

CPSC 467b: Cryptography and Computer Security

Michael J. Fischer

Lecture 19
April 4, 2013

Non-interactive Interactive Proofs

Feige-Fiat-Shamir Signatures

Pseudorandom Sequence Generation

Quadratic Residues Revisited

- The Legendre symbol

- Jacobi symbol

- Computing the Jacobi symbol

Non-interactive Interactive Proofs

Eliminating interaction from interactive proofs

Going from serial composition to parallel composition reduces communication overhead but may sacrifice of zero knowledge.

Rather surprisingly, one can go a step further and eliminate the interaction from interactive proofs altogether.

The idea is that Alice will provide Bob with a **trace of a pretend execution** of an interactive proof of herself interacting with Bob.

Bob will check that the trace is a valid execution of the protocol.

Of course, that isn't enough to convince Bob that Alice isn't cheating, for how does he ensure that Alice simulates random query bits b_i for him, and how does he ensure that Alice chooses her x_i 's before knowing the b_i 's?

Keeping Alice from cheating

The solution is to make the b_i 's depend in an unpredictable way on the x_i 's.

We base the b_i 's on the value of a “random-looking” hash function H applied to the concatenation of the x_i 's.

A non-interacting version of FFS

Here's how it works in, say, the parallel composition of t copies of the simplified FFS protocol.

- ▶ The honest Alice chooses x_1, \dots, x_t according to the protocol.
- ▶ Next she chooses $b_1 \dots b_t$ to be the first t bits of $H(x_1 \cdots x_t)$.
- ▶ Finally, she computes y_1, \dots, y_t , again according to the protocol.
- ▶ She sends Bob a single message consisting of $x_1, \dots, x_t, y_1, \dots, y_t$.
- ▶ Bob computes $b_1 \dots b_t$ to be the first t bits of $H(x_1 \cdots x_t)$ and then performs each of the t checks of the FFS protocol, accepting Alice's proof only if all checks succeed.

Why can't Alice cheat?

A cheating Alice can choose y_i arbitrarily and then compute a valid x_i for a given b_i .

If she chooses the b_i 's first, the x_i 's she computes are unlikely to hash to a string that begins with $b_1 \dots b_t$.¹

¹This assumes that the hash function “looks like” a random function. We have already seen artificial examples of hash functions that do not have this property.

Why can't Alice cheat?

If some b_i does not agree with the corresponding bit of the hash function, she can either change b_i and try to find a new y_i that works with the given x_i , or she can change x_i to try to get the i^{th} bit of the hash value to change.

However, neither of these approaches works. The former may require knowledge of Alice's secret; the latter will cause the bits of the hash function to change “randomly”.

Brute force cheating

One way Alice can attempt to cheat is to use a brute-force attack.

For example, she could generate all of the x_i 's to be squares of the y_i with the hopes that the hash of the x_i 's will make all $b_i = 0$.

But that is likely to require 2^{t-1} attempts on average.

If t is chosen large enough (say $t = 80$), the number of trials Alice would have to do in order to have a significant probability of success is prohibitive.

Of course, these observations are not a proof that Alice can't cheat; only that the obvious strategies don't work.

Nevertheless, it is plausible that a cheating Alice not knowing Alice's secret, really wouldn't be able to find a valid such "non-interactive interactive proof".

Contrast with true interactive proofs

With a true zero-knowledge interactive proof, Bob does not learn anything about Alice's secret, nor can Bob impersonate Alice to Carol after Alice has authenticated herself to Bob.

On the other hand, if Alice sends Bob a valid non-interactive proof, then Bob can in turn send it on to Carol.

Even though Bob couldn't have produced it on his own, it is still valid.

So here we have the curious situation that Alice needs her secret in order to produce the non-interactive proof string π , and Bob can't learn Alice's secret from π , but now Bob can use π itself in an attempt to impersonate Alice to Carol.

Feige-Fiat-Shamir Signatures

Similarity between signature scheme and non-interactive IP

A signature scheme has a lot in common with the “non-interactive interactive” proofs.

In both cases, there is only a one-way communication from Alice to Bob.

