CPSC 467: Cryptography and Computer Security

Michael J. Fischer

Lecture 16
October 28, 2013

Euler criterion

Testing for membership in QR_p

Theorem (Euler Criterion)

An integer a is a non-trivial[1] quadratic residue modulo an odd prime p iff

$$a^{(p-1)/2} \equiv 1 \pmod{p}.$$

Proof in forward direction.

Let $a \equiv b^2 \pmod{p}$ for some $b \not\equiv 0 \pmod{p}$. Then

$$a^{(p-1)/2} \equiv (b^2)^{(p-1)/2} \equiv b^{p-1} \equiv 1 \pmod{p}$$

by Euler's theorem, as desired. □

[1] A non-trivial quadratic residue is one that is not equivalent to 0 (mod p).

Proof of Euler Criterion

Proof in reverse direction.

Suppose $a^{(p-1)/2} \equiv 1 \pmod{p}$. Clearly $a \not\equiv 0 \pmod{p}$. We find a square root b of a modulo p.

Let g be a primitive root of p. Choose k so that $a \equiv g^k \pmod{p}$, and let $\ell = (p-1)k/2$. Then

$$g^\ell \equiv g^{(p-1)k/2} \equiv (g^k)^{(p-1)/2} \equiv a^{(p-1)/2} \equiv 1 \pmod{p}.$$

Since g is a primitive root, $(p-1) \mid \ell$. Hence, $2|k$ and $k/2$ is an integer.

Let $b = g^{k/2}$. Then $b^2 \equiv g^k \equiv a \pmod{p}$, so b is a non-trivial square root of a modulo p, as desired. □

Finding Square Roots

Finding square roots modulo prime $p \equiv 3 \pmod 4$

The Euler criterion lets us test membership in QR_p for prime p, but it doesn't tell us how to quickly find square roots. They are easily found in the special case when $p \equiv 3 \pmod 4$.

Theorem

Let $p \equiv 3 \pmod 4$, $a \in \mathrm{QR}_p$. Then $b = a^{(p+1)/4} \in \sqrt{a} \pmod p$.

Proof.

$p + 1$ is divisible by 4, so $(p + 1)/4$ is an integer. Then

$$b^2 \equiv (a^{(p+1)/4})^2 \equiv a^{(p+1)/2} \equiv a^{1+(p-1)/2} \equiv a \cdot 1 \equiv a \pmod p$$

by the Euler Criterion. □

Finding square roots for general primes

We now present an algorithm due to D. Shanks[2] that finds square roots of quadratic residues modulo any odd prime p.

[2]Shanks's algorithm appeared in his paper, "Five number-theoretic algorithms", in Proceedings of the Second Manitoba Conference on Numerical Mathematics, Congressus Numerantium, No. VII, 1973, 51–70. Our treatment is taken from the paper by Jan-Christoph Schlage-Puchta", "On Shank's Algorithm for Modular Square Roots", *Applied Mathematics E-Notes, 5* (2005), 84–88.

Shank's algorithm

Let p be an odd prime. Write $\phi(p) = p - 1 = 2^s t$, where t is odd. (Recall: s is # trailing 0's in the binary expansion of $p - 1$.)

Because p is odd, $p - 1$ is even, so $s \geq 1$.

A special case

In the special case when $s = 1$, then $p - 1 = 2t$, so $p = 2t + 1$.

Writing the odd number t as $2\ell + 1$ for some integer ℓ, we have

$$p = 2(2\ell + 1) + 1 = 4\ell + 3,$$

so $p \equiv 3 \pmod 4$.

This is exactly the case that we handled above.

Overall structure of Shank's algorithm

Let $p - 1 = 2^s t$ be as above, where p is an odd prime.

Assume $a \in \mathrm{QR}_p$ is a quadratic residue and $u \in \mathrm{QNR}_p$ is a quadratic non-residue.

We can easily find u by choosing random elements of $\mathbf{Z}_p^*$ and applying the Euler Criterion.

The goal is to find x such that $x^2 \equiv a \pmod{p}$.

Shanks's algorithm

1. Let s, t satisfy $p - 1 = 2^s t$ and t odd.
2. Let $u \in \mathrm{QNR}_p$.
3. $k = s$
4. $z = u^t \bmod p$
5. $x = a^{(t+1)/2} \bmod p$
6. $b = a^t \bmod p$
7. while $(b \not\equiv 1 \pmod p)$ {
8. let m be the least integer with $b^{2^m} \equiv 1 \pmod p$
9. $y = z^{2^{k-m-1}} \bmod p$
10. $z = y^2 \bmod p$
11. $b = bz \bmod p$
12. $x = xy \bmod p$
13. $k = m$
14. }
15. return x

Loop invariant

The congruence

$$x^2 \equiv ab \pmod{p}$$

is easily shown to be a loop invariant.

