

Lecture 1

1 Overview of Cryptography

This section gives an overview of the various branches of cryptography. ¹

1.1 Brief History

While cryptography has been in use since the times of the Romans, it has only recently been formalized into a science.

- Claude Shannon (1940s)
 - Formalized "Perfect Security" while working on the ENIGMA project
 - Defined the private key communication model
- Whitfield Diffie and Martin Hellman (1970s)
 - Laid foundations for cryptographic techniques based on complexity theoretical assumptions
 - Defined Public-Key Cryptography
- Shafi Goldwasser and Silvio Micali (1980s)
 - Provided formal and first satisfactory definition of encryption

1.2 Some example problem spaces in Cryptography

Secrecy in Communication

Two parties (Alice and Bob) want to be able to communicate with each other in some manner such that an adversary (Eve) intercepting their communication will not be able to get any information about the message being sent.

¹Scribe notes by Ratzin, Oliveria, Draganova, Chandran, Moriarty; 1/10 - 1/12, 2005

Non-Malleability

Alice and Bob want to be able to communicate with each other in some manner such that Eve intercepting their communication will not be able to modify any message of the communication to a related message.

Authentication/Data Integrity

In this problem Alice would like to send Bob a message in such a way that Eve cannot replace the message with his own message without Bob's knowledge.

1.3 Cryptographic Primitives

Cryptographic primitives are used as building blocks to build more advanced security protocols and crypto-systems.

Zero Knowledge Proofs

In this problem, there exists two parties, a prover P and a verifier V . P has a statement that he wishes to prove to V . P and V would like to run a protocol at the end of which V is convinced that the statement is true but gains no more knowledge than the validity of the statement. P , even though being infinitely powerful, should not be able to prove false statements to V .

Coin Flipping Protocol

A coin flipping protocol is one in which Alice and Bob, who don't trust each other, run a protocol at the end of which, they agree on a fair coin flip such that neither of them can cheat.

Commitment Protocol

A commitment protocol is a two-round protocol between Alice and Bob in which Alice first commits to a message and sends the commitment to Bob. Bob cannot get any information about the committed message from the commitment and Alice cannot later change the message that he first committed to. In the second round, Alice sends the decommitment to the message.

Secure two-party Computation

In this protocol Alice and Bob have inputs x_a and x_b and they wish to compute some function $f(x_a, x_b)$. At the end of the protocol, Alice (Bob) should learn nothing more than what can be computed given x_a (x_b) and $f(x_a, x_b)$.

2 Background in Complexity Theory

2.1 Basing Cryptography on Complexity Theory

Modern cryptography is based on assumptions on complexity theory. One of these assumptions is the existence of a one way function. Informally a one way function is a function that is easy to compute on an input, but given the output it is hard to come up with an inverse. Computation is computationally “easy” if it takes polynomial time in the size of the input. Computation is computationally “hard” if it takes super-polynomial time in the size of the input.

It is easy to see that if $P=NP$ then no one way function can exist. Thus we must make the assumption that $P \neq NP$. While there is no proof that these complexity assumptions are true many researchers believe them to be true.

Once we have made some basic complexity assumptions we can then build protocols on these assumptions by showing that if the protocol is broken we can break the underlining complexity assumption. In other words if there is a program that can break the protocol we can use this program as a subroutine for breaking the underlining complexity assumption. Thus the protocol is at least as strong as the underlining complexity assumption.

For example one complexity assumption is that it is hard to factor the product of two large primes. We can show that if we assume this to be hard $\Rightarrow \exists$ A One-way Function $\Rightarrow \exists$ A pseudo-random generator $\Rightarrow \exists$ A secure commitment scheme $\Rightarrow \exists$ A coin-flipping protocol.

Thus if there exists an adversary that can break the coin flipping protocol we can use this adversary to factor large primes which would violate the complexity assumption.

2.2 Complexity Classes

A **language** is a subset of a universal set of alphabet, that adheres to some rules. An **algorithm** is a procedure that decides if a given input is in the language.

The class of languages **P** is the set of all languages that can be recognized in deterministic polynomial time.

The class of languages **NP** is the set of all languages that can be recognized in non-deterministic polynomial time.

A **Probabilistic Turing Machine** is a Turing Machine which has an additional input tape that is a special read-once tape that contains randomness. A Probabilistic Turing Machine may give its output with certain error probabilities. Thus we can define complexity classes based on these error probabilities.

A function $g(n) : N \rightarrow N$ is **negligible** if $\forall c, \exists N_c$ such that $\forall n > N_c, g(n) < \frac{1}{n^c}$. Otherwise, $g(n)$ is said to be **non-negligible**.

Randomized Polytime (RP) is the set of all languages that can be recognized in polynomial time with the following properties:

$$\begin{aligned} Pr_w[M_w(x) = Yes|x \in L] &> \frac{1}{h(n)} \\ Pr_w[M_w(x) = Yes|x \notin L] &= 0 \end{aligned}$$

where w is the coin tosses of Turing machine M . The probability is taken over all coin tosses of M . $h(n)$ is any function of n such that $h(n)$ is non-negligible. Note that by running M polynomial many times using fresh randomness on every run, we can boost the probability of success to $1 - \epsilon(n)$ where $\epsilon(n)$ is negligible. If M outputs “No” on any input, we can conclude that $x \notin L$. If M outputs “Yes” on all runs, we can conclude that $x \in L$ with probability $1 - \epsilon(n)$.

Co-RP is the set of all languages that can be recognized in polynomial time with the following properties:

$$\begin{aligned} Pr_w[M_w(x) = Yes|x \in L] &= 1 \\ Pr_w[M_w(x) = Yes|x \notin L] &< 1 - \frac{1}{h(n)} \end{aligned}$$

where $h(n)$ is any non-negligible function.

