So far in this course, we have shown that

PRF $\Rightarrow$ CPA-secure encryption $\Rightarrow$ authenticated encryption
$\Rightarrow$ secure MAC

conceptually "simpler" object

From HW1, we saw how to construct a PRF from a (length-doubling) PRG

can be built from any PRG with 1-bit stretch

Question: Can we distill this further? Can we base symmetric cryptography on an even simpler primitive?

- Cryptography is about exploiting some kind of asymmetry: we want an operation that is "easy" for honest users, but hard for adversaries
- Suggests a notion of "hard to invert": (cannot recover seed from PRG, cannot decrypt without knowledge of secret, etc.)

Definition. A function $f: \mathcal{X} \to \mathcal{Y}$ is one-way if

1. f is efficiently computable
2. for all efficient adversaries $\mathcal{A}$:

$$\Pr[x \xleftarrow{R} \mathcal{X}, y \leftarrow \mathcal{A}(f(x)) : f(x) = f(y)] = \text{negl}(\lambda)$$

"Function is hard to invert on average"

Technically, $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda \in \mathbb{N}}$ and $\mathcal{Y} = \{\mathcal{Y}_\lambda\}_{\lambda \in \mathbb{N}}$ are sets indexed by a security parameter λ and $|\mathcal{X}_\lambda| = \text{poly}(\lambda)$.

Theorem (Håstad-Impagliazzo-Levin-Luby). OWF $\Rightarrow$ PRG [implies OWF is sufficient (and necessary) for symmetric cryptography]

We will consider a weaker statement: one-way permutation $\Rightarrow$ PRG

Definition. A function $f: \mathcal{X} \to \mathcal{X}$ is a one-way permutation if

1. f is one-way
2. f is a permutation

Goal: given a OWP $f: \mathcal{X} \to \mathcal{X}$, can we construct a PRG with one-bit stretch.

Idea: if $x \xleftarrow{R} \mathcal{X}$, then $f(x)$ is uniformly random

moreover, given $f(x)$, should be difficult to recover (all of) x $\leftarrow$ leverage this to get 1 pseudorandom bit

Definition. Let $f: \mathcal{X} \to \mathcal{Y}$ be a one-way function. Then $h: \mathcal{X} \to \mathcal{R}$ is a hard-core predicate for f if no efficient adversary can distinguish the following distributions:

$$D_0 : \{x \xleftarrow{R} \mathcal{X} : (f(x), h(x))\}$$
$$D_1 : \{x \xleftarrow{R} \mathcal{X}, r \xleftarrow{R} \mathcal{R} : (f(x), r)\}$$

If a OWP has a hard-core predicate, that immediately implies a PRG:

$$\text{PRG}(s) := f(s) \,\|\, h(s)$$

Typically, we will consider hard-core bits (i.e., $\mathcal{R} = \{0,1\}$)

Lemma. Let $f: \mathcal{X} \to \mathcal{Y}$ be a one-way function. Suppose $h: \mathcal{X} \to \{0,1\}$ is unpredictable in the following sense: for all efficient adversaries $\mathcal{A}$:

$$\left|\Pr[x \xleftarrow{R} \mathcal{X} : \mathcal{A}(f(x)) = h(x)] - \tfrac{1}{2}\right| = \text{negl}(\lambda)$$

If h is unpredictable, then it is a hard-core bit. [Note: Converse of this is immediate]

Proof. Suppose there exists an efficient A that can distinguish between $(f(x),h(x))$ and $(f(x),r)$ for $x \xleftarrow{R} X$ and $b \xleftarrow{R} \{0,1\}$ with advantage ε. We use A to build a predictor B:

1. On input $f(x)$, sample $b \xleftarrow{R} \{0,1\}$ and run A on input $((f(x), b))$.
2. If A outputs 1, then output b. Otherwise, output $1-b$.

Intuition: Suppose A is more likely to output 1 given inputs from the "hard-core bit distribution". This means that A outputs 1 if we "guess correctly."

