

Welcome to CS 355! - Spring 2026

"Advanced" Topics in Cryptography

↑ a.k.a. we have prereqs: - CS 255 or equiv.
- interest and enthusiasm :)

not to scare you!

Clarification: CS 355 is a "topics" class. Normally, taught by Dan Boneh's PhD students
Last year was an exception: Dan himself taught a course on ZK proofs.
This year: back to normal...

2025-26 Instructors

- Trisha Datta
- Aditi Partap
- Ro Nema ← me!

TA: Jasdeep Sidhu

Today's Agenda

- About CS 355
- Course Logistics
- One-Way Functions (OWFs)
- Pseudorandom Generators (PRGs)
- Hardcore Bits

Course Goals

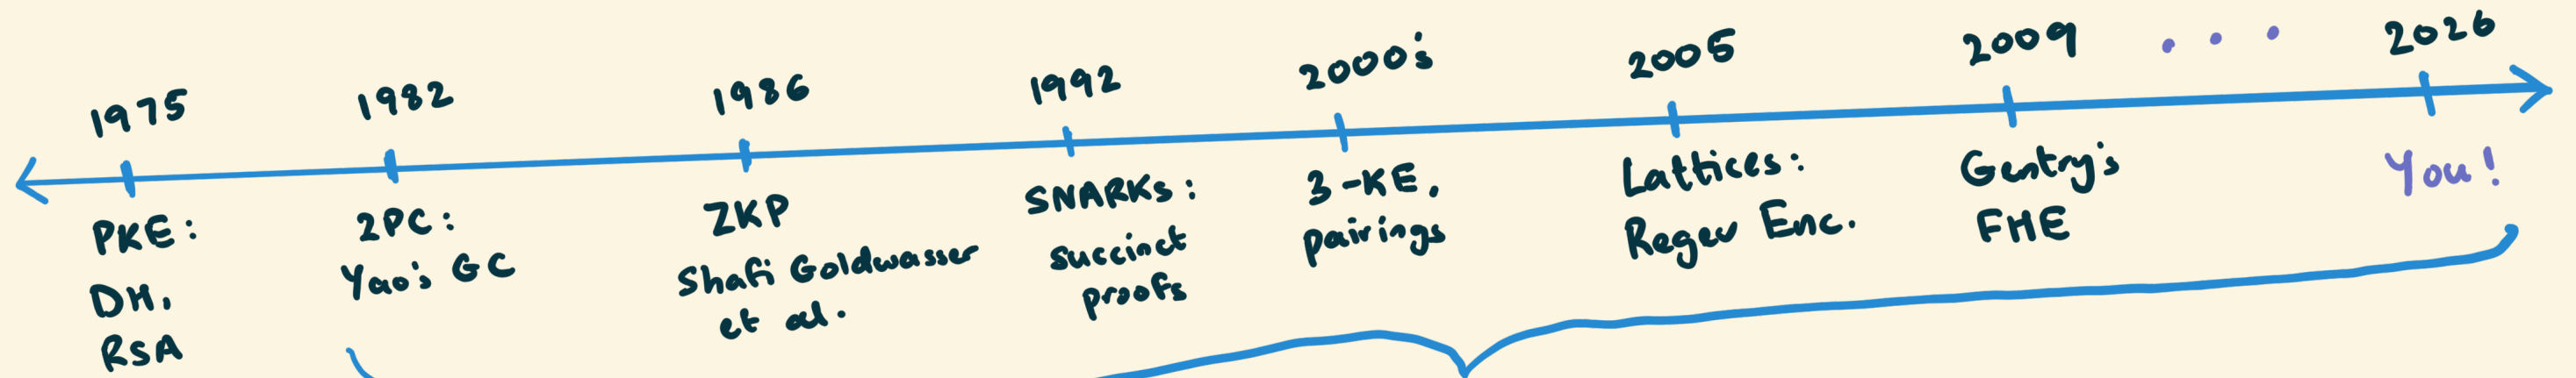
First advanced crypto course:

- learn techniques and formalisms
- understand open research problems
- prepare you for research!

Or ... last advanced crypto course:

- applications of cutting-edge crypto
- digest new papers
- prepare you to use crypto to "change the world"!