Bounded Probabilistic Polytime (BPP or PPT) is the set of all languages that can be recognized in polynomial time with the following properties:

$$\begin{aligned} Pr_w[M_w(x) = Yes|x \in L] &> 1/2 + h(n) \\ Pr_w[M_w(x) = Yes|x \notin L] &< 1/2 - g(n) \end{aligned}$$

where $h(n)$ and $g(n)$ are non-negligible functions.

We can reduce error probability in BPP by running M polynomial many times with fresh randomness each time. We can reduce the error probability to a negligible function. Below we show how to do so with $h(n) = g(n) = 1/6$.

The **Chernoff Bound** states that given n independent random variables X_1, X_2, \dots, X_n with identical probability distributions, if $X = \sum_{i=1}^n X_i$, then

$$Pr[X \geq (1 + \beta)E(X)] < e^{-\beta^2 E(X)/2}$$

where $E(X)$ is the expectation of X .

The new machine M' runs M, k times. Let $X_i = 1$ if M' makes a mistake on the i^{th} run of machine M . Since each run of M is independent, $Pr[X_i = 1] = (1 - 2/3) = 1/3$. If we run the machine k times, then $E(X) = k/3$. M' 's output will be wrong if more than half the outcomes of M are incorrect. Thus M' 's output will be incorrect if $\sum_{i=1}^n X_i \geq k/2$. Thus setting β to $1/2$ in the Chernoff bound gives us an upper bound on this probability. Thus we get that

$$Pr[X \geq \frac{3}{2} \cdot \frac{k}{3}] < e^{-k/24}$$

Notice that the probability of M' making a mistake can be made negligible.

It is the case that $P \subseteq RP \subseteq NP$. Clearly any problem in P can be solved in RP by simply ignoring all randomness. Any problem in RP can be solved in NP by guessing the randomness.

3 Comparison of Uniform and Non-Uniform Complexity Classes

Uniform algorithms are those that use Turing Machines to make their decision. The class of problems that can be solved by these algorithms is known as Uniform Complexity class.

Non-Uniform algorithms are a family of polynomial size circuits one circuit for each input size. They can also be thought of as Turing machines which get a single polynomial length advise string for all inputs of the same length. These definitions can be shown to be equivalent. The class of problems that can be solved by these algorithms is known as Non-Uniform Complexity class denoted by **P / poly**.

3.1 P/Poly \subseteq NP?

The following is a wrong proof for the claim that $P/Poly \subseteq NP$.

Construct an NP machine that guesses the advice that was given to a P/Poly machine. Now, any problem in P / Poly is in NP. The problem with this proof is that although the NP machine uses the advice to make a decision, it does not know if the advice is correct for every instance of size $|x|$.

To give an idea of the power of P/poly we will in fact show that it can recognize undecidable languages. Consider some standard enumeration of the Turing machines M_1, M_2, \dots, M_n . Now consider the following language L: $x \in L$ iff machine number $|x|$ halts (Machine number in Godel numbering). Since we are allowed magic advice for each input length, the magic advice could tell us which machines halt and which do not in the above enumeration. Hence a P/Poly algorithm can solve the above problem which is undecidable.

3.2 BPP \subseteq P/Poly

P/poly circuits are so powerful that they do not require coin flips to make decisions. In fact, Leonard Adleman proved that $BPP \subseteq P/Poly$.

Theorem 1 *BPP \subseteq P/Poly. [Adleman]*

Proof Given a BPP machine M with $2/3$ probability of falsely rejecting an input and $1/3$ probability of falsely accepting an input, we can easily create another BPP machine M' that runs M many times on the input (with new randomness every time) to reduce the probability of error to a very small value [see previous subsection]. Let this new machine M' on input x be characterized in the following way:

$$Pr_w[M'(x) = yes|x \in L] > (1 - 2^{-(n+1)})$$

$$Pr_w[M'(x) = yes|x \notin L] < 2^{-(n+1)}$$

where $n = |x|$. Let r be the number of coin flips used by the machine.

Now let us construct a matrix with 2^n rows and 2^r columns, where the rows are all possible inputs and columns are all possible coin-flip sequences. This matrix is the space of all possible inputs (including coin flips) to the machine M' . Put in a cell of the matrix a value 1 when there is an error corresponding to the input and the coin flips of the cell's position, and a value 0 when there is no error (see Figure 1).

Now, the probability of a 1 in a cell, is the probability that machine M' makes a mistake $< 2^{-(n+1)}$. Hence, the expected number of ones in the matrix $< 2^n \cdot 2^r \cdot 2^{-(n+1)} = 2^{r-1}$. The number of columns in the matrix = 2^r . By pigeonhole principle, at least half the columns in the matrix will have no ones in them. In other words, at least half the coin flip sequences will give the correct result on all 2^n inputs. Choose one of these coin flip sequences and




Figure 1: Input space for a BPP machine M' with r coin flips and input string of size n

hardwire the P/poly circuit with this set of coin flips. By doing this, the P/poly circuit will be able to solve any problem in BPP without making use of any randomness. ■

4 Introduction to One-Way Functions

4.1 Overview

In secure cryptosystems the user should have the ability to encrypt its data in probabilistic polynomial time; but the adversary should fail to decrypt the encrypt data in probabilistic polynomial time.

As a motivating example, suppose we wish to construct a cryptosystem consisting of an encryption E and a decryption D , both which are poly-time. The encryption takes a clear-text message x and some random bits r , and gives $y = E(x, r)$. A polynomial-time adversary has access only to the cipher-text y . For the cryptosystem to be secure, it should be hard for the adversary to recover the clear-text, i.e. a poly-time adversary who is given $E(x, r)$




Figure 2: An illustration for a secure cryptosystem

should not be able to figure out any information about x .




