

# 5

## Pseudorandom Generators

One-time pad requires a key that's as long as the plaintext. Let's forget that we know about this limitation. Suppose Alice & Bob share only a short  $\lambda$ -bit secret  $k$ , but they want to encrypt a  $2\lambda$ -bit plaintext  $m$ . They don't know that (perfect) one-time secrecy is impossible in this setting ([Exercise 2.11](#)), so they try to get it to work anyway using the following reasoning:

- ▶ The only encryption scheme they know about is one-time pad, so they decide that the ciphertext will have the form  $c = m \oplus ??$ . This means that the unknown value  $??$  must be  $2\lambda$  bits long.
- ▶ In order for the security of one-time pad to apply, the unknown value  $??$  should be uniformly distributed.
- ▶ The process of obtaining the unknown value  $??$  from the shared key  $k$  should be *deterministic*, so that the sender and receiver compute the same value and decryption works correctly.

Let  $G$  denote the process that transforms the key  $k$  into this mystery value. Then  $G : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{2\lambda}$ , and the encryption scheme is  $\text{Enc}(k, m) = m \oplus G(k)$ .


It is not hard to see that if  $G$  is a deterministic function, then there are only  $2^{2\lambda}$  possible outputs of  $G$ , so the distribution of  $G(k)$  cannot be uniform in  $\{0, 1\}^{2\lambda}$ . We therefore cannot argue that the scheme is secure in the same way as one-time pad.

However, what if the distribution of  $G(k)$  values is not perfectly uniform but only "close enough" to uniform? Suppose no polynomial-time algorithm can distinguish the distribution of  $G(k)$  values from the uniform distribution. Then surely this ought to be "close enough" to uniform for practical purposes. This is exactly the idea of **pseudorandomness**. It turns out that if  $G$  has a pseudorandomness property, then the encryption scheme described above is actually secure (against polynomial-time adversaries, in the sense discussed in the previous chapter).

### 5.1 Definitions


A **pseudorandom generator (PRG)** is a deterministic function  $G$  whose outputs are longer than its inputs. When the input to  $G$  is chosen uniformly at random, it induces a certain distribution over the possible output. As discussed above, this output distribution cannot be uniform. However, the distribution is *pseudorandom* if it is **indistinguishable from the uniform distribution**. More formally:

Definition 5.1 (PRG security) Let  $G : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{\lambda+\ell}$  be a deterministic function with  $\ell > 0$ . We say that  $G$  is a **secure pseudorandom generator (PRG)** if  $\mathcal{L}_{\text{prg-real}}^G \approx \mathcal{L}_{\text{prg-rand}}^G$ , where:



The value  $\ell$  is called the **stretch** of the PRG. The input to the PRG is typically called a **seed**.

Below is an illustration of the distributions sampled by these libraries, for a **length-doubling** ( $\ell = \lambda$ ) PRG (not drawn to scale):



$\mathcal{L}_{\text{prg-real}}$  samples from distribution of red dots, by first sampling a uniform element of  $\{0, 1\}^\lambda$  and performing the action of  $G$  on that value to get a red result in  $\{0, 1\}^{2\lambda}$ . The other library  $\mathcal{L}_{\text{prg-rand}}$  directly samples the uniform distribution on  $\{0, 1\}^{2\lambda}$  (in green above).

To understand PRGs, you must simultaneously appreciate two ways to compare the PRG's output distribution with the uniform distribution:

- ▶ From a *relative* perspective, the PRG's output distribution is tiny. Out of the  $2^{2\lambda}$  strings in  $\{0, 1\}^{2\lambda}$ , only  $2^\lambda$  are possible outputs of  $G$ . These strings make up a  $2^\lambda / 2^{2\lambda} = 1/2^\lambda$  fraction of  $\{0, 1\}^{2\lambda}$  — a **negligible fraction!**
- ▶ From an *absolute* perspective, the PRG's output distribution is huge. There are  $2^\lambda$  possible outputs of  $G$ , which is an **exponential amount!**

The illustration above only captures the *relative* perspective (comparing the red dots to the entire extent of  $\{0, 1\}^{2\lambda}$ ), so it can lead to some misunderstanding. Just looking at this picture, it is hard to imagine how the two distributions could be indistinguishable. How could a calling program could *not* notice whether it's seeing the whole set or just a negligible fraction of the whole set? Well, if you run in polynomial-time in  $\lambda$ , then  $2^\lambda$  and  $2^{2\lambda}$  are both so enormous that it doesn't really matter that one is vastly bigger than the other. The relative *sizes* of the distribution don't really help distinguish, since it is not a viable strategy for the distinguisher to "measure" the size of the distribution it's sampling.

Consider: there are about  $2^{75}$  molecules in a teaspoon of water, and about  $2^{2 \cdot 75}$  molecules of water in Earth’s oceans. Suppose you dump a teaspoon of water into the ocean and let things mix for a few thousand years. Even though the teaspoon accounts for only  $1/2^{75}$  of the ocean’s contents, that doesn’t make it easy to keep track of all  $2^{75}$  water molecules that originated in the teaspoon! If you are small enough to see individual water molecules, then a teaspoon of water looks as big as the ocean.

### Discussion & Pitfalls

- ▶ Do not confuse the interface of a PRG (it takes in a seed as input) with the interface of the security libraries  $\mathcal{L}_{\text{prg-}\star}$  (their `QUERY` subroutine doesn’t take any input)! A PRG is indeed an algorithm into which you can feed any string you like. However, **security is only guaranteed** when the PRG is being used exactly the way that is described in the security libraries — in particular, when the seed is chosen uniformly/secretly and not used for anything else.

Nothing prevents a user from putting an adversarially-chosen  $s$  into a PRG, or revealing a PRG seed to an adversary, etc. You just get no security guarantee from doing it, since it’s not the situation reflected in the PRG security libraries.

- ▶ It doesn’t really make sense to say that “0010110110 a random string” or “0000000001 is a pseudorandom string.” Randomness and pseudorandomness are **properties of the process used to generate a string**, not properties of the individual strings themselves. When we have a value  $z = G(s)$  where  $G$  is a PRG and  $s$  is chosen uniformly, you could say that  $z$  was “chosen pseudorandomly.” You could say that the output of some process is a “pseudorandom distribution.” But it is slightly sloppy (although common) to say that a string  $z$  “is pseudorandom”.
- ▶ There are common statistical tests you can run, which check whether some data has various properties that you would expect from a uniform distribution.<sup>1</sup> For example, are there roughly an equal number of 0s and 1s? Does the substring 01010 occur with roughly the frequency I would expect? If I interpret the string as a series of points in the unit square  $[0, 1)^2$ , is it true that roughly  $\pi/4$  of them are within Euclidean distance 1 of the origin?

