## Lecture Notes 19

## 93 Composing Zero-Knowledge Proofs

One round of the simplified FFS protocol has probability 0.5 of error. That is, Mallory could fool Bob half the time. This is unacceptably high for most applications. By repeating the protocol $t$ times, we reduce this error probability to $1 / 2^{t}$. Taking $t=20$, for example, reduces the probability of error to less than on in a million. The downside of such serial repetition is that it also requires $t$ round trip messages between Alice and Bob (plus a final message from Alice to Bob).

## 94 Parallel Execution of Zero-Knowledge Proofs

One could also imagine running the $t$ executions of the protocol in parallel. Let $\left(x_{i}, b_{i}, y_{i}\right)$ be the messages exchanged during the $i^{\text {th }}$ execution of the simplified FFS protocol, $1 \leq i \leq t$. In a protocol execution, Alice sends $\left(x_{1}, \ldots, x_{t}\right)$ to Bob, then Bob sends $\left(b_{1}, \ldots, b_{t}\right)$ to Alice, then Alice sends $\left(y_{1}, \ldots, y_{t}\right)$ to Bob, and finally Bob checks the $t$ sets of values he has received, accepting only if all checks pass.

A parallel execution is certainly attractive in practice, for it reduces the number of round-trip messages between Alice and Bob to only $11 / 2$. The downside is that the resulting protocol may not be zero knowledge by our definition. Intuitively, the important difference between the serial and parallel versions of the protocol is that in the latter, Bob gets to know all of the $x_{i}$ 's before choosing the $b_{i}$ 's. While it seems implausible that this would actually help a cheating Bob to compromise Alice secret, the simulation proof used to show that a protocol is zero knowledge no longer works. First Sam would have to guess $\left(\hat{b}_{1}, \ldots \hat{b}_{t}\right)$. Then he can construct the $x_{i}$ 's and $y_{i}$ 's as before. But when he finally gets to the point in Mallory's program that Mallory generates the $b_{i}$ 's, the chance is very high that his initial guesses were wrong and he will be forced to start over again. Indeed, the probability that all $t$ initial guesses are correct is only $1 / 2^{t}$.

## 95 Full Feige-Fiat-Shamir Authentication Protocol

The full Feige-Fiat-Shamir Authentication Protocol combines ideas of serial and parallel execution to get a protocol that exhibits some of the properties of both.

A Blum prime is a prime $p$ such that $p \equiv 3(\bmod 4)$. A Blum integer is a number $n=p q$, were $p$ and $q$ are both Blum primes. A special property of Blum integers is that if $a$ is a quadratic residue modulo $n$, then exactly one of $a$ 's four square roots modulo $n$ is itself a quadratic residue.

To see this, note that $a$ is a quadratic residue modulo both $p$ and $q$. Claim 3, section 64, of lecture notes 12, shows that one of $a$ 's two square roots is a quadratic residue modulo $p$, say $b_{p}$. Then $-b_{p}$ is not a quadratic residue since -1 is not a quadratic residue modulo the Blum prime $p$ by the Euler Criterion of lecture notes 12, section 63. Similar, exactly one of $a$ 's two square roots modulo $q$ is a quadratic residue. Applying the Chinese Remainder Theorem, it follows that exactly one of $a$ 's four square roots modulo $n$ is a quadratic residue.

To generate the public and private keys of the full FFS protocol, Alice chooses a Blum integer $n$. Next she chooses random numbers $s_{1}, \ldots, s_{k} \in \mathbf{Z}_{n}^{*}$ and random bits $c_{1}, \ldots, c_{k} \in\{0,1\}$. From these, she computes $v_{i}=(-1)^{c_{i}} s_{i}^{-2} \bmod n$, for $i=1, \ldots, k$. She makes $\left(n, v_{1}, \ldots, v_{k}\right)$ public and keeps $\left(n, s_{1}, \ldots, s_{k}\right)$ private.