Figure 3: An example for a cryptosystem

This brings up two questions: what assumptions do we need to design such a cryptosystem, and what is meant by security of the cryptosystem? In this lecture we will answer only the first question.

4.2 Necessary assumptions

P and NP

For 1-way functions to exist, our first assumption must be that $P \neq NP$. If we allow the adversary to use non-deterministic polynomial time machines, then A can always break our encryption. If $P \neq NP$, then we know that there is some function f such that $f^{-1}(y)$ is hard to compute by a poly-time adversary A on some number of instances of y . All we can say about the number of those hard instances is that they are infinitely many. If there were only finitely many hard instances, then one can create a look-up table for those hard pairs, and thus get a polynomial time algorithm that easily inverts f in all instances.

BPP and NP

Next, we want to assume that the adversary A is as ‘smart’ as possible. Since we are allowing our machine to flip coins, we want to assume that A can do that as well, and that A can occasionally make mistakes. Thus, let A be any BPP (bounded probabilistic polynomial time) machine. Our next assumption becomes $\text{BPP} \neq \text{NP}$. Since $\text{P} \subseteq \text{NP}$, this assumption is at least as strong as the assumption $\text{P} \neq \text{NP}$. Now we are guaranteed that there is some f such that a BPP adversary fails to invert f on some infinite number of instances. By fails to invert we mean that the probability that A finds some z so that $f(z) = y$ is very small. We make this notion precise in the next section.

Ave-P and Ave-NP

It is important to note that the assumption $\text{BPP} \neq \text{NP}$, while necessary, is not sufficient to guarantee the existence of 1-way functions. Even though we know that infinitely many hard instances exist, we may not have a way of actually finding them. Consider the function f defined below, that illustrates the well-known NP-complete problem of determining whether a given graph G has a Hamiltonian cycle.

$$f(G, H) = \begin{cases} G & \text{if } H \text{ is a hamiltonian cycle of } G, \\ \underbrace{00 \dots 0}_{|G|} & \text{otherwise.} \end{cases}$$

If, $\text{BPP} \neq \text{NP}$, there may be infinitely many pairs (G, H) for which f is hard to invert, but no poly-time algorithm is known that can generate them.

For a 1-way function to exist, we want hard instances to be easy to find. One way is to ask that ‘most’ instances are hard. Maybe, for a sufficiently large input size $|x|$, it is hard to invert f for most $f(x)$ ². Is assuming ave-P \neq ave-NP enough to guarantee the existence of 1-way functions? It is not clear, since it could be the case that whenever we pick y at random, and try to find $f^{-1}(y)$ it is hard, but whenever we pick x at random and ask our enemy to invert $f^{-1}(f(x))$ it is easy. The reason for this is that the distribution of $f(x)$, if we start from the uniform distribution on x , maybe far from uniform. Thus, we need to assume not only that ave-P \neq ave-NP, but also that there are functions which are hard to invert on the uniform distribution of the inputs.

Still, we are only talking about the existence of hard ‘unsolved’ problems. Given an output y , we want $f^{-1}(y)$ to be hard to compute. For a 1-way function to exist, however, we need hard ‘solved’ problems. We want an output y for which $f^{-1}(y)$ is hard for the adversary to compute, but we know an answer x such that $f(x) = y$.

In summary, if 1-way functions exist, than it must be that $\text{P} \neq \text{NP}$, $\text{BPP} \neq \text{NP}$, and ave-P \neq ave-NP; i.e. our assumptions are necessary. It is not known whether they are sufficient.

²This idea was formalized by Levin as an average case analog of the P vs NP question

4.3 Negligible and noticeable functions

When we talk about 1-way functions f , we do not require that f is one-to-one. The same output y can be produced by more than one output x . That is, we can have $f(x) = f(x')$ while $x \neq x'$. We consider the adversary A successful in inverting $y = f(x)$ if A produces some x' such that $f(x') = y$. Formally, A inverts $f(x)$ if

$$A(f(x)) \subseteq f^{-1}(f(x)).$$

$A(f(x))$ is the value x' that A produces given $f(x)$, and $f^{-1}(f(x))$ are all those inputs z for which $f(z) = f(x)$. A is successful if it is able to produce one inverse of $f(x)$. Certainly, if f is one-to-one, then each $f(x)$ has a unique inverse and if A is successful in this case, then A was able to recover x . For simplicity, we sometimes write $\Pr_{x,w}[A \text{ inverts } f(x)]$ instead of $\Pr_{x,w}[A(f(x)) \subseteq f^{-1}(f(x))]$.

Next, we want to give a formal definition of what it means for A to fail. We recall negligible functions.

Definition 2 A function $\epsilon : \mathbb{N} \rightarrow \mathbb{R}$ is **negligible** if for all $c > 0$, there exists an integer N_c so that $\forall n \geq N_c$

$$\epsilon(n) < \frac{1}{n^c}.$$

A negligible function is a function that vanishes faster than the inverse of any polynomial. We say that A fails to invert $f(x)$ if

$$\Pr_{x,w}[A(f(x)) \subseteq f^{-1}(f(x))] < \epsilon(n)$$

for some negligible function $\epsilon(n)$, where $n = |x|$.

Here it is worth mentioning that if A has a negligible probability of success, then even if A attempts to invert f a polynomial number of times, its probability of success will not amplify but will remain negligible.

We defined negligible probability of success as occurring with probability smaller than any polynomial fraction. A polynomial probability of success makes a function noticeable (non-negligible).