Historical Perspective



don't worry, stuff happened here too

CS 255
focus on secure communication

CS 355
focus on secure computation

About CS355

Topics

1. Foundations
2. Cryptanalysis
3. Elliptic Curve Cryptography
4. Zero Knowledge Proofs
5. Multiparty Computation
6. Lattice-based Cryptography

Why Take CS355?

- Cryptography is about the balance of power and privacy
- Widely deployed
- Beautiful math!

Why NOT Take CS 355?

- We assume mathematical maturity
 - Should not be your first proof-based course
 - possible prereqs. : CS 255, 265, 254 ; MATH 120, ...
- can be time-consuming

Course Logistics

Website: cs355.stanford.edu (we don't use Canvas)

Communication: ↙ don't spoil problems!

- Ed: Q&A, announcements
- Anonymous Feedback Form: welcome throughout the quarter :)
- cs355@cs.stanford.edu: email all 3 instructors
↖ you can also privately message us on Ed.

Lectures: • in-person, not recorded

- attendance highly encouraged
- notes available after each class
- no textbook; supplemental readings available
↖ but there is Boneh & Shoup: cryptobook.us

Office Hours: • see course calendar: both TA and Instructors have OH

- group work environment

Problem Sets :

- due every 2 weeks (5 total : each 20% of total grade)
- due on Fridays, 6 pm on Gradescope
- LaTeX mandatory. Template available. Overleaf.com works well!
- PSet 1 is out! Due Apr 10
- 3 late days total, at most 1 day for Pset 5
- No exams :)

Collaboration, AI, Academic Honesty

- encouraged to work together BUT write solutions by yourself
- list collaborators (obviously, no penalty for collaborating)
- AI/LLMs, in this class, are considered knowledgeable (if fallible) "friends" outside of this class (they are not collaborators). Thus, you are allowed to interact with them the same way you would interact with such friends

OK: build intuition, get ideas, learn a topic, et cetera

NOT OK: solve an entire problem, write up the solution, rob you of the joy of learning :(

Hardness & Reductions

- hardness and reductions are asymptotic
- honest parties are "efficient" (some polynomial in λ)
- attacks are "inefficient" (super-polynomial or even exponential in λ)
 - e.g. decryption with secret key is efficient
BUT intractable/inefficient without it

tunable security parameter



• Security ~~Proofs~~ Reductions!

Assumption.

Problem is hard \Rightarrow Cryptosystem is secure

Problem is easy \Leftarrow (efficient) Attack on cryptosystem
(i.e., efficient solution to problem)



e.g. factoring, discrete log, CDH, DDH,
"AES is a block cipher"

This week: Build crypto from one of the most minimal assumptions
"One-Way Functions"

We'll show: OWFs \Rightarrow PRGs \Rightarrow PRFs \Rightarrow Block ciphers

Today: OWF \Rightarrow PRG₊₁: PRGs with only 1-bit stretch

One-Way Functions (OWFs)

Let $X := \{X_\lambda\}_{\lambda \in \mathbb{N}}$ and $Y := \{Y_\lambda\}_{\lambda \in \mathbb{N}}$ be families of finite sets indexed by sec. param. λ
 \leftarrow input space \leftarrow output space natural numbers

e.g. $X_\lambda = \{0, 1\}^{\lambda}$
 $Y_\lambda = \{0, 1\}^{P(\lambda)}$
for some poly. P

Family of functions: $F = \{f_\lambda : X_\lambda \rightarrow Y_\lambda\}_{\lambda \in \mathbb{N}}$

intuitively, f_λ is "easy to compute"
BUT "hard to invert"

Formally,

we're leaving some stuff out to simplify the presentation for this course

Def. F is a one-way function family if

(1) for all λ , f_λ is deterministic and computable in polynomial time
"easy to compute"

(2) for all probabilistic polynomial-time (PPT) algorithms A ,
efficient
there exists a "negligible" function $\text{negl}(\cdot)$ such that

$$\Pr_{x \leftarrow \mathcal{X}_\lambda} [f_\lambda(A(f_\lambda(x))) = f_\lambda(x)] \leq \text{negl}(\lambda)$$

A 's randomness \rightarrow $x \leftarrow \mathcal{X}_\lambda$

$f_\lambda(A(f_\lambda(x))) = f_\lambda(x)$ "hard to invert" (on average)

$\text{negl}(\lambda)$ what does this mean??