It's clearly true initially since $x^2 \equiv a^{t+1}$ and $b \equiv a^t \pmod{p}$.

Each time through the loop, a is unchanged, b gets multiplied by y^2 (lines 10 and 11), and x gets multiplied by y (line 12); hence the invariant remains true regardless of the value of y.

If the program terminates, we have $b \equiv 1 \pmod{p}$, so $x^2 \equiv a$, and x is a square root of $a \pmod{p}$.

Termination proof (sketch)

The algorithm terminates after at most $s - 1$ iterations of the loop.

To see why, we look at the orders[3] of b and $z \pmod{p}$ and show the following loop invariant:

> *At the start of each loop iteration (before line 8),* $\text{ord}(b)$ *is a power of 2 and* $\text{ord}(b) < \text{ord}(z) = 2^k$.

After line 8, $m < k$ since $2^m = \text{ord}(b) < 2^k$. Line 13 sets $k = m$ for the next iteration, so k decreases on each iteration.

The loop terminates when $b \equiv 1 \pmod{p}$. Then $\text{ord}(b) = 1 < 2^k$, so $k \geq 1$. Hence, the loop is executed at most $s - 1$ times.

[3] Recall that the order of an element g modulo p is the least positive integer k such that $g^k \equiv 1 \pmod{p}$.

QR Probabilistic Cryptosystem

A hard problem associated with quadratic residues

Let $n = pq$, where p and q are distinct odd primes.

Recall that each $a \in \mathrm{QR}_n$ has 4 square roots, and $1/4$ of the elements in $\mathbf{Z}_n^*$ are quadratic residues.

Some elements of $\mathbf{Z}_n^*$ are easily recognized as non-residues, but there is a subset of non-residues (which we denote as Q_n^{00}) that are *hard to distinguish* from quadratic residues without knowing p and q.

This allows for public key encryption of single bits: A random element of QR_n encrypts 1; a random element of Q_n^{00} encrypts 0.

Quadratic residues modulo $n = pq$

Let $n = pq$, p, q distinct odd primes.

We divide the numbers in $\mathbf{Z}_n^*$ into four classes depending on their membership in QR_p and QR_q.[4]

- Let $Q_n^{11} = \{a \in \mathbf{Z}_n^* \mid a \in \mathrm{QR}_p \cap \mathrm{QR}_q\}$.
- Let $Q_n^{10} = \{a \in \mathbf{Z}_n^* \mid a \in \mathrm{QR}_p \cap \mathrm{QNR}_q\}$.
- Let $Q_n^{01} = \{a \in \mathbf{Z}_n^* \mid a \in \mathrm{QNR}_p \cap \mathrm{QR}_q\}$.
- Let $Q_n^{00} = \{a \in \mathbf{Z}_n^* \mid a \in \mathrm{QNR}_p \cap \mathrm{QNR}_q\}$.

Under these definitions,

$$\mathrm{QR}_n = Q_n^{11}$$

$$\mathrm{QNR}_n = Q_n^{00} \cup Q_n^{01} \cup Q_n^{10}$$

[4] To be strictly formal, we classify $a \in \mathbf{Z}_n^*$ according to whether or not $(a \bmod p) \in \mathrm{QR}_p$ and whether or not $(a \bmod q) \in \mathrm{QR}_q$.

Quadratic residuosity problem

Definition (Quadratic residuosity problem)

The *quadratic residuosity problem* is to decide, given $a \in Q_n^{00} \cup Q_n^{11}$, whether or not $a \in Q_n^{11}$.

Fact

There is no known feasible algorithm for solving the quadratic residuosity problem that gives the correct answer significantly more than 1/2 the time for uniformly distributed random $a \in Q_n^{00} \cup Q_n^{11}$, unless the factorization of n is known.

Goldwasser-Micali probabilistic cryptosystem

The Goldwasser-Micali cryptosystem is based on the assumed hardness of the quadratic residuosity problem.

The public key consist of a pair $e = (n, y)$, where $n = pq$ for distinct odd primes p, q, and y is any member of Q_n^{00}.
The private key consists of p.
The message space is $\mathcal{M} = \{0, 1\}$. (Single bits!)

To encrypt $m \in \mathcal{M}$, Alice chooses a random $r \in \mathbf{Z}_n^*$ and sets $a = r^2 \bmod p$. The result a is a random element of $\mathrm{QR}_n = Q_n^{11}$.

