Proof of Theorem 3

Let $P = \{P_k\}_{k \in \mathbb{N}}$ be a family of unapproximable trapdoor predicates. Let $q(k)$ be a polynomial upper bound on the number of random bits used by any $p \in P_k$. When used with security parameter k , we view the oracle as a function $\mathcal{O} : \{0, 1\}^* \rightarrow \{0, 1\}^{kq(k)}$.

We define a family $F = \{F_k\}_{k \in \mathbb{N}}$ of trapdoor functions in the random oracle model as follows: We associate to any $p \in [P_k]$ the function f defined on input $b_1 \dots b_k \in \{0, 1\}^k$ by

$$f(b_1 \dots b_k) = p(b_1; r_1) \parallel \dots \parallel p(b_k; r_k),$$

where

$$r_1 \parallel \dots \parallel r_k = \mathcal{O}(b_1 \dots b_k), \quad r_i \in \{0, 1\}^{q(k)}.$$

The generator $F\text{-Gen}$ takes input 1^k , runs $(p, tp) \leftarrow P\text{-Gen}(1^k)$ and outputs (f, tp) where f is as defined above. It is clear that f can be evaluated in polynomial time using the evaluator $P\text{-Eval}$ for p .

Notice that f can be inverted given the trapdoor information. Given f, tp , and $y_1 \parallel \dots \parallel y_k = f(b_1 \dots b_k)$, inverter $F\text{-Inv}$ computes $b_i = P\text{-Inv}(p, tp, y_i)$ for $i = 1, \dots, k$, and outputs $b_1 \dots b_k$. Furthermore, f is injective because P has zero decryption error: in this inversion process, $P\text{-Inv}$ correctly returns b_i , so we correctly recover the full input. It remains to show that F is one-way.

Claim. F is one-way.

We prove this claim by describing several probabilistic experiments, modifying the role of the oracle with each experiment. The first arises from the definition of a family of one-way functions in the random oracle model. Let A be any PPT, let k be any positive integer, and let $q = q(k)$.

Experiment 1.

- (1) Choose a random oracle $\mathcal{O} : \{0, 1\}^* \rightarrow \{0, 1\}^{kq(k)}$.
- (2) Choose $p \leftarrow P_k$

- (3) Select b_1, \dots, b_k uniformly and independently from $\{0, 1\}$.
- (4) Let $r_1 \| \dots \| r_k = \mathcal{O}(b_1 \dots b_k)$, where $|r_i| = q(k)$ for each i .
- (5) Let $x = p(b_1; r_1) \| \dots \| p(b_k; r_k)$.
- (6) Compute $z \leftarrow A^{\mathcal{O}}(1^k, p, x)$.

We need to prove the following:

Claim. For every PPT A , the probability that $z = b_1 \dots b_k$ in Experiment 1 is a negligible function of k .

To prove Claim 5.1, we first analyze what happens when the r_i 's are chosen independently of the oracle, as in the following experiment: Let A be any PPT, let k be any positive integer, and let $q = q(k)$.

Experiment 2.

- (1)–(3) As in Experiment 1.
- (4) Select r_1, \dots, r_k uniformly and independently from $\{0, 1\}^q$.
- (5)–(6) As in Experiment 1.

Claim. For every PPT A , the probability that $z = b_1 \dots b_k$ in Experiment 2 is a negligible function of k .

Claim 5.1 follows from standard direct product arguments [Ya,GNW]. Specifically, Claim 5.1 is a special case of the uniform complexity version of the Concatenation Lemma in [GNW, Lemma 10].

Claim. For every PPT A , the probability that \mathcal{O} is queried at point $b_1 \dots b_k$ during the execution of $A^{\mathcal{O}}(1^k, p, x)$ in Step 6 of Experiment 2 is a negligible function of k .

Proof. Suppose that the probability that \mathcal{O} is queried at point $b_1 \dots b_k$ was greater than $1/s(k)$ for infinitely many k , where s is a polynomial. Then we could obtain a PPT A' that violates Claim 5.1 as follows. Let $t(k)$ be a polynomial bound on the running time of A . A' does the following on input $(1^k, p, x)$:

- (1) Select i uniformly from $\{1, \dots, t(k)\}$.
- (2) Simulate A on input $(1^k, p, x)$, with the following changes:
 - (1) Replace the oracle responses with strings randomly selected on-line, with the condition that multiple queries at the same point give the same answer.
 - (2) Halt the simulation at the i 'th oracle query and let w be this query.
- (3) Output w .