Definition 3 A function $\nu : \mathbb{N} \rightarrow \mathbb{R}$ is **noticeable (non-negligible)** if there exists $c > 0$, and there exists an integer N_c so that $\forall n \geq N_c$

$$\nu(n) > \frac{1}{n^c}.$$

Noticeable and negligible functions are not perfect negations of each other. There are functions that are neither noticeable nor negligible. For example, the function $f : \mathbb{N} \rightarrow \mathbb{R}$ given by

$$f(n) = \begin{cases} n & \text{if } n \text{ is odd,} \\ \frac{1}{2^n} & \text{if } n \text{ is even} \end{cases}$$

is negligible on the even lengths and noticeable on the odd lengths, so overall, f is neither.

5 One-Way Functions

5.1 Informal Definition of One-way Function

The 1-way function problem can be described as a game between a Challenger C (a P-time machine) and an Adversary A (a BPP machine):

1. The Challenger chooses an input length n for a one-way function, which he hopes is “large enough”. He then picks x such that $|x| = n$ and computes $y = f(x)$, giving the result y to A .
2. A tries to compute $f^{-1}(y)$ during a polynomial amount of time in the length of $|f(x)|$, and sends its guess z back to the Challenger.
3. A wins if $f(x) = f(z)$, otherwise the Challenger wins. f is a 1-way function if the probability of all BPP adversaries to win is negligible, for a sufficiently big n .

Taking in account the above perspective, we can define a 1-way function f informally:

1. f can be computed in deterministic polynomial time.
2. f is “hard to invert” for all PPT adversaries A .
3. f has polynomially-related input/output.

5.2 Uniform and Non-Uniform One-way Functions

Definition 4 *A function f is said to be a uniform strong one-way function if the following conditions hold:*

1. f is polynomial-time computable.
2. f is hard to invert for a random input: $\forall c > 0 \forall A \in PPT \exists N_c$ such that $\forall n > N_c$:

$$\Pr_{\{x,w\}}[A \text{ inverts } f(x)] < \frac{1}{n^c}$$

where “PPT” stands for probabilistic polynomial time, $|x| = n$ and w are coin-flips of the probabilistic algorithm A .

3. I/O length of f is polynomially related: $\exists c_1, c_2$ such that $|x|^{c_1} < |f(x)| < |x|^{c_2}$.

Condition 3 is necessary to assure that both the Challenger and the Adversary do polynomially related work as a function of their input. Note that the Adversary still has two trivial ways of attacking $f(x)$:

- (1) A can always try to invert $f(x)$ by simply guessing what the inputs are and
- (2) A can use a huge table to store pairs $(x, f(x))$, sorted, say by the value of $f(x)$.

Neither (1) nor (2) are good strategies for attack since (1) is successful only with a negligible probability and (2) is avoided by requiring that A is of polynomial size.

A *non-uniform* 1-way function is defined exactly as above, except that the adversary is formulated not as a PPT machine, but as a family of poly-size circuits. Recall that a family of poly-size circuits is a set of circuits, one for each input length n , such that the size of each circuit is polynomially related (in size) to the length of the input.




Figure 4: Uniform and non-uniform 1-way functions

Definition 5 A function f is said to be a non-uniform strong one-way function if the following conditions hold:

1. f is polynomial-time computable.
2. f is hard to invert: $\forall c > 0 \forall$ non-uniform poly-size families A of circuits $\exists N_c$ such that $\forall n > N_c$:

$$Pr_x[A \text{ inverts } f(x)] < \frac{1}{n^c}$$

where $|x| = n$.

3. I/O length of f is polynomially related: $\exists \epsilon, c$ such that $|x|^\epsilon < |f(x)| < |x|^c$.

Note that the probability of A inverting $f(x)$ is taken only over all x and not over any coin flips. This is because, since A is a family of poly-size circuits, the optimum coin flip for A to invert $f(x)$ can be hardwired into A .

For simplicity, we say that f is a 1-way function when f is a uniform strong 1-way function. We now prove that if a function f is *non-uniform* then it is also *uniform*, hence we have the scenario in Figure 4.

Theorem 6 *If f is a non-uniform one-way function, then f is also a uniform one-way function.*

Proof We will prove the contrapositive, i.e., instead of proving $A \Rightarrow B$, we will prove $\neg B \Rightarrow \neg A$. Suppose that f is not a uniform one-way function. Then there exists a constant $c > 0$, and a PPT adversary A such that for an infinite number of integers n , for all strings x of length n ,

$$Pr_{\{x,w\}}[A \text{ inverts } f(x)] > \frac{1}{n^c}$$

where w are coin-flips of the adversary. Our objective is to find a poly-size collection of circuits A' to substitute our PPT adversary A . Let $\epsilon(n) = \frac{1}{n^c}$ and define the set $GOOD$ to be

$$GOOD = \{x | Pr_w[A \text{ inverts } f(x)] > \frac{\epsilon(n)}{2}\}.$$

By conditioning on whether $x \in GOOD$ we get

$$\begin{aligned} Pr_{x,w}[A \text{ inverts } f(x)] &= Pr_{x,w}[A \text{ inverts } f(x) | x \in GOOD] \cdot Pr_x[x \in GOOD] \\ &\quad + Pr_{x,w}[A \text{ inverts } f(x) | x \notin GOOD] \cdot Pr_x[x \notin GOOD] \end{aligned} \quad (1)$$

Hence,

$$\begin{aligned} Pr_x[x \in GOOD] &= \frac{Pr_{x,w}[A \text{ inverts } f(x)] - Pr_{x,w}[A \text{ inverts } f(x) | x \notin GOOD] \cdot Pr_x[x \notin GOOD]}{Pr_{x,w}[A \text{ inverts } f(x) | x \in GOOD]} \\ &> Pr_{x,w}[A \text{ inverts } f(x)] - Pr_{x,w}[A \text{ inverts } f(x) | x \notin GOOD] \cdot Pr_x[x \notin GOOD] \\ &> Pr_{x,w}[A \text{ inverts } f(x)] - Pr_{x,w}[A \text{ inverts } f(x) | x \notin GOOD] \\ &\geq \epsilon(n) - \frac{\epsilon(n)}{2} \\ &= \frac{\epsilon(n)}{2} \end{aligned} \quad (2)$$

