CPSC 467: Cryptography and Computer Security

Michael J. Fischer

Lecture 22
November 27, 2017

BBS Pseudorandom Sequence Generator

Secret Splitting

Shamir's Secret Splitting Scheme
- Secret Splitting with Dishonest Parties

Appendix: Security of BBS

BBS Pseudorandom Sequence Generator

Blum primes and integers

A *Blum prime* is a prime p such that $p \equiv 3 \pmod 4$.

A *Blum integer* is a number $n = pq$, where p and q are Blum primes.

If p is a Blum prime, then $-1 \in \mathrm{QNR}_p$. This follows from the Euler criterion, since $\frac{p-1}{2}$ is odd. By definition of the Legendre symbol, $\left(\frac{-1}{p}\right) = -1$.

If n is a Blum integer, then $-1 \in \mathrm{QNR}_n$, but now

$$\left(\frac{-1}{n}\right) = \left(\frac{-1}{p}\right)\left(\frac{-1}{q}\right) = (-1)(-1) = 1.$$

Square roots of Blum primes

Theorem

Let p be a Blum prime, $a \in QR_p$, and $\{b, -b\} = \sqrt{a}$ be the two square roots of a. Then exactly one of b and $-b$ is itself a quadratic residue.

Proof.

$(-b)^{(p-1)/2} \neq b^{(p-1)/2}$ since

$$(-b)^{(p-1)/2} = (-1)^{(p-1)/2} b^{(p-1)/2} = (-1) b^{(p-1)/2}.$$

Both $(-b)^{(p-1)/2}$ and $b^{(p-1)/2}$ are in $\sqrt{1} = \{\pm 1\}$, so it follows from the Euler criterion that one of b, $-b$ is a quadratic residue and the other is not. □

Square roots of Blum integers

Theorem (QR square root)

Let $n = pq$ be a Blum integer and $a \in \mathrm{QR}_n$. Exactly one of a's four square roots modulo n is a quadratic residue.

Proof of QR square root theorem

Consider $\mathbf{Z}_p^*$ and $\mathbf{Z}_q^*$. $a \in \mathrm{QR}_p$ and $a \in \mathrm{QR}_q$.

Let $\{b, -b\} \in \sqrt{a} \pmod{p}$. By the previous theorem, exactly one of these numbers is in QR_p. Call that number b_p.

Similarly, one of the square roots of $a \pmod{q}$ is in QR_q, say b_q.

Applying the Chinese Remainder Theorem, it follows that exactly one of a's four square roots modulo n is in QR_n.

A cryptographically secure PRSG

We present a cryptographically secure pseudorandom sequence generator due to Blum, Blum, and Shub (BBS).

BBS is defined by a Blum integer $n = pq$ and an integer ℓ.

It maps strings in $\mathbf{Z}_n^*$ to strings in $\{0,1\}^\ell$.

Given a seed $s_0 \in \mathbf{Z}_n^*$, we define a sequence $s_1, s_2, s_3, \ldots, s_\ell$, where $s_i = s_{i-1}^2 \bmod n$ for $i = 1, \ldots, \ell$.

The ℓ-bit output sequence BBS(s_0) is $b_1, b_2, b_3, \ldots, b_\ell$, where $b_i = \mathrm{lsb}(s_i)$ is the least significant bit of s_i.

QR assumption and Blum integers

The security of BBS is based on the assumed difficulty of determining, for a given a with Jacobi symbol 1, whether or not a is a quadratic residue, i.e., whether or not $a \in \mathrm{QR}_n$.

We just showed that Blum primes and Blum integers have the important property that every quadratic residue a has exactly one square root y which is itself a quadratic residue.

Call such a y the *principal square root* of a and denote it (ambiguously) by $\sqrt{a} \pmod{n}$ or simply by $\sqrt{a}$ when it is clear that mod n is intended.

Security of BBS

We show in the appendix that BBS is cryptographically secure.

The proof reduces the problem of predicting the output of BBS to the quadratic residuosity problem for numbers with Jacobi symbol 1 over Blum integers.

To do this reduction, we show that if there is a judge J that successfully distinguishes BBS(S) from U, then there is a feasible algorithm A for distinguishing quadratic residues from non-residues with Jacobi symbol 1, contradicting the above version of the QR hardness assumption.

Secret Splitting

Two-key locks

There are many situations in which one wants to grant access to a resource only if a sufficiently large group of agents cooperate.

For example, the office safe of a supermarket might require both the manager's key and the armored car driver's key in order to be opened.

This protects the store against a dishonest manager or armored car driver, and it also prevents an armed robber from coercing the manager into opening the safe.

A similar 2-key system is used for safe deposit boxes in banks.

Two-part secret splitting

We might like to achieve the same properties for cryptographic keys or other secrets. (This concept was introduced in Lecture 19.)