The definition of pseudorandomness is kind of a “master” definition that encompasses all of these statistical tests and more. After all, what is a statistical test, but a polynomial-time procedure that obtains samples from a distribution and outputs a yes/no decision? Pseudorandomness means that *every* statistical test that “passes” uniform data will also “pass” pseudorandomly generated data.

## 5.2 Pseudorandom Generators in Practice

You are probably expecting to now see at least one example of a secure PRG. Unfortunately, things are not so simple. We have no examples of secure PRGs! If it were possible to prove

<sup>1</sup>For one list of such tests, see <http://csrc.nist.gov/publications/nistpubs/800-22-rev1a/SP800-22rev1a.pdf>.

that some function  $G$  is a secure PRG, **it would resolve the famous P vs NP problem** — the most famous unsolved problem in computer science (and arguably, all of mathematics).

The next best thing that cryptographic research can offer are **candidate PRGs**, which are *conjectured* to be secure. The best examples of such PRGs are the ones that have been subjected to significant public scrutiny and resisted all attempts at attacks so far.

In fact, the entire rest of this book is based on cryptography that is only *conjectured* to be secure. How is this possible, given the book’s stated focus on *provable security*? As you progress through the book, pay attention to how all of the provable security claims are *conditional* — if  $X$  is secure then  $Y$  is secure. You will be able to trace back through this web of implications and discover that there are only a small number of underlying cryptographic primitives whose security is merely *conjectured* (PRGs are one example of such a primitive). Everything else builds on these primitives in a provably secure way.

With that disclaimer out of the way, surely *now* you can be shown an example of a conjectured secure PRG, right? There are indeed some conjectured PRGs that are simple enough to show you at this point, but you won’t find them in the book. The problem is that none of these PRG candidates are really used in practice. When you really need a PRG in practice, you would actually use a PRG that is built from something called a block cipher (which we won’t see until [Chapter 6](#)). A block cipher is *conceptually* more complicated than a PRG, and can even be built from a PRG (in principle). That explains why this book starts with PRGs. In practice, a block cipher is just a more useful object, so that is what you would find easily available (even implemented with specialized CPU instructions in most CPUs). When we introduce block ciphers (and pseudorandom functions), we will discuss how they can be used to construct PRGs.

## How NOT to Build a PRG

We can appreciate the challenges involved in building a PRG “from scratch” by first looking at an obvious idea for a PRG and understanding why it’s insecure.

**Example** *Let’s focus on the case of a length-doubling PRG. It should take in  $\lambda$  bits and output  $2\lambda$  bits. The output should look random when the input is sampled uniformly. A natural idea is for the candidate PRG to simply repeat the input twice. After all, if the input  $s$  is random, then  $s||s$  is also random, too, right?*

$G(s) :$ <hr style="width: 50%; margin: 0;"/> return $s  s$
--

*To understand why this PRG is insecure, first let me ask you whether the following strings look like they were sampled uniformly from  $\{0, 1\}^8$ :*

**11011101, 01010101, 01110111, 01000100, ...**

*Do you see any patterns? Every string has its first half equal to its second half. That is a conspicuous pattern because it is relatively rare for a uniformly chosen string to have this property.*

Of course, this is exactly what is wrong with this simplistic PRG  $G$  defined above. Every output of  $G$  has equal first/second halves. But it is rare for uniformly sampled strings to have this property. We can formalize this observation as an attack against the PRG-security of  $G$ :

$\mathcal{A}$
$x  y := \text{QUERY}()$
return $x \stackrel{?}{=} y$

The first line means to obtain the result of  $\text{QUERY}$  and set its first half to be the string  $x$  and its second half to be  $y$ . This calling program simply checks whether the output of  $\text{QUERY}$  has equal halves.

To complete the attack, we must show that this calling program has non-negligible bias distinguishing the  $\mathcal{L}_{\text{prg-}\star}$  libraries.

- ▶ When linked to  $\mathcal{L}_{\text{prg-real}}$ , the calling program receives outputs of  $G$ , which always have matching first/second halves. So  $\Pr[\mathcal{A} \diamond \mathcal{L}_{\text{prg-real}}^G \Rightarrow 1] = 1$ . Below we have filled in  $\mathcal{L}_{\text{prg-real}}$  with the details of our  $G$  algorithm:

<table border="1" style="width: 100%;"> <tr><th style="text-align: center;"><math>\mathcal{A}</math></th></tr> <tr><td><math>x  y := \text{QUERY}()</math></td></tr> <tr><td>return <math>x \stackrel{?}{=} y</math></td></tr> </table>	$\mathcal{A}$	$x  y := \text{QUERY}()$	return $x \stackrel{?}{=} y$	◇	<table border="1" style="width: 100%;"> <tr><th style="text-align: center;"><math>\mathcal{L}_{\text{prg-real}}^G</math></th></tr> <tr><td><math>\text{QUERY}():</math></td></tr> <tr><td><math>s \leftarrow \{0, 1\}^\lambda</math></td></tr> <tr><td>return <math>s  s</math></td></tr> </table>	$\mathcal{L}_{\text{prg-real}}^G$	$\text{QUERY}():$	$s \leftarrow \{0, 1\}^\lambda$	return $s  s$
$\mathcal{A}$									
$x  y := \text{QUERY}()$									
return $x \stackrel{?}{=} y$									
$\mathcal{L}_{\text{prg-real}}^G$									
$\text{QUERY}():$									
$s \leftarrow \{0, 1\}^\lambda$									
return $s  s$									

- ▶ When linked to  $\mathcal{L}_{\text{prg-rand}}$ , the calling program receives uniform samples from  $\{0, 1\}^{2\lambda}$ .

