

Lecture 4

Lecturer: Yevgeniy Dodis

Scribe: David C. Kandathil

This lecture will study the notion of *hardcore bit* for a given OWF f . Intuitively, such a hardcore bit $h(x)$ is easy to compute from x , but almost impossible to even *guess* well from $f(x)$. We will see specific examples of hardcore bits for modular exponentiation, RSA and Rabin's squaring function. Next we will show a groundbreaking result of Goldreich-Levin, that (more or less) shows a general hardcore bit for *any* OWF. We will then consider two natural applications of hardcore bits to the problem of encryption. Firstly, we will show an intuitively good public-key encryption for one bit, and then a “plausible” secret-key encryption which encrypts $k+1$ bits with k -bit key (thus beating the Shannon information-theoretic bound). We will then try to extend the hardcore bit construction to extracting many “pseudorandom bits”, but analogy to the S/Key system. We will notice that our many-bit construction seems to satisfy a very special notion of security, which we call “next-bit unpredictability”. We then will make a formal definition of a *pseudorandom generator*, which seems to be a more relevant primitive for our encryption applications. We stop by asking the question of whether our “next-bit secure” construction is indeed a pseudorandom generator.

1 HARDCORE BITS

Last lecture we addressed some of the criticism over straightforward usages of OWF's, OWP's and TDP's. Specifically, our main criticism was the fact that a OWF $f(x)$ could reveal a lot of partial information about x (remember generic example $f(x_1, x_2) = (f'(x_1), x_2)$ that reveals half of its input bits, or more realistic one of exponentiation $f(x) = g^x \bmod p$ that reveals the $LSB(x) = y^{(p-1)/2}$).

The obvious solution seemed to try to *completely* hide all information about x , given $f(x)$. However, this leads to a vicious circle, since it is really equivalent to the problem of secure encryption that we started from. Instead, we explore the idea of *completely* hiding not *all*, but only a *specific and carefully chosen partial information about x* , when given $f(x)$. The first preliminary step towards this goal is to determine how to completely hide exactly just *one* bit of information about the plaintext x . This leads us to the following definition:

Definition 1 (Hardcore bit) A function $h : \{0, 1\}^* \rightarrow \{0, 1\}$ is called a *hardcore bit* for a function f if

- $h(x)$ is polynomial time computable (from x):

$$(\exists \text{ poly-time } H)(\forall x)[H(x) = h(x)]$$

- No PPT algorithm that can predict $h(x)$ given $f(x)$ better than flipping a coin:

$$(\forall \text{ PPT } A) \Pr[A(f(x)) = h(x) \mid x \leftarrow^r \{0, 1\}^k] \leq \frac{1}{2} + \text{negl}(k)$$

Remark 2 Notice, we compare the success of A to a $\frac{1}{2} + \text{negl}(k)$ rather than $\text{negl}(k)$, as we did for OWF's. The reason is that the output of h is now only one bit, so A can always guess if with probability $\frac{1}{2}$ by flipping a coin. Of course, we could have said that h is difficult to compute by saying that A succeeds with probability at most $1 - \epsilon$, where ϵ is non-negligible. However, we really want to say much more (and that is why we started with $h(x)$ being just 1 bit for now): not only is h hard to compute, it is even hard to predict.

Thus, a hardcore bit $h(x)$ pinpoints an aspect of x that is truly hidden given $f(x)$. Namely, the knowledge of $f(x)$ does not allow us to predict $h(x)$ any better than *without it* (i.e., by flipping a coin), so $h(x)$ looks random given $f(x)$. It is to be noted that we need not require that $h(x)$ be a bit that is selected from the string x itself, but, in general, may depend on x in more complex (but efficiently computable) ways. This is not inconsistent with the idea that $h(x)$ is supposed to represent some general information about x . We might want to attempt the construction in two ways:

1. Taking as hypothesis that a *concrete* function is OWF, exhibit a hardcore bit for that function. (This is a useful, but not very general construction.)
2. Taking as hypothesis that an *arbitrary* function is OWF, exhibit a hardcore bit for that function. (This is the strongest construction we can hope for.)

It is to be noted that these approaches are analogous to those we adopted for constructing what might, antonymously, be called “softcore” bits.