If $m = 0$, set $c = a$ (which is in Q_n^{11}).
If $m = 1$, set $c = ay \bmod n$ (which is in Q_n^{00}).

The problem of finding m given c is equivalent to the problem of testing if $c \in \mathrm{QR}_n (= Q_n^{11})$, given that $c \in Q_n^{00} \cup Q_n^{11}$.

Decryption in Goldwasser-Micali encryption

Bob, knowing the private key p, can use the Euler Criterion to quickly determine whether or not $c \in \mathrm{QR}_p$ and hence whether $c \in Q_n^{11}$ or $c \in Q_n^{00}$, thereby determining m.

Eve's problem of determining whether c encrypts 0 or 1 is the same as the problem of distinguishing between membership in Q_n^{00} and Q_n^{11}, which is just the quadratic residuosity problem, assuming the ciphertexts are uniformly distributed.

One can show that every element of Q_n^{11} is equally likely to be chosen as the ciphertext c in case $m = 0$, and every element of Q_n^{00} is equally likely to be chosen as the ciphertext c in case $m = 1$. If the messages are also uniformly distributed, then any element of $Q_n^{00} \cup Q_n^{11}$ is equally likely to be the ciphertext.

Important facts about quadratic residues

1. If p is odd prime, then $|\mathrm{QR}_p| = |\mathbf{Z}_p^*|/2$, and for each $a \in \mathrm{QR}_p$, $|\sqrt{a}| = 2$.
2. If $n = pq$, $p \neq q$ odd primes, then $|\mathrm{QR}_n| = |\mathbf{Z}_n^*|/4$, and for each $a \in \mathrm{QR}_n$, $|\sqrt{a}| = 4$.
3. Euler criterion: $a \in \mathrm{QR}_p$ iff $a^{(p-1)/2} \equiv 1 \pmod{p}$, p odd prime.
4. If n is odd prime, $a \in \mathrm{QR}_n$, can feasibly find $y \in \sqrt{a}$.
5. If $n = pq$, $p \neq q$ odd primes, then distinguishing Q_n^{00} from Q_n^{11} is believed to be infeasible. Hence, infeasible to find $y \in \sqrt{a}$. Why?
   If not, one could attempt to find $y \in \sqrt{a}$, check that $y^2 \equiv a \pmod{n}$, and conclude that $a \in Q^{11}$ if successful.

Authentication While Preventing Impersonation

Preventing impersonation

A fundamental problem with all of the password authentication schemes discussed so far is that Alice reveals her secret to Bob every time she authenticates herself.

This is fine when Alice trusts Bob but not otherwise.

After authenticating herself once to Bob, then Bob can masquerade as Alice and impersonate her to others.

Authentication requirement

When neither Alice nor Bob trust each other, there are two requirements that must be met:

1. Bob wants to make sure that an impostor cannot successfully masquerade as Alice.
2. Alice wants to make sure that her secret remains secure.

At first sight these seem contradictory, but there are ways for Alice to prove her identity to Bob without compromising her secret.

Challenge-Response Authentication Protocols

Challenge-response authentication protocols

In a challenge-response protocol, Bob presents Alice with a challenge that only the true Alice (or someone knowing Alice's secret) can answer.

Alice answers the challenge and sends her answer to Bob, who verifies that it is correct.

Bob learns the response to his challenge but Alice never reveals her secret.

If the protocol is properly designed, it will be hard for Bob to determine Alice's secret, even if he chooses the challenges with that end in mind.

Challenge-response protocol from a signature scheme

A challenge-response protocol can be built from a digital signature scheme (S_A, V_A).

(The same protocol can also be implemented using a symmetric cryptosystem with shared key k.)

	Alice		Bob
1.		$\xleftarrow{r}$	Choose random string r.
2.	Compute $s = S_A(r)$	$\xrightarrow{s}$	Check $V_A(r, s)$.

Requirements on underlying signature scheme

This protocol exposes Alice's signature scheme to a chosen plaintext attack.

A malicious Bob can get Alice to sign any message of his choosing.

Alice had better have a different signing key for use with this protocol than she uses to sign contracts.

While we hope our cryptosystems are resistant to chosen plaintext attacks, such attacks are very powerful and are not easy to defend against.

Anything we can do to limit exposure to such attacks can only improve the security of the system.

Limiting exposure to chosen plaintext attack: try 1

We explore some ways that Alice might limit Bob's ability to carry out a chosen plaintext attack.

