

CS 6501 Week 8: Zero-Knowledge Proof Systems

In a zero-knowledge proof system, a prover can convince the verifier that some statement x is true (without revealing anything more about x).

In many cases, we want a stronger property: the prover actually "knows" why a statement is true (e.g., it knows a "witness")

For instance, consider the following language:

$$\mathcal{L} = \{h \in \mathbb{G} \mid \exists x \in \mathbb{Z}_p: h = g^x\} = \mathbb{G}$$

\uparrow group of order p \uparrow generator of \mathbb{G}

Note: this definition of \mathcal{L} implicitly defines an NP relation R :
 $R(h, x) = 1 \iff h = g^x \in \mathbb{G}$

In this case, all statements in \mathbb{G} are true (i.e., contained in \mathcal{L}), but we can still consider a notion of proving knowledge of the discrete log of an element $h \in \mathbb{G}$ — conceptually stronger property than proof of membership

Philosophical question: What does it mean to "know" something?

If a prover is able to convince an honest verifier that it "knows" something, then it should be possible to extract that quantity from the prover.

Definition. An interactive proof system (P, V) is a proof of knowledge for an NP relation R if there exists an efficient extractor \tilde{E} such that for any x and any prover P^* proof of knowledge is parameterized by a specific relation R (as opposed to the language \mathcal{L})

$$\Pr[w \leftarrow \tilde{E}^{P^*}(x) : R(x, w) = 1] \geq \Pr[\langle P^*, V \rangle(x) = 1] - \epsilon$$

\uparrow more generally, could be polynomially smaller \uparrow knowledge error

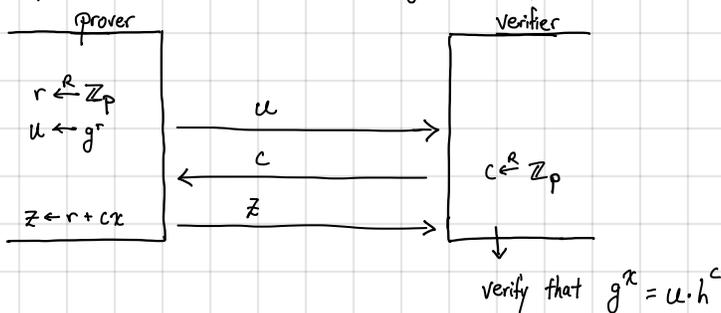
Trivial proof of knowledge: prover sends witness in the clear to the verifier
 \hookrightarrow In most applications, we additionally require zero-knowledge

Note: knowledge is a strictly stronger property than soundness

\hookrightarrow if protocol has knowledge error $\epsilon \Rightarrow$ it also has soundness error ϵ (i.e. a dishonest prover convinces an honest verifier of a false statement with probability at most ϵ)

Proving knowledge of discrete log (Schnorr's protocol)

Suppose prover wants to prove it knows x such that $h = g^x$ (i.e. prover demonstrates knowledge of discrete log of h base g)



Completeness: if $z = r + cx$, then

$$g^z = g^{r+cx} = g^r g^{cx} = u \cdot h^c$$

zero knowledge only required to hold against an honest verifier (e.g., view of the honest verifier can be simulated)

Honest-Verifier Zero-Knowledge: build a simulator as follows (familiar strategy: run the protocol in "reverse"):

on input (g, h) :

1. sample $z \xleftarrow{R} \mathbb{Z}_p$

2. sample $c \xleftarrow{R} \mathbb{Z}_p$

3. set $u = g^z / h^c$ and output (u, c, z)

uniformly random group element since z is uniformly random

uniformly random challenge

chosen so that $g^z = u \cdot h^c$

(relation satisfied by a valid proof)

simulated transcript is identically distributed as the real transcript with an honest verifier

What goes wrong if the challenge is not sampled uniformly at random (i.e., if the verifier is dishonest)

Above simulation no longer works (since we cannot sample z first)

↳ To get general zero knowledge, we require that the verifier first commit to its challenge (using a statistically hiding commitment)

Knowledge: Suppose P^* is (possibly malicious) prover that convinces honest verifier with probability 1. We construct an extractor as follows:

for simplicity, we assume P^* succeeds with probability 1

1. Run the prover P^* to obtain an initial message u .

2. Send a challenge $c_1 \xleftarrow{R} \mathbb{Z}_p$ to P^* . The prover replies with a response z_1 .

3. "Rewind" the prover P^* so its internal state is the same as it was at the end of Step 1. Then, send another challenge $c_2 \xleftarrow{R} \mathbb{Z}_p$ to P^* . Let z_2 be the response of P^* .

