

# CS 6501 Week 1: Definitions and Foundations


Instructor: David Wu (dwu4@virginia.edu)

TA: Suya (fs5xz@virginia.edu)

What this course is all about:

Zero-knowledge proofs: How do you convince someone that a statement is true but not reveal anything more about the statement?  
What does it even mean to "know" something?

Fully homomorphic encryption: Can we compute on encrypted data (e.g., preserve full confidentiality of the data and yet, perform arbitrary computations over the hidden data?)



Multiparty computation: "Anything that can be computed with the help of a trusted party can be computed without one!"

Public-key cryptography:

How can two users who have never met communicate securely over a public and untrusted network?

Cryptography started as a means of protecting communication. Has now evolved to also protect computation.

We are at an "inflection" point in the development of cryptography/

1. Real, large-scale deployment of "fancy" cryptography (i.e. beyond public-key cryptography)

↳ Systems like Zcash (relies on pairings-based cryptography and zero-knowledge proof systems)

2. Potential threat of quantum computers requires re-thinking much of the existing public-key cryptography

- Google recently ran pilot project implementing ring-LWE based key-exchange into Chrome (alongside traditional Diffie-Hellman)

- Ongoing NIST competition to standardize new post-quantum cryptography (expect 5-7 years to converge on new standards)

Course objectives: 1. Provide broad survey of modern cryptography (from a theory-focused lens) ↳ so we will be starting with foundations  
2. Prepare you for research in cryptography ↳ Will likely offer an introduction to applied cryptography course next year

↳ Course will be fast-paced at the beginning, but should be self-contained.

More thorough treatment of symmetric crypto and public-key crypto (especially how to use them properly in systems) will be covered in the applied crypto course.

Logistics and administrative: Course website: <https://www.cs.virginia.edu/dwu4/courses/sp19> ← refer here for late day policy, homework submission instructions, office hours, notes...  
See Piazza for announcements, Gradescope for homework submission  
Anonymous feedback form also available ← please provide feedback!  
Course consists of 5 problem sets, weighted equally; no exams or projects  
Course TA: Suya [see website for contact information]

# Foundations of Modern Cryptography


Modern cryptography is study of hardness: some task should be "easy" for an honest user, but "difficult" for an adversary  
 e.g., if Alice encrypts a message to Bob, it should be easy for Alice to encrypt the message and for Bob to decrypt; however, should be difficult for eavesdropper who intercepts the message to decrypt

How do we model this mathematically?

basic building block of symmetric cryptography (will see how to use PRGs/PRFs/PRPs to build encryption schemes, message authentication codes next week)

Consider the notion of a pseudorandom generator (PRG):

- A PRG takes a short seed  $s$  and expands it to a much longer random string  $G(s)$ .
- We will see (next week) that PRGs (and related primitives) are very useful for constructing encryption schemes, message integrity mechanisms, and more.



**Definition.** A function  $G: \{0,1\}^{\lambda} \rightarrow \{0,1\}^{\ell(\lambda)}$  is a secure pseudorandom generator (PRG) if for all efficient algorithms  $A$ :

$$\left| \Pr[s \xleftarrow{\$} \{0,1\}^{\lambda} : A(G(s)) = 1] - \Pr[t \xleftarrow{\$} \{0,1\}^{\ell(\lambda)} : A(t) = 1] \right| = \text{negl}(\lambda)$$

(Annotations:  $\lambda$  is security parameter,  $\ell(\lambda)$  is stretch of the PRG, running time is polynomial in the input length, without loss of generality, assume  $A$  just outputs a single bit, a function  $f(\lambda)$  is negligible if  $f(\lambda) = o(1/\lambda^c)$  for all  $c \in \mathbb{N}$  (i.e., the function  $f$  is smaller than all polynomials in  $\lambda$ ), we view  $A$  as a "distinguisher": on input a string of length  $\ell(\lambda)$ , guesses whether it is the output of a PRG or if it is truly random - notice that typically  $\ell(\lambda) \gg \lambda$  so  $2^{\lambda} \ll 2^{\ell(\lambda)}$  (remarkable that PRGs plausibly exist!))