<table border="1" style="width: 100%;"> <tr><th style="text-align: center;"><math>\mathcal{A}</math></th></tr> <tr><td><math>x  y := \text{QUERY}()</math></td></tr> <tr><td>return <math>x \stackrel{?}{=} y</math></td></tr> </table>	$\mathcal{A}$	$x  y := \text{QUERY}()$	return $x \stackrel{?}{=} y$	◇	<table border="1" style="width: 100%;"> <tr><th style="text-align: center;"><math>\mathcal{L}_{\text{prg-rand}}^G</math></th></tr> <tr><td><math>\text{QUERY}():</math></td></tr> <tr><td><math>r \leftarrow \{0, 1\}^{2\lambda}</math></td></tr> <tr><td>return <math>r</math></td></tr> </table>	$\mathcal{L}_{\text{prg-rand}}^G$	$\text{QUERY}():$	$r \leftarrow \{0, 1\}^{2\lambda}$	return $r$
$\mathcal{A}$									
$x  y := \text{QUERY}()$									
return $x \stackrel{?}{=} y$									
$\mathcal{L}_{\text{prg-rand}}^G$									
$\text{QUERY}():$									
$r \leftarrow \{0, 1\}^{2\lambda}$									
return $r$									

$\mathcal{A}$  outputs 1 whenever we sample a string from  $\{0, 1\}^{2\lambda}$  with equal first/second halves. What exactly is the probability of this happening? There are several ways to see that the probability is  $1/2^\lambda$  (this is like asking the probability of rolling doubles with two dice, but each die has  $2^\lambda$  sides instead of 6). Therefore,  $\Pr[\mathcal{A} \diamond \mathcal{L}_{\text{prg-rand}}^G \Rightarrow 1] = 1/2^\lambda$ .

The advantage of this adversary is  $1 - 1/2^\lambda$  which is certainly non-negligible — it does not even approach 0 as  $\lambda$  grows. This shows that  $G$  is not a secure PRG.

This example illustrates how randomness/pseudorandomness is a property of the *entire process*, not of individual strings. If you take a string of **1**s and concatenate it with another string of **1**s, you get a long string of **1**s. “Containing only **1**s” is a property of individual strings. If you take a “random string” and concatenate it with another “random string,” you might not get a “random long string.” Being random is not a property of an individual string, but of the entire process that generates it.

Outputs from this  $G$  have equal first/second halves, which is an obvious pattern. The challenge of designing a secure PRG is that its outputs must have *no discernable pattern!* Any pattern will lead to an attack similar to the one shown above.

### Related Concept: Random Number Generation

The security of a PRG requires the seed to be chosen uniformly. In practice, the seed has to come from somewhere. Generally a source of “randomness” is provided by the hardware or operating system, and the process that generates these random bits is (confusingly) called a random *number* generator (RNG).

In this course we won’t cover low-level random *number* generation, but merely point out what makes it different than the PRGs that we study:

- ▶ The job of a PRG is to take a small amount of “ideal” (in other words, uniform) randomness and extend it.
- ▶ By contrast, an RNG usually takes many inputs over time and maintains an internal state. These inputs are often from physical/hardware sources. While these inputs are “noisy” in some sense, it is hard to imagine that they would be statistically *uniform*. So the job of the RNG is to “refine” (sometimes many) sources of noisy data into uniform outputs.

## 5.3 Application: Shorter Keys in One-Time-Secret Encryption

We revisit the motivating example from the beginning of this chapter. Alice & Bob share only a  $\lambda$ -bit key but want to encrypt a message of length  $\lambda + \ell$ . The main idea is to expand the key  $k$  into a longer string using a PRG  $G$ , and use the result as a one-time pad on the (longer) plaintext. More precisely, let  $G : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{\lambda+\ell}$  be a PRG, and define the following encryption scheme:

Construction 5.2  
(Pseudo-OTP)

$\mathcal{K} = \{0, 1\}^\lambda$	<u>KeyGen:</u>	<u>Enc(<math>k, m</math>):</u>	<u>Dec(<math>k, c</math>):</u>
$\mathcal{M} = \{0, 1\}^{\lambda+\ell}$	$k \leftarrow \mathcal{K}$	return $G(k) \oplus m$	return $G(k) \oplus c$
$\mathcal{C} = \{0, 1\}^{\lambda+\ell}$	return $k$		

The resulting scheme will not have (perfect) one-time secrecy. That is, encryptions of  $m_L$  and  $m_R$  will not be identically distributed in general. However, the distributions will be *indistinguishable* if  $G$  is a secure PRG. The precise flavor of security obtained by this construction is the following.


Definition 5.3 *Let  $\Sigma$  be an encryption scheme, and let  $\mathcal{L}_{\text{ots-L}}^\Sigma$  and  $\mathcal{L}_{\text{ots-R}}^\Sigma$  be defined as in Definition 2.8 (and repeated below for convenience). Then  $\Sigma$  has **(computational) one-time secrecy** if  $\mathcal{L}_{\text{ots-L}}^\Sigma \approx \mathcal{L}_{\text{ots-R}}^\Sigma$ . That is, if for all polynomial-time distinguishers  $\mathcal{A}$ , we have  $\Pr[\mathcal{A} \diamond \mathcal{L}_{\text{ots-L}}^\Sigma \Rightarrow 1] \approx \Pr[\mathcal{A} \diamond \mathcal{L}_{\text{ots-R}}^\Sigma \Rightarrow 1]$ .*

$\mathcal{L}_{\text{ots-L}}^\Sigma$	$\mathcal{L}_{\text{ots-R}}^\Sigma$
EAVESDROP( $m_L, m_R \in \Sigma.\mathcal{M}$ ):	EAVESDROP( $m_L, m_R \in \Sigma.\mathcal{M}$ ):
$k \leftarrow \Sigma.\text{KeyGen}$	$k \leftarrow \Sigma.\text{KeyGen}$
$c \leftarrow \Sigma.\text{Enc}(k, m_L)$	$c \leftarrow \Sigma.\text{Enc}(k, m_R)$
return $c$	return $c$


This is essentially the same as [Definition 2.8](#), except we are using  $\approx$  (indistinguishability) instead of  $\equiv$  (interchangeability).

**Claim 5.4** *Let pOTP denote [Construction 5.2](#). If pOTP is instantiated using a secure PRG  $G$  then pOTP has computational one-time secrecy.*


**Proof** We must show that  $\mathcal{L}_{\text{ots-L}}^{\text{pOTP}} \approx \mathcal{L}_{\text{ots-R}}^{\text{pOTP}}$ . As usual, we will proceed using a sequence of hybrids that begins at  $\mathcal{L}_{\text{ots-L}}^{\text{pOTP}}$  and ends at  $\mathcal{L}_{\text{ots-R}}^{\text{pOTP}}$ . For each hybrid library, we will demonstrate that it is indistinguishable from the previous one. Note that we are allowed to use the fact that  $G$  is a secure PRG. In practical terms, this means that if we can express some hybrid library in terms of  $\mathcal{L}_{\text{prg-real}}^G$  (one of the libraries in the PRG security definition), we can replace it with its counterpart  $\mathcal{L}_{\text{prg-rand}}^G$  (or vice-versa). The PRG security of  $G$  says that such a change will be indistinguishable.




The starting point is  $\mathcal{L}_{\text{ots-L}}^{\text{pOTP}}$ , shown here with the details of pOTP filled in.