Q: this formulation vs. $x = A(f_\lambda(x))$.
What's the difference and why?

For the remainder of the course:

we'll drop the λ index when it's clear...

λ means λ is "fixed" for the sake of presentation

Negligible Function: formalizing "very small" in cryptography when dealing with asymptotics

A function g is negligible if

$$g(x) \in o(x^{-c}) \quad \forall c \in \mathbb{N}$$

← this is little- o not Big- O ! although in this case it doesn't matter...

"asymptotically smaller than the inverse of any polynomial"

Examples. $2^{-\lambda}$, $2^{-\lambda/2}$, $2^{-\sqrt{\lambda}}$ are all negligible.

$$\parallel \lambda^{\log \lambda (2^{-\lambda})} = \lambda^{-\lambda \log \lambda (2)} = \lambda^{-\lambda \cdot \frac{\log_2(2)}{\log_2(\lambda)}} = \lambda^{-\frac{\lambda}{\log_2(\lambda)}}$$

← super-constant

$$\Rightarrow o(\lambda^{-c}) \quad \forall c \in \mathbb{N},$$

grows slower than any polynomial

What does " $f(\lambda) \leq \text{negl}(\lambda)$ " mean?

↳ abuse of notation... just means f is negligible in λ

Helpful identities. (1) $f \leq \text{negl}$, $g \leq \text{poly} \Rightarrow f \cdot g \leq \text{negl}$

(2) $f > \text{negl}$, $g \leq \text{poly} \Rightarrow f/g > \text{negl}$

← not negl

(2) "Hard to invert" as a game:

$$f_\lambda: X_\lambda \rightarrow Y_\lambda$$

Challenger

Adversary A

$$x \xleftarrow{\$} X_\lambda$$

$$\xrightarrow{f_\lambda(x)}$$

$$\xleftarrow{\hat{x}}$$

$$\text{OWFAdv}[A, F](\lambda) = \Pr_A [f_\lambda(x) = f_\lambda(\hat{x})]$$

F is a one-way function family if \forall PPT A ,
there exists negligible function negl s.t.

$$\text{OWFAdv}[A, F](\lambda) \leq \text{negl}(\lambda)$$

Q: Can F be one-way if $X_\lambda = \{0, 1\}^{\log(\lambda)}$?

A: No! Regardless of f_λ , consider (PPT) A that outputs $\hat{x} \xleftarrow{\$} X_\lambda$.

It has advantage $\geq 1/\lambda \notin \text{negl}(\lambda)$

Candidate OWFs

(1) $X_\lambda = (\kappa_1, \kappa_2)$ where κ_1, κ_2 are λ -bit primes

$$Y_\lambda = \mathbb{N}$$

$$f_\lambda(\kappa_1, \kappa_2) := \kappa_1 \cdot \kappa_2$$

Factoring

(2) $f_{p,g}(x) := g^x \pmod{p}$

p is λ -bit prime,
 $g \in \mathbb{Z}/p^*$, $x \in \mathbb{Z}/p-1$

Modular Exponentiation or
DLOG

(3) $f(x_1, \dots, x_n, S \subseteq [n]) = (x_1, \dots, x_n, \sum_{i \in S} x_i)$

Subset Sum

(4) "Levin's One-Way Function": f_L is OWF $\Leftrightarrow \exists$ OWF

Pseudo-Random Generators

$G := \{ G_\lambda : \{0,1\}^\lambda \rightarrow \{0,1\}^{\ell(\lambda)} \}$ Takes random input and generates longer input that "looks random"



How do we formalize "looks random"?

The notion of (Computational) Indistinguishability!

Def. G is a PRG family if \forall PPT A ,

$$\left| \Pr_A [A(G_\lambda(s)) = 1 : s \leftarrow \{0,1\}^\lambda] - \Pr_A [A(r) = 1 : r \leftarrow \{0,1\}^{\ell(\lambda)}] \right| \leq \text{negl}(\lambda)$$

PRGAdv^{!!}[A, G](λ)

Cannot distinguish between $G(s)$ and random $\ell(\lambda)$ -bit string
↑
random λ -bit string

OWFs \Rightarrow PRG with $\ell(\lambda) = \lambda + 1$

Only one-bit stretch.