2 HARDCORE BITS FOR CONCRETE OWF'S

The concrete function that we consider is the exponentiation function mod p , $f(x) = y = g^x \bmod p$, where g is the generator of \mathbb{Z}_p^* . We define the most significant bit of x , MSB , as follows:

$$MSB(x) = \begin{cases} 0, & \text{if } x < \frac{p-1}{2} \\ 1, & \text{if } x \geq \frac{p-1}{2} \end{cases}$$

Remark 3 $MSB(x)$ not defined to be simply x_1 in order to make it unbiased, since the prime p is not a perfect power of 2.

Theorem 4 If $f(x) = (g^x \bmod p)$ is a OWF, then $MSB(x)$ is a hardcore bit for f .

Proof: A rigorous proof for the above theorem exists. It is our usual proof by contradiction, which takes as hypothesis that MSB is not an hardcore bit for f , and proves that f is not OWF. The proof is constructive, and explicitly transforms any PPT that can compute $MSB(x)$ from $f(x)$ with non-negligible advantage, into a PPT that can compute the Discrete Log with non-negligible probability. This proof is, however, somewhat technical, so we settled for a simpler, but quite representative result stated below. \square

Lemma 5 If there exists a PPT that can always compute $MSB(x)$ from $f(x)$, then there is a PPT that can always invert $f(x)$ (i.e., compute the discrete log of $y = f(x)$).

Proof: The idea of the algorithm is very simple. We know that $LSB(x) = x_k$ is easy to compute given $y = g^x \bmod p$. This way we determine x_k . Now we can transform y into $g^{[x_1 \dots x_{k-1} 0]} = (g^{[x_1 \dots x_{k-1}]})^2 \bmod p$ (by dividing y by g if $x_k = 1$). We also know how to extract square roots modulo p . So it seems like we can compute $g^{[x_1 \dots x_{k-1}]}$, and keep going the same way (take LSB , extract square root, etc.) until we get all the bits of x . However, there is a problem. The problem is that $g^{[x_1 \dots x_{k-1} 0]}$ has two square roots: $y_0 = g^{[x_1 \dots x_{k-1}]}$ (the one we want) and $y_1 = (-g^{[x_1 \dots x_{k-1}]}) = g^{\frac{p-1}{2} + [x_1 \dots x_{k-1}]}$ (here we used the fact that $-1 = g^{(p-1)/2}$). So after we compute the square roots y_0 and y_1 , how do we know which root is really $y_0 = g^{[x_1 \dots x_{k-1}]}$? Well, this is exactly what the hardcore bit MSB tells us! Namely, $MSB(Dlog(y_0)) = 0$ and $MSB(Dlog(y_1)) = 1$. The complete algorithm follows.

```

 $i = k;$ 
while ( $y \geq 1$ ) do /*  $y = f(x)$  */
begin
    output (  $x_i = LSB(Dlog(y))$  );
    /* Assertion:  $x = [x_1 x_2 \dots x_i]$  */
    if (  $x_i == 1$  ) then  $y := y/g$ ;
    /* Assertion:  $y = g^{[x_1 x_2 \dots x_{i-1} 0]} = (g^{[x_1 x_2 \dots x_{i-1}]})^2$  */
    Let  $y_0$  and  $y_1$  be square roots of  $y$ ;
    If (  $MSB((Dlog(y_1)) == 0$  )
        then  $y := y_1$ 
        else  $y := y_2$ 
     $i := i - 1$ ;
end

```

To summarize, the value of MSB plays a critical role in distinguishing which square root of y corresponds to $x/2$. This enables us to use the LSB iteratively, so that the process is continued to extract all the bits of x . \square

It turns out that the other OWF's we study have natural hardcore bits as well:

1. LSB and MSB are hardcore bits for Rabin's Squaring Function.
2. All the bits of x are hardcore bits for RSA .

3 CONSTRUCTION OF A HARDCORE BIT FOR ARBITRARY OWF

Looking at the previous examples, we would now like to see if any OWF f has some easy and natural hardcore bits. Specifically, it would be great if at least one the following two statements was true:

1. Any OWF has some particular bit x_i in x (i depends on f) which is hardcore for f .
2. A concrete boolean function h (which is not necessarily an input bit) is a hardcore bit for all OWF's f .

Unfortunately, both of these hopes are false in general.