Instead of letting Bob choose the string r for Alice to sign, r is constructed from two parts, r_1 and r_2.

r_1 is chosen by Alice; r_2 is chosen by Bob. Alice chooses first.

	Alice		Bob
1.	Choose random string r_1	$\xrightarrow{r_1}$	
2.		$\xleftarrow{r_2}$	Choose random string r_2.
3.	Compute $r = r_1 \oplus r_2$		Compute $r = r_1 \oplus r_2$
4.	Compute $s = S_A(r)$	$\xrightarrow{s}$	Check $V_A(r, s)$.

Problem with try 1

The idea is that neither party should be able to control r.

Unfortunately, that idea does not work here because Bob gets r_1 before choosing r_2.

Instead of choosing r_2 randomly, a cheating Bob can choose $r_2 = r \oplus r_1$, where r is the string that he wants Alice to sign.

Thus, try 1 is no more secure against chosen plaintext attack than the original protocol.

Limiting exposure to chosen plaintext attack: try 2

Another possibility is to choose the random strings in the other order—Bob chooses first.

	Alice		Bob
1.		$\xleftarrow{r_2}$	Choose random string r_2.
2.	Choose random string r_1	$\xrightarrow{r_1}$	
3.	Compute $r = r_1 \oplus r_2$		Compute $r = r_1 \oplus r_2$
4.	Compute $s = S_A(r)$	$\xrightarrow{s}$	Check $V_A(r, s)$.

Try 2 stops chosen plaintext attack

Now Alice has complete control over r.

No matter how Bob chooses r_2, Alice's choice of a random string r_1 ensures that r is also random.

This thwarts Bob's chosen plaintext attack since r is completely random.

Thus, Alice only signs random messages.

Problem with try 2

Unfortunately, try 2 is totally insecure against active eavesdroppers. Why?

Suppose Mallory listens to a legitimate execution of the protocol between Alice and Bob.

From this, he easily acquires a valid signed message (r_0, s_0). How does this help Mallory?

Mallory sends $r_1 = r_0 \oplus r_2$ in step 2 and $s = s_0$ in step 4.

Bob computes $r = r_1 \oplus r_2 = r_0$ in step 3, so his verification in step 4 succeeds.

Thus, Mallory can successfully impersonate Alice to Bob.

Further improvements

Possible improvements to both protocols.

1. Let $r = r_1 \cdot r_2$ (concatenation).
2. Let $r = h(r_1 \cdot r_2)$, where h is a cryptographic hash function.

In both cases, neither party now has full control over r.

This weakens Bob's ability to launch a chosen plaintext attack if Alice chooses first.

This weakens Mallory's ability to impersonate Alice if Bob chooses first.

Feige-Fiat-Shamir Authentication Protocol

Concept of zero knowledge

In all of the challenge-response protocols above, Alice releases some partial information about her secret by producing signatures that Bob could not compute by himself.

The Feige-Fiat-Shamir protocol allows Alice to prove knowledge of her secret *without revealing any information about the secret itself*.

Such protocols are called *zero knowledge*, which we will discuss shortly.

Feige-Fiat-Shamir protocol: overview

Alice authenticates herself by successfully completing several rounds of a protocol that requires knowledge of a secret s.

In a single round, protocol, Bob has at least a 50% chance of catching an impostor Mallory.

By repeating the protocol t times, the error probability (that is, the probability that Bob fails to catch Mallory) drops to $1/2^t$.

This can be made acceptably low by choosing t to be large enough.

For example, if $t = 20$, then Mallory has only one chance in a million of successfully impersonating Alice.

Feige-Fiat-Shamir protocol: preparation

The Feige-Fiat-Shamir protocol is based on the difficulty of computing square roots modulo composite numbers.

- Alice chooses $n = pq$, where p and q are distinct large primes.
- Next she picks a quadratic residue $v \in \mathrm{QR}_n$ (which she can easily do by choosing a random element $u \in \mathbf{Z}_n^*$ and letting $v = u^2 \bmod n$).
- Finally, she chooses s to be the smallest square root of $v^{-1} \pmod n$.[5] She can do this since she knows the factorization of n.

She makes n and v public and keeps s private.

[5] Note that if v is a quadratic residue, then so is $v^{-1} \pmod n$.

A simplified one-round FFS protocol

Here's a simplified one-round version.

	Alice		Bob
1.	Choose random $r \in \mathbf{Z}_n$.		
	Compute $x = r^2 \bmod n$.	$\xrightarrow{x}$	
2.		$\xleftarrow{b}$	Choose random $b \in \{0, 1\}$.
3.	Compute $y = rs^b \bmod n$.	$\xrightarrow{y}$	Check $x = y^2 v^b \bmod n$.

