

Lecture 10: Cryptography (Modern)

*Week 8**UCSB 2014*

We talked about a few historical methods in which people have encrypted messages in our last talk. Today, we're going to talk about the methods people use today! Specifically, we're going to talk about the **Diffie-Hellman key exchange system!**

1 Key Exchanges: A Puzzle

First, before we start, consider the following puzzle:

Puzzle. The Untrustworthy Ferryman.

- Good news! You've found yourself in the possession of a shoebox full of rubies, diamonds, gold, etc. You also have a padlock and key, which you can use to lock the outside of the box.
- Bad news! You're trapped on a desert island in the middle of nowhere.
- Good news! Every day at noon, a ferryman approaches your island. His boat is just big enough to contain him and a small box. If you could get the box delivered to your friend on the mainland, in a way where he could open the box, they could use a handful of the jewels to rescue you and set you free!
- Bad news! The ferryman will steal and not deliver anything that is not a locked box. Locked boxes he will take and deliver to your friend on the mainland for the price of one jewel. But anything else he will steal!
- Possibly useful news? Your friend also has a box and padlock and key, and the ferryman is willing to take packages back to you from your friend as part of the deal. But again, he will steal anything that is not a locked box.

How can you get him the treasure?

The answer is on the next page; try to solve it before continuing!

Answer. Here's the answer!

- You lock up your box and send it to your friend via the ferryman. Because the box is locked, the ferryman is honest and delivers the box.
- Your friend receives the box, and **locks the box with his padlock as well!** I.e. now the box has two locks on it.
- The doubly-locked box is sent back to you. You remove your lock, and send it back to your friend.
- Your friend receives the box with just his lock on it! He removes his lock, takes the jewels, rescues you, and you live your life in comfort and/or splendor.

This idea, roughly speaking, is how modern cryptographic systems work!

2 Public Key Exchange Systems

In general, a public key exchange system, roughly speaking, is any cryptographic system that works like our puzzle above — i.e. where there are two parties that are communicating, who do so in the following manner:

- Party A takes a message and encrypts it with a secret key, known only to them. Then they send this message to party B .
- Party B takes that encrypted message and encrypts it a second time with their secret key, and sends it back to A .
- Party A then takes this doubly-encrypted message and decrypts it with their key. If our encryption and decryption functions **commute** — in other words, if encrypting and then decrypting is the same thing as decrypting and then encrypting — then this “undoes” the first layer of encryption, leaving a message encrypted with only B 's key. This is sent back to B .
- Party B then takes that message and decrypts it; it is now readable by party B !

This is the idea for how most modern cryptographic systems work. With the rest of this class, we're going to explore one specific method of this encryption scheme: the **Diffie-Hellman** key exchange system!

First, here's a fun property:

Theorem 1. Fermat's Little Theorem. *Let p be a prime number. Take any $a \neq 0$ in $\mathbb{Z}/p\mathbb{Z}$. Then*

$$a^{p-1} \equiv 1 \pmod{p}.$$

Proof. On the HW! □

Why do we care about this? Because it lets us create one of the first (and one of the most fundamental) modern cryptographic systems!

3 Diffie-Hellman

Algorithm. The Diffie-Hellman key exchange algorithm works as follows: suppose that we have two people, Alice and Bob, that want to communicate over a public network.

Ahead of time, Alice and Bob agree on a public base prime p to work in, as well as a public seed g . Selecting this value p and seed g requires some thought in practice due to some complicated mathematical tricks people can perform. However, if you just pick a big prime p (like 300+ digits) this usually works pretty well. Choices for g are a little harder to pin down when they are “good,” but usually most anything works — in practice, g is usually 2, 3 or 5.

From there, Alice and Bob each pick a “secret key” a, b from $\mathbb{Z}/p\mathbb{Z}$. Then, they do the following:

1. Alice sends the number g^a publicly to Bob.
2. Bob then takes this received number g^a and uses his key to raise it to the b -th power, which gives him g^{ab} .
3. Bob then sends the number g^b publicly to Alice.
4. Alice then takes this received number g^b and uses her key to raise it to the a -th power, which gives her g^{ab} .
5. Bob and Alice are now both in possession of the same secret key g^{ab} . Any eavesdroppers will only have heard g^a and g^b , and thus have no obvious way to figure out g^{ab} .
6. To communicate a message $m \in \mathbb{Z}/p\mathbb{Z}$, then, either Alice or Bob can send to the other party $m \cdot g^{ab}$. No other parties know the secret key, so they cannot decode this message.
7. Moreover, because both Alice and Bob know the secret key g^{ab} , along with the values g^a, g^b , they can use Fermat’s little theorem to calculate its inverse! Specifically: notice that Bob, knowing b, p , and g^a , can calculate

$$(g^a)^{p-1-b} = g^{a(p-1)-ab} = g^{a(p-1)} \cdot g^{-ab} = (g^{p-1})^a \cdot g^{-ab}.$$

By Fermat’s little theorem, we know that

$$g^{p-1} \equiv 1 \pmod{p},$$

and therefore that

$$(g^{p-1})^a \cdot g^{-ab} \equiv g^{-ab} \pmod{p}.$$

Therefore, Bob can calculate g^{-ab} and multiply it by their received encrypted message to get the original message back. Alice, using a, p , and g^b , can do the same trick as well; this allows both people to read their messages!

Why does this work? This is because the **discrete logarithm problem**, described below, is considered a “hard” problem to answer:

Problem. The **discrete logarithm problem** is the following task: Suppose you’re given a pair of numbers a, b in $\mathbb{Z} \bmod p\mathbb{Z}$. The discrete logarithm problem asks the user for a value k such that

$$a^k \equiv b \pmod{p},$$

if some such value exists.

This problem is surprisingly hard to calculate. One naive algorithm that you might be tempted to try is simply raising a to higher and higher powers, and then reducing mod p each time until you got something that is equal to b . However, this takes a hideously long time to run: if we just take values of k starting at 1 and going up, we’re effectively looking through the elements of $\mathbb{Z}/p\mathbb{Z}$ one at a time, in the order $a, a^2, a^3 \dots$. If we assume that b is chosen more or less at random, then we’d expect to find b after going through about half of $\mathbb{Z}/p\mathbb{Z}$ (and in the worst-case after going through all of $\mathbb{Z}/p\mathbb{Z}$.) In either case, we have runtime that is linear in the size of the group! Therefore, if our group has say 300 digits (like we have for the size of prime requested in the algorithm above,) our algorithm has runtime that’s absolutely awful (i.e. exponential in the number of digits in the size of the group.)

(It bears noting that this problem is an interesting one: it’s both in NP, because it can be checked easily, but **unlike** many other problems it’s not yet been shown to be NP-complete. So there’s an odd chance that it might be in P after all; or at least it’s more likely to be in P than things like the traveling salesman problem?)

As a consequence, Alice and Bob in the discussion above can be reasonably sure that (unless someone comes up with an algorithm that’s much faster than the one above) their communications are going by un-eavesdropped, because everything they send needs the discrete log problem to be decoded without keys!