Then A' , when used in Experiment 2, outputs $b_1 \dots b_k$ with probability greater than $1/(s(k)t(k))$ for infinitely many k , which contradicts Claim 5.1. \square

In order to deduce Claim 5.1 from Claims 5.1 and 5.1, we give an equivalent reformulation of Experiment 1: Let A be any PPT, let k be any positive integer, and let $q = q(k)$.

Experiment 3.

- (1)–(3) As in Experiment 1.
- (4) Select r_1, \dots, r_k uniformly and independently from $\{0, 1\}^q$.
- (5) Let $x = p(b_1; r_1) \parallel \dots \parallel p(b_k; r_k)$.
- (6) Modify \mathcal{O} at location $b_1 \dots b_k$ to have value $r_1 \parallel \dots \parallel r_k$.
- (7) Compute $z \leftarrow A^{\mathcal{O}}(1^k, p, x)$.

We now argue that Experiment 3 is equivalent to Experiment 1. In Experiment 1, r_1, \dots, r_k are uniformly and independently distributed in $\{0, 1\}^q$ and after Step 5 of Experiment 1 the only information about the oracle that has been used is that $r_1 \parallel \dots \parallel r_k = \mathcal{O}(b_1 \dots b_k)$. Thus, the final distribution on all random variables are identical in the two experiments and it suffices to prove Claim 5.1 for Experiment 3 rather than Experiment 1.

Proof. Let E be the event that $z = b_1 \dots b_k$ in Experiment 3. Let F be the event that \mathcal{O} is queried at point $b_1 \dots b_k$ during the execution of $A^{\mathcal{O}}(p, x)$ in Step 7 of Experiment 3. To show that E occurs with negligible probability, it suffices to argue that both F and $E \wedge \overline{F}$ occur with negligible probability.

First we show that F occurs with negligible probability. Notice that whether or not $A^{\mathcal{O}}$ queries \mathcal{O} at $b_1 \dots b_k$ in Experiment 3 will not change if Step 6 is removed. This is because its behavior cannot be affected by the change in $\mathcal{O}(b_1 \dots b_k)$ until it has already queried that position of the oracle. If Step 6 is removed from Experiment 3, we obtain Experiment 2. Hence, the probability of F is negligible by Claim 5.1.

Similarly, the probability that $[z = b_1 \dots b_k$ and $A^{\mathcal{O}}$ never queries the oracle at $b_1 \dots b_k]$ will not change if Step 6 is removed. Thus, the probability of $E \cap \overline{F}$ is bounded above by the probability that $z = b_1 \dots b_k$ in Experiment 2, which is negligible by Claim 5.1. \square

Remark 5. If the family of unapproximable trapdoor predicates we start with has negligible decryption error, then the family of trapdoor functions we construct will in general also have negligible decryption error and may fail to be injective with some small probability.

By first reducing the decryption error of the predicate family to $\exp(-\Omega(k^3))$ as in the proof of Theorem 2 and then using the oracle to derandomize the inversion algorithm, one can produce an *injective* family that has *zero* decryption error with probability $1 - 2^{-k}$ (where the probability is just taken over the choice of the oracle).

Acknowledgments

The first author was supported by a 1996 Packard Foundation Fellowship in Science and Engineering, and by NSF CAREER Award CCR-9624439. The third and fourth authors were supported by DOD/NDSEG Graduate Fellowships and partially by DARPA grant DABT-96-C-0018.

The starting point of this research was a question posed to us by Shafi Goldwasser, namely whether trapdoor permutations could be built from the assumptions underlying the Ajtai-Dwork cryptosystem.

Thanks to Oded Goldreich and the members of the Crypto 98 program committee for their comments on the paper.