Assume OWF $f: \{0,1\}^\lambda \rightarrow \{0,1\}^\lambda$

First Attempt.

Just leak the first bit of input

$$g(x) = x_1 \parallel f(x)$$

$x_1 \parallel x_2 \parallel \dots \parallel x_\lambda$ \parallel is concatenation

$$1 \parallel 0 = 10$$

From now on, we're sweeping the family formalization under the rug. We'll drop the λ subscript as well.

g is not necessarily a PRG...

Turns out it's still one-way! **why?**

Pf. By contradiction.

Suppose g is not one-way.

Then \exists PPT A s.t.

$$\Pr [g(A(g(x))) = g(x)] > \text{negl}(\lambda)$$

We will use A to construct PPT B s.t.

$$\Pr [f(B(f(x))) = f(x)] > \text{negl}(\lambda)$$

$B(y)$: Sample $b \leftarrow \{0,1\}$;
Output $A(b \parallel y)$.

Intuition. b = first bit of input with prob. $1/2$.

If that's the case, A outputs inverse with non-negl. prob and B "wins".

Formally, $\Pr [f(B(f(x))) = f(x)]$

$$\geq \Pr [f(B(f(x))) = f(x) \wedge b = x_1]$$

$$= \Pr [f(B(f(x))) = f(x) \mid b = x_1] \cdot \Pr [b = x_1]$$

$$= \Pr [f(A(x_1 \parallel f(x))) = f(x)] \cdot 1/2$$

$$\geq \Pr [g(A(g(x))) = g(x)] \cdot 1/2 > \text{negl}(\lambda)$$

$> \text{negl}(\lambda) \Rightarrow f$ is not one-way. \downarrow

What went wrong? x_1 can be easy to guess given the output...
Can we derive a bit that is hard to guess?

We would like a predicate on the input that hides it even when given the output of the OWF!

Def. An efficient predicate $B: \{0,1\}^\lambda \rightarrow \{0,1\}$ is "hard-core" for a OWF $f: \{0,1\}^\lambda \rightarrow Y$ if for all PPT A ,

$$\Pr[A(f(x)) = B(x)] \leq \frac{1}{2} + \text{negl}(\lambda)$$



not much better than random guessing

$$\left| \Pr[A(f(x)) = B(x)] - \frac{1}{2} \right| \leq \text{negl}(\lambda)$$

!!
HCBAdv $[A, B, f]$

We also call B a Hard-core Bit (HCB) of f

Theorem [Goldreich-Levin 89]. "Every OWF can be transformed into another OWF with a HCB"

Idea. Take a random linear combination of the bits of x

Where is the randomness coming from? From the input!

Let f be a OWF.

Define $f'(x, r) := (f(x), r)$ $|r| = |x| = \lambda$

f' is also a OWF. Why?

Claim. $B(x, r) := \sum_{i \in \lambda} x_i \cdot r_i \pmod{2}$ is a HCB of f' .

Proof won't be covered in class. Resources will be available.

Involves non-trivial probability arguments. Proof usually built up by relaxing assumptions on the error in reduction.

In class: given HCB we can construct a PRG with 1-bit stretch!

Okay, I lied... We'll show something similar but very close.

We'll also assume our OWF is a permutation a.k.a. OWP

"one-way permutation"

Given OWP f and its HCB B^* ,

Construct PRG,

$$G(x) = f(x) || B(x)$$

* if f is OWP,
then f' is also OWP
as defined in Goldreich-Levin Thm.

Thm. $f: \{0,1\}^\lambda \rightarrow \{0,1\}^\lambda$ OWP with HCB B

\Downarrow
 $G: \{0,1\}^\lambda \rightarrow \{0,1\}^{\lambda+1}$ is a PRG.

Pf. Suppose not.

