Instructor: Gorjan Alagic (<u>galagic@umd.edu</u>); ATL 3102, office hours: by appointment **Textbook:** *Introduction to Modern Cryptography*, Katz and Lindell;

Webpage: <u>alagic.org/cmsc-456-cryptography-spring-2020/</u> (slides, reading);
Piazza: piazza.com/umd/spring2020/cmsc456
ELMS: active, slides and reading posted there, assignments will be as well.
Gradescope: active, access through ELMS.

<u>Check these setups asap, and let me know if you run into issues!</u>

TAs (Our spot: shared open area across from IRB 5234)

- Elijah Grubb (egrubb@cs.umd.edu) 11am-12pm TuTh (Iribe);
- Justin Hontz (jhontz@terpmail.umd.edu) 1pm-2pm MW (Iribe);

Additional help:

- Chen Bai (cbai1@terpmail.umd.edu) 3:30-5:30pm Tu (2115 ATL, starting Feb 4)
- Bibhusa Rawal (bibhusa@terpmail.umd.edu) 3:30-5:30pm Th (2115 ATL, starting Feb 6)

RECAP: THE BIG IDEA

Shannon's theorem: if you want perfect secrecy, one-time pad is as good as it gets. Limits of one-time pad:

- can only send one message;
- message cannot be longer than the shared key;
- what if you can't share key in advance?
- what if Eve is allowed to *change* messages?

solution: computationally-secure encryption

later in the course

Obviously:

- the crypto we use every day does not suffer from these drawbacks;
- ... (Shannon's theorem) it follows that we must relax perfect secrecy; What to do?
 - 1. allow adversaries to succeed with exponentially-small probability (roughly);
 - 2. allow adversaries to succeed after exponential time (roughly);

Claim: if done right, no "real" security loss, but huge gain in features!



(To decrypt: reverse the arrows from ciphertext to plaintext.)

PRG scheme

Key generation :	sample $k \leftarrow \{0,1\}^n$;
Encryption :	$\mathbf{Enc}_k(m) = m \oplus \mathbf{G}(k);$
Decryption :	$\mathbf{Dec}_k(c) = c \oplus \mathbf{G}(k)$.

RECAP: COMPUTATIONALLY-SECURE ENCRYPTION

Idea: use "good enough" randomness in OTP (instead of perfect.)

How to make this formal?

- 1. define pseudorandomness;
- 2. define notion of "efficient" and "inefficient";
- 3. define a relaxed notion of security;
- 4. prove that this scheme works.



Efficient algorithms: probabilistic, polynomial-time (PPT.)

• running time: at most polynomial in the input size;

AND

• success probability: at least inverse-polynomial in input size;

Think: "Achieves a noticeable success probability, in a reasonable amount of time."

Inefficient algorithms:

• running time: superpolynomial in input size (i.e., exponential or nearly so.);

OR

• success probability: at most inverse-superpolynomial (i.e., negligibile.)

Think: "can achieve noticeable success probability <u>only</u> by spending an UNreasonable amount of time."



Cryptographic pseudorandomness

Definition. A **pseudorandom generator** is a deterministic, polynomial-time algorithm *G* satisfying the following:

- 1. (expansion) $G: \{0,1\}^n \to \{0,1\}^{\ell(n)}$ for some fixed polynomial ℓ satisfying $\ell(n) > n$ for all n.
- 2. (pseudorandomness) for every PPT algorithm **D**,

$$\left|\Pr_{s \leftarrow \{0,1\}^n} \left[\boldsymbol{D}(\boldsymbol{G}(s)) = 1 \right] - \Pr_{r \leftarrow \{0,1\}^{\ell(n)}} \left[\boldsymbol{D}(r) = 1 \right] \right| \le \operatorname{negl}(n).$$



RECAP. "IND' SECRECY

Indistinguishability experiment (IND).