References

- AjDw. M. AJTAI AND C. DWORK. A public-key cryptosystem with worst-case / average-case equivalence. *Proceedings of the 29th Annual Symposium on the Theory of Computing*, ACM, 1997. 287
- AMM. ADLEMAN, MANDERS AND MILLER. On taking roots in finite fields. *Proceedings of the 18th Symposium on Foundations of Computer Science*, IEEE, 1977.
- BHSV. M. BELLARE, S. HALEVI, A. SAHAI, AND S. VADHAN. Many-to-one trapdoor functions and their relation to public-key cryptosystems. Full version of this paper, available via <http://www-cse.ucsd.edu/users/mihir>.
- BeRo. M. BELLARE AND P. ROGAWAY. Random oracles are practical: a paradigm for designing efficient protocols. *Proceedings of the First Annual Conference on Computer and Communications Security*, ACM, 1993. 286, 287, 294
- Be. E. BERLEKAMP. Factoring polynomials over large finite fields. *Mathematics of Computation*, Vol. 24, 1970, pp. 713–735.
- BIMi. M. BLUM AND S. MICALI. How to generate cryptographically strong sequences of pseudo-random bits, *SIAM Journal on Computing*, Vol. 13, No. 4, 850-864, November 1984. 286
- Ca. R. CANETTI. Towards realizing random oracles: Hash functions that hide all partial information. *Advances in Cryptology – Crypto 97 Proceedings*, Lecture Notes in Computer Science Vol. 1294, B. Kaliski ed., Springer-Verlag, 1997. 286
- CGH. R. CANETTI, O. GOLDREICH AND S. HALEVI. The random oracle model, revisited. *Proceedings of the 30th Annual Symposium on the Theory of Computing*, ACM, 1998. 286
- DiHe. W. DIFFIE AND M. HELLMAN. New directions in cryptography. *IEEE Trans. Info. Theory*, Vol. IT-22, No. 6, November 1976, pp. 644–654. 284, 284, 284, 287
- DDN. D. DOLEV, C. DWORK, AND M. NAOR. Non-Malleable Cryptography. *Proceedings of the 23rd Annual Symposium on the Theory of Computing*, ACM, 1991.
- ElG. T. EL GAMAL. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Trans. Inform. Theory*, Vol. 31, 1985, pp. 469–472. 287
- GoLe. O. GOLDREICH AND L. LEVIN. A hard predicate for all one-way functions. *Proceedings of the 21st Annual Symposium on the Theory of Computing*, ACM, 1989. 291, 292, 292, 293
- GoMi. S. GOLDWASSER AND S. MICALI. Probabilistic Encryption. *Journal of Computer and System Sciences*, Vol. 28, April 1984, pp. 270–299. 284, 284, 284, 284, 286, 286, 287, 289, 289, 291, 293, 293
- GNW. O. GOLDREICH, N. NISAN, AND A. WIGDERSON. On Yao’s XOR Lemma. *Electronic Colloquium on Computational Complexity*, TR95-050. March 1995. <http://www.eccc.uni-trier.de/eccc/> 296, 296

- HILL. J. HÅSTAD, R. IMPAGLIAZZO, L. LEVIN AND M. LUBY. Construction of a pseudo-random generator from any one-way function. Manuscript. Earlier versions in STOC 89 and STOC 90.
- ImLu. R. IMPAGLIAZZO AND M. LUBY. One-way Functions are Essential for Complexity-Based Cryptography. *Proceedings of the 30th Symposium on Foundations of Computer Science*, IEEE, 1989. 293, 293
- ImRu. R. IMPAGLIAZZO AND S. RUDICH. Limits on the provable consequences of one-way permutations. *Proceedings of the 21st Annual Symposium on the Theory of Computing*, ACM, 1989. 285, 285, 286, 287, 287, 291, 291
- NaYu. M. NAOR AND M. YUNG. Public-Key Cryptosystems Provably Secure against Chosen Ciphertext Attacks. *Proceedings of the 22nd Annual Symposium on the Theory of Computing*, ACM, 1990.
- Rab. M. RABIN. Digitalized Signatures and Public Key Functions as Intractable as Factoring. *MIT/LCS/TR-212*, 1979. 289
- Ya. A. YAO. Theory and applications of trapdoor functions. *Proceedings of the 23rd Symposium on Foundations of Computer Science*, IEEE, 1982. 284, 286, 291, 293, 293, 294, 296