

Focus thus far in the course: protecting communication (e.g., message confidentiality and message integrity)

Remainder of course: protecting computations

Zero-knowledge: a defining idea at the heart of theoretical cryptography

- ↳ Idea will seem very counter-intuitive, but surprisingly powerful
- ↳ Showcases the importance and power of definitions (e.g., "What does it mean to know something?")

with surprising implications
(DSA/ECDSA signatures based on ZK!)

We begin by introducing the notion of a "proof system"

- Goal: A prover wants to convince a verifier that some statement is true

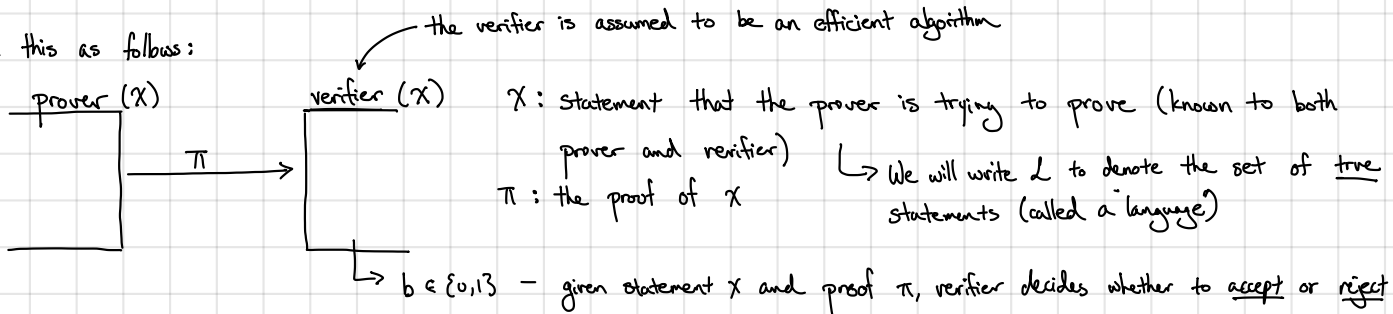
e.g., "This Sudoku puzzle has a unique solution"

"The number N is a product of two prime numbers p and q "

"I know the discrete log of h base g "

} these are all examples of statements

We model this as follows:



Properties we care about:

- Completeness: Honest prover should be able to convince honest verifier of true statements

$$\forall x \in L : \Pr[\pi \leftarrow P(x) : V(x, \pi) = 1] = 1$$

- Soundness: Dishonest prover cannot convince honest verifier of false statement

$$\forall x \notin L : \Pr[\pi \leftarrow P(x) : V(x, \pi) = 1] < 1/3$$

[Could relax requirement to allow for some error]

Important: We are not restricting to efficient provers (for now)

Typically, proofs are "one-shot" (i.e., single message from prover to verifier) and the verifier's decision algorithm is deterministic

↳ Languages with these types of proof systems precisely coincide with NP (proof of statement x is to send NP witness w)


Recall that NP is the class of languages where there is a deterministic solution-checker:

$$L \in NP \iff \exists \text{ efficiently-computable relation } R \text{ s.t.}$$

$$x \in L \iff \exists w \in \{0,1\}^{|x|} : R(x, w) = 1$$

\uparrow \uparrow \uparrow \uparrow
 statement language witness NP relation

Proof system for NP:




Perfect completeness + soundness

Going beyond NP: we augment the model as follows

- Add randomness: the verifier can be a randomized algorithm
- Add interaction: verifier can ask "questions" to the prover

Interactive proof systems [Goldwasser-Micali-Rackoff]:



Verifier randomness is critical. Otherwise, class of languages that can be recognized collapses to NP. (See HW5).

Interactive proof should satisfy completeness + soundness (as defined earlier)

We define $IP[k]$ to denote class of languages where there is an interactive proof with k messages.

We write $IP = IP[\text{poly}(n)]$ where n is the statement length

(i.e., IP is the class of languages with an interactive proof with polynomially-many rounds)


What is the power of IP ?

- For constant number of messages, seems comparable to NP
- Going from constant to polynomial number of rounds is significant!

Theorem. (Lund-Fortnow-Karloff-Nisan '90, Shamir '90) $IP = PSPACE$.

✓ the set of languages that can be checked in polynomial space

Consider following example: Suppose prover wants to convince verifier that $N = pq$ where p, q are prime (and secret).



Proof is certainly complete and sound, but now verifier also learned the factorization of N ... (may not be desirable if prover was trying to convince verifier that N is a proper RSA modulus (for a cryptographic scheme) without revealing factorization in the process

↳ In some sense, this proof conveys information to the verifier [i.e., verifier learns something it did not know before seeing the proof]

Zero-knowledge: ensure that verifier does not learn anything (other than the fact that the statement is true)

How do we define "zero-knowledge"? We will introduce a notion of a "simulator."

for a language L

Definition. An interactive proof system $\langle P, V \rangle$ is zero-knowledge if for all efficient (and possibly malicious) verifiers V^* , there exists an efficient simulator S such that for all $x \in L$:

$$\text{View}_{V^*}(\langle P, V \rangle(x)) \approx S(x)$$

random variable denoting the set of messages sent and received by V^* when interacting with the prover P on input x

What does this definition mean?