First, using many attempts of the adversary to invert using fresh coin-flips each time, we can amplify $\Pr_w[A \text{ inverts } f(x) | x \in \text{GOOD}]$. Then, using the same technique as in the proof of $BPP \subseteq P/poly$, it is possible to show that there are sequences of coin flips r such that A correctly inverts all elements of GOOD on r . Therefore we can hardwire r and build a circuit family A' which inverts at least $\frac{\epsilon(n)}{2}$ of all strings x . Therefore f is not a non-uniform one-way function. ■

6 Number Theory

6.1 Modular Arithmetic

Informally, we can think of modular arithmetic as arithmetic as usual over the integers, except that if we are working modulo n , then every result x is replaced by the element of $0, 1, \dots, n-1$ that is equivalent to x , modulo n (that is, x is replaced by $x \bmod n$). This informal model is sufficient if we stick to the operations of addition, subtraction, and multiplication. A more formal model for modular arithmetic, which we now give, is best described within the framework of group theory.

Definition 7 A group (S, \oplus) is a set S together with a binary operation \oplus defined on S for which the following properties hold:

1. Closure: For all $a, b \in S$, we have $a \oplus b \in S$.
2. Identity: There is an element $e \in S$, called the identity of the group, such that $e \oplus a = a \oplus e = a$ for all $a \in S$.
3. Associativity: For all $a, b, c \in S$, we have $(a \oplus b) \oplus c = a \oplus (b \oplus c)$.
4. Inverses: For each $a \in S$, there exists a unique element $b \in S$, called the inverse of a , such that $a \oplus b = b \oplus a = e$.

As an example, consider the familiar group $(\mathbb{Z}, +)$ of the integers \mathbb{Z} under the operation of addition: 0 is the identity, and the inverse of a is $-a$. If a group (S, \oplus) satisfies the commutative law $a \oplus b = b \oplus a$ for all $a, b \in S$, then it is an *abelian group*. If a group (S, \oplus) satisfies $|S| < \infty$, then it is a *finite group*.

We give two small facts about finite groups.

Lemma 8 For any finite group (G, \cdot) , $g^m = 1$ for any nonzero $g \in G$ and $m = |G|$.

Lemma 9 For any finite group (G, \cdot) , $g^x = g^{x \bmod m}$ for any nonzero $g \in G$, $m = |G|$, and $x \in \mathbb{Z}$.

Proof Let $x = x' \bmod m \Rightarrow x = km + x' \Rightarrow g^x = g^{km+x'} \Rightarrow g^x = g^{km} g^{x'} = 1^k g^{x'} = g^{x'}$.
 ■

6.2 The multiplicative group Z_N^*

For any positive integer n , let Z_n stand for the set $\{0, 1, 2, \dots, n-1\}$ of n elements. Define

$$Z_N^* = \{x | 1 \leq x \leq N \text{ and } \gcd(x, N) = 1\}$$

i.e., Z_N^* contains all positive integers less than and relatively prime to N . Z_N^* is a group under multiplication modulo N . The function $\varphi(N) : \mathbf{Z} \rightarrow \mathbf{Z}$ defined by

$$\varphi(N) = |Z_N^*|$$

is the so called Euler phi function.

When N is a prime number, say p ,

$$Z_p^* = \{1, 2, \dots, p-1\}$$

and $\varphi(p) = |Z_p^*| = p-1$.

We are mostly interested in those integers N that are the product of two distinct primes p and q . For $N = pq$,

$$\varphi(N) = \varphi(pq) = |Z_N^*| = (p-1)(q-1).$$

6.3 The Chinese Remainder Theorem for the case $N = pq$

Let $N = pq$ where p and q are distinct prime numbers. The Chinese Remainder Theorem allows us to understand the structure of Z_N^* by considering the 'easier' to work with Z_p^* and Z_q^* .

Theorem 10 Let $N = pq$ where p and q are distinct prime numbers. Then, the map from Z_N^* to $Z_p^* \times Z_q^*$ given by

$$x \mapsto (x \pmod{p}, x \pmod{q})$$

is one-to-one and onto.

By the above theorem, every x in Z_N^* can be written as (x_p, x_q) where $x_p \in Z_p^*$ and $x_q \in Z_q^*$. Conversely, for every element (a, b) in $Z_p^* \times Z_q^*$, there is a unique x in Z_N^* so that $x \equiv a \pmod{p}$ and $x \equiv b \pmod{q}$.

Consider, for example, $N = 10$, $p = 2$, and $q = 5$. Then, $|Z_N^*| = (5 - 1)(2 - 1) = 4$ and, in particular, $Z_N^* = \{1, 3, 7, 9\}$ while $Z_2^* = \{1\}$ and $Z_5^* = \{1, 2, 3, 4\}$. The bijection is

$$\begin{aligned} 1 &\longmapsto (1, 1) \\ 3 &\longmapsto (1, 3) \\ 7 &\longmapsto (1, 2) \\ 9 &\longmapsto (1, 4). \end{aligned}$$

In fact, if we know the factorization of N we can simplify computations that have to be performed modulo N into computations modulo p and modulo q . Say that we are trying to multiply two elements x and y of Z_N^* and that p and q are two k -bit primes. Rather than first computing xy and reducing it modulo the $2k$ -bit number N , we can instead multiply the corresponding (x_p, x_q) and (y_p, y_q) and thus perform two multiplications modulo k -bit numbers.