- 1. A outputs two messages m_0, m_1 with $|m_0| = |m_1|$;
- 2. We sample a key $k \leftarrow \text{KeyGen}$, and a coin $b \leftarrow \{0,1\}$; then we give A the ciphertext $c \leftarrow \text{Enc}_k(m_b)$;
- 3. A outputs a bit b'.

We say **A** wins if b = b'.

 $A \xrightarrow{m_0} \operatorname{Enc}_k \xrightarrow{c} A \xrightarrow{b'} b'$ $A \xrightarrow{m_1} \xrightarrow{c} A \xrightarrow{b'} b'$ $A \xrightarrow{m_1} \xrightarrow{c} A \xrightarrow{b'} b'$

Definition. An encryption scheme (**KeyGen**, **Enc**, **Dec**) has **indistinguishable ciphertexts** if, for every PPT adversary *A*,

$$\Pr[A \text{ wins IND}] \le \frac{1}{2} + \operatorname{negl}(n).$$

The PRG scheme is secure.

Theorem. The PRG scheme has indistinguishable ciphertexts.

Proof.

Proof by contradiction: "if PRG scheme is broken, then the underlying PRG is broken."

More concretely:

- Let G be a PRG, and let $\Pi(G)$ be the PRG scheme using G;
- Given a PPT \boldsymbol{A} that wins IND game against $\Pi(\boldsymbol{G})$...
- ... we build a PPT **D** who distinguishes the output of **G** from random:



PRG ENCRYPTION SECURITY PROOF

"If there's an attacker **A** that can win the IND game, then there's a distinguisher **D** against **G**."



assumption!

Let's analyze D.

Two cases:

- (1.) r is uniformly random in $\{0,1\}^{\ell(n)}$.
 - Then **D** is an <u>exact</u> simulation of this IND game:
 - A plays against the one-time pad with keylength $\ell(n)$;
 - by perfect secrecy of OTP, A loses: Pr[b = b'] = 1/2;
 - it follows that Pr[z = 1] = 1/2.

(2.) $r = \mathbf{G}(s)$ for uniformly random $s \in \{0,1\}^n$.

- Then **D** is an <u>exact</u> simulation of this IND game:
- *A* plays against the PRG scheme with PRG *G*;
- by assumption, A wins noticeably, i.e. $Pr[b = b'] \ge 1/2 + 1/p(n)$ for some polynomial p;
- it follows that $\Pr[z = 1] = 1/2 + 1/p(n)$.

$$\left|\Pr[\boldsymbol{D}(\boldsymbol{G}(s)) = 1] - \Pr[\boldsymbol{D}(r) = 1]\right| = \left|\left(\frac{1}{2} + \frac{1}{p(n)}\right) - \frac{1}{2}\right| = \frac{1}{p(n)}$$

IV. PSELDORANDOM FUNCTIONS

Reading: p.71-95

PRGs enable:

- *fixed-length* encryption with poly-size messages;
- with stateful schemes, allows multiple messages, up to a total $\ell(n)$ bits.



- but how do you decrypt? What if the ciphertexts arrive out of order?
- and what if you don't want to keep state? (a potential attack avenue.)
- and what if you want to send *arbitrarily many* messages? 🔆

MORE POWERFUL ATTACKS

So far...

- our model still grants adversary very little power;
- they are only a passive observer;
- in real world, they can do much more!

For example: they can interrogate systems.

- try to connect to some authorized system;
- guess passwords and see what happens;
- send transmissions and see if they decrypt to something;
- use real world power over parties to get them to send encrypted messages.

How do we capture things like this in our framework? Oracle algorithms.



Oracle algorithm:

- same as a regular algorithm, but can "invoke" a special subroutine;
- this subroutine simply evaluates some function, and has no other effect;
- the subroutine behaves like a "black box" or an "oracle";
- it costs the algorithm only one timestep to query the oracle.

Think: invoking a compiled library method/function when programming.