The first hybrid step is to factor out the computations involving  $G$ , in terms of the  $\mathcal{L}_{\text{prg-real}}^G$  library.



From the PRG security of  $G$ , we may replace the instance of  $\mathcal{L}_{\text{prg-real}}^G$  with  $\mathcal{L}_{\text{prg-rand}}^G$ . The resulting hybrid library  $\mathcal{L}_{\text{hyb-2}}$  is indistinguishable from the previous one.



A subroutine has been inlined. Note that the resulting library is precisely  $\mathcal{L}_{\text{ots-L}}^{\text{OTP}}$  involving **standard one-time pad** on plaintexts of size  $\lambda + \ell$ . We have essentially proven that pOTP is indistinguishable from standard OTP, and therefore we can apply the security of OTP.



$\mathcal{L}_{\text{hyb-4}}$	$\mathcal{L}_{\text{ots-R}}^{\text{OTP}}$
	$\text{EAVESDROP}(m_L, m_R):$ $z \leftarrow \{0, 1\}^{\lambda+\ell}$ $c := z \oplus m_R$ return $c$

The (perfect) one-time secrecy of  $r\text{OTP}$  allows us to replace  $\mathcal{L}_{\text{ots-L}}^{\text{OTP}}$  with  $\mathcal{L}_{\text{ots-R}}^{\text{OTP}}$ ; they are interchangeable.

The rest of the proof is essentially a “mirror image” of the previous steps, in which we perform the same steps but in reverse (and with  $m_R$  being used instead of  $m_L$ ).

$\mathcal{L}_{\text{hyb-5}}$	$\mathcal{L}_{\text{prg-rand}}^G$	$\diamond$
	$\text{EAVESDROP}(m_L, m_R):$ $z \leftarrow \text{QUERY}()$ $c := z \oplus m_R$ return $c$	

A statement has been factored out into a subroutine, which happens to exactly match  $\mathcal{L}_{\text{prg-rand}}^G$ .

$\mathcal{L}_{\text{hyb-6}}$	$\mathcal{L}_{\text{prg-real}}^G$	$\diamond$
	$\text{EAVESDROP}(m_L, m_R):$ $z \leftarrow \text{QUERY}()$ $c := z \oplus m_R$ return $c$	

From the PRG security of  $G$ , we can replace  $\mathcal{L}_{\text{prg-rand}}^G$  with  $\mathcal{L}_{\text{prg-real}}^G$ . The resulting library is indistinguishable from the previous one.

$\mathcal{L}_{\text{ots-R}}^{\text{pOTP}}$	$\mathcal{L}_{\text{ots-R}}^{\text{pOTP}}$
	$\text{EAVESDROP}(m_L, m_R):$ $k \leftarrow \{0, 1\}^\lambda$ $c := G(k) \oplus m_R$ return $c$

A subroutine has been inlined. The result is  $\mathcal{L}_{\text{ots-R}}^{\text{pOTP}}$ .

Summarizing, we showed a sequence of hybrid libraries satisfying the following:

$$\mathcal{L}_{\text{ots-L}}^{\text{pOTP}} \equiv \mathcal{L}_{\text{hyb-1}} \approx \mathcal{L}_{\text{hyb-2}} \equiv \mathcal{L}_{\text{hyb-3}} \equiv \mathcal{L}_{\text{hyb-4}} \equiv \mathcal{L}_{\text{hyb-5}} \approx \mathcal{L}_{\text{hyb-6}} \equiv \mathcal{L}_{\text{ots-R}}^{\text{pOTP}}$$

Hence,  $\mathcal{L}_{\text{ots-L}}^{\text{pOTP}} \approx \mathcal{L}_{\text{ots-R}}^{\text{pOTP}}$ , and pOTP has (computational) one-time secrecy. ■

## ★ 5.4 Contrapositive Point of View on Security Proofs

We just proved the statement “if  $G$  is a secure PRG, then pOTP has one-time secrecy,” but let’s also think about the contrapositive of that statement:

If the pOTP scheme is **not** one-time secret, then  $G$  is **not** a secure PRG.

If the pOTP scheme is not secure, then there is some distinguisher  $\mathcal{A}$  that can distinguish the two  $\mathcal{L}_{\text{ots-}\star}$  libraries with better than negligible advantage. Knowing that such an  $\mathcal{A}$  exists, can we indeed break the security of  $G$ ?



Imagine going through the sequence of hybrid libraries from the proof, using this hypothetical  $\mathcal{A}$  as the calling program. We know that one of the steps of the proof must break down since  $\mathcal{A}$  successfully distinguishes between the endpoints of the hybrid sequence. Some of the steps of the proof were unconditional; for example, factoring out and inlining subroutines *never* has an effect on the calling program. These steps of the proof always hold; they are the steps where we write  $\mathcal{L}_{\text{hyb-}i} \equiv \mathcal{L}_{\text{hyb-}(i+1)}$ .

The steps where we write  $\mathcal{L}_{\text{hyb-}i} \approx \mathcal{L}_{\text{hyb-}(i+1)}$  are *conditional*. In our proof, the steps  $\mathcal{L}_{\text{hyb-}1} \approx \mathcal{L}_{\text{hyb-}2}$  and  $\mathcal{L}_{\text{ots-R}}^{\text{OTP}} \approx \mathcal{L}_{\text{hyb-}4}$  relied on  $G$  being a secure PRG. So if a hypothetical  $\mathcal{A}$  was able to break the security of pOTP, then that same  $\mathcal{A}$  *must* also successfully distinguish between  $\mathcal{L}_{\text{hyb-}1}$  and  $\mathcal{L}_{\text{hyb-}2}$ , or between  $\mathcal{L}_{\text{hyb-}5}$  and  $\mathcal{L}_{\text{hyb-}6}$  – the only conditional steps of the proof.