Clearly, it is easy to find (x_p, x_q) if we are given x by computing x modulo p and x modulo q . For the above simplification to work, we should also be able to convert back using a polynomial time algorithm. To do so, we first find integers $s, t < N$ so that

- 1) $s \equiv 1 \pmod{p}$ and $s \equiv 0 \pmod{q}$ and
- 2) $t \equiv 0 \pmod{p}$ and $t \equiv 1 \pmod{q}$

i.e., we can informally think of s as $(1, 0)$ and of t as $(0, 1)$ ³. Then, given any (a, b) in $Z_p^* \times Z_q^*$ compute $x = as + bt \pmod{N}$. The value x is the unique element of Z_N^* that gets mapped to (a, b) .

For example, consider Z_{10}^* again. We get $s = 5$ since $5 \equiv 1 \pmod{2}$ and $5 \equiv 0 \pmod{5}$, and $t = 6$ since $6 \equiv 0 \pmod{2}$ and $6 \equiv 1 \pmod{5}$. To convert, say, $(1, 4)$ from an element of $Z_2^* \times Z_5^*$ into an element of Z_{10}^* , we compute

$$1 \cdot s + 4 \cdot t = 1 \cdot 5 + 4 \cdot 6 \pmod{10} = 9.$$

6.4 Quadratic Residues and quadratic non-residues

We call an element $a \in Z_N^*$ a *quadratic residue (QR) modulo N* if there exists an $x \in Z_N^*$ such that $x^2 \equiv a \pmod{N}$. Informally, we refer to a as a square in Z_N^* and we call x its

³Computing s and t is done in polynomial time by using the generalized Euclidean algorithm.

square root. If a is not a square in Z_N^* , we call a a *quadratic non-residue modulo N* . We let QR_N denote the set of all quadratic residues in Z_N^* . For example, in Z_{10}^* , $QR_N = \{1, 9\}$.

In Z_p^* , where p is an odd prime, exactly half of the elements of Z_p^* are quadratic residues. This fact follows from the following lemma.

Lemma 11 *If p is an odd prime and $a \in Z_p^*$, then a has either 0 square roots or 2 distinct square roots in Z_p^* .*

Proof Take any $a \in Z_p^*$. If a is a quadratic non-residue modulo p , then a has no square roots in Z_p^* and we are done. Otherwise, there is some $x \in Z_p^*$ so that $x^2 \equiv a \pmod{p}$. Then, $p-x$ is also in Z_p^* and $(p-x)^2 = p^2 - 2px + x^2$ so that $(p-x)^2 \equiv a \pmod{p}$. Thus, both x and $p-x$ are square roots of a and they are distinct since $x = p-x$ contradicts the fact that p is odd.

Now, if y is yet another square root of a , then $x^2 \equiv y^2 \pmod{p}$ and $p|x^2 - y^2 = (x-y)(x+y)$. Since p is a prime, this implies that either $p|x-y$ or $p|x+y$. That is, $y \equiv x \pmod{p}$ or $y \equiv -x \pmod{p}$ leading to $y = x$ or $y = p-x$. ■

We turn again to the case $N = pq$ where p and q are distinct primes. We require, in addition, that both p and q are odd. It turns out that exactly $\frac{1}{4}$ of the elements of Z_N^* are quadratic residues. Note that this is not true for Z_{10}^* where $|QR_{10}| = 2$ while $|Z_{10}^*| = 4$.

Lemma 12 *If $N = pq$ where p and q are distinct odd primes, and if $a \in Z_N^*$, then a has either 0 square roots or 4 distinct square roots in Z_N^* .*

Proof Again, if a has no square roots, there is nothing to prove. Assume that there exists some $x \in Z_N^*$ so that

$$x^2 \equiv a \pmod{N}. \quad (3)$$

By the Chinese Remainder Theorem, we can write x as (x_p, x_q) and a as (a_p, a_q) in $Z_p^* \times Z_q^*$. But then (3) implies that

$$x_p^2 \equiv a_p \pmod{p} \text{ and } x_q^2 \equiv a_q \pmod{q}$$

i.e., x_p is a square root of a_p in Z_p^* and x_q is a square root of a_q in Z_q^* . By Lemma 11, the only possibilities for x are (x_p, x_q) , $(p-x_p, x_q)$, $(x_p, q-x_q)$, and $(p-x_p, q-x_q)$. All four of those are distinct since both p and q are odd. Similar argument using Lemma 11 shows that those are the only square roots of a . Thus, a has exactly 4 distinct square roots in Z_N^* . ■

⁴We are using the fact that if p is a prime number that divides the product ab , then either p divides a or p divides b .

It is known that computing square roots in Z_p^* can be done in polynomial-time. If we are given the factorization of N as pq , then using the bijection given by the Chinese Remainder Theorem we will see that we can also compute square roots in Z_N^* in polynomial-time. We will show, however, that when the factorization of N is not known, then it is as ‘hard’ to compute square roots modulo N as it is to factor N .

6.5 The Legendre symbol and the Jacobi symbol

Quadratic residues are important enough to prompt the definition of further notation that allows dealing with them. For any prime p , the *Legendre symbol*, $L_p(y)$ is defined to be

$$L_p(y) = \begin{cases} 1 & \text{if } y \text{ is a quadratic residue modulo } p, \\ -1 & \text{otherwise.} \end{cases}$$

For any $N = pq$, where p and q are distinct primes, the *Jacobi symbol*, $J_N(y)$ is defined to be

$$J_N(y) = L_p(y)L_q(y).$$

The Jacobi symbol provides a generalization of the Legendre symbol and can further be defined for any integer. Note that it is not true that $J_N(y) = 1$ implies that y is a quadratic residue modulo N . It could be that $L_p(y) = L_q(y) = -1$ and therefore y is not a quadratic residue modulo N .