1. From arbitrary OWF f , it is possible to construct another OWF g , such that none of the bits of x are hardcore for g . (cf. Handout).
2. For any boolean function h and OWF f , if we let $g(x) = f(x) \circ h(x)$, then: (1) g is also a OWF; (2) h is not hardcore for g . Part (1) follows from the fact that an inverter A for g would imply the one for f . Indeed, given $y = f(x)$, we can ask the inverter A for g to invert both $f(x) \circ 0$ and $f(x) \circ 1$, and see if at least one of them succeeds. Part (2) is obvious. Thus, no “universal” h exists.

Despite these negative news, it turns out that we nevertheless have a very simple hardcore bit for an arbitrary OWF. This is the celebrated Goldreich Levin construction.

4 GOLDREICH LEVIN CONSTRUCTION

We will begin with a definition that generalizes the concept of selecting a specific bit from a binary string.

Definition 6 (Parity) *If $x = x_1x_2\dots x_k \in \{0,1\}^k$, and $r = r_1r_2\dots r_k \in \{0,1\}^k$, then $h(x, r) = r_1x_1 \oplus r_2x_2 \dots \oplus r_kx_k = (r_1x_1 + r_2x_2 + \dots + r_kx_k \bmod 2)$ is called the parity of x with respect to r .*

Notice, r can be viewed as a selector for the bits of x to be included in the computation of parity. Further, the expression for the notation for the inner product, \cdot , can be used profitably, i.e., $h(x, r) = (r \cdot x)$ can be viewed as the inner product of binary vectors x and r modulo 2. The “basis strings” e_i with exactly one bit $r_i = 1$, give the usual specific bit selection x_i .

The Goldreich-Levin theorem essentially says that “if f is a OWF, then most parity functions $h(x, r)$ are hardcore bits for f .” To make the above statement more precise, it is convenient to introduce an auxiliary function $g_f(x, r) = f(x) \circ r$ (the concatenation of $f(x)$ and r), where $|x| = |r|$. Notice, a random input for g_f samples both r and x at random from $\{0,1\}^k$, so $h(x, r) = x \cdot r$ indeed computes a “random parity for (randomly selected) x ”. Notice also, that if f is a OWF/OWP/TDP, then so is g_f (in particular, g_f is a permutation if f is). Now, the Goldreich-Levin theorem states that

Theorem 7 (Goldreich-Levin Bit) *f is a OWF, then $h(x, r)$ is a hardcore bit for g_f . More formally:*

$$(\forall \text{ PPT } A) \quad \Pr[A(f(x), r) = (x \cdot r) \mid x, r \leftarrow^r \{0,1\}^k] < \frac{1}{2} + \text{negl}(k)$$

Remark 8 *A hardcore bit for g_f is as useful as the one for f , since f is really computationally equivalent to g_f (since r is part of the output, inverting g_f exactly reduces to inverting f). Thus, we will often abuse the terminology and say “since f is a OWF, let us take its hardcore bit $h(x)$ ”. But that we really mean that in case we are not aware of some simple hardcore bit for a specific f , we can always take the Goldreich-Levin bit for g_f and use it instead. Similar parsing should be given to statement of the form “every OWF has a hardcore bit”. Again, in the worst case always use the Goldreich-Levin bit for g_f . Finally, another popular interpretation of this result is that “most parities of f are hardcore”.*

Proof: A rigorous proof for the above theorem exists. It is an indirect proof by contradiction, which takes as hypothesis that $h(x, r)$ is *not* a hardcore bit for g_f , and proves that f is not a OWF. The proof is constructive, and explicitly transforms any PPT that can compute $h(x, r)$ with non negligible probability for most values of r , given $f(x)$ and r , into a PPT that can compute x with non negligible probability, given $f(x)$. However and despite the simplicity of theorem statement, the full proof is extremely technical. Therefore, we will again only give a good intuition of why it works. For simplicity, we will assume that f (and thus g_f) are *permutations*, so that $(x \cdot r)$ is uniquely defined given $f(x) \circ r$.

First, assume that we are given a PPT A that *always* computes $(x \cdot r)$ given $f(r) \circ r$. Well, then everything is extremely simple. Given $y = f(x)$ that we need to invert, we already observed that $x \cdot e_i = x_i$ is the i -th bit of x , where e_i is a vector with 1 only at position i . Thus, asking $A(y, e_i)$ will give us x_i , so we can perfectly learn x bit by bit.

