

# COS433/Math 473: Cryptography

Mark Zhandry

Princeton University

Fall 2020

# Announcements/Reminders

HW5 due Nov 10


PR2 due Dec 5

Previously on COS 433...

# Digital Signatures

(aka public key MACs)

# Message Integrity in Public Key Setting



Goal: If Eve changed  $m$ , Bob should reject

# Digital Signatures

Algorithms:

- **Gen()**  $\rightarrow$  (sk,pk)
- **Sign(sk,m)**  $\rightarrow \sigma$
- **Ver(pk,m, $\sigma$ )**  $\rightarrow$  0/1

Correctness:

$$\Pr[\text{Ver}(\text{pk}, m, \text{Sign}(\text{sk}, m)) = 1 : (\text{sk}, \text{pk}) \leftarrow \text{Gen}()] = 1$$

# Building Digital Signatures

Non-trivial to construct with provable security

Most efficient constructions have heuristic security

# Signatures from TDPs

$$\mathbf{Gen}_{\text{sig}}() = \mathbf{Gen}()$$

$$\mathbf{Sign}(\text{sk}, m) = F^{-1}(\text{sk}, H(m))$$

$$\mathbf{Ver}(\text{pk}, m, \sigma): F(\text{pk}, \sigma) == H(m)$$

**Theorem:** If  $(\mathbf{Gen}, F, F^{-1})$  is a secure TDP, and  $H$  is “modeled as a random oracle”, then  $(\mathbf{Gen}_{\text{sig}}, \mathbf{Sign}, \mathbf{Ver})$  is (strongly) CMA-secure



# Basic Rabin Signatures

**Gen<sub>sig</sub>()**: let **p,q** be random large primes  
**sk** = (**p,q**), **pk** = **N** = **pq**

**Sign(sk,m)**: Solve equation  $\sigma^2 = H(m) \bmod N$   
using factors **p,q**

- Output  $\sigma$

**Ver(pk,m, $\sigma$ )**:  $\sigma^2 \bmod N == H(m)$

# Today

Signatures cont.

Identification protocols

# Schnorr Signatures

$$\text{sk} = w$$

$$\text{pk} = h := g^w$$

**Sign(sk,m):**

- $r \leftarrow \mathbb{Z}_p$
- $a \leftarrow g^r$
- $b \leftarrow H(m,a)$
- $c \leftarrow r + wb$
- Output **(a,c)**

**Ver(h,m,(a,c)):**

$$b \leftarrow H(m,a)$$

$$a \times h^b == g^c?$$

**Theorem:** If Dlog is hard and **H** is modeled as a random oracle, then Schnorr signatures are strongly CMA secure

# What's the Smallest Signature?

RSA Hash-and-Sign: 2 kilobits

ECDSA (variant of Schnorr using “elliptic curves”):  
around 512 bits

BLS: 256 bits


Are 128-bit signatures possible?

- No fundamental reason for impossibility, but all (practical) schemes require 256 bits or more


# Digital Signatures and the Public Key Infrastructure




# Digital Signatures and the Public Key Infrastructure



# Digital Signatures and the Public Key Infrastructure



# Digital Signatures and the Public Key Infrastructure






# Takeaway


Need some authenticated channel to ensure distribution of public keys

But how to authenticate channel in the first place without being able to distribute public keys?

# Solution: Certificate Authorities



# Solution: Certificate Authorities



$$Cert_{CA \rightarrow B} = \text{Sign}(sk_{CA}, \text{"Bob's public key is } pk_B \text{"})$$

# Solution: Certificate Authorities

Bob is typically some website

- Obtains **Cert** by, say, sending someone in person to CA with **pk<sub>B</sub>**
- Only needs to be done once

If Alice trusts CA, then Alice will be convinced that **pk<sub>B</sub>** belongs to Bob

Alice typically gets **pk<sub>CA</sub>** bundled in browser

# Limitations

Everyone must trust same CA

- May have different standards for issuing certs

Single point of failure: if  $\mathbf{sk}_{CA}$  is compromised, whole system is compromised

Single CA must handle all verification

# Multiple CAs

There are actually many CA's,  $CA_1$ ,  $CA_2$ ,...

Bob obtains cert from all of them, sends all the certs with his public key

As long as Alice trusts one of the CA's, she will be convinced about Bob's public key

# Certificate Chaining

CA issues **Cert**<sub>CA→B</sub> for Bob

Bob can now use his signing key to issue **Cert**<sub>B→D</sub> to Donald

Donald can now prove his public key by sending  
**(Cert**<sub>CA→B</sub>, **Cert**<sub>B→D</sub>)

- Proves that CA authenticated Bob, and Bob authenticated Donald

# Certificate Chaining

For Bob to issue his own certificates, a standard cert should be insufficient

- CA knows who Bob is, but does not trust him to issue certs on its behalf

Therefore, Bob should have a stronger cert:

**$\text{Cert}_{CA \rightarrow B} = \text{Sign}(\text{sk}_{CA}, \text{"Bob's public key is } \mathbf{pk}_B \text{ and he can issue certificates on behalf of CA"})$**



# Certificate Chaining

One root CA

Many second level CAs  $CA_1, CA_2, \dots$

- Each has **Cert**<sub>CA→CA<sub>i</sub></sub>

Advantage: eases burden on root

Disadvantage: now multiple points of failure

# Invalidating Certificates

Sometimes, need to invalidate certificates

- Private key stolen
- User leaves company
- Etc

Options:

- Expiration
- Explicit revocation

# Identification Protocols

# Identification



# Identification



# Identification

To identify yourself, you need something the adversary doesn't have

Typical factors:

- What you **are**: biometrics (fingerprints, iris scans,...)
- What you **have**: Smart cards, SIM cards, etc
- What you **know**: Passwords, PINs, secret keys

Today

# Types of Identification Protocols

Secret key:



Public Key:



# Types of Attacks


Direct Attack:






# Types of Attacks

Eavesdropping/passive:




# Types of Attacks

Eavesdropping/passive:




# Types of Attacks

Man-in-the-Middle/Active:



# Types of Attacks

Man-in-the-Middle/Active:



# Basic Password Protocol

Never ever (ever ever...) use




**sk == vk?**

# Problem with Basic Pwd Protocol

**vk** must be kept secret at all costs

Issue:



User	Pwd
Alice	$\text{pwd}_A$
Bob	$\text{pwd}_B$
Charlie	$\text{pwd}_C$
...	...

# Problem with Basic Pwd Protocol

**vk** must be kept secret at all costs


Issue:



**pw**



**d**







# Slightly Better Version

STILL never ever (ever ever...) use

Let **H** be a hash function



**sk=pwd**

**sk**




**vk=H(pwd)**

**H(sk) == vk?**

# Slightly Better Version

STILL never ever (ever ever...) use

Let **H** be a hash function



User	Pwd
Alice	$H(\text{pwd}_A)$
Bob	$H(\text{pwd}_B)$
Charlie	$H(\text{pwd}_C)$
...	...

# Slightly Better Version

STILL never ever (ever ever...) use

Advantage of hashing:

- Now if pwd database is leaks, adversary only gets hashes passwords
- For identification protocol, need actual password
- Therefore, adversary needs to invert hash function to break protocol
- Presumed hard

# Weak Passwords

Data from 10M passwords leaked in 2016:

**17%**



RANK	PASSWORD	9.	123123	18.	654321
1.	123456	10.	987654321	19.	555555
2.	123456789	11.	qwertyuiop	20.	3rjs1la7qe
3.	qwerty	12.	myn0ob	21.	google
4.	12345678	13.	123321	22.	1q2w3e4r5t
5.	111111	14.	666666	23.	123qwe
6.	1234567890	15.	18atcskd2w	24.	zxcvbnm
7.	1234567	16.	7777777	25.	1q2w3e
8.	password	17.	1q2w3e4r		



50% of available passwords

# Weak Passwords

Of course, pwds that have been leaked are likely the particularly common ones

Even so, 360M pwds covers about 25% of all users

# Online Dictionary Attacks

Suppose attacker gets list of usernames

Attacker tries logging in to each with **pwd** = '123456'

5-17% of accounts will be compromised

# Online Dictionary Attacks

How to slow down attacker?

- Lock out after several unsuccessful attempts
  - Honest users may get locked out too
- Slow down response after each unsuccessful attempt
  - 1s after 1<sup>st</sup>, 2s after 2<sup>nd</sup>, 4s after 3<sup>rd</sup>, etc

# Offline Dictionary Attacks

Suppose attacker gets hashed password  **$vk = H(pwd)$**

Attack:

- Assemble dictionary of 360M common passwords
- Hash each, and check if you get  **$vk$**
- If so, you have just found  **$pwd$** !

On modern hardware, takes a few seconds to recover a password 25% of the time



# Offline Dictionary Attacks

Now consider what happens when adversary gets entire hashed password database

- Hash dictionary once:  $O(|D|)$
- Index dictionary by hashes
- Lookup each database entry in dictionary:  $O(|L|)$


To get 25% of passwords takes  $O(|D|+|L|)$  time

- Amortize cost of hashing dictionary over many passwords

# Salting

Let **H** be a hash function

**s<sub>i</sub>** random



User	Salt	Pwd
Alice	$s_A$	$H(s_A, \text{pwd}_A)$
Bob	$s_B$	$H(s_B, \text{pwd}_B)$
Charlie	$s_C$	$H(s_C, \text{pwd}_C)$
...	...	...

# Salting

Salt length? Enough to make each user's salt unique

- At least 64 bits

Salting kills amortization:

- To recover Alice's key, adversary must hash entire dictionary with  $\mathbf{s}_A$
- To recover Bob's key, adversary must hash entire dictionary with  $\mathbf{s}_B$
- Must hash entire dictionary again for each user

Running time:  $O(|D| \times |L|)$

# Unique Passwords

Different websites may employ different standards for password security

- Some may store passwords in clear, some may hash without salt, some may salt

If you use the same password at a bank (high security) and your high school reunion (low security), could end up with your password stolen

# Unique Passwords

Solutions:

- Password managers
- Salt master password to generate website-specific password (e.g. pwdhash):

Master password: **pwd**

Pwd for abcdefg.com: **H(abcdefg.com,pwd)**

# What Hash Function to Use

In LinkedIn leak (using Sha1), 90% of passwords were recovered within a week

Problem: Sha1 is very fast!

To make hashing harder, want hash function that is just slow enough to be unnoticeable to user

# What Hash Function to Use

Examples: PBKDF2, bcrypt

- Iterate hash function many times:

$$H'(x) = H(H(H(\dots H(x)\dots)))$$

- Set #iterations to get desired hashing time

Still problem:

- Adversary may have special purpose hardware  
⇒ Can eval much faster than you can (50,000x)

# What Hash Function to Use

Memory-hard functions: functions that require a lot of memory to compute

- As far as we know, no special purpose memory
- Attacker doesn't gain advantage using special purpose hardware

Examples: Scrypt, Argon2i