Let's examine the two cases:

- Suppose the hypothetical  $\mathcal{A}$  successfully distinguishes between  $\mathcal{L}_{\text{hyb-}1}$  and  $\mathcal{L}_{\text{hyb-}2}$ . Let's recall what these libraries actually look like:

$$\mathcal{L}_{\text{hyb-}1} = \begin{array}{|l} \text{EAVESDROP}(m_L, m_R): \\ \hline z \leftarrow \text{QUERY}() \\ c = z \oplus m_L \\ \text{return } c \end{array} \diamond \mathcal{L}_{\text{prg-real}}^G;$$

$$\mathcal{L}_{\text{hyb-}2} = \begin{array}{|l} \text{EAVESDROP}(m_L, m_R): \\ \hline z \leftarrow \text{QUERY}() \\ c = z \oplus m_L \\ \text{return } c \end{array} \diamond \mathcal{L}_{\text{prg-rand}}^G.$$

Interestingly, these two libraries share a common component that is linked to either  $\mathcal{L}_{\text{prg-real}}$  or  $\mathcal{L}_{\text{prg-rand}}$ . (This is no coincidence!) Let's call that common library  $\mathcal{L}^*$  and write

$$\mathcal{L}_{\text{hyb-}1} = \mathcal{L}^* \diamond \mathcal{L}_{\text{prg-real}}^G; \quad \mathcal{L}_{\text{hyb-}2} = \mathcal{L}^* \diamond \mathcal{L}_{\text{prg-rand}}^G.$$

Since  $\mathcal{A}$  successfully distinguishes between  $\mathcal{L}_{\text{hyb-}1}$  and  $\mathcal{L}_{\text{hyb-}2}$ , the following advantage is non-negligible:


$$\left| \Pr[\mathcal{A} \diamond \mathcal{L}^* \diamond \mathcal{L}_{\text{prg-real}}^G \Rightarrow 1] - \Pr[\mathcal{A} \diamond \mathcal{L}^* \diamond \mathcal{L}_{\text{prg-rand}}^G \Rightarrow 1] \right|.$$

But with a change of perspective, this means that  $\mathcal{A} \diamond \mathcal{L}^*$  is a calling program that successfully distinguishes  $\mathcal{L}_{\text{prg-real}}^G$  from  $\mathcal{L}_{\text{prg-rand}}^G$ . In other words,  $\mathcal{A} \diamond \mathcal{L}^*$  **breaks the PRG security of  $G$ !**

- Suppose the hypothetical  $\mathcal{A}$  only distinguishes  $\mathcal{L}_{\text{hyb-}5}$  from  $\mathcal{L}_{\text{hyb-}6}$ . Going through the reasoning above, we will reach a similar conclusion but with a different  $\mathcal{L}^*$  than before (in fact, it will be mostly the same library but with  $m_L$  replaced with  $m_R$ ).

So you can think of our security proof as a very roundabout way of saying the following:

If you give me an adversary/distinguisher  $\mathcal{A}$  that breaks the one-time secrecy of pOTP, then I can use it to build an adversary that breaks the PRG security of  $G$ . More specifically,  $G$  is guaranteed to be broken by (at least) one of the two distinguishers:<sup>2</sup>



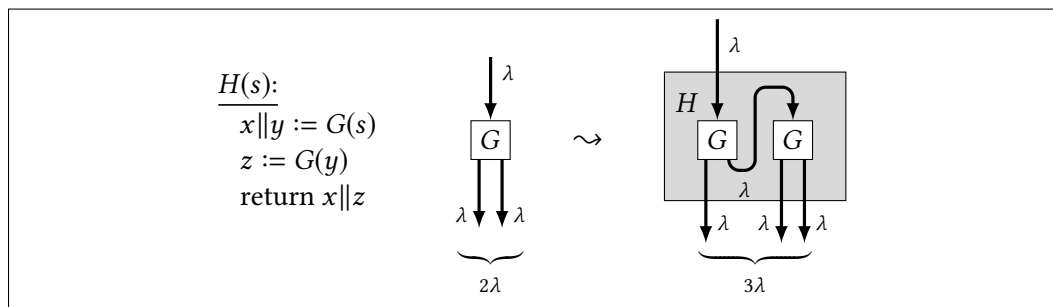
In fact, this would be the “traditional” method for proving security that you might see in many other cryptography textbooks. You would suppose you are given an adversary  $\mathcal{A}$  breaking pOTP, then you would demonstrate why at least one of the adversaries given above breaks the PRG. Unfortunately, it is quite difficult to explain to students how one is supposed to *come up with* these PRG-adversaries in terms of the one-time-secrecy adversaries. For that reason, we will reason about proofs in the “if X is secure then Y is secure” realm, rather than the contrapositive “if Y is insecure then X is insecure.”

## 5.5 Extending the Stretch of a PRG

Recall that the *stretch* of a PRG measures how much longer its output is than its input. A PRG with very long stretch seems much more useful than one with small stretch — at least, such a PRG can be used to encrypt a longer plaintext using the pseudo-OTP construction. Is there a limit to the stretch of a PRG?

In this section we will see that once you can extend a PRG a little bit, you can also extend it a lot. This is a curious fact about pseudorandomness that is not true about a truly uniform distributions. We will demonstrate the concept by extending a PRG with stretch  $\lambda$  into one with stretch  $2\lambda$ , but the idea can be used to increase the stretch of any PRG indefinitely, as we will see in the next section.

**Construction 5.5 (PRG feedback)** Let  $G : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{2\lambda}$  be a length-doubling PRG (i.e., a PRG with stretch  $\lambda$ ). When we write  $a||b := G(s)$ , it means that  $a$  is the first  $\lambda$  bits of  $G$ 's output and  $b$  is the second  $\lambda$  bits. Define the length-tripling function  $H : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{3\lambda}$  as follows:

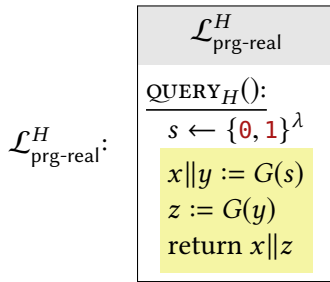


<sup>2</sup>The statement is somewhat non-constructive, since we don't know for sure which of the two distinguishers will be the one that actually works. But a way around this is to consider a single distinguisher that flips a coin and acts like the first one with probability 1/2 and acts like the second one with probability 1/2.


Claim 5.6 *If  $G$  is a secure length-doubling PRG, then  $H$  (defined above) is a secure length-tripling PRG.*

Proof We want to show that  $\mathcal{L}_{\text{prg-real}}^H \approx \mathcal{L}_{\text{prg-rand}}^H$ . As usual, we do so with a hybrid sequence. Since we assume that  $G$  is a secure PRG, we are allowed to use the fact that  $\mathcal{L}_{\text{prg-real}}^G \approx \mathcal{L}_{\text{prg-rand}}^G$ . In this proof, we will use the fact twice: once for each occurrence of  $G$  in the code of  $H$ .


Since  $\mathcal{L}_{\text{prg-}\star}^H$  and  $\mathcal{L}_{\text{prg-}\star}^G$  will both appear in this proof, and both libraries/interfaces have a subroutine named “QUERY”, we will rename these subroutines  $\text{QUERY}_H$  and  $\text{QUERY}_G$  to reduce confusion.




The starting point is  $\mathcal{L}_{\text{prg-real}}^H$ , shown here with the details of  $H$  filled in.