Unfortunately, our assumption on A is too strong. In reality we do not know that it succeeds for *random* r 's. In particular, maybe it always refuses to work for “basis” $r = e_i$. So let us be more reasonable and assume that

$$\Pr[A(f(x), r) = (x \cdot r) \mid x, r \leftarrow^r \{0, 1\}^k] > \frac{3}{4} + \epsilon$$

where ϵ is non-negligible. Here we use $3/4$ instead of $1/2$ for the reason to be clear in a second. Still we cannot ask $r = e_i$ even in this case. The idea is to notice that for any $r \in \{0, 1\}^k$

$$(x \cdot r) \oplus (x \cdot (r \oplus e_i)) = \left(\sum_{j \neq i} r_j x_j + r_i x_i \bmod 2 \right) \oplus \left(\sum_{j \neq i} r_j x_j + (1 - r_i)x_i \bmod 2 \right) = x_i$$

Moreover both r and $(r \oplus e_i)$ are *individually random* when r is chosen at random. Hence, for any fixed index i ,

$$\begin{aligned} \Pr[A(y, r) \neq (x \cdot r) \mid x, r \leftarrow^r \{0, 1\}^k] &< \frac{1}{4} - \epsilon \\ \Pr[A(y, r \oplus e_i) \neq (x \cdot (r \oplus e_i)) \mid x, r \leftarrow^r \{0, 1\}^k] &< \frac{1}{4} - \epsilon \end{aligned}$$

Thus, with probability at least $1 - 2(\frac{1}{4} - \epsilon) = \frac{1}{2} + 2\epsilon$, A will be correct in both cases, and hence we correctly recover x_i with probability $\frac{1}{2} + 2\epsilon$. Well, this probability is greater than a half, so repeating the same experiment roughly $t = O(\log k / \epsilon^2)$ times (each time picking a brand new r ; notice also that t is polynomial in k by assumption on ϵ) and taking the majority of the answers, we determine each x_i correctly with probability $1 - 1/k^2$. We now repeat the whole procedure for all indices i . We get that the probability that *at least one x_i is wrong* is at most $k/k^2 = 1/k$, so we recover the entire x correctly with probability $1 - 1/k$, which is certainly non-negligible, contradicting the one-wayness of f .

Still, our assumption about the success of A with probability $\frac{3}{4} + \epsilon$ is too much. We can only assume $\frac{1}{2} + \epsilon$. As we said, a very careful choice of random r 's allows us to extend the proof above even to this case, but this is too technical. \square

We will now look at how to make use of what we have at hand. We will not be terribly rigorous for the time being, and will proceed with the understanding that speculative adventures are acceptable. This time we will hand-wave a little, but will return to the problematic sections in the next lecture.

5 PUBLIC KEY CRYPTOSYSTEM FOR ONE BIT

It seems intuitive that, if we have a hardcore bit, we *should* be able to send *one* bit of information with complete security in the Public Key setting. And we show exactly that.

- **Scenario**

Bob(B) wants to send a bit b to Alice(A). Eve(E) tries to get b . Alice has a public key PK and a secret key SK hidden from everybody.

- **Required Primitives**

1. TDP f will be the public key PK and its trapdoor information t will be Alice's secret key SK .
2. Hardcore bit h for f . If needed, can apply Goldreich-Levin to get it.

- **Protocol**

B selects a random $x \in \{0, 1\}^k$ and sends A the ciphertext $c = \langle f(x), h(x) \oplus b \rangle$.

- **Knowledge of the Concerned Parties before Decryption**

B : b, x, c, f, h .

E : c, f, h .

A : c, t, f, h .

- **Decryption by A**

x is obtained from $f(x)$ using the trapdoor t ;

$h(x)$ is computed from x ;

b is obtained from $(h(x) \oplus b)$ using $h(x)$.

- **Security from E**

Intuitively, to learn anything about b , E must learn something about $h(x)$. But E only knows $f(x)$. Since h is hardcore, E cannot predict $h(x)$ better than flipping a coin, so b is completely hidden.

We note that this scheme is grossly inefficient. Even though it *seems* like we are using an elephant to kill an ant, look what we accomplished: we constructed the first secure public-key cryptosystem!

Having said this, we would still like to improve the efficiency of this scheme. Can we send arbitrary number of bits by using a single x above? More generally, can we extract a lot of random looking bits from a single x ?