- ▶ Alice signs a message and sends it to Bob.
- ▶ Bob verifies it without further interaction with Alice.
- ▶ If Bob hands the message to Carol, then Carol can also verify that it was signed by Alice.

Not surprisingly, the “non-interactive interactive proof” ideas can be used to turn the Feige-Fiat-Shamir authentication protocol into a signature scheme.

Signature scheme from non-interactive IP

We present a signature scheme based on a slightly simplified version of the full FFS authentication protocol in which all of the v_i 's in the public key are quadratic residues, and n is not required to be a Blum integer, only a product of two distinct odd primes.

The public verification key is (n, v_1, \dots, v_k) , and the private signing key is (n, s_1, \dots, s_k) , where $v_j = s_j^{-2} \pmod n$ ($1 \leq j \leq k$).

Signing algorithm

To sign a message m , Alice simulates t parallel rounds of FFS.

- ▶ She first chooses random $r_1, \dots, r_t \in \mathbf{Z}_n - \{0\}$ and computes

$$x_i = r_i^2 \bmod n \quad (1 \leq i \leq t).$$

- ▶ She computes $u = H(mx_1 \cdots x_t)$, where H is a suitable cryptographic hash function.
- ▶ She chooses $b_{1,1}, \dots, b_{t,k}$ according to the first tk bits of u :

$$b_{i,j} = u_{(i-1)*k+j} \quad (1 \leq i \leq t, 1 \leq j \leq k).$$

- ▶ Finally, she computes

$$y_i = r_i s_1^{b_{i,1}} \cdots s_k^{b_{i,k}} \bmod n \quad (1 \leq i \leq t).$$

The signature is

$$s = (b_{1,1}, \dots, b_{t,k}, y_1, \dots, y_t).$$

Verification algorithm

To verify the signed message (m, s) , Bob computes

$$z_i = y_i^2 v_1^{b_{i,1}} \cdots v_k^{b_{i,k}} \bmod n \quad (1 \leq i \leq t).$$

Bob checks that each $z_i \neq 0$ and that $b_{1,1}, \dots, b_{t,k}$ are equal to the first tk bits of $H(mz_1 \cdots z_t)$.

When both Alice and Bob are honest, it is easily verified that $z_i = x_i$ ($1 \leq i \leq t$). In that case, Bob's checks all succeed since $x_i \neq 0$ and $H(mz_1 \cdots z_t) = H(mx_1 \cdots x_t)$.

Forgery

To forge Alice's signature, an impostor must find $b_{i,j}$'s and y_i 's that satisfy the equation

$$b_{1,1} \dots b_{t,k} \preceq H(m(y_1^2 v_1^{b_{1,1}} \dots v_k^{b_{1,k}} \text{ mod } n) \dots (y_t^2 v_1^{b_{t,1}} \dots v_k^{b_{t,k}} \text{ mod } n)).$$

where “ \preceq ” means string prefix. It is not obvious how to solve such an equation without knowing a square root of each of the v_i^{-1} 's and following essentially Alice's procedure.

Pseudorandom Sequence Generation

Pseudorandom sequence generators revisited

Cryptographically strong pseudorandom sequence generators were introduced in Lecture 6 in connection with stream ciphers.

We next define carefully what it means for a pseudorandom sequence generator (PRSG) to be *cryptographically strong*.

We then show how to build one that is provably secure. It is based on the quadratic residuosity assumption (Lecture 17) on which the Goldwasser-Micali probabilistic cryptosystem is based.

Desired properties of a PRSG

A pseudorandom sequence generator (PRSG) maps a “short” random seed to a “long” pseudorandom bit string.

We want a PRSG to be *cryptographically strong*, that is, it must be **difficult to correctly predict any generated bit**, even knowing all of the other bits of the output sequence.

In particular, it must also be difficult to find the seed given the output sequence, since otherwise the whole sequence is easily generated.

Thus, a PRSG is a one-way function and more.

Note: While a hash function might generate hash values of the form yy and still be strongly collision-free, such a function could not be a PRSG since it would be possible to predict the second half of the output knowing the first half.

Expansion amount

I am being intentionally vague about how much **expansion** we expect from a PRSG that maps a “short” seed to a “long” pseudorandom sequence.