The first invocation of  $G$  has been factored out into a subroutine. The resulting hybrid library includes an instance of  $\mathcal{L}_{\text{prg-real}}^G$ .




From the PRG security of  $G$ , we can replace the instance of  $\mathcal{L}_{\text{prg-real}}^G$  with  $\mathcal{L}_{\text{prg-rand}}^G$ . The resulting hybrid library is indistinguishable.




A subroutine has been inlined.




Choosing  $2\lambda$  uniformly random bits and then splitting them into two halves has exactly the same effect as choosing  $\lambda$  uniformly random bits and independently choosing  $\lambda$  more.



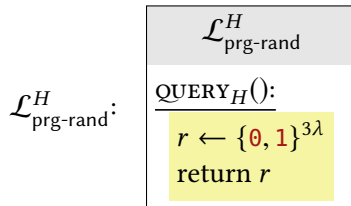
The remaining appearance of  $G$  has been factored out into a subroutine. Now  $\mathcal{L}_{\text{prg-real}}^G$  makes its second appearance.



Again, the PRG security of  $G$  lets us replace  $\mathcal{L}_{\text{prg-real}}^G$  with  $\mathcal{L}_{\text{prg-rand}}^G$ . The resulting hybrid library is indistinguishable.



A subroutine has been inlined.



Similar to above, concatenating  $\lambda$  uniform bits with  $2\lambda$  independently uniform bits has the same effect as sampling  $3\lambda$  uniform bits. The result of this change is  $\mathcal{L}_{\text{prg-rand}}^H$ .

Through this sequence of hybrid libraries, we showed that:


$$\mathcal{L}_{\text{prg-real}}^H \equiv \mathcal{L}_{\text{hyb-1}} \approx \mathcal{L}_{\text{hyb-2}} \equiv \mathcal{L}_{\text{hyb-3}} \equiv \mathcal{L}_{\text{hyb-4}} \equiv \mathcal{L}_{\text{hyb-5}} \approx \mathcal{L}_{\text{hyb-6}} \equiv \mathcal{L}_{\text{hyb-7}} \equiv \mathcal{L}_{\text{prg-rand}}^H.$$

Hence,  $H$  is a secure PRG. ■

## ★ 5.6 Applications: Stream Cipher & Symmetric Ratchet

The PRG-feedback construction can be generalized in a natural way, by continuing to feed part of  $G$ 's output into  $G$  again. The proof works in the same way as for the previous construction — the security of  $G$  is applied one at a time to each application of  $G$ .

**Claim 5.7** *If  $G$  is a secure length-doubling PRG, then for any  $n$  (polynomial function of  $\lambda$ ) the following construction  $H_n$  is a secure PRG with stretch  $n\lambda$ :*



The fact that this chain of PRGs can be extended indefinitely gives another useful functionality:


Definition 5.8  
(Stream cipher)

A **stream cipher** is an algorithm  $G$  that takes a seed  $s$  and length  $\ell$  as input, and outputs a string. It should satisfy the following requirements:

1.  $G(s, \ell)$  is a string of length  $\ell$ .
2. If  $i < j$ , then  $G(s, i)$  is a prefix of  $G(s, j)$ .
3. For each  $n$ , the function  $G(\cdot, n)$  is a secure PRG.

You can think of a stream cipher as a special PRG that generates a *pseudorandom stream of unbounded length*, but that only returns a finite prefix of that stream (as much as you request with the parameter  $\ell$ ).

The PRG-feedback construction can be used to construct a secure stream cipher in the natural way: given seed  $s$  and length  $\ell$ , keep iterating the PRG-feedback main loop until  $\ell$  bits have been generated.



### Symmetric Ratchet

Suppose Alice & Bob share a symmetric key  $k$  and are using a secure messaging app to exchange messages over a long period of time. Later in the course we will see techniques that Alice & Bob could use to securely encrypt many messages using a single key. However, suppose Bob's device is compromised and an attacker learns  $k$ . Then the attacker can decrypt all past, present, and future ciphertexts that it saw!

Alice & Bob can protect against such a key compromise by using the PRG-feedback stream cipher to constantly “update” their shared key. Suppose they do the following, starting with their shared key  $k$ :

- ▶ They use  $k$  to seed a chain of length-doubling PRGs, and both obtain the same stream of pseudorandom keys  $t_1, t_2, \dots$


- ▶ They use  $t_i$  as a key to send/receive the  $i$ th message. The details of the encryption are not relevant to this example.
- ▶ After making a call to the PRG, they erase the PRG input from memory, and only remember the PRG's output. After using  $t_i$  to send/receive a message, they also erase it from memory.

This way of using and forgetting a sequence of keys is called a **symmetric ratchet**.

Construction 5.9  
(Symm Ratchet)

```

 $s_0 = k$ 
for  $i = 1$  to  $\infty$ :
   $s_i || t_i := G(s_{i-1})$ 
  erase  $s_{i-1}$  from memory
  use  $t_i$  to encrypt/decrypt the  $i$ th message
  erase  $t_i$  from memory
    
```



Suppose that an attacker compromises Bob's device after  $n$  ciphertexts have been sent. The only value residing in memory is  $s_n$ , which the attacker learns. Since  $G$  is deterministic, the attacker can now compute  $t_{n+1}, t_{n+2}, \dots$  in the usual way and decrypt all future ciphertexts that are sent.

However, we can show that the attacker learns no information about  $t_1, \dots, t_n$  from  $s_n$ , which implies that the previous ciphertexts remain safe. By compromising the key  $s_n$ , the adversary only compromises the security of *future* messages, but not *past* messages. Sometimes this property is called **forward secrecy**, meaning that messages in the present are protected against a key-compromise that happens in the future.

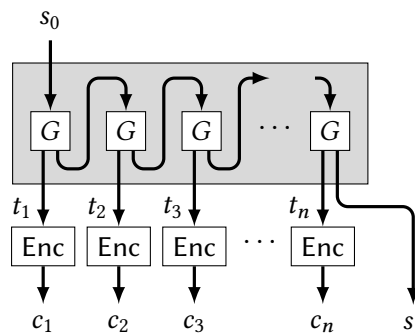
This construction is called a **ratchet**, since it is easy to advance the key sequence in the forward direction (from  $s_n$  to  $s_{n+1}$ ) but hard to reverse it (from  $s_{n+1}$  to  $s_n$ ). The exercises explore the problem of explicitly reversing the ratchet, but the more relevant property for us is whether the attacker learns anything about the ciphertexts that were generated before the compromise.