Example:

- recall the class NP and polynomial-time reductions;
- there is a poly-time algorithm which, given a SAT oracle, solves the Traveling Salesman Problem;





MORE POWERFUL ATTACKS

Oracle algorithms for us:

- allow adversaries to "query" cryptosystem in various ways;
- they can use whatever they learn to try to devise a better attack.

Why give away more power?

- model real situations more accurately;
- eliminate unnecessary weaknesses in system;
- explore ultimate limits of what crypto can do.



MORE POWERFUL ATTACKS

Oracle algorithms for us:

- allow adversaries to "query" cryptosystem in various ways;
- they can use whatever they learn to try to devise a better attack.

Why give away more power?

- model real situations more accurately;
- eliminate unnecessary weaknesses in system;
- explore ultimate limits of what crypto can do.

For example:

- give Eve access to encryption!
- can we still have secrecy?



PSELDORANDOM FUNCTIONS

A more powerful primitive: pseudorandom functions.

PRG

- public algorithm **G**;
- if you plug in a random string...
- ... you get back a longer, pseudorandom string.

```
s \leftarrow \{0,1\}^n
```



PRF

- keyed algorithm **F** (kind of like encryption);
- if you plug in *any* string..
- ... you get back a pseudorandom string.

$$k \leftarrow \{0,1\}^n$$



- query anywhere, as many times as you want...
- ... and still the output looks random!

PSELDORANDOM FUNCTIONS



PSELDORANDOM FUNCTIONS

What does a PRF do?



<u>"Real" experiment:</u>

- pick a key $k \leftarrow \{0,1\}^n$;
- "put function F_k in a box";
- give box to an adversary;
- adversary cannot open box...
- ... but can plug in any input, and get output.

<u>"Ideal" experiment:</u>

- pick a completely random function **R**;
- "put function **R** in a box";
- give box to an adversary;
- adversary cannot open box...
- ... but can plug in any input, and get output.



Formal definition: "oracle adversaries can't tell it apart from totally random"





Output of PRF looks random!

- it's indistinguishable from uniformly random;
- we know from OTP that uniformly random ciphertexts are good;
- so let's build encryption!
- to encrypt, just apply the PRF: $\mathbf{Enc}_k(m) = \mathbf{F}_k(m)$.

Easy to check: adversary can't distinguish output from random.

Does this work?

PRFs vs PRGs

PRFs vs PRGs.

Can you build a PRG from a PRF?

Easy:

- Let $F: \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n$ be a PRF;
- Build a $G: \{0,1\}^n \rightarrow \{0,1\}^{2n}$ like this:
- $\boldsymbol{G}(s) \coloneqq \boldsymbol{F}_s(00 \cdots 00) || \boldsymbol{F}_s(00 \cdots 01)$
- Easy to extend to arbitrary-length **G**.

Note:

- in PRG setting, seed is uniformly random...
- so we can use it for the key.

So are PRFs stronger than PRGs?

PRFs vs PRGs

Can you build a PRF from a PRG?

Let $G: \{0,1\}^n \to \{0,1\}^{2n}$ be a PRG, and define: $G_0: \{0,1\}^n \to \{0,1\}^n$ by $G_0(x) = G(x)|_1^n$ "apply G, take left half" $G_1: \{0,1\}^n \to \{0,1\}^n$ by $G_0(x) = G(x)|_{n+1}^{2n}$ "apply G, take right half"

Define \mathbf{F}_k : $\{0,1\}^n \times \{0,1\}^n \rightarrow \{0,1\}^n$ by

$$\boldsymbol{F}_k(\boldsymbol{x}) = G_{\boldsymbol{x}_n}(G_{\boldsymbol{x}_n}(\cdots(G_{\boldsymbol{x}_1}(G_{\boldsymbol{x}_0}(k))\cdots))$$

Theorem (Goldwasser, Goldreich, Micali '88): F_k is a pseudorandom function.