Intuitively: the outputs of a PRG on a random seed looks indistinguishable from a truly random string (assuming adversary does not see the seed)

Alternative characterization in terms of distributions: for an adversary  $A$ , define the following random variables:

$W_0: \Pr[s \xleftarrow{\$} \{0,1\}^{\lambda} : A(G(s)) = 1]$  "pseudorandom" distribution

$W_1: \Pr[t \xleftarrow{\$} \{0,1\}^{\ell(\lambda)} : A(t) = 1]$  "truly random" distribution

$\text{PRGAdv}[A, G] := |W_0 - W_1|$

**Definition.** A function  $G: \{0,1\}^{\lambda} \rightarrow \{0,1\}^{\ell(\lambda)}$  is a secure PRG if for all efficient adversaries  $A$ ,  $\text{PRGAdv}[A, G] = \text{negl}(\lambda)$ .

PRG security definition requires that two distribution ensembles look indistinguishable to any computationally-bounded adversary. More generally, we can write:

**Definition.** Let  $\lambda \in \mathbb{N}$  be a security parameter. Let  $D_1 = \{D_{1,\lambda}\}_{\lambda \in \mathbb{N}}$  and  $D_2 = \{D_{2,\lambda}\}_{\lambda \in \mathbb{N}}$  be two collections (ensembles) of distributions indexed by  $\lambda$ . Then,

$D_1$  and  $D_2$  are computationally indistinguishable (denoted  $D_1 \stackrel{c}{\approx} D_2$ ) if for all efficient adversaries  $A$ ,

$$\left| \Pr[x_1 \leftarrow D_{1,\lambda} : A(1^{\lambda}, x_1) = 1] - \Pr[x_2 \leftarrow D_{2,\lambda} : A(1^{\lambda}, x_2) = 1] \right| = \text{negl}(\lambda)$$

(Annotations: the adversary is given  $1^{\lambda}$  (the all-ones string of length  $\lambda$ ); this allows the running time of the algorithm to be  $\text{poly}(\lambda)$  (the draws from  $D_{1,\lambda}$  and  $D_{2,\lambda}$  may be much shorter than  $\lambda$ ))

Intuitively: no efficient adversary can tell a sample from  $D_1$  from a sample from  $D_2$

PRG security definition: A function  $G: \{0,1\}^{\lambda} \rightarrow \{0,1\}^{\ell(\lambda)}$  is a secure PRG if  $\{s \xleftarrow{\$} \{0,1\}^{\lambda} : G(s)\}_{\lambda \in \mathbb{N}} \stackrel{c}{\approx} \{t \xleftarrow{\$} \{0,1\}^{\ell(\lambda)} : t\}$ .

Do PRGs exist? We don't know! More difficult than resolving P vs. NP!


However, it is not hard to see that if PRGs exist, then  $P \neq NP$ . [Try proving this yourself]

↳ What we can say is that if one-way functions (OWF) exist, then there exists a PRG where  $l(\lambda) = \lambda + 1$  (i.e., a PRG with a 1-bit stretch)

[We will explore this more thoroughly on HW1]

But what if we want PRGs with longer stretch? For example, can we build PRGs with stretch  $l(\lambda) = \text{poly}(\lambda)$  for arbitrary polynomials?

Blum-Micali PRG: Suppose  $G: \{0,1\}^\lambda \rightarrow \{0,1\}^{\lambda+1}$  is a secure PRG. We build a PRG with stretch  $l(\lambda) = \text{poly}(\lambda)$  as follows:



Why is this constructing a secure PRG?

↳ Intuitively, if  $s_0$  is uniformly random, then  $G(s_0) = (b_1, s_1)$  is uniformly random so we can feed  $s_1$  into the PRG and take  $b_1$  as the first output bit of the PRG  $\Rightarrow$  iterate until we have  $l$  output bits




Theorem. If  $G: \{0,1\}^\lambda \rightarrow \{0,1\}^{\lambda+1}$  is a secure PRG, then the Blum-Micali generator  $G^{(l)}: \{0,1\}^\lambda \rightarrow \{0,1\}^{l(\lambda)}$  is also a secure PRG for all  $l = \text{poly}(\lambda)$ .