Formally:
$$\begin{aligned}
\Pr[B(f(x)) = h(x)] &= \Pr[A(f(x), b) = h(x)] \\
&= \Pr[A(f(x), b) = 1 \mid b = h(x)]\underbrace{\Pr[b = h(x)]}_{=1/2} + \underbrace{\Pr[A(f(x), b) = 0 \mid b = 1-h(x)]}_{=1-\Pr[A(f(x),b)=1 \mid b = 1-h(x)]}\underbrace{\Pr[b = 1-h(x)]}_{=1/2} \\
&= \frac{1}{2} + \frac{1}{2}\Big(\Pr[A(f(x),b) = 1 \mid b = h(x)] - \Pr[A(f(x),b) = 1 \mid b = 1-h(x)]\Big) \\
&= \frac{1}{2} + \frac{1}{2}\Big(\underbrace{\Pr[A(f(x),h(x)) = 1]}_{\alpha} - \underbrace{\Pr[A(f(x), b) = 1 \mid b = 1-h(x)]}_{\beta}\Big)
\end{aligned}$$

Now,
$$\begin{aligned}
\varepsilon &= \big|\Pr[A(f(x),h(x)) = 1] - \Pr[A(f(x),b) = 1]\big| \\
&= \big|\alpha - \Pr[A(f(x), b) = 1 \mid b = h(x)]\Pr[b = h(x)] \\
&\qquad - \Pr[A(f(x),b) = 1 \mid b = 1-h(x)]\Pr[b = h(x)]\big| \\
&= \left|\alpha - \tfrac{1}{2}(\alpha + \beta)\right| \\
&= \tfrac{1}{2}|\alpha - \beta|
\end{aligned}$$

$$\begin{aligned}
\left|\Pr[B(f(x)) = h(x)] - \tfrac{1}{2}\right| &= \left|\tfrac{1}{2}(\alpha - \beta)\right| \\
&= \varepsilon
\end{aligned}$$

Theorem (Goldreich-Levin). Let $f: \{0,1\}^n \to \{0,1\}^m$ be a one-way function. For a string $r \in \{0,1\}^n$, define the function $h_r: \{0,1\}^n \to \{0,1\}$ where $h_r(x) = \sum r_i x_i \pmod 2$. Then the function $g(x,r) := (f(x), r)$ is one-way and h_r is a hard-core predicate for g.

Observe that if f is a OWP, then so is g

Proof Idea. One-wayness of g immediately follows from one-wayness of f. Suffices to show that h_r is hard-core. Suppose that h_r is not a hard-core predicate for g. This means that there is an adversary A that can predict h_r given $(f(x), r)$ with probability $\frac{1}{2} + \varepsilon$. We will use g to construct an adversary B that can invert f (and thus g).

Here: Suppose A succeeds with probability 1:
$$\Pr[A(g(x,r)) = h_r(x)] = 1 \quad (\text{for } x, r \xleftarrow{R} \{0,1\}^n)$$
Given $y = f(x)$, run A on inputs $\underbrace{(y, e_1), \ldots, (y, e_n)}$ where e_i is the i^{th} basis vector
$$\begin{aligned}
h_{e_i}(x) &= \langle e_i, x\rangle \bmod 2 \\
&= x_i \in \{0,1\}
\end{aligned}$$

Suppose now that A succeeds with probability $3/4 + \varepsilon$ for constant $\varepsilon > 0$:

Evaluating at $e_1, \ldots, e_n$ not guarenteed to work since A could be wrong on all of these inputs

Analysis proceeds in two steps:

integers modulo 2

1. Fix an $x \in \mathbb{Z}_2^n$. Suppose we have a function $t: \mathbb{Z}_2^n \to \{0,1\}$ where
$$\Pr[r \xleftarrow{R} \mathbb{Z}_2^n : t(r) = \langle x, r \rangle] \geq \frac{3}{4} + \varepsilon$$
We show that we can learn x by evaluating t on carefully-chosen points.
Similar to before, t could be wrong on $e_1, \dots, e_n$. Need evaluation points to be random.

Sample $r \xleftarrow{R} \mathbb{Z}_2^n$. and evaluate t at r and $e_1 + r$.
By assumption: $\Pr[t(r) = \langle x, r \rangle] \geq \frac{3}{4} + \varepsilon$
$\Pr[t(r+e_1) = \langle x, r + e_1 \rangle] \geq \frac{3}{4} + \varepsilon$ (since $r + e_1$ with $r \xleftarrow{R} \mathbb{Z}_2^n$ is uniform)
But these events are not independent: inputs are correlated!
Consider the complements: $\Pr[t(r) \neq \langle x, r \rangle] < \frac{1}{4} - \varepsilon$
$\Pr[t(r+e_1) \neq \langle x, r + e_1 \rangle] < \frac{1}{4} - \varepsilon$
$\Rightarrow$ By union bound:
$\Pr[t(r) \neq \langle x, r \rangle \text{ or } t(r+e_1) \neq \langle x, r+e_1 \rangle]$
$< \frac{1}{2} - 2\varepsilon < \frac{1}{2}$