Claim 5.10 *If the symmetric ratchet (Construction 5.9) is used with a secure PRG  $G$  and an encryption scheme  $\Sigma$  that has uniform ciphertexts (and  $\Sigma.\mathcal{K} = \{0, 1\}^\lambda$ ), then the first  $n$  ciphertexts are pseudorandom, even to an eavesdropper who compromises the key  $s_n$ .*

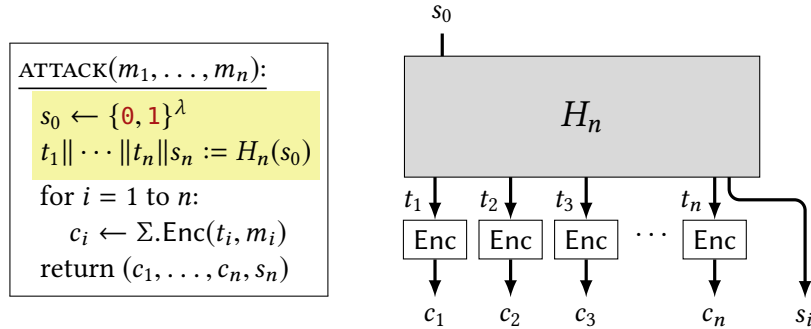
Proof We are considering an attack scenario in which  $n$  plaintexts are encrypted, and the adversary sees their ciphertexts as well as the ratchet-key  $s_n$ . This situation is captured by the following library:

```


ATTACK( $m_1, \dots, m_n$ ):
   $s_0 \leftarrow \{0, 1\}^\lambda$ 
  for  $i = 1$  to  $n$ :
     $s_i || t_i := G(s_{i-1})$ 
     $c_i \leftarrow \Sigma.\text{Enc}(t_i, m_i)$ 
  return ( $c_1, \dots, c_n, s_n$ )
    
```




As we have seen, the shaded box (the process that computes  $t_1, \dots, t_n$  from  $s_0$ ) is actually a PRG. Let us rewrite the library in terms of this PRG  $H_n$ :




Now, we can apply the PRG security of  $H_n$  and instead choose  $t_1, \dots, t_n$  and  $s_n$  uniformly. This change is indistinguishable, by the security of the PRG. Note that we have not written out the standard explicit steps (factor out the first two lines of `ATTACK` in terms of  $\mathcal{L}_{\text{prg-real}}$ , replace with  $\mathcal{L}_{\text{prg-rand}}$ , and inline).



At this point, the encryption scheme is being used “as intended,” meaning that we generate its keys  $t_i$  uniformly/independently, and use each key only for one encryption and nothing else. Formally speaking, this means we can factor out the body of the for-loop in terms of  $\mathcal{L}_{\text{ots-real}}$ :



We can now replace  $\mathcal{L}_{\text{ots-real}}$  with  $\mathcal{L}_{\text{ots-rand}}$  and inline the subroutine (without showing the intermediate library). The result is:





This final library is indistinguishable from the first one. As promised, we showed that the attacker cannot distinguish the first  $n$  ciphertexts from random values, even when seeing  $s_n$ . ■

This proof used the uniform-ciphertexts property, but the same logic applies to basically any encryption property you care about — just imagine factoring out the encryption steps in terms of a different library than  $\mathcal{L}_{\text{ots}\$-real}$ .

## Exercises

- 5.1. Let  $G : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{\lambda+\ell}$  be an injective (i.e., 1-to-1) PRG. Consider the following distinguisher:

$\mathcal{A}$
$x := \text{QUERY}()$ for all $s' \in \{0, 1\}^\lambda$ : if $G(s') = x$ then return 1 return 0

- (a) What is the advantage of  $\mathcal{A}$  in distinguishing  $\mathcal{L}_{\text{prg-real}}^G$  and  $\mathcal{L}_{\text{prg-rand}}^G$ ? Is it negligible?  
 (b) Does this contradict the fact that  $G$  is a PRG? Why or why not?  
 (c) What happens to the advantage if  $G$  is not injective?
- 5.2. Let  $G : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{\lambda+\ell}$  be an injective PRG, and consider the following distinguisher:

$\mathcal{A}$
$x := \text{QUERY}()$ $s' \leftarrow \{0, 1\}^\lambda$ return $G(s') \stackrel{?}{=} x$

What is the advantage of  $\mathcal{A}$  in distinguishing  $\mathcal{L}_{\text{prg-real}}^G$  from  $\mathcal{L}_{\text{prg-rand}}^G$ ?

*Hint:* When computing  $\Pr[\mathcal{A} \diamond \mathcal{L}_{\text{prg-rand}}^G \text{ outputs } 1]$ , separate the probabilities based on whether  $x \in \text{im}(G)$  or not. ( $\text{im}(G)$  is defined in a previous problem)

- 5.3. For any PRG  $G : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{\lambda+\ell}$  there will be many strings in  $\{0, 1\}^{\lambda+\ell}$  that are impossible to get as output of  $G$ . Let  $S$  be any such set of impossible  $G$ -outputs, and consider the following adversary that has  $S$  hard-coded:

$\mathcal{A}$
$x := \text{QUERY}()$ return $x \in S$

What is the advantage of  $\mathcal{A}$  in distinguishing  $\mathcal{L}_{\text{prg-real}}^G$  from  $\mathcal{L}_{\text{prg-rand}}^G$ ? Why does an adversary like this one not automatically break every PRG?

- 5.4. Show that the scheme from Section 5.3 does not have *perfect* one-time secrecy, by showing that there must exist two messages  $m_1$  and  $m_2$  whose ciphertext distributions differ.

*Hint:* There must exist strings  $s_1, s_2 \in \{0, 1\}^{2\lambda}$  where  $s_1 \in \text{im}(G)$ , and  $s_2 \notin \text{im}(G)$ . Use these two strings to find two messages  $m_1$  and  $m_2$  whose ciphertext distributions assign different probabilities to  $s_1$  and  $s_2$ . Note that it is legitimate for an attacker to “know”  $s_1$  and  $s_2$ , as these are properties of  $G$  alone, and do not depend on the random choices made “at runtime” — when the library executes the encryption algorithms.

- 5.5. In the PRG feedback construction (Construction 5.5), there are two calls to  $G$ . The security proof applies the PRG security rule to both of them, starting with the first. Describe what happens when you try to apply the PRG security of  $G$  to these two calls in the opposite order. Does the proof still work, or does it work only in the order that was presented?
- 5.6. Let  $\ell' > \ell > 0$ . Extend the “PRG feedback” construction to transform any PRG of stretch  $\ell$  into a PRG of stretch  $\ell'$ . Formally define the new PRG and prove its security using the security of the underlying PRG.
- 5.7. Prove that if  $G$  is a secure PRG, then so is the function  $H(s) = G(\bar{s})$ .
- 5.8. Let  $G : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{3\lambda}$  be a secure length-**tripling** PRG. For each function below, state whether it is also a secure PRG. If the function is a secure PRG, give a proof. If not, then describe a successful distinguisher and explicitly compute its advantage. When we write  $a\|b\|c := G(s)$ , each of  $a, b, c$  have length  $\lambda$ .