PRFs vs PRGs

Can you build a PRF from a PRG? "GGM PRF" Let $G: \{0,1\}^n \rightarrow \{0,1\}^{2n}$ be a PRG, and define: $G_0: \{0,1\}^n \to \{0,1\}^n$ by $G_0(x) = G(x)|_1^n$ $\boldsymbol{F}_{k}(\boldsymbol{x}) = G_{\boldsymbol{x}_{n}}(G_{\boldsymbol{x}_{n}}(\cdots(G_{\boldsymbol{x}_{1}}(G_{\boldsymbol{x}_{0}}(\boldsymbol{k}))\cdots)$ $G_1: \{0,1\}^n \to \{0,1\}^n$ by $G_0(x) = G(x)|_{n+1}^{2n}$ k **Example:** suppose n=3 G • compute $F_k(101)$. G $x_0 = 1$ $x_1 = 0$ G $x_2 = 1$

PRF ENCRYPTION

What's a PRF good for?

Lots of things! Like really powerful encryption:

Construction (PRF encryption). Let $F: \{0,1\}^n \times \{0,1\}^m \rightarrow \{0,1\}^\ell$ be a PRF. Define a scheme:

- **KeyGen**: sample a PRF key $k \leftarrow \{0,1\}^n$;
- Enc: on input a message $m \in \{0,1\}^{\ell}$, sample $r \leftarrow \{0,1\}^m$ and output $(r, F_k(r) \oplus m)$;
- **Dec**: on input a ciphertext (r, c), output $c \oplus F_k(r)$.



Some properties

- at its core, there's still OTP
- can send arbitrarily-many messages!
- encryption is now a *randomized* algorithm



CPAATTACKS

What about security?

It turns out, we get a big upgrade there too.

Chosen plaintext attacks (CPA):

- at anytime during their attack...
- ... adversary can ask for an encryption of any message!
- "chosen plaintext";



Examples.

- against OTP: query $\operatorname{Enc}_k(0^n) = 0^n \oplus k = k$. Complete key recovery with one query!
- against PRG scheme: can only encrypt a limited number of times...
- ... so adversary can just use them all up!

RECALL: "IND' SECRECY

Indistinguishability experiment (IND).

- 1. A outputs two messages m_0, m_1 with $|m_0| = |m_1|$;
- 2. We sample a key $k \leftarrow \text{KeyGen}$, and a coin $b \leftarrow \{0,1\}$; then we give A the ciphertext $c \leftarrow \text{Enc}_k(m_b)$;
- 3. A outputs a bit b'.

We say **A** wins if b = b'.

 $A \xrightarrow{m_0} \operatorname{Enc}_k \xrightarrow{c} A \xrightarrow{b'} b'$ $A \xrightarrow{m_1} \xrightarrow{c} A \xrightarrow{b'} b'$ $A \xrightarrow{m_1} \xrightarrow{c} A \xrightarrow{b'} b'$

Definition. An encryption scheme (**KeyGen**, **Enc**, **Dec**) has **indistinguishable ciphertexts** if, for every PPT adversary **A**,

$$\Pr[A \text{ wins IND}] \le \frac{1}{2} + \operatorname{negl}(n).$$

IND-CPA

Indistinguishability under Chosen Plaintext Attack. INDCPA experiment:

- 1. Sample a key $k \leftarrow$ KeyGen;
- 2. Give adversary oracle access to Enc_k ;
- 3. A^{Enc_k} outputs two messages m_0, m_1 with $|m_0| = |m_1|$;
- 4. Sample a coin $b \leftarrow \{0,1\}$; give **A** ciphertext $c \leftarrow \mathbf{Enc}_k(m_b)$;
- 5. A^{Enc_k} outputs a bit b'.

We say **A** wins if b = b'.



Definition. An encryption scheme (KeyGen, Enc, Dec) is IND-CPA if, for every PPT adversary A,

$$\Pr[A \text{ wins INDCPA experiment}] \leq \frac{1}{2} + \operatorname{negl}(n).$$

End of Lecture 3.