4. Compute and output $x = (z_1 - z_2)(c_1 - c_2)^{-1} \in \mathbb{Z}_p$.

Since P^* succeeds with probability 1 and the extractor perfectly simulates the honest verifier's behavior, with probability 1, both (u, c_1, z_1) and (u, c_2, z_2) are both accepting transcripts. This means that

$$g^{z_1} = u \cdot h^{c_1} \quad \text{and} \quad g^{z_2} = u \cdot h^{c_2}$$

$$\Rightarrow \frac{g^{z_1}}{h^{c_1}} = \frac{g^{z_2}}{h^{c_2}} \Rightarrow g^{z_1 + c_2 x} = g^{z_2 + c_1 x}$$

$$\Rightarrow x = (z_1 - z_2)(c_1 - c_2)^{-1} \in \mathbb{Z}_p \quad \text{with overwhelming probability, } c_1 \neq c_2$$

Thus, extractor succeeds with overwhelming probability.

(Boneh-Shoup, Lemma 19.2)

If P^* succeeds with probability ϵ , then need to rely on "Rewinding Lemma" to argue that extractor obtains two accepting transcripts with probability at least $\epsilon^2 - 1/p$.

How can a prover both prove knowledge and yet be zero-knowledge at the same time?

- ↳ Extractor operates by "rewinding" the prover (if the prover has good success probability, it can answer most challenges correctly).
- ↳ But in the real (actual) protocol, verifier cannot rewind (i.e., verifier only sees prover on fresh protocol executions), which can provide zero-knowledge.

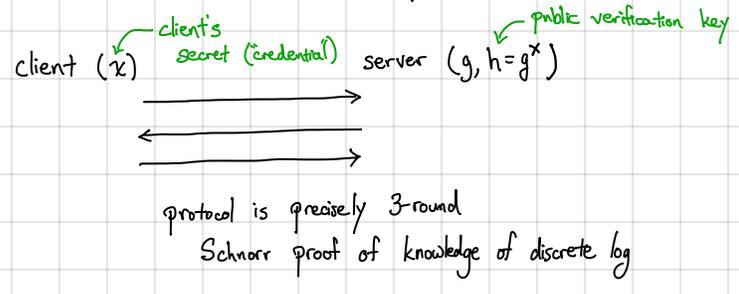
Identification Protocol from Discrete Log

Suppose a client wants to authenticate to the server

↳ Goal: security against active adversaries (adversary sees contents of the server and can interact arbitrarily with the client) this setting

vanilla password based authentication does not provide active security in this setting

Can directly build such a scheme from Schnorr's protocol:



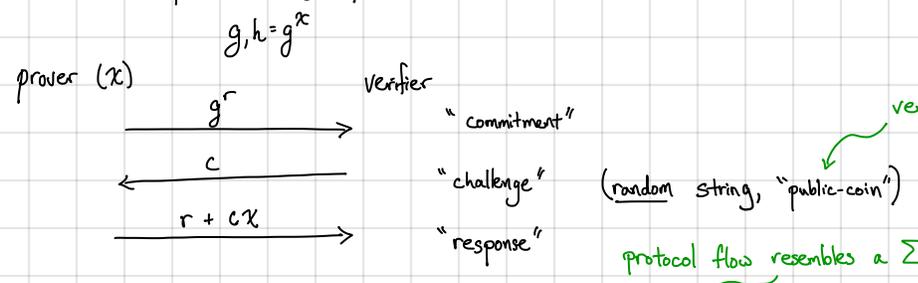
Essentially, the discrete log of h (base g) is the client's "password" and instead of sending the password in the clear to the server, the client proves in zero-knowledge that it knows x

Correctness of this protocol follows from completeness of Schnorr's protocol

(Active) security follows from knowledge property and zero-knowledge

- ↳ Intuitively: knowledge says that any client that successfully authenticates must know secret x
- zero-knowledge says that interactions with honest client (i.e., the prover) do not reveal anything about x
- (for active security, require protocol that provides general zero-knowledge rather than just HVZK)

More general view: Σ -protocols (Sigma protocols)



verifier has no secret randomness (Arthur-Merlin proofs)

- Properties:

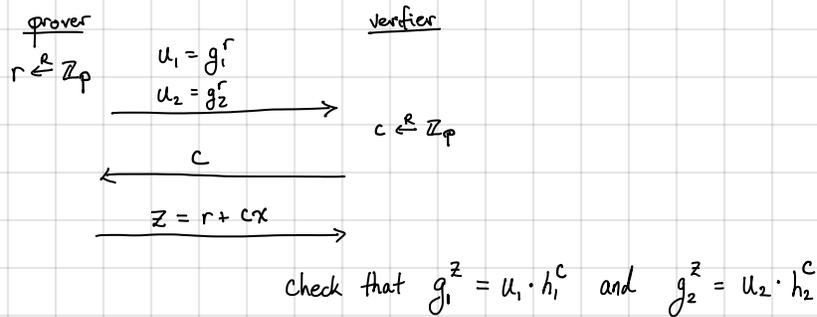
 1. Completeness
 2. Honest-Verifier Zero-Knowledge
 3. Proof of Knowledge

Protocols with this structure (commitment-challenge-response) are called Σ -protocols (Sigma protocols)

Many variants of Schnorr protocols: can be used to prove knowledge of statements like:

- Common discrete log: x such that $h_1 = g_1^x$ and $h_2 = g_2^x$ (useful for building a verifiable random function)
- DDH tuple: (g, u, v, w) is a DDH tuple - namely, that $u = g^\alpha$, $v = g^\beta$, and $w = g^{\alpha\beta}$ for $\alpha, \beta \in \mathbb{Z}_p$
 - ↳ useful for proving relations on ElGamal ciphertexts (e.g., that a particular ElGamal ciphertext encrypts either 0 or 1)
 - ↳ useful building block in constructions of DDH-based oblivious transfer (OT) protocols - Naor-Pinkas (more details next lecture)
 - ↳ Reduces to proving common discrete log: (g, u, v, w) is a DDH tuple if and only if there is an x such that $v = g^x$ and $w = u^x$

Showing that $h_1 = g_1^x$ and $h_2 = g_2^x$:



Completeness and HVZK follows as in Schnorr's protocol.

Knowledge: Two scenarios:

- If prover uses inconsistent commitment (i.e., $u_1 = g_1^{r_1}$ and $u_2 = g_2^{r_2}$ where $r_1 \neq r_2$), then over choice of honest verifier's randomness, then prover can only succeed with probability at most $1/p$:

$$z = r_1 + x_1 c = r_2 + x_2 c \quad (\text{if verifier accepts})$$

Diagram showing the derivation of the equation above:

- $u_1 = g_1^{r_1}$ and $h_1 = g_1^{x_1}$ are linked by a green arrow.
- $u_2 = g_2^{r_2}$ and $h_2 = g_2^{x_2}$ are linked by a green arrow.
- A green arrow points from $u_2 = g_2^{r_2}$ to $h_2 = g_2^{x_2}$.

This means that

$$(r_1 - r_2) = t(x_2 - x_1)$$

If $r_1 \neq r_2$, there is at most 1 $c \in \mathbb{Z}_p$ where this relation holds. Since c is uniform over \mathbb{Z}_p , the verifier accepts with probability at most $1/p$.

- If prover succeeds with $1/\text{poly}(\lambda)$ probability, then it must use a "consistent" commitment. Can build extractor just as in Schnorr's protocol. Knowledge error larger by additive $1/p$ term (from above analysis).

If we want to prove the AND of many statements, then we can prove each one in sequence. What if we want to prove the OR of many statements. The difficulty is not revealing which statement is true (or in the case of proof of knowledge, which witness the prover knows).

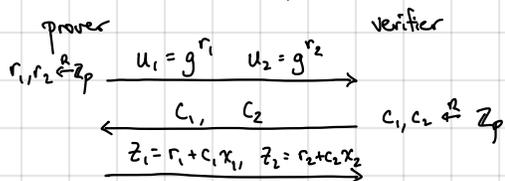
We will work with the following: Prover wants to show that it knows either x_1 or x_2 such that $h_1 = g^{x_1}$ or $h_2 = g^{x_2}$

↳ Statement: (g, h_1, h_2)

Witness: x_1 or x_2 where $h_1 = g^{x_1}$, $h_2 = g^{x_2}$

"1 out of 2 discrete logs"

Starting point: Run Schnorr protocol in parallel:

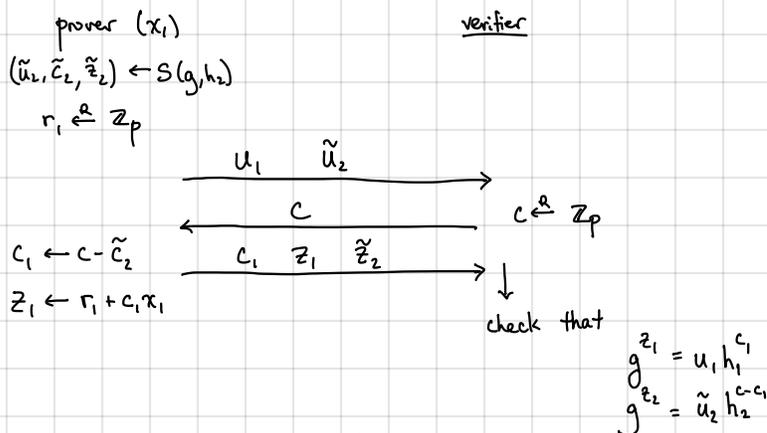


Problem: Honest prover only knows one of x_1 or x_2 so it cannot correctly answer both challenges (unless it knew both x_1 and x_2)

Key idea: Prover will simulate the transcript it does not know.

Suppose prover knows x_1 . Then, it will first run the Schnorr simulator on input (g, h_2) to obtain transcript $(\tilde{u}_2, \tilde{c}_2, \tilde{z}_2)$.

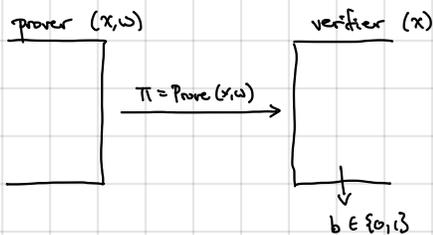
↳ But challenge c_2 may not match \tilde{c}_2 ... To address this, we will have the verifier send a single challenge $c \in \mathbb{Z}_p$ and the prover can pick c_1 and c_2 such that $c_1 + c_2 = c \in \mathbb{Z}_p$



Completeness, HVZK and proof of knowledge follow very similarly as in the proof of Schnorr's protocol.

(NIZK)

Non-interactive zero-knowledge: Can we construct a zero-knowledge proof system where the proof is a single message from the prover to the verifier?



Why do we care? Interaction in practice is expensive!

Unfortunately, NIZKs are only possible for sufficiently-easy languages (i.e., languages in BPP).

languages that can be decided by a randomized polynomial-time algorithm (w.h.p.)

↳ The simulator (for ZK property) can essentially be used to decide the language

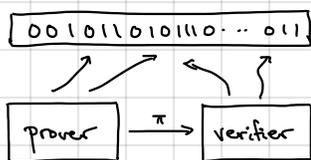
if $x \in L : S(x) \rightarrow \pi$ and π should be accepted by the verifier (by ZK)

if $x \notin L : S(x) \rightarrow \pi$ but π should not be accepted by verifier (by soundness)

} NIZK impossible for NP unless $NP \in BPP$ (wildly!)

Impossibility results tell us where to look! If we cannot succeed in the "plain" model, then move to a different one:

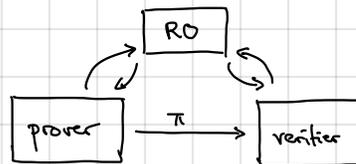
Common random/reference string (CRS) model:



prover and verifier have access to shared randomness (could be a uniformly random string or a structured string)

in this model, simulator is allowed to choose (i.e., simulate) the CRS in conjunction with the proof, but soundness is defined with respect to an honestly-generated CRS [asymmetry between the capabilities of the real prover and the simulator]

random oracle model:

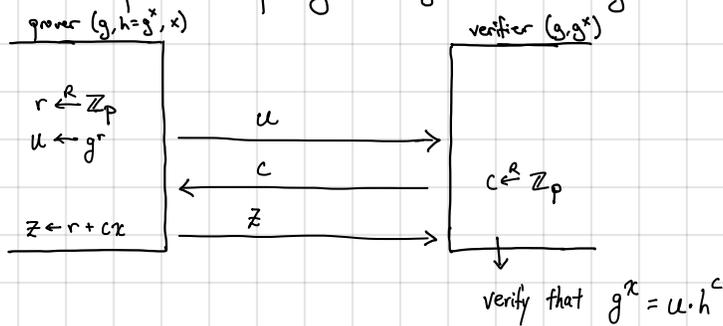


in this model, simulator can "program" the random oracle (again, asymmetry between real prover and the simulator)

⇒ In both cases, simulator has additional "power" than the real prover, which is critical for enabling NIZK constructions for NP.