(a) 
$$\begin{array}{l} \overline{H(s):} \\ x\|y\|z := G(s) \\ \text{return } G(x)\|G(z) \end{array}$$

(b) 
$$\begin{array}{l} \overline{H(s):} \\ x\|y\|z := G(s) \\ \text{return } x\|y \end{array}$$

(c) 
$$\begin{array}{l} \overline{H(s):} \\ x := G(s) \\ y := G(s) \\ \text{return } x\|y \end{array}$$

(d) 
$$\begin{array}{l} \overline{H(s):} \\ x := G(s) \\ y := G(0^\lambda) \\ \text{return } x\|y \end{array}$$

(e) 
$$\begin{array}{l} \overline{H(s):} \\ x := G(s) \\ y := G(0^\lambda) \\ \text{return } x \oplus y \end{array}$$

(f) 
$$\begin{array}{l} // H : \{0, 1\}^{2\lambda} \rightarrow \{0, 1\}^{3\lambda} \\ \overline{H(s_L\|s_R):} \\ x := G(s_L) \\ y := G(s_R) \\ \text{return } x \oplus y \end{array}$$

(g) 
$$\begin{array}{l} // H : \{0, 1\}^{2\lambda} \rightarrow \{0, 1\}^{6\lambda} \\ \overline{H(s_L\|s_R):} \\ x := G(s_L) \\ y := G(s_R) \\ \text{return } x\|y \end{array}$$

- 5.9. Let  $G : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{3\lambda}$  be a secure length-**tripling** PRG. Prove that each of the following functions is also a secure PRG:

(a) 
$$\begin{array}{l} // H : \{0, 1\}^{2\lambda} \rightarrow \{0, 1\}^{4\lambda} \\ \hline H(s_L || s_R): \\ y := G(s_R) \\ \text{return } s_L || y \end{array}$$

Note that  $H$  includes half of its input directly in the output. How do you reconcile this fact with the conclusion of Exercise 5.12(b)?

(b) 
$$\begin{array}{l} // H : \{0, 1\}^{2\lambda} \rightarrow \{0, 1\}^{3\lambda} \\ \hline H(s_L || s_R): \\ \text{return } G(s_L) \end{array}$$

- 5.10. A frequently asked question in cryptography forums is whether it's possible to determine which PRG implementation was used by looking at output samples.

Let  $G_1$  and  $G_2$  be two PRGs with matching input/output lengths. Define two libraries  $\mathcal{L}_{\text{which-prg}}^{G_1}$  and  $\mathcal{L}_{\text{which-prg}}^{G_2}$  as follows:

$\mathcal{L}_{\text{which-prg}}^{G_1}$	$\mathcal{L}_{\text{which-prg}}^{G_2}$
$\begin{array}{l} \text{QUERY}(): \\ s \leftarrow \{0, 1\}^\lambda \\ \text{return } G_1(s) \end{array}$	$\begin{array}{l} \text{QUERY}(): \\ s \leftarrow \{0, 1\}^\lambda \\ \text{return } G_2(s) \end{array}$

Prove that if  $G_1$  and  $G_2$  are both secure PRGs, then  $\mathcal{L}_{\text{which-prg}}^{G_1} \approx \mathcal{L}_{\text{which-prg}}^{G_2}$  — that is, it is infeasible to distinguish which PRG was used simply by receiving output samples.

- ★ 5.11. Prove that if PRGs exist, then  $P \neq NP$ .

*Hint:*  $\{y \mid \exists s : G(s) = y\} \in NP$ . Prove the contrapositive! Use the powerful assumption that  $P = NP$  to construct an efficient adversary to attack any candidate PRG.

- 5.12. (a) Let  $f$  be any function. Show that the following function  $G$  is **not** a secure PRG, no matter what  $f$  is. Describe a successful distinguisher and explicitly compute its advantage:

$$\begin{array}{l} G(s): \\ \hline \text{return } s || f(s) \end{array}$$

- (b) Let  $G : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{\lambda+\ell}$  be a candidate PRG. Suppose there is a polynomial-time algorithm  $V$  with the property that it inverts  $G$  with non-negligible probability. That is,

$$\Pr_{s \leftarrow \{0, 1\}^\lambda} [V(G(s)) = s] \text{ is non-negligible.}$$

Show that if an algorithm  $V$  exists with this property, then  $G$  is not a secure PRG. In other words, construct a distinguisher contradicting the PRG-security of  $G$  and show that it achieves non-negligible distinguishing advantage.

*Note:* Don't assume anything about the output of  $V$  other than the property shown above. In particular,  $V$  might very frequently output the "wrong" thing.

5.13. Let  $s_0, s_1, \dots$  and  $t_1, t_2, \dots$  be defined as in the symmetric ratchet (Construction 5.9).

- (a) Prove that if  $G$  is a secure PRG then the following two libraries are indistinguishable, for any polynomial-time algorithm  $\mathcal{A}$ :

$\mathcal{L}_{\text{left}}$	$\mathcal{L}_{\text{right}}$
$\text{TEST}():$ $s_{n-1} \leftarrow \{0, 1\}^\lambda$ $s_n \  t_n := G(s_{n-1})$ $\tilde{t} = \mathcal{A}(s_n)$ $\text{return } \tilde{t} \stackrel{?}{=} t_n$	$\text{TEST}():$ $s_{n-1} \leftarrow \{0, 1\}^\lambda$ $\tilde{t} = \mathcal{A}(s_n)$ $t_n \leftarrow \{0, 1\}^\lambda$ $\text{return } \tilde{t} \stackrel{?}{=} t_n$

- (b) What is  $\Pr[\text{TEST outputs true}]$  in  $\mathcal{L}_{\text{right}}$ ?
- (c) Prove that for any polynomial-time algorithm  $\mathcal{A}$ ,  $\Pr[\mathcal{A}(s_n) = t_n]$  is negligible, where  $s_n, t_n$  are generated as in the symmetric ratchet construction.
- (d) Prove that for any polynomial-time algorithm  $\mathcal{A}$ ,  $\Pr[\mathcal{A}(s_n) = s_{n-1}]$  is negligible. In other words, “turning the ratchet backwards” is a hard problem.  
*Hint:* the proof should be a few lines, a direct corollary of part (c).