Let k be the key for a symmetric cryptosystem. One might wish to split k into two *shares* k_1 and k_2 so that by themselves, neither k_1 nor k_2 by itself reveals any information about k, but when suitably combined, k can be recovered.

A simple way to do this is to choose k_1 uniformly at random and then let $k_2 = k \oplus k_1$.

Both k_1 and k_2 are uniformly distributed over the key space and hence give no information about k.

However, combined with XOR, they reveal k, since $k = k_1 \oplus k_2$.

Comparison with one-time pad

Indeed, the one-time pad cryptosystem in the appendix of Lecture 3 can be viewed as an instance of secret splitting.

Here, Alice's secret is her message m.

The two shares are the ciphertext c and the key k.

Neither by themselves gives any information about m, but together they reveal $m = k \oplus c$.

Multi-share secret splitting

Secret splitting generalizes to more than two shares.

Imagine a large company that restricts access to important company secrets to only its five top executives, say the president, vice-president, treasurer, CEO, and CIO.

They don't want any executive to be able to access the data alone since they are concerned that an executive might be blackmailed into giving confidential data to a competitor.

Multi-share secret splitting (cont.)

On the other hand, they also don't want to require that all five executives get together to access their data because

- this would be cumbersome;
- they worry about the death or incapacitation of any single individual.

They decide as a compromise that any three of them should be able to access the secret data, but one or two of them operating alone should not have access.

Shamir's Secret Splitting Scheme

(τ, k) threshold secret spitting scheme

A (τ, k) *threshold secret splitting scheme* splits a secret s into *shares* $s_1, \ldots, s_k$.

Any subset of τ or more shares allows s to be recovered, but no subset of shares of size less than τ gives any information about s.

The executives of the previous example thus want a $(3, 5)$ threshold secret splitting scheme: The secret key is to be split into 5 shares, any 3 of which allow the secret to be recovered.

A threshold scheme based on polynomials

Shamir proposed a threshold scheme based on polynomials.

A *polynomial of degree d* is an expression

$$f(x) = a_0 + a_1 x + a_2 x^2 + \ldots + a_d x^d,$$

where $a_d \neq 0$.

The numbers $a_0, \ldots, a_d$ are called the *coefficients* of f.

A polynomial can be simultaneously regarded as a function and as an object determined by its vector of coefficients.

Interpolation

Interpolation is the process of finding a polynomial that goes through a given set of points.

Fact

Let $(x_1, y_1), \ldots, (x_k, y_k)$ *be points, where all of the* x_i*'s are distinct. There is a unique polynomial* $f(x)$ *of degree at most* $k-1$ *that passes through all* k *points, that is, for which* $f(x_i) = y_i$ $(1 \leq 1 \leq k)$.

f can be found using Lagrangian interpolation. This statement generalizes the familiar statement from high school geometry that two points determine a line.

Lagrangian interpolation method

One way to understand Lagrangian interpolation is to consider the polynomial

$$\delta_i(x) = \frac{(x - x_1)(x - x_2)\dots(x - x_{i-1}) \cdot (x - x_{i+1})\dots(x - x_k)}{(x_i - x_1)(x_i - x_2)\dots(x_i - x_{i-1}) \cdot (x_i - x_{i+1})\dots(x_i - x_k)}$$

Although this looks at first like a rational function, it is actually just a polynomial in x since the denominator contains only the x-values of the given points and not the variable x.

$\delta_i(x)$ has the easily-checked property that $\delta_i(x_i) = 1$, and $\delta_i(x_j) = 0$ for $j \neq i$.

Lagrangian interpolation method (cont.)

Using $\delta_i(x)$, the polynomial

$$p(x) = \sum_{i=1}^{k} y_i \, \delta_i(x)$$

is the desired interpolating polynomial, since $p(x_i) = y_i$ for $i = 1, \ldots, k$.

To actually find the coefficients of $p(x)$ when written as

$$p(x) = \sum_{i=0}^{k} a_i x^i,$$

it is necessary to expand $p(x)$ by multiplying out the factors and collect like terms.

Interpolation over finite fields

Interpolation also works over finite fields such as $\mathbf{Z}_p$ for prime p.

It is still true that any k points with distinct x coordinates determine a unique polynomial of degree at most $k-1$ over $\mathbf{Z}_p$.

Of course, we must have $k \leq p$ since $\mathbf{Z}_p$ has only p distinct coordinate values in all.

Shamir's secret splitting scheme

Here's how Shamir's (τ, k) secret splitting scheme works.

Let Alice (also called the *dealer*) have secret s.

She first chooses a prime $p > k$ and announces it to all players.