Intuitively, “**short**” is a length like we use for cryptographic keys—long enough to prevent brute-force attacks, but generally much shorter than the data we want to deal with. Typical seed lengths might range from 128 to 2048.

By “**long**”, we mean much larger sizes, perhaps thousands or even millions of bits, but polynomially related to the seed length.

Incremental generators

In practice, the output length is usually variable. We can request as many output bits from the generator as we like (within limits), and it will deliver them.

In this case, “long” refers to the maximum number of bits that can be delivered while still maintaining security.

Also, in practice, the bits are generally delivered a few at a time rather than all at once, so we don't need to announce in advance how many bits we want but can go back as needed to get more.

Notation for PRSG's

In a little more detail, a *pseudorandom sequence generator* G is a function from a domain of *seeds* \mathcal{S} to a domain of strings \mathcal{X} .

We will generally assume that all of the seeds in \mathcal{S} have the same length n and that \mathcal{X} is the set of all binary strings of length $\ell = \ell(n)$, where $\ell(\cdot)$ is a polynomial and $n \ll \ell(n)$.

$\ell(\cdot)$ is called the *expansion factor* of G .

What does it mean for a string to look random?

Intuitively, we want the strings $G(s)$ to “look random”.
But what does it mean to “look random”?

Chaitin and Kolmogorov defined a string to be “random” if its shortest description is almost as long as the string itself.

By this definition, most strings are random by a simple counting argument.

For example, `011011011011011011011011011` is easily described as the pattern `011` repeated 9 times. On the other hand, `101110100010100101001000001` has no obvious short description.

While philosophically very interesting, these notions are somewhat different than the statistical notions that most people mean by randomness and do not seem to be useful for cryptography.

Randomness based on probability theory

We take a different tack.

We assume that the seeds are chosen uniformly at random from \mathcal{S} .

Let S be a uniformly distributed random variable over \mathcal{S} .

Then $X \in \mathcal{X}$ is a derived random variable defined by $X = G(S)$.

For $x \in \mathcal{X}$,

$$\Pr[X = x] = \frac{|\{s \in \mathcal{S} \mid G(s) = x\}|}{|\mathcal{S}|}.$$

Thus, $\Pr[X = x]$ is the probability of obtaining x as the output of the PRSG for a randomly chosen seed.

Randomness amplifier

We think of $G(\cdot)$ as a *randomness amplifier*.

We start with a short truly random seed and obtain a long string that “looks like” a random string, even though we know it’s not uniformly distributed.

In fact, the distribution $G(S)$ is very much non-uniform.

Because $|\mathcal{S}| \leq 2^n$, $|\mathcal{X}| = 2^\ell$, and $n \ll \ell$, **most strings in \mathcal{X} are not in the range of G** and hence have probability 0.

For the uniform distribution U over \mathcal{X} , all strings have the same non-zero probability $1/2^\ell$.

U is what we usually mean by a *truly random* variable on ℓ -bit strings.

Computational indistinguishability

We have just seen that the probability distributions of $X = G(S)$ and U are quite different.

Nevertheless, it may be the case that all feasible probabilistic algorithms behave essentially the same whether given a sample chosen according to X or a sample chosen according to U .

If that is the case, we say that X and U are *computationally indistinguishable* and that G is a *cryptographically strong* pseudorandom sequence generator.

Some implications of computational indistinguishability

Before going further, let me describe some functions G for which $G(S)$ is readily distinguished from U .

Suppose every string $x = G(s)$ has the form $b_1b_1b_2b_2b_3b_3\dots$, for example 0011111100001100110000....

Algorithm $A(x)$ outputs “G” if x is of the special form above, and it outputs “U” otherwise.

A will always output “G” for inputs from $G(S)$. For inputs from U , A will output “G” with probability only

$$\frac{2^{\ell/2}}{2^\ell} = \frac{1}{2^{\ell/2}}.$$

How many strings of length ℓ have the special form above?

Judges

Formally, a *judge* is a probabilistic polynomial-time algorithm J that takes an ℓ -bit input string x and outputs a single bit b .

Thus, it defines a *random function* from \mathcal{X} to $\{0, 1\}$.