Fiat-Shamir heuristic: from Σ -protocols to NIZK in RO model

Recall Schnorr's protocol for proving knowledge of discrete log:



In this protocol, verifier's message is uniformly random (and in fact, is "public coin" - the verifier has no secrets)

Key idea: Replace the verifier's challenge with a hash function $H: \{0,1\}^* \rightarrow \mathbb{Z}_p$

Namely, instead of sampling $c \xleftarrow{R} \mathbb{Z}_p$, we sample $c \leftarrow H(g, h, u)$. ← prover can now compute this quantity on its own!

Security of Fiat-Shamir:

- Completeness: Same as Schnorr's protocol
 - Zero-Knowledge: Same as in Schnorr's protocol; namely, simulator samples $c \xleftarrow{R} \mathbb{Z}_p$, $z \xleftarrow{R} \mathbb{Z}_p$, computes u , and then programs RO at (g, h, u) to c .
 - Knowledge: Construct extractor as follows: given (possibly malicious) prover P^* :
 - Run P^* to obtain proof $\pi = (u, z)$ where challenge $c = H(g, h, u)$ at verification time
 - ↳ Note that at some point, P^* must have queried the random oracle on input (g, h, u)
 - Run P^* again, but when it queries RO, use different responses
 - ↳ Can extract discrete log as before
- need to argue that with high probability, P^* will output proof with same committed value u (follows by "forking lemma")

Signatures from discrete log in RO model (Schnorr):

- Verification key is $(g, h = g^x)$ and signing key is x
- To sign a message m , signer proves knowledge of x (discrete log of h) using Fiat-Shamir (and where challenge is derived from message): e.g., $c \leftarrow H(g, h, u, m)$.
- Security essentially follows from security of Schnorr's identification protocol (together with Fiat-Shamir)
 - ↳ specifically, challenger answers signing queries using the ZK simulator (programming RO as needed for consistency)
 - ↳ forged signature on a new message m is a proof of knowledge of the discrete log (can be extracted from adversary)

More generally, any Σ -protocol can be used to build a signature scheme using the Fiat-Shamir heuristic (by using the message to derive the challenge via RO)

Length of Schnorr's signature: vk: $(g, h = g^x)$ sk: x σ : $(g^r, \underbrace{c = H(g, h, g^r, m)}_{\text{can be computed given other components, so do not need to include}}, z = r + cx)$ verification checks that $g^z = g^r h^c$

$\Rightarrow |\sigma| = 2 \cdot |G|$ [512 bits if $|G| = 2^{256}$]

But, can do better... observe that challenge c only needs to be 128-bits (the knowledge error of Schnorr is $1/|C|$ where C is the set of possible challenges), so we can sample a 128-bit challenge rather than 256-bit challenge. Thus instead of sending (g^r, z) , instead send (c, z) and compute $g^r = g^z/h^c$ and that $c = H(g, h, g^r, m)$. Then resulting signatures are 384 bits

128 bit challenge ↙
+
256 bit group element

In practice, ^{↙ TLS protocol} we use a variant of Schnorr's signature scheme called DSA/ECDSA ^{digital signature algorithm / elliptic-curve DSA}

↳ larger signatures (2 group elements - 512 bits) and proof only in "generic group" model [but we use it because Schnorr was patented ... until 2008]

Important note: Schnorr signatures (and DSA/ECDSA) are randomized, and security relies on having good randomness

↳ What happens if randomness is reused for two different signatures?

Then, we have

$$\left. \begin{aligned} \sigma_1 &= (g^r, c_1 = H(g, h, g^r, m_1), z_1 = r + c_1 x) \\ \sigma_2 &= (g^r, c_2 = H(g, h, g^r, m_2), z_2 = r + c_2 x) \end{aligned} \right\} z_1 - z_2 = (c_1 - c_2)x \Rightarrow x = (c_1 - c_2)^{-1}(z_1 - z_2)$$

This is precisely the set of relations the knowledge extractor uses to recover the discrete log x (i.e., the signing key)!

↙ and in some bad Bitcoin wallets

↳ in PlayStation 3, the randomness was a fixed constant! Enables hackers to deliver arbitrary firmware updates to device!

Deterministic Schnorr: We want to replace the random value $r \leftarrow \mathbb{Z}_p$ with one that is deterministic, but which does not compromise security

↳ Derive randomness from message using a PRF. In particular, signing key includes a secret PRF key k , and signing algorithm computes $r \leftarrow F(k, m)$ and $\sigma \leftarrow \text{Sign}(sk, m; r)$.

↳ Avoids randomness reuse/misuse vulnerabilities.