6 PUBLIC KEY CRYPTOSYSTEM FOR ARBITRARY NUMBER OF BITS

We will try to generalize the system in a manner analogous to what we did with the S/Key system. Remember, there we published the value $y_0 = f^T(x)$, and kept giving the server the successive preimages of y_0 . So maybe we can do the same thing here, except we will use *hardcore bits of successive preimages to make them into a good one-time pad!* Notice also that publishing $f^t(x)$ will still allow Alice to get back all the way to x since she has the trapdoor.

- **Scenario**

Bob(B) wants to send a string $m = m_1 \dots m_n$ to Alice(A). Eve(E) tries to get “some information” about m . Alice has a public key PK and a secret key SK hidden from everybody.

- **Required Primitives**

1. As before, TDP f will be the public key PK and its trapdoor information t will be Alice’s secret key SK .
2. Hardcore bit h for f . If needed, can apply Goldreich-Levin to get it.

- **Protocol**

B selects a random $x \in \{0, 1\}^k$ and sends A the ciphertext $c = \langle f^n(x), G'(x) \oplus m \rangle$, where

$$G'(x) = h(f^{n-1}(x)) \circ h(f^{n-2}(x)) \circ \dots \circ h(x) \quad (1)$$

Notice, $G'(x)$ really serves as a “computational one-time pad”.

- **Knowledge of the Concerned Parties before Decryption**

B : m, x, c, f, h .

E : c, f, h .

A : c, t, f, h .

- **Decryption by A**

x is obtained from $f^n(x)$ using the trapdoor t by going “backwards” n times;

$G'(x)$ is computed from x by using h and f in the “forward” direction n times;

m is obtained from $(G'(x) \oplus m)$ using $G'(x)$.

- **Security from E**

The intuition about the security is no longer that straightforward. Intuitively, if we let $p_1 = h(f^{n-1}(x)), \dots, p_n = h(x)$ be the one-time pad bits, it seems like p_1 looks random given $f^n(x)$, so m_1 is secure for now. On the other hand, $p_2 = h(f^{n-2}(x))$ looks secure even given $f^n(x)$ and p_1 (since both can be computed from $f^{n-1}(x)$ and p_2 is secure even if $f^{n-1}(x)$ is completely known). And so on. So we get a very strange kind of “security”. Even if the adversary knows $f^n(x)$ (which he does as it is part

of c), and even if he somehow learns m_1, \dots, m_{i-1} , which would give it p_1, \dots, p_{i-1} , he still cannot predict p_i , and therefore, m_i is still secure. So we get this “next-bit security”: given first $(i - 1)$ bits, E cannot predict the i -th one. It is completely unclear if

- “Next-bit security” is what we really want from a good encryption (we certainly want at least this security, but does suffice?) For example, does our system satisfy analogously defined “previous-bit security”?
- Our scheme satisfies some more “reasonable” notion of security.

The answers to these questions will come soon.

At this point, we turn our attention to Secret Key Cryptography based on symmetric keys. We would like to contemplate whether we could transcend Shannon’s Theorem on key lengths.

7 SECRET KEY CRYPTOSYSTEMS

Recall, our main question in secret key encryption was to break the Shannon bound. Namely, we would like to encrypt a message of length n using a secret key of a much smaller size k , i.e. $k < n$ (and hopefully, $k \ll n$). We right away propose a possible solution by looking at the corresponding public-key example for encrypting many bits.

Recall, in the public key setting we used $G'(x)$ (see Equation (1)) as a “computational one-time pad” for m , where x was chosen at random by Bob. Well, now we can do the same thing, but make x the shared secret! Notice also that now we no longer need the trapdoor, so making f OWP suffices.

- **Scenario**

Bob(B) wants to send a string $m = m_1 \dots m_n$ to Alice(A). Eve(E) tries to get “some information” about m . Alice and Bob share a random k -bit key x which is hidden from Eve.

- **Required Primitives**

1. OWP f which is known to everyone.
2. Hardcore bit h for f . If needed, can apply Goldreich-Levin to get it.

- **Protocol**

B sends A the ciphertext $c = G'(x) \oplus m$, where $G'(x)$ is same as in Equation (1).

$$G'(x) = h(f^{n-1}(x)) \circ h(f^{n-2}(x)) \circ \dots \circ h(x) \quad (2)$$

As before, $G'(x)$ really serves as a “computational one-time pad”.

- **Knowledge of the Concerned Parties before Decryption**

B: m, x, c, f, h .