Proof. Our goal is to show that the following distributions are computationally indistinguishable:

$$D_0 = \{s_0 \xleftarrow{R} \{0,1\}^\lambda : G^{(l)}(s_0)\}$$

$$D_1 = \{t \xleftarrow{R} \{0,1\}^{l(\lambda)} : t\}$$

We will use a "hybrid" argument. Specifically, we first define a sequence of intermediate distributions:



⋮

$D_1 = D_{0,l}$        $b_1 \xleftarrow{R} \{0,1\}$        $b_2 \xleftarrow{R} \{0,1\}$        $\dots$        $b_l \xleftarrow{R} \{0,1\}$

Basic idea: in distribution  $D_{0,i}$ , the first  $i$  bits of output are generated uniformly at random while the remaining bits are generated using the Blum-Micali generator

Let  $A$  be an efficient distinguisher. Define  $p_i = \Pr[x \xleftarrow{R} D_{0,i} : A(x) = 1]$ .

Our objective is to show that  $\text{PRGAdv}[A, G] = |p_0 - p_l| = \text{negl}(\lambda)$ .

Proof (cont.). We use the triangle inequality:

$$\begin{aligned} \text{PRGAdv}[A, G] &= |p_0 - p_\ell| = |p_0 - p_1 + p_1 - p_2 + \dots + p_{\ell-1} - p_\ell| \\ &\leq |p_0 - p_1| + |p_1 - p_2| + \dots + |p_{\ell-1} - p_\ell| \end{aligned}$$


Claim. If  $G$  is a secure PRG, then for all efficient adversaries  $A$ ,  $|p_i - p_{i+1}| = \text{negl}(\lambda)$ .

Proof. We will show the contrapositive: if  $A$  can distinguish distributions  $D_{0,i}$  and  $D_{0,i+1}$ , then  $A$  can break pseudorandomness of  $G$ .


Suppose  $|p_i - p_{i+1}| = \epsilon$ . We use  $A$  to build a distinguisher  $B$  for  $G$ . Algorithm  $B$  works as follows:

1. On input a string  $z \in \{0,1\}^{\lambda+1}$ , algorithm  $B$  parses  $z$  as  $(b_{i+1}, s_{i+1})$  where  $b_{i+1} \in \{0,1\}$  and  $s_{i+1} \in \{0,1\}^\lambda$
2. Sample  $b_1, \dots, b_i \stackrel{R}{\leftarrow} \{0,1\}$ .
3. Compute  $b_{i+2}, \dots, b_\ell$  using Blum-Micali with seed  $s_{i+1}$ . Give  $b_1 \dots b_\ell$  to  $A$  and output whatever  $A$  outputs.

In pictures:




Two possibilities: 1. Suppose  $z = G(s_i)$  for some  $s_i \stackrel{R}{\leftarrow} \{0,1\}^\lambda$ . Then, above picture looks like this:



In this case,  $b_1, \dots, b_\ell$  is distributed exactly as in distribution  $D_{0,i}$  and so  $A$  outputs 1 with prob.  $p_i$ .

2. Suppose  $z \stackrel{R}{\leftarrow} \{0,1\}^{\lambda+1}$ . Then above picture looks like this:



In this case,  $b_1, \dots, b_\ell$  is distributed exactly as in distribution  $D_{0,i+1}$  and so  $A$  outputs 1 with prob.  $p_{i+1}$ .

$$\begin{aligned} \text{Thus, PRGAdv}[B, G] &= |\Pr[s \stackrel{R}{\leftarrow} \{0,1\}^\lambda : B(G(s)) = 1] - \Pr[z \stackrel{R}{\leftarrow} \{0,1\}^{\lambda+1} : B(z) = 1]| \\ &= |p_i - p_{i+1}| \\ &= \epsilon \end{aligned}$$

Since  $B$  outputs whatever  $A$  outputs

Since  $B$  is efficient (assuming  $A$  is efficient), by security of  $G$ ,  $\text{PRGAdv}[B, G] = \text{negl}(\lambda)$ . Thus,  $\epsilon = |p_i - p_{i+1}| = \text{negl}(\lambda)$ , and the claim follows. ■

To complete the proof of the main theorem, we have that

$$\begin{aligned} |p_0 - p_\ell| &\leq |p_0 - p_1| + \dots + |p_{\ell-1} - p_\ell| \\ &\leq l \cdot \text{negl}(\lambda) \\ &= \text{negl}(\lambda) \quad \text{since } l = \text{poly}(\lambda). \quad \blacksquare \end{aligned}$$

Proof strategy recap: 1. Hybrid arguments: to argue indistinguishability of a pair of distributions, begin by identifying a simple set of intermediate distributions, and argue that each pair of adjacent distributions is indistinguishable

2. Security reduction (proof by contrapositive): To show a statement of the form "If  $X$  is secure, then  $Y$  is secure," show instead the statement "If  $Y$  is not secure, then  $X$  is not secure." In the proof, show that if there exists an adversary for  $Y$  (i.e.  $Y$  is not secure), then there exists an adversary for  $X$ .