We can compute $L_p(y)$ in polynomial-time. If $N = pq$ we can also compute $J_N(y)$ in polynomial-time even if we are given N but not p and q . Even if we have computed that $J_N(y) = 1$, however, no polynomial-time algorithm is known that can determine whether y is a quadratic residue modulo N .

7 The Rabin candidate for a 1-way function

Based on the Number theory background we introduced so far, we can consider the function $f_N : Z_N^* \rightarrow QR_N$ given by

$$f_N(x) \equiv x^2 \pmod{N}$$

where $N = pq$ as before. Note that f_N is not one-to-one but, in fact, is 4-to-1 as shown by Lemma 12.

In 1979, Michael Rabin was the first to show that f_N would be a 1-way function if factorization is ‘hard’. In order to prove this, we first show a small fact.

Lemma 13 Let $N = pq$ where p and q are distinct odd primes. If $x, y \in Z_N^*$ are such that $x \neq \pm y$ and

$$x^2 \equiv y^2 \pmod{N}$$

then given x, y , and N , we can efficiently determine p and q ; i.e. factor N .

Proof Since $x^2 \equiv y^2 \pmod{N}$, $N | x^2 - y^2 = (x - y)(x + y)$. On the other hand, $x \neq \pm y$ implies that $x - y \not\equiv 0 \pmod{N}$ and $x + y \not\equiv 0 \pmod{N}$. The prime $p | (x - y)(x + y)$ so it must be that $p | x - y$ or $p | x + y$ while N does not divide $x - y$ nor $x + y$. Thus, $\gcd(x - y, N) = p$ or $\gcd(x + y, N) = p$ (It is known that the gcd of two numbers can be computed in poly time). We are done since we were able to find a factor of N . ■

Theorem 14 (Rabin, 1979) Let $N = pq$ where p and q are distinct odd primes. The function $f_N(x) \equiv x^2 \pmod{N}$ is a 1-way function if and only if factoring N cannot be done in polynomial time.

Proof (can factor \Rightarrow can invert) If N can be factored efficiently, given an output $y \in QR_N$, compute p and q , and then, find the representation (y_p, y_q) in $Z_p^* \times Z_q^*$. Also using a poly-time algorithm we can then find a square root z_p of y_p in Z_p^* and z_q of y_q in Z_q^* . We have produced a square root (z_p, z_q) of (y_p, y_q) in $Z_p^* \times Z_q^*$. Converting (z_p, z_q) back to some $z \in Z_N^*$, we get $f_N(z) = y$. We were able to describe a polynomial time algorithm that inverts f_N which contradicts our assumption that f_N is a 1-way function.

(can invert \Rightarrow can factor) The converse statement holds the essence of the theorem and requires more work. Assume that F_N is not a 1-way function. Formally, this means that there exists some constant $c > 0$ and some PPT adversary A such that for any integer M , there is some input length $n \geq M$ such that

$$\Pr_{x,w}[A \text{ inverts } f(x)] > \frac{1}{n^c} \tag{4}$$

for inputs x of length n . Let $\epsilon(n) = \frac{1}{n^c}$. We will use A to find another PPT algorithm A' for which

$$\Pr_{N,w'}[A' \text{ factors } N] > \frac{\epsilon(n)}{2}.$$

The algorithm A' can be described as follows:

- (1) Given N , choose some $y \in Z_N^*$, compute $z = y^2 \pmod{N}$, and give z and N to A .
- (2) Take the output $x = A(z, N)$ produced by A and check whether $x^2 = z$ and whether $x \neq \pm y$.

(3) If both of the above are true, use Lemma 13 to factor N . Otherwise, give up.

Certainly, A' runs in polynomial time but we need to consider the probability of success for A' .

$$\begin{aligned} Pr_{N,w}[A' \text{ factors } N] &= Pr_{x,w}[A \text{ inverts } f_N] \cdot Pr_x[x \neq \pm y | A \text{ inverts } f_N] \\ &> \epsilon(n) \cdot \frac{1}{2} = \frac{\epsilon(n)}{2} \end{aligned}$$

This follows from (4) and from the fact that z has two square roots $\pm y$ and two other square roots $\pm y'$. So, if A gives a square root x of z , then with probability $\frac{1}{2}$ $x \neq \pm y$. Thus, assuming that f_N is not a 1-way function we were able to show that N can then be factored in polynomial time which is a contradiction and we are done. ■

8 Weak One-Way Functions

A one-way function, also called a strong one-way function, is a function that one cannot invert successfully in polynomial time except with negligible probability. A weak one-way function is a function that one cannot invert successfully in polynomial time with noticeable probability.

A motivating example for weak one way function: The problem of factoring $N = pq$ when p , and q are very big prime numbers (about the same number of bits) is widely believed to be a hard problem. What if we define a function $f(x, y) = x \cdot y$ where x and y are big (k bits) random integers. Is f a 1-way function?

Let A be the following poly-time algorithm:

- (1) A receives z and checks if $\frac{z}{2}$ is an integer.
- (2) If it is, A outputs $(2, \frac{z}{2})$. Otherwise, A gives up.

Since each of x and y are even with probability a half, then the probability that z is even is $\frac{3}{4}$. With certainty, f is not a 1-way function.

What we really refer to when saying that factoring is hard is that the function f above is hard to invert, by density of primes, on a particular part of its domain. The probability for a k -bit integer to be a prime number is $\frac{1}{k}$, making the probability that $f(x, y)$ is a product of two primes $\frac{1}{k^2}$. In this case, it is believed hard to invert f with probability greater than $\frac{1}{n^2}$.