E: c, f, h .

A: c, x, f, h .

- **Decryption by A**

$G'(x)$ is computed from secret key x by using h and f in the “forward” direction n times;

m is obtained from $(G'(x) \oplus m)$ using $G'(x)$.

- **Security from E**

Again, the intuition about the security is not straightforward. Similar to the public-key example, it seems like we get what we called “next-bit security”: given first $(i - 1)$ bits of m (or of $G'(x)$), *E* cannot predict the i -th bit of m (or $G'(x)$). It is completely unclear if

- “Next-bit security” is what we really want from a good encryption (we certainly want at least this security, but does suffice?) For example, does our system satisfy analogously defined “previous-bit security”?
- Our scheme satisfies some more “reasonable” notion of security.

Again, the answers to these questions will come soon,

Before moving on, we also make several more observations. First, notice that we could make n very large (in particular, much larger than k). Also, we can make the following optimization. Recall that in the public key scenario we also send $f^n(x)$ to Alice so that she can recover x . Now, it seems like there is no natural way to use it, so we really computed it almost for nothing... But wait, we could use $f^n(x)$ to make our one-time pad longer! Namely, define

$$G(x) = f^n(x) \circ G'(x) = f^n(x) \circ h(f^{n-1}(x)) \circ h(f^{n-2}(x)) \circ \dots \circ h(x) \quad (3)$$

Now we can use $G(x)$ as the one-time pad for messages of length $n + k$, which is *always* greater than our key size k , even if $n = 1$! Indeed, we claim that our intuitively defined “next-bit security” holds for $G(x)$ as well. Indeed, for $i < k$, $f^n(x)$ is *completely random* (since x is random), so predicting the i -th bit based on the first $(i - 1)$ bits is hopeless. While for $i > k$ our informal argument anyway assumed that Eve knows $f^n(x)$ (it was part of the encryption). We will discuss later if it really pays off in the long run to use this efficiency improvement (can you think of a reason why it might be good to keep the same x for encrypting more than one message?).

However, using either $G(x)$ or $G'(x)$ as our one-time pads still has its problems that we mentioned above. Intuitively, what we really want from a “computational one-time pad” is that it really *looks completely random to Eve*. We now formalize what it means, by defining an extremely important concept of a *pseudorandom number generator*.

8 PSEUDO RANDOM GENERATORS

Intuitively, a *pseudorandom number generator* (PRG) stretches a short random seed $x \in \{0, 1\}^k$ into a longer output $G(x)$ of length $p(k) > k$ which nevertheless “looks” like a random $p(k)$ -bit strings to any computationally bounded adversary. For clear reasons, the adversary here is called a *distinguisher*.

Definition 9 (Pseudo Random Generator) A deterministic polynomial-time computable function $G : \{0, 1\}^k \rightarrow \{0, 1\}^{p(k)}$ (defined for all $k > 0$) is called a pseudorandom number generator (PRG) if

1. $p(k) > k$ (it should be stretching).
2. There exists no PPT distinguishing algorithm D which can tell $G(x)$ apart from a truly random string $R \in \{0, 1\}^{p(k)}$. To define this formally, let 1 encode “pseudorandom” and 0 encode “random”. Now we say that for any PPT D

$$|\Pr(D(G(x)) = 1 | x \leftarrow^r \{0, 1\}^k) - \Pr(D(R) = 1 | R \leftarrow^r \{0, 1\}^{p(k)})| < \text{negl}(k)$$

We observe that we require the length of x be less than the output length of $G(x)$. This is done since otherwise an identity function will be a trivial (and useless) PRG. It should not be that easy! On the other hand, requiring $p(k) > k$ makes this cryptographic primitive quite non-primitive to construct (no pun intended).

Secondly, we are not creating pseudorandomness from the thin air. We are taking a *truly random seed* x , and stretch it to “computationally random” output $G(x)$. In other words, $G(x)$ is computationally indistinguishable from a random sequence (i.e., looks random), only provided that (much shorter seed) x is random.

9 POINTS TO PONDER

We would like to conclude with some questions which are food for thought.

1. Is our public-key encryption really good?
2. What about the secret-key encryption?
3. Are $G'(x)$ and $G(x)$ (see Equations (2) and (3)) pseudorandom generators?
4. Can we output the bits of $G'(x)$ in “forward” order?
5. Is “next-bit security” enough to imply a true PRG?