This means that for every input x , the output is 1 with some probability p_x , and the output is 0 with probability $1 - p_x$.

If the input string is a random variable X , then the probability that the output is 1 is the weighted sum of p_x over all possible inputs x , where the weight is the probability $\Pr[X = x]$ of input x occurring.

Thus, the output value is itself a random variable $J(X)$, where

$$\Pr[J(X) = 1] = \sum_{x \in \mathcal{X}} \Pr[X = x] \cdot p_x.$$

Formal definition of indistinguishability

Two random variables X and Y are *ϵ -indistinguishable by judge J* if

$$|\Pr[J(X) = 1] - \Pr[J(Y) = 1]| < \epsilon.$$

Intuitively, we say that G is *cryptographically strong* if $G(S)$ and U are ϵ -indistinguishable for suitably small ϵ by all judges that do not run for too long.

A careful mathematical treatment of the concept of indistinguishability must relate the length parameters n and ℓ , the error parameter ϵ , and the allowed running time of the judges

Further formal details may be found in Goldwasser and Bellare and in [handout 9](#).

Quadratic Residues Revisited

Legendre symbol

Let p be an odd prime, a an integer. The *Legendre symbol* $\left(\frac{a}{p}\right)$ is a number in $\{-1, 0, +1\}$, defined as follows:

$$\left(\frac{a}{p}\right) = \begin{cases} +1 & \text{if } a \text{ is a non-trivial quadratic residue modulo } p \\ 0 & \text{if } a \equiv 0 \pmod{p} \\ -1 & \text{if } a \text{ is not a quadratic residue modulo } p \end{cases}$$

By the Euler Criterion, we have

Theorem

Let p be an odd prime. Then

$$\left(\frac{a}{p}\right) \equiv a^{\left(\frac{p-1}{2}\right)} \pmod{p}$$

Note that this theorem holds even when $p \mid a$.

Properties of the Legendre symbol

The Legendre symbol satisfies the following *multiplicative property*:

Fact

Let p be an odd prime. Then

$$\left(\frac{a_1 a_2}{p}\right) = \left(\frac{a_1}{p}\right) \left(\frac{a_2}{p}\right)$$

Not surprisingly, if a_1 and a_2 are both non-trivial quadratic residues, then so is $a_1 a_2$. Hence, the fact holds when

$$\left(\frac{a_1}{p}\right) = \left(\frac{a_2}{p}\right) = 1.$$

Product of two non-residues

Suppose $a_1 \notin \text{QR}_p$, $a_2 \notin \text{QR}_p$. The above fact asserts that **the product $a_1 a_2$ is a quadratic residue** since

$$\left(\frac{a_1 a_2}{p}\right) = \left(\frac{a_1}{p}\right) \left(\frac{a_2}{p}\right) = (-1)(-1) = 1.$$

Here's why.

- ▶ Let g be a primitive root of p .
- ▶ Write $a_1 \equiv g^{k_1} \pmod{p}$ and $a_2 \equiv g^{k_2} \pmod{p}$.
- ▶ Both k_1 and k_2 are odd since $a_1, a_2 \notin \text{QR}_p$.
- ▶ But then $k_1 + k_2$ is even.
- ▶ Hence, $g^{(k_1+k_2)/2}$ is a square root of $a_1 a_2 \equiv g^{k_1+k_2} \pmod{p}$, so $a_1 a_2$ is a quadratic residue.

The Jacobi symbol

The *Jacobi symbol* extends the Legendre symbol to the case where the “denominator” is an arbitrary odd positive number n .

Let n be an odd positive integer with prime factorization $\prod_{i=1}^k p_i^{e_i}$. We define the *Jacobi symbol* by

$$\left(\frac{a}{n}\right) = \prod_{i=1}^k \left(\frac{a}{p_i}\right)^{e_i} \quad (1)$$

The symbol on the left is the Jacobi symbol, and the symbol on the right is the Legendre symbol.

(By convention, this product is 1 when $k = 0$, so $\left(\frac{a}{1}\right) = 1$.)

The Jacobi symbol extends the Legendre symbol since the two definitions coincide when n is an odd prime.

Meaning of Jacobi symbol

What does the Jacobi symbol mean when n is not prime?