When both parties are honest, Bob accepts Alice because

$$x = y^2 v^b \bmod n.$$

This holds because

$$y^2 v^b \equiv (rs^b)^2 v^b \equiv r^2 (s^2 v)^b \equiv x(v^{-1} v)^b \equiv x \pmod{n}.$$

A dishonest Alice

We now turn to the security properties of the protocol when "Alice" is dishonest, that is, when Mallory is attempting to impersonate the real Alice.

Theorem

Suppose Mallory doesn't know a square root of v^{-1}. Then Bob's verification will fail with probability at least 1/2.

Proof that Mallory can't successfully cheat

Proof.

In order for Mallory to successfully fool Bob, he must come up with x in step 1 and y in step 3 satisfying

$$x = y^2 v^b \bmod n.$$

Mallory sends x in step 1 before Bob chooses b, so he does not know which value of b to expect.

When Mallory receives b, he responds by sending a value y_b to Bob.

We consider two cases.

(continued...)

Proof: case 1

Proof (continued).

Case 1: There is at least one $b \in \{0, 1\}$ for which y_b fails to satisfy

$$x = y^2 v^b \bmod n.$$

Since $b = 0$ and $b = 1$ each occur with probability $1/2$, this means that Bob's verification will fail with probability at least $1/2$, as desired.

(continued. . .)

Proof: case 2

Proof (continued).

Case 2: y_0 and y_1 both satisfy the verification equation, so $x = y_0^2 \bmod n$ and $x = y_1^2 v \bmod n$.

We can solve these equations for v^{-1} to get

$$v^{-1} \equiv y_1^2 x^{-1} \equiv y_1^2 y_0^{-2} \pmod{n}$$

But then $y_1 y_0^{-1} \bmod n$ is a square root of v^{-1}.

Since Mallory was able to compute both y_1 and y_0, then he was also able to compute a square root of v^{-1}, contradicting the assumption that he doesn't "know" a square root of v^{-1}. □

Successful cheating with probability 1/2

We remark that it *is* possible for Mallory to cheat with success probability 1/2.

- He guesses the bit b that Bob will send him in step 2 and generates a pair (x, y).
- If he guesses $b = 0$, then he chooses $x = r^2 \bmod n$ and $y = r \bmod n$, just as Alice would have done.
- If he guesses $b = 1$, then he chooses y arbitrarily and $x = y^2 v \bmod n$.

He proceeds to send x in step 1 and y in step 3.

The pair (x, y) is accepted by Bob Mallory's guess of b turns out to be correct, which will happen with probability 1/2.

A dishonest Bob

We now consider the case of a dishonest Mallory impersonating Bob, or simply a dishonest Bob who wants to capture Alice's secret.

Alice would like assurance that her secret is protected if she follows the protocol,regardless of what Mallory (Bob) does.

Consider what Mallory knows at the end of the protocol.

Mallory sends $b = 0$

Suppose Mallory sends $b = 0$ in step 2.

Then he ends up with a pair (x, y), where y is a random number and x is its square modulo n.

Neither of these numbers depend in any way on Alice secret s, so Mallory gets no direct information about s.

It's also of no conceivable use to Mallory in trying to find s by other means, for he can compute such pairs by himself without involving Alice.

If having such pairs would allow him find a square root of v^{-1}, then he was already able to compute square roots, contrary to the assumption that finding square roots modulo n is difficult.

Mallory sends $b = 1$

Suppose Mallory sends $b = 1$ in step 2.

Now he ends up with the pair (x, y), where $x = r^2 \bmod n$ and $y = rs \bmod n$.

While y might seem to give information about s, observe that y itself is just a random element of $\mathbf{Z}_n$. This is because r is random, and the mapping $r \to rs \bmod n$ is one-to-one for all $s \in \mathbf{Z}_n^*$. Hence, as r ranges through all possible values, so does $rs \bmod n$.

What does Mallory learn from x?

Nothing that he could not have computed himself knowing y, for $x = y^2 v \bmod n$.

Again, all he ends up with is a random number (y in this case) and a quadratic residue x that he can compute knowing y.

Mallory learns nothing from (x, y)

In both cases, Mallory ends up with information that he could have computed without interacting with Alice.

Hence, if he could have discovered Alice's secret by talking to Alice, then he could have also done so on his own, contradicting the hardness assumption for computing square roots.

This is the sense in which Alice's protocol releases zero knowledge about her secret.