Pseudorandom functions (PRFs) - the "workhorse" of symmetric cryptography (will see applications next week)


↳ PRG "compresses" a long random string into a short seed

↳ PRF "compresses" a random function into a short key

Definition. A pseudorandom function (PRF) with key-space  $K = \{K_\lambda\}_{\lambda \in \mathbb{N}}$ , domain  $X = \{X_\lambda\}_{\lambda \in \mathbb{N}}$ , and range  $Y = \{Y_\lambda\}_{\lambda \in \mathbb{N}}$  is a family of efficiently-computable functions  $F = \{F_\lambda\}_{\lambda \in \mathbb{N}}$  where  $F_\lambda : K_\lambda \times X_\lambda \rightarrow Y_\lambda$ . We say that  $F$  is secure if for all efficient adversaries  $A$ ,

$$\text{PrAdv}[A, F] = |\Pr[W_b = 1] - \Pr[W_c = 1]| = \text{negl}(\lambda),$$

where for  $b \in \{0, 1\}$ , we define  $W_b$  to be the output of  $A$  in the following experiment:



Intuitively: A PRF is secure if no efficient adversary can distinguish the input/output behavior of the PRF (with a random and secret) key from the outputs of a truly random function (on the same domain and range).

↳ Very surprising: truly random functions do not have efficient representations, and yet we can mimic the behavior of such a function using a very short seed

↳ Will see many applications next week.

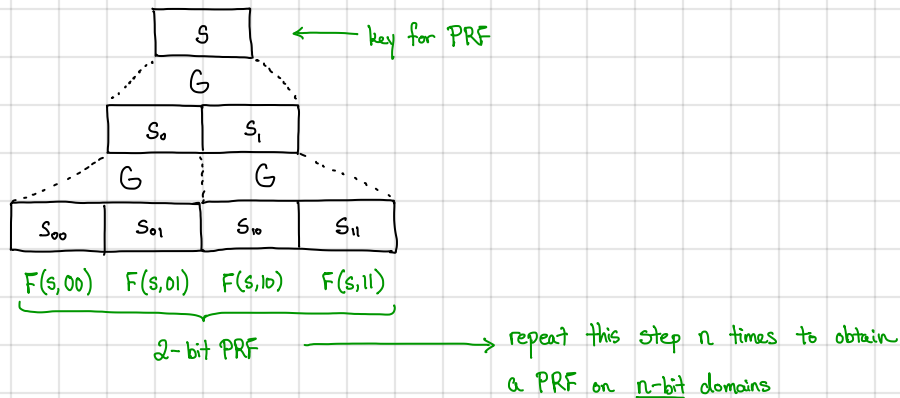
PRF  $\Rightarrow$  PRG. Suppose we have a PRF  $F : \{0, 1\}^\lambda \times \{0, 1\}^\lambda \rightarrow \{0, 1\}$ . We can construct a PRG  $G : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{\ell(\lambda)}$  for any  $\ell(\lambda) = \text{poly}(\lambda)$  as follows:

$$G(s) := \underbrace{F(s, 1) \parallel F(s, 1) \parallel \dots \parallel F(s, \ell)}_{\substack{\text{evaluate the PRF on inputs} \\ 1 \text{ through } \ell \\ \text{(known as "counter mode")}}}$$

Security proof (sketch): When  $s$  is sampled uniformly at random, then  $F(s, \cdot)$  is computationally indistinguishable from random function  $f(\cdot)$ . Now,  $f(1) \parallel f(2) \parallel \dots \parallel f(\ell)$  is uniformly random string.


Formally, should set up a reduction.

PRG  $\Rightarrow$  PRF. (Goldreich-Goldwasser-Micali). Suppose we have a length-doubling PRG  $G : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{2\lambda}$ .