$\text{View}_{V^*}(P \leftrightarrow V^*(x))$: this is what V^* sees in the interactive proof protocol with P

$S(x)$: this is a function that only depends on the statement x , which V^* already has

If these two distributions are indistinguishable, then anything that V^* could have learned by talking to P , it could have learned just by invoking the simulator itself, and the simulator output only depends on x , which V^* already knows

↳ In other words, anything V^* could have learned (i.e., computed) after interacting with P , it could have learned without ever talking to P !


Very remarkable definition!

↖ can in fact be constructed from OWFs

More remarkable: Using cryptographic commitments, then every language $L \in \text{IP}$ has a zero-knowledge proof system.

↳ Namely, anything that can be proved can be proved in zero-knowledge!

We will show this theorem for NP languages. Here it suffices to construct a single zero-knowledge proof system for an NP-complete language. We will consider the language of graph 3-colorability.



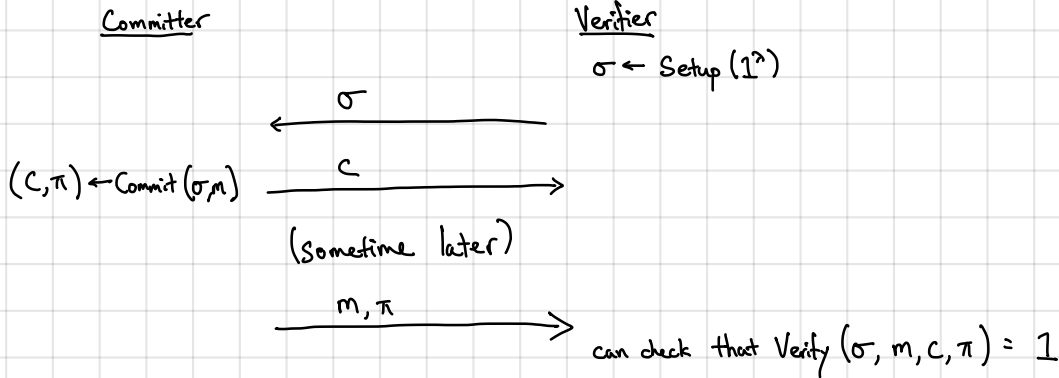
3-coloring: given a graph G , can you color the vertices so that no adjacent nodes have the same color?

cryptographic analog of a sealed "envelope"

We will need a commitment scheme. A (non-interactive) commitment scheme consists of three algorithms (Setup, Commit, Open):

- Setup(1^λ) $\rightarrow \sigma$: Outputs a common reference string (used to generate/validate commitments) σ
- Commit(σ, m) $\rightarrow (c, \pi)$: Takes the CRS σ and message m and outputs a commitment c and opening π
- Verify(σ, m, c, π) $\rightarrow 0/1$: Checks if c is a valid commitment to m (given π)

Typical setup:

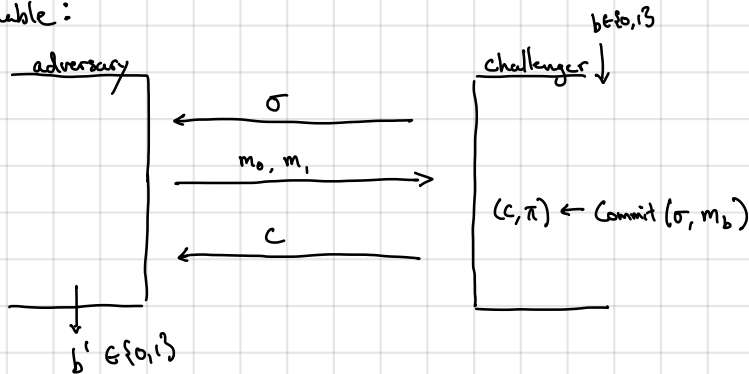


Requirements:

- Correctness: for all messages m :

$$\Pr[\sigma \leftarrow \text{Setup}(1^\lambda); (c, \pi) \leftarrow \text{Commit}(\sigma, m); \text{Verify}(\sigma, c, m, \pi) = 1] = 1$$

- Hiding: for all common reference strings $\sigma \in \{0, 1\}^n$ and all efficient A , following distributions are computationally indistinguishable:



$$|\Pr[b' = 1 \mid b = 0] - \Pr[b' = 1 \mid b = 1]| = \text{negl}(\lambda)$$

- Binding: for all adversaries A , if $\sigma \leftarrow \text{Setup}(1^\lambda)$, then

$$\Pr[(m_0, m_1, c, \pi_0, \pi_1) \leftarrow A : m_0 \neq m_1 \text{ and } \text{Verify}(\sigma, c, m_0, \pi_0) = 1 = \text{Verify}(\sigma, c, m_1, \pi_1)] = \text{negl}(\lambda)$$