- ▶ If $\left(\frac{a}{n}\right) = +1$, a **might or might not be** a quadratic residue.
- ▶ If $\left(\frac{a}{n}\right) = 0$, then $\gcd(a, n) \neq 1$.
- ▶ If $\left(\frac{a}{n}\right) = -1$ then a is **definitely not** a quadratic residue.

Jacobi symbol = +1 for $n = pq$

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

$$\left(\frac{a}{n}\right) = \left(\frac{a}{p}\right) \left(\frac{a}{q}\right) \quad (2)$$

there are two cases that result in $\left(\frac{a}{n}\right) = 1$:

1. $\left(\frac{a}{p}\right) = \left(\frac{a}{q}\right) = +1$, or
2. $\left(\frac{a}{p}\right) = \left(\frac{a}{q}\right) = -1$.

Case of both Jacobi symbols = +1

If $\left(\frac{a}{p}\right) = \left(\frac{a}{q}\right) = +1$, then $a \in \text{QR}_p \cap \text{QR}_q = \text{QR}_n^{11}$.

It follows by the Chinese Remainder Theorem that $a \in \text{QR}_n$.

This fact was implicitly used in the proof sketch that $|\sqrt{a}| = 4$.

Case of both Jacobi symbols = -1

If $\left(\frac{a}{p}\right) = \left(\frac{a}{q}\right) = -1$, then $a \in \text{QNR}_p \cap \text{QNR}_q = \text{Q}_n^{00}$.

In this case, a is *not* a quadratic residue modulo n .

Such numbers a are sometimes called “pseudo-squares” since they have Jacobi symbol 1 but are not quadratic residues.

Computing the Jacobi symbol

The Jacobi symbol $\left(\frac{a}{n}\right)$ is easily computed from its definition (equation 1) and the Euler Criterion, given the factorization of n .

Similarly, $\gcd(u, v)$ is easily computed without resort to the Euclidean algorithm given the factorizations of u and v .

The remarkable fact about the Euclidean algorithm is that it lets us compute $\gcd(u, v)$ efficiently, without knowing the factors of u and v .

A similar algorithm allows us to compute the Jacobi symbol $\left(\frac{a}{n}\right)$ efficiently, without knowing the factorization of a or n .

Identities involving the Jacobi symbol

The algorithm is based on identities satisfied by the Jacobi symbol:

$$1. \left(\frac{0}{n}\right) = \begin{cases} 1 & \text{if } n = 1 \\ 0 & \text{if } n \neq 1; \end{cases}$$

$$2. \left(\frac{2}{n}\right) = \begin{cases} 1 & \text{if } n \equiv \pm 1 \pmod{8} \\ -1 & \text{if } n \equiv \pm 3 \pmod{8}; \end{cases}$$

$$3. \left(\frac{a_1}{n}\right) = \left(\frac{a_2}{n}\right) \text{ if } a_1 \equiv a_2 \pmod{n};$$

$$4. \left(\frac{2a}{n}\right) = \left(\frac{2}{n}\right) \cdot \left(\frac{a}{n}\right);$$

$$5. \left(\frac{a}{n}\right) = \begin{cases} \left(\frac{n}{a}\right) & \text{if } a, n \text{ odd and } \neg(a \equiv n \equiv 3 \pmod{4}) \\ -\left(\frac{n}{a}\right) & \text{if } a, n \text{ odd and } a \equiv n \equiv 3 \pmod{4}. \end{cases}$$

A recursive algorithm for computing Jacobi symbol

```
/* Precondition: a, n >= 0; n is odd */
int jacobi(int a, int n) {
    if (a == 0) /* identity 1 */
        return (n==1) ? 1 : 0;
    if (a == 2) /* identity 2 */
        switch (n%8) {
            case 1: case 7: return 1;
            case 3: case 5: return -1;
        }
    if ( a >= n ) /* identity 3 */
        return jacobi(a%n, n);
    if (a%2 == 0) /* identity 4 */
        return jacobi(2,n)*jacobi(a/2, n);
    /* a is odd */ /* identity 5 */
    return (a%4 == 3 && n%4 == 3) ? -jacobi(n,a) : jacobi(n,a);
}
```