Definition 15 Weak One-Way Functions

f is a **weak one-way function** if:

1. f is polynomial-time computable.
2. $\exists c > 0 \forall$ probabilistic polynomial-time $A, \exists N_c$ such that $\forall n > N_c$

$$Pr_{w,x}[A_w(f(x)) \notin f^{-1}(f(x))] > \frac{1}{n^c} = \epsilon(n)$$

where $|x| = n$, w are coin-flips of A , and $A_w(f(x)) \notin f^{-1}(f(x))$ means "A does not invert $f(x)$ ".

3. I/O size is polynomially related.

Let's now prove the main result about weak and strong 1-way functions [Yao]:

Theorem 16 *There exists a weak one-way function if and only if there exists a strong one-way function.*

Proof First, let us show a trivial direction, i.e., the existence of a strong 1-way function implies that of a weak 1-way function: condition 2 of a strong 1-way function can be re-written as follows: $\forall c > 0 \forall A \in PPT \exists N_c$ s.t. $\forall n > N_c$:

$$Pr_{\{x,w\}}[A \text{ does not invert } f(x)] > 1 - \frac{1}{n^c} > \frac{1}{n^c}$$

which implies condition 2. of a weak 1-way function.

We now prove the converse: given a weak 1-way function f_0 , we will construct a strong 1-way function f_1 . We will demonstrate that f_1 is a strong 1-way function by contradiction: we assume an adversary A_1 for f_1 and then demonstrate an effective adversary A_0 for f_0 . We can assume that f_0 is length-preserving and maps m bits to m bits. Condition 2. of a weak one-way function f_0 can be restated as:

- $\exists \epsilon_{f_0} > 0 \forall A_0 \in PPT \exists M_{\epsilon_{f_0}}$ s.t. $\forall m > M_{\epsilon_{f_0}}$:

$$Pr_{\{x,w\}}[A_0 \text{ inverts } f_0(x)] \leq 1 - \frac{1}{m^{\epsilon_{f_0}}} = 1 - \epsilon_0(m)$$

where $|x| = m$; w are coin-flips of A ; and $\epsilon_0(m) \triangleq \frac{1}{m^{\epsilon_{f_0}}}$.

To construct f_1 , we amplify the "hardness" of weak 1-way function f_0 by applying f_0 in parallel $q \triangleq \frac{2m}{\epsilon_0(m)}$ times:

$$f_1(x_1, \dots, x_q) \triangleq f_0(x_1), \dots, f_0(x_q).$$

where each x_i , $1 \leq i \leq q$ is a uniformly and independently chosen m -bit input to f_0 . Notice that our f_1 maps $n = \frac{2m^2}{\epsilon_0(m)}$ bits to n bits. We claim that f_1 is a strong 1-way function. The proof is by contradiction. Suppose f_1 is not a strong 1-way function. Then $\exists A_1, \exists c$ s.t. for infinitely many inputs of length n ,

$$Pr_{\{\vec{x}, w\}}[A_1 \text{ inverts } f_1(\vec{x})] > \frac{1}{n^c} \triangleq \epsilon_1(n) \triangleq \epsilon_2(m)$$

Notice that we can redefine $\epsilon_1(n)$ in terms of $\epsilon_2(m)$ since m and n are polynomially related. If we can show how to construct A_0 (using above A_1 as a subroutine) such that A_0 will invert f_0 with probability (over x and w) greater than $1 - \epsilon_0(m)$ we will achieve a contradiction with weak one-wayness of f_0 . Our algorithm $A_0(f_0(x))$ is as follows:

Algorithm $A_0(y)$:
 repeat procedure $Q(y)$ at most $\frac{4m^2}{\epsilon_2(m)\epsilon_0(m)}$ times;
 stop whenever $Q(y)$ succeeds and output $f_0^{-1}(y)$,
 otherwise output “fail to invert”.

Procedure $Q(y)$:
 for i from 1 to $q = \frac{2m}{\epsilon_0(m)}$ do:
 STEP1: pick $x_0, \dots, x_{i-1}, x_{i+1}, \dots, x_q$
 (where each x_j is independently chosen m -bit number)
 STEP2: call $A_1(f_0(x_0), \dots, f_0(x_{i-1}), y, f_0(x_{i+1}), \dots, f_0(x_q))$
 (procedure $Q(y)$ succeeds if A_1 above inverts)

We must estimate the success probability of $A_0(f_0(x))$, where the probability is over x and coin-flips w of A_0 . Define x (of length m) to be BAD if

$$Pr_w[Q(f(x)) \text{ succeeds}] < \frac{\epsilon_2(m)\epsilon_0(m)}{4m}$$

We claim that:

$$Pr_x[x \text{ is BAD}] < \frac{\epsilon_0(m)}{2}$$

To show this we assume (towards the contradiction) that $Pr_x[x \text{ is BAD}] \geq \frac{\epsilon_0(m)}{2}$. Then

$$\begin{aligned}
Pr_{\{\vec{x},w\}}[A_1 \text{ inverts } f_1(\vec{x})] &= Pr_{\{\vec{x},w\}}[A_1 \text{ inverts } f_1(\vec{x}) | \text{some } x_i \in BAD] \cdot Pr_{\vec{x}}[\text{some } x_i \in BAD] \\
&+ Pr_{\{\vec{x},w\}}[A_1 \text{ inverts } f_1(\vec{x}) | \forall i, x_i \notin BAD] \cdot Pr_{\vec{x}}[\forall i, x_i \notin BAD] \\
&\leq \sum_{i=1}^{\frac{\epsilon_0(m)}{2m}} [Pr_{\{\vec{x},w\}}[A_1 \text{ inverts } f_1(\vec{x}) | x_i \in BAD]] \cdot Pr_{\vec{x}}[\forall i, x_i \notin BAD] \\
&+ Pr_{\{