Thus, with prob. at least $\frac{1}{2} + 2\varepsilon$, $t(r) = \langle x, r \rangle$ and $t(r+e_1) = \langle x, r + e_1 \rangle$
Set $z = t(r) + t(r + e_1)$
If $t(r) = \langle x, r \rangle$ and $t(r+e_1) = \langle x, r + e_1 \rangle$,
$$t(r+e_1) - t(r) = \langle x, r + e_1 \rangle - \langle x, r \rangle = \langle x, e_1 \rangle = x_1$$
Idea: Sample k independent pairs $(r_i, r_i \oplus e_1)$ for $r_i \xleftarrow{R} \mathbb{Z}_2^n$ and compute estimates $z_1, \dots, z_k$
Take the first bit $\hat{x}_1$ to be $\text{Majority}(z_1, \dots, z_k)$
Repeat this procedure to obtain estimates $\hat{x}_2, \dots, \hat{x}_n$. Output $\hat{x}_1 \hat{x}_2 \cdots \hat{x}_n$.

Analysis will use a Chernoff bound. Simple version for our setting:
Let $X_1, \dots, X_k \in \{0,1\}$ be independent random variables where $\Pr[X_i = Y] \geq \frac{1}{2} + \varepsilon$. Then,
$$\Pr[\text{Majority}(X_1, \dots, X_k) \neq Y] \leq 2e^{-2\varepsilon^2 k}$$
In particular, if $\varepsilon = O(1)$, $\Pr[\text{Majority}(X_1, \dots, X_k) \neq Y] \leq 2^{-O(k)}$

(when $k = O(n)$) for each bit of x

By the Chernoff bound, $\hat{x}_1 = x_1$ with probability $1 - \text{negl}(n)$. Repeating this n times yields the desired result
Total evaluations of t: $O(n^2)$

Can also set $k = O(\log n)$. Each bit correct with prob. $1 - \frac{1}{n}$, so correct with constant probability.

2. Our setting is not quite this:
$$\Pr[x, r \xleftarrow{R} \{0,1\}^n : A(f(x), r) = \langle x, r \rangle] \geq \frac{3}{4} + \varepsilon$$
randomness taken over both x and r while above analysis only looks at r.

Let's say an x is "good" if
$$\Pr[r \xleftarrow{R} \{0,1\}^n : A(f(x), r) = \langle x, r \rangle] \geq \frac{3}{4} + \frac{\varepsilon}{2}$$
If x is "good", then can recover x using above algorithm.

How many x's are good? If $\Pr[x \xleftarrow{R} \{0,1\}^n : x \text{ is "good"}]$ is non-negligible, then we have proven the claim. Algorithm B runs above decoder on A and recovers x whenever x is good, which happens with non-negligible probability.

If A succeeds on $(\frac{3}{4} + \varepsilon)$-fraction of x's, cannot have "too many" bad x's. (Averaging argument).
Suppose δ fraction of x's are bad. Then, probability of A succeeding over choice of $x, r \xleftarrow{R} \{0,1\}^n$ is at most

$$\delta\left(\frac{3}{4}+\frac{\varepsilon}{2}\right)+(1-\delta)$$
$$= 1-\frac{\delta}{4}+\frac{\delta\varepsilon}{2}$$

constant for all $\varepsilon > 0$, $\varepsilon \leq \frac{1}{4}$

Require that $1-\frac{\delta}{4}+\frac{\delta\varepsilon}{2} \geq \frac{3}{4}+\varepsilon \implies 1-\delta+2\delta\varepsilon \geq 4\varepsilon$
$\implies \delta(1-2\varepsilon) \leq 1-4\varepsilon \implies \delta \leq \frac{1-4\varepsilon}{1-2\varepsilon}$

Conclusion: At most constant fraction is "bad" so inversion will succeed on constant fraction of inputs.

HW3: Show how to go from $\frac{3}{4}+\varepsilon$ to $\frac{1}{2}+\varepsilon$ for constant $\varepsilon > 0$