Constructing the polynomial

She next constructs a polynomial

$$f = a_0 + a_1 x + a_2 x^2 \ldots a_{\tau-1} x^{\tau-1}$$

of degree at most $\tau - 1$ as follows:

- She sets $a_0 = s$ (the secret).
- She chooses $a_1, \ldots, a_{\tau-1} \in Z_p$ at random.

Constructing the shares

She constructs the k shares as follows:

- She chooses $x_i = i$. $\quad (1 \leq i \leq k)$
- She chooses $y_i = f(i)$. $\quad (1 \leq i \leq k)$[1]
- Share $s_i = (x_i, y_i) = (i, f(i))$.

[1] $f(i)$ is the result of evaluating the polynomial f at the value $x = i$. All arithmetic is over the field $\mathbf{Z}_p$, so we omit explicit mention of mod p.

Recovering the secret

Theorem

s can be reconstructed from any set T of τ or more shares.

Proof.

Suppose $s_{i_1}, \ldots, s_{i_\tau}$ are τ distinct shares in T.

By interpolation, there is a unique polynomial $g(x)$ of degree $d \leq \tau - 1$ that passes through these shares.

By construction of the shares, $f(x)$ also passes through these same shares; hence $g = f$ as polynomials.

In particular, $g(0) = f(0) = s$ is the secret. □

Protection from unauthorized disclosure

Theorem

For any set T' of fewer than τ shares and any possible secret $\hat{s}$, there is a polynomial $\hat{f}$ that interprets those shares and reveals $\hat{s}$.

Proof.

Let $T' = \{s_{i_1}, \ldots, s_{i_r}\}$ be a set of $r < \tau$ shares.

In particular, for each $s' \in \mathbf{Z}_p$, there is a polynomial $g_{s'}$ that interpolates the shares in $T' \cup \{(0, s')\}$.

Each of these polynomials passes through all of the shares in T', so each is a plausible candidate for f. Moreover, $g_{s'}(0) = s'$, so each s' is a plausible candidate for the secret s. □

No information about secret

One can show further that the number of polynomials that interpolate $T' \cup \{(0, s')\}$ is the same for each $s' \in \mathbf{Z}_p$, so each possible candidate s' is equally likely to be s.

Hence, the shares in T' give no information at all about s.

Secret splitting with semi-honest parties

Shamir's scheme is an example of a protocol that works assuming *semi-honest* parties.

These are players that follow the protocol but additionally may collude in an attempt to discover secret information.

We just saw that no coalition of fewer than τ players could learn anything about the dealer's secret, even if they pooled all of their shares.

Secret splitting with dishonest dealer

In practice, either the dealer or some of the players (or both) may be dishonest and fail to follow the protocol. The honest players would like some guarantees even in such situations.

A dishonest dealer can always lie about the true value of her secret. Even so, the honest players want assurance that their shares do in fact encode a unique secret, that is, all sets of τ shares reconstruct the same secret s.

Failure of Shamir's scheme with dishonest dealer

Shamir's (τ, k) threshold scheme assumes that all k shares lie on a single polynomial of degree at most $\tau - 1$.

This might not hold if the dealer is dishonest and gives bad shares to some of the players.

The players have no way to discover that they have bad shares until later when they try to reconstruct s, and maybe not even then.

Verifiable secret sharing

In *verifiable secret sharing*, the sharing phase is an active protocol involving the dealer and all of the players.

At the end of this phase, either the dealer is exposed as being dishonest, or all of the players end up with shares that are consistent with a single secret.

Needless to say, protocols for verifiable secret sharing are quite complicated.

Dishonest players

Dishonest players present another kind of problem. These are players that fail to follow the protocol. During the reconstruction phase, they may fail to supply their share, or they may present a (possibly different) corrupted share to each other player.

With Shamir's scheme, a share that just disappears does not prevent the secret from being reconstructed, as long as enough valid shares remain.

But a player who lies about his share during the reconstruction phase can cause other players to reconstruct incorrect values for the secret.

Fault-tolerance in secret sharing protocols

A *fault-tolerant secret sharing scheme* should allow the secret to be correctly reconstructed, even in the face of a certain number of corrupted shares.

Of course, it may be desirable to have schemes that can tolerate dishonesty in both dealer and a limited number of players.

The interested reader is encouraged to explore the extensive literature on this subject.

Appendix: Security of BBS

Blum integers and the Jacobi symbol

Fact

Let n be a Blum integer and $a \in \mathrm{QR}_n$. Then $\left(\frac{a}{n}\right) = \left(\frac{-a}{n}\right) = 1$.

Proof.