Notice that every $v_{i}$ is either a quadratic residue or the negation of a quadratic residue. It is easily shown that all of the $v_{i}$ have Jacobi symbol 1 modulo $n$. A round of the protocol itself is shown in Figure 1. The protocol is repeated for a total of $t$ rounds.

$$
\text { Alice } \quad \text { Bob }
$$

1. Choose random $r \in \mathbf{Z}_{n}-\{0\}$.

Choose random $c \in\{0,1\}$.
Compute $x=(-1)^{c} r^{2} \bmod n \quad \xrightarrow{x}$
2.
3. Compute $y=r s_{1}^{b_{1}} \cdots s_{k}^{b_{k}} \bmod n$.
$\stackrel{b_{1}, \ldots, b_{k}}{\longleftrightarrow}$. Choose random $b_{1}, \ldots, b_{k} \in\{0,1\}$.
4. Compute $z=y^{2} v_{1}^{b_{1}} \cdots v_{k}^{b_{k}} \bmod n$. Check $z \equiv \pm x(\bmod n)$ and $z \neq 0$.

Figure 1: One round of the full Feige-Fiat-Shamir authentication protocol.
To see that the protocol works when both Alice and Bob are honest, one needs to verify that Bob's checks will succeed. Plugging in for $y$, we get that

$$
z=r^{2}\left(s_{1}^{2 b_{1}} \cdots s_{k}^{2 b_{k}}\right)\left(v_{1}^{b_{1}} \cdots v_{k}^{b_{k}}\right) \bmod n
$$

Since $v_{i}=(-1)^{c_{i}} s_{k}^{-2}$, it follows that $s_{i}^{2} v_{i}=(-1)^{c_{i}}$. Hence,

$$
z \equiv r^{2}\left(s_{1}^{2} v_{1}\right)^{b_{1}} \cdots\left(s_{k}^{2} v_{k}\right)^{b_{k}} \equiv x(-1)^{c}(-1)^{c_{1} b_{1}} \cdots(-1)^{c_{k} b_{k}} \equiv \pm x(\bmod n)
$$

Moreover, since $x \neq 0$, then also $z \neq 0$. Hence, Bob's checks succeed.
The chance that a bad Alice can fool Bob is only $1 / 2^{k t}$. The authors recommend $k=5$ and $t=4$ for a failure probability of $1 / 2^{20}$.

## 96 Non-interactive Interactive Proofs

We have seen that going from a serial composition of interactive proofs to a parallel version reduces communication overhead but possibly at the sacrifice of zero knowledge. Rather surprisingly, one can go a step further to eliminate the interaction from interactive proofs altogether.

The idea here is that Alice will provide Bob with a trace of an execution of herself in an interactive protocol with Bob. Bob will check the trace to make sure that it is a possible record of an interaction between the two of them when both are following 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?

The solution to both of these puzzles is to make the $b_{i}$ 's depend in an unpredictable way on the $x_{i}$ 's. We do this by choosing the $b_{i}$ 's from the value of a hash function applied to the concatenation of the $x_{i}$ 's. Here's how it works in, say, the parallel composition of $t$ copies of the simplified FFS protocol. The honest Alice chooses $x_{1}, \ldots, x_{t}$ according to the protocol. Next she chooses
$b_{1} \ldots b_{t}$ to be the first $t$ bits of $H\left(x_{1} \cdots x_{t}\right)$. Finally, she computes $y_{1}, \ldots, y_{t}$, again according to the protocol. She sends Bob a single message consisting of $x_{1}, \ldots, x_{t}, y_{1}, \ldots, y_{t}$. Bob computes $b_{1} \ldots b_{t}$ to be the first $t$ bits of $H\left(x_{1} \cdots x_{t}\right)$ and then performs each of the $t$ checks of the FFT protocol, accepting Alice's proof only if all checks succeed.