Security proof: naive approach is to replace all the leaf nodes with uniformly random values exponentially many leaf nodes, so this requires exponentially-many hybrids, which is problematic

Alternative proof strategy is to proceed query-by-query: for each query, replace value of all intermediate nodes with random string



Final hybrid: adversary get random string for all of its queries

↳ Total number of hybrids is  $Q \cdot n = \text{poly}(\lambda)$   
 # PRF queries ↗ depth of tree ↖

Pseudorandom permutations (PRPs): Similar to PRFs, except replace the function with a permutation

Definition. A function family  $F_\lambda: \mathcal{K}_\lambda \times \mathcal{X}_\lambda \rightarrow \mathcal{X}_\lambda$  is a secure PRP if

1. For all keys  $k \in \mathcal{K}_\lambda$ ,  $F_\lambda(k, \cdot)$  is a permutation on  $\mathcal{X}_\lambda$
2. The function family  $F_\lambda$  is pseudorandom ( $F_\lambda(k, \cdot)$  is computationally indistinguishable from  $f(\cdot)$  when  $k \in \mathcal{K}_\lambda$  and  $f \in \text{Perm}[\mathcal{X}_\lambda]$ )

set of all permutations on  $\mathcal{X}_\lambda$  ↖

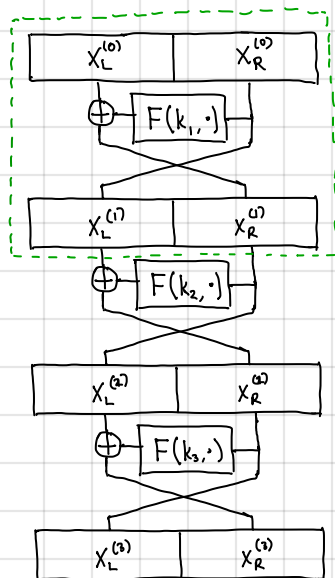
PRP  $\Rightarrow$  PRF. (PRF switching lemma). Suppose  $F_\lambda: \mathcal{K}_\lambda \times \mathcal{X}_\lambda \rightarrow \mathcal{X}_\lambda$  is a secure PRP. Then, for any efficient adversary making at most  $Q$  queries,

$$|\text{PRPAdv}[A, F] - \text{PRFAdv}[A, F]| \leq \frac{Q^2}{|\mathcal{X}_\lambda|}$$

Intuition. If a PRP is over a large (i.e., super-polynomial) domain, then a PRP is as good as a PRF. Follows from fact that a random permutation behaves like a random function (up to the birthday bound - e.g., when a collision occurs)

PRF  $\Rightarrow$  PRP. (Luby-Rackoff). We can construct a PRP from any PRF using the 3-round Feistel construction:

Input:  $X = X_L^{(0)} \parallel X_R^{(0)}$  (viewed as bit-strings)



Feistel round

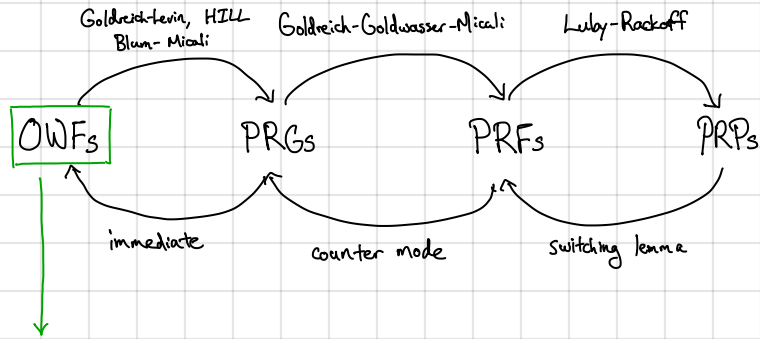
Luby-Rackoff (Informal). If  $F$  is a secure PRF, then the 3-round Feistel network (shown to the left) is a secure PRP.

Important: independent "round" keys  $(k_1, k_2, k_3)$  used in the Feistel network

This construction is invertible.  
 (basic design of DES block cipher)

The story so far: Symmetric primitives: OWFs, PRGs, PRFs, PRPs

All of these primitives are closely related:



"simplest" notion that captures cryptographic hardness (often considered the fundamental primitive in symmetric cryptography)