This follows from the fact that if a is a quadratic residue modulo a Blum prime, then $-a$ is a quadratic non-residue. Hence,

$$\left(\frac{a}{p}\right) = -\left(\frac{-a}{p}\right) \text{ and } \left(\frac{a}{q}\right) = -\left(\frac{-a}{q}\right), \text{ so}$$

$$\left(\frac{a}{n}\right) = \left(\frac{a}{p}\right) \cdot \left(\frac{a}{q}\right) = \left(-\left(\frac{-a}{p}\right)\right) \cdot \left(-\left(\frac{-a}{q}\right)\right) = \left(\frac{-a}{n}\right).$$

□

Blum integers and the least significant bit

The low-order bits of x mod n and $(-x)$ mod n always differ when n is odd.

Let $\mathrm{lsb}(x) = (x \bmod 2)$ be the least significant bit of integer x.

Fact

If n is odd, then $\mathrm{lsb}(x \bmod n) \oplus \mathrm{lsb}((-x) \bmod n) = 1$.

First-bit prediction

A *first-bit predictor with advantage* ϵ is a probabilistic polynomial time algorithm A that, given $b_2, \ldots, b_\ell$, correctly predicts b_1 with probability at least $1/2 + \epsilon$.

This is not sufficient to establish that the pseudorandom sequence BBS(S) is indistinguishable from the uniform random sequence U, but if it did not hold, there certainly would exist a distinguishing judge.

Namely, the judge that outputs 1 if $b_1 = A(b_2, \ldots, b_\ell)$ and 0 otherwise would output 1 with probability greater than $1/2 + \epsilon$ in the case that the sequence came from BBS(S) and would output 1 with probability exactly $1/2$ in the case that the sequence was truly random.

BBS has no first-bit predictor under the QR assumption

If BBS has a first-bit predictor A with advantage ϵ, then there is a probabilistic polynomial time algorithm Q for testing quadratic residuosity with the same accuracy.

Thus, if quadratic-residue-testing is "hard", then so is first-bit prediction for BBS.

Theorem

Let A be a first-bit predictor for BBS(S) with advantage ϵ. Then we can find an algorithm Q for testing whether a number x with Jacobi symbol 1 is a quadratic residue, and Q will be correct with probability at least $1/2 + \epsilon$.

Construction of Q

Assume that A predicts b_1 given $b_2, \ldots, b_\ell$.

Algorithm $Q(x)$ tests whether or not a number x with Jacobi symbol 1 is a quadratic residue modulo n.

It outputs 1 to mean $x \in \mathrm{QR}_n$ and 0 to mean $x \notin \mathrm{QR}_n$.

To $Q(x)$:

1. Let $\hat{s}_2 = x^2 \bmod n$.
2. Let $\hat{s}_i = \hat{s}_{i-1}^2 \bmod n$, for $i = 3, \ldots, \ell$.
3. Let $\hat{b}_1 = \mathrm{lsb}(x)$.
4. Let $\hat{b}_i = \mathrm{lsb}(\hat{s}_i)$, for $i = 2, \ldots, \ell$.
5. Let $c = A(\hat{b}_2, \ldots, \hat{b}_\ell)$.
6. If $c = \hat{b}_1$ then output 1; else output 0.

Why Q works

Since $\left(\frac{x}{n}\right) = 1$, then either x or $-x$ is a quadratic residue. Let s_0 be the principal square root of x or $-x$. Let $s_1, \ldots, s_\ell$ be the state sequence and $b_1, \ldots, b_\ell$ the corresponding output bits of BBS(s_0).

We have two cases.

Case 1: $x \in \mathrm{QR}_n$. Then $s_1 = x$, so the state sequence of BBS(s_0) is

$$s_1, s_2, \ldots, s_\ell = x, \hat{s}_2, \ldots, \hat{s}_\ell,$$

and the corresponding output sequence is

$$b_1, b_2, \ldots, b_\ell = \hat{b}_1, \hat{b}_2, \ldots, \hat{b}_\ell.$$

Since $\hat{b}_1 = b_1$, $Q(x)$ correctly outputs 1 whenever A correctly predicts b_1. This happens with probability at least $1/2 + \epsilon$.

Why Q works (cont.)

Case 2: $x \in \mathrm{QNR}_n$, so $-x \in \mathrm{QR}_n$. Then $s_1 = -x$, so the state sequence of BBS(s_0) is

$$s_1, s_2, \ldots, s_\ell = -x, \hat{s}_2, \ldots, \hat{s}_\ell,$$

and the corresponding output sequence is

$$b_1, b_2, \ldots, b_\ell = \neg\hat{b}_1, \hat{b}_2, \ldots, \hat{b}_\ell.$$

Since $\hat{b}_1 = \neg b_1$, $Q(x)$ correctly outputs 0 whenever A correctly predicts b_1. This happens with probability at least $1/2 + \epsilon$.
In both cases, $Q(x)$ gives the correct output with probability at least $1/2 + \epsilon$.