A cheating Alice can choose $y_{i}$ arbitrarily and then compute a valid $x_{i}$ for a given $b_{i}$. But if she chooses the $b_{i}$ 's first, the chance that the $x_{i}$ 's she then computes will hash to a string that begins with $b_{1} \ldots b_{t}$ is only $1 / 2^{t}{ }^{1}$ 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 $i^{\text {th }}$ bit of the hash value to change. However, neither of these approaches works. The former requires knowledge of Alice's secret; the latter will cause all of the bits of the hash function to change "randomly".

Note that 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".

Even if this is so, there is an important difference between the true interactive proofs we have been discussing and this kind of non-interactive proof. 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 $\pi$, and Bob can't learn Alice's secret from $\pi$, but now Bob can use $\pi$ itself in an attempt to impersonate Alice to Carol.

## 97 Feige-Fiat-Shamir Signatures

A signature scheme has a lot in common with the "non-interactive interactive" proofs introduced in section 96 . In both cases, there is only a one-way communication from Alice to Bob. Alice signs a message and sends it to Bob. Bob then 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 of section 95 into a signature scheme. The signature scheme we present here is based on a slightly simplified version of the aforementioned 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 $\left(n, v_{1}, \ldots, v_{k}\right)$, and the private signing key is $\left(n, s_{1}, \ldots, s_{k}\right)$, where $v_{j}=s_{j}^{-2} \bmod n(1 \leq j \leq k)$.

To sign a message $m$, Alice simulates $t$ rounds of the protocol in parallel. She first chooses random $r_{1}, \ldots, r_{t} \in \mathbf{Z}_{n}-\{0\}$ and computes

$$
x_{i}=r_{i}^{2} \bmod n(1 \leq i \leq t)
$$

[^0]Next she computes $u=H\left(m x_{1} \cdots x_{t}\right)$, where $H$ is a suitable cryptographic hash function. She chooses $b_{1,1}, \ldots, b_{t, k}$ according to the first $t k$ bits of $u$, that is,

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

Finally, she computes

$$
y_{i}=r s_{1}^{b_{i, 1}} \cdots s_{k}^{b_{i, k}} \bmod n(1 \leq i \leq t) .
$$

The signature is

$$
s=\left(b_{1,1}, \ldots, b_{t, k}, y_{1}, \ldots, y_{t}\right) .
$$

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

$$
z_{i}=y_{i}^{2} v_{1}^{b_{i, 1}} \cdots v_{k}^{b_{i, k}} \bmod n(1 \leq i \leq t)
$$

Bob checks that each $z_{i} \neq 0$ and that $b_{1,1}, \ldots, b_{t, k}$ are equal to the first $t k$ bits of $H\left(m z_{1} \cdots z_{t}\right)$.
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 $H\left(m z_{1} \cdots z_{t}\right)=H\left(m x_{1} \cdots x_{t}\right)$.

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

$$
b_{1,1} \ldots b_{t, k} \preceq H\left(m\left(y_{1}^{2} v_{1}^{b_{1,1}} \cdots v_{k}^{b_{1, k}} \bmod n\right) \ldots\left(y_{t}^{2} v_{1}^{b_{t, 1}} \cdots v_{k}^{b_{t, k}} \bmod n\right)\right) .
$$

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.

## 98 Two-Part Secret Splitting

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, a store safe 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.

We might like to achieve the same properties for cryptographic keys or other secrets. For example, if $k$ is the secret decryption key for a cryptosystem, one might wish to split $k$ into two shares $k_{1}$ and $k_{2}$. By themselves, neither $k_{1}$ nor $k_{2}$ 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 $k=k_{1} \oplus k_{2}$.

Indeed, the one-time pad cryptosystem of lecture notes 3, section 14, 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 $m=k \oplus c$.

## 99 Multi-Part 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.

On the other hand, they also don't want to require that all five executives get together to access their data, both because this would be cumbersome and also because 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 not one or two of them operating alone.

A $(\tau, k)$ threshold secret splitting scheme splits a secret $s$ into shares $s_{1}, \ldots, s_{k}$. Any subset of $\tau$ or more shares allows $s$ to be recovered, but no subset of shares of size less than $\tau$ gives any information about $s$.


[^0]:    ${ }^{1}$ This 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.