$\Rightarrow \exists$ PPT (Adv.) A s.t. $\text{PRGAdv}[A, G] > \text{negl}(\lambda)$

i.e., $\textcircled{1} \left| \Pr_{x,A} [A(f(x) || B(x)) = 1] - \Pr_{r,A} [A(r) = 1] \right| > \epsilon(\lambda)$ where $\epsilon(\lambda)$ is a non-negl function.

in words, A can distinguish between random and $f(x) || B(x)$

How do we use this? Notice that we can write $r \in \{0,1\}^{\lambda+1}$ as

$$r = y || b \quad \text{where } y \in \{0,1\}^\lambda, b \in \{0,1\}$$

Since f is a permutation,

$$r = f(x) || b \quad \text{for some } x.$$

Moreover $x \leftarrow \{0,1\}^\lambda, b \leftarrow \{0,1\} \Rightarrow r$ is uniformly random in $\{0,1\}^{\lambda+1}$

$$\text{So, } \Pr_{r, A} [A(r) = 1] = \Pr_{x, b, A} [A(f(x) || b) = 1]$$

← when $b = B(x)$, this is just the first term

$$\begin{aligned} \text{We can simplify: } \Pr [A(f(x) || b) = 1] &= \Pr [b = B(x)] \cdot \Pr [A(f(x) || b) = 1] \\ &+ \Pr [b = \overline{B(x)}] \cdot \Pr [A(f(x) || b) = 1] \\ &= \frac{1}{2} \cdot \Pr [A(f(x) || B(x)) = 1] + \frac{1}{2} \cdot \Pr [A(f(x) || \overline{B(x)}) = 1] \end{aligned}$$

Plugging this in ①,

$$\textcircled{1}: \left| \Pr [A(f(x) || B(x)) = 1] - \frac{1}{2} \Pr [A(f(x) || B(x)) = 1] - \frac{1}{2} \Pr [A(f(x) || \overline{B(x)}) = 1] \right| > \epsilon(\lambda)$$

$$\Rightarrow \left| \Pr [A(f(x) || B(x)) = 1] - \Pr [A(f(x) || \overline{B(x)}) = 1] \right| > 2\epsilon(\lambda) \leftarrow \text{still non-negl.}$$

What does this tell us? We can use A to distinguish if the HCB is flipped!

Construct adv. for HCB, β as follows:

$B(y)$: Sample $b \leftarrow \{0, 1\}$
 IF $A(y || b) = 1$: output b
 Else: output \overline{b}

Claim. \mathcal{B} "guesses" HCB with non-negl. probability

PF. $\left| \Pr[\mathcal{B}(f(x)) = \mathcal{B}(x)] - \frac{1}{2} \right| = \left| \Pr[\mathcal{B}(f(x)) = \mathcal{B}(x) \mid b = \mathcal{B}(x)] \cdot \Pr[b = \mathcal{B}(x)] + \Pr[\mathcal{B}(f(x)) = \mathcal{B}(x) \mid b = \overline{\mathcal{B}(x)}] \cdot \Pr[b = \overline{\mathcal{B}(x)}] - \frac{1}{2} \right|$

\nearrow
 need to show this is non-negl.

$$= \left| \Pr[\mathcal{A}(f(x) \parallel \overset{b}{\mathcal{B}(x)}) = 1] \cdot \frac{1}{2} + \Pr[\mathcal{A}(f(x) \parallel \overset{b}{\overline{\mathcal{B}(x)}}) = 0] \cdot \frac{1}{2} - \frac{1}{2} \right|$$

(conditionally)

$$= \frac{1}{2} \left| \Pr[\mathcal{A}(f(x) \parallel \mathcal{B}(x)) = 1] + (1 - \Pr[\mathcal{A}(f(x) \parallel \overline{\mathcal{B}(x)}) = 1]) - 1 \right|$$

$$= \frac{1}{2} \left| \Pr[\mathcal{A}(f(x) \parallel \mathcal{B}(x)) = 1] - \Pr[\mathcal{A}(f(x) \parallel \overline{\mathcal{B}(x)}) = 1] \right|$$

$$> \frac{1}{2} \cdot (2\varepsilon(\lambda)) = \varepsilon(\lambda)$$

Therefore, $\text{HCBAdv}[\mathcal{B}, \mathcal{B}, f] > \varepsilon(\lambda)$, non-negl. \Downarrow

Using Goldreich-Levin transformation,

$$G(x, r) := \underbrace{f(x)}_{f'(x, r) \text{ is still OWF}} \parallel r \parallel \mathcal{B}(x, r)$$

\swarrow HCB of f'

$$(x, r) \in \{0, 1\}^{2\lambda}$$

G : 2λ bits to $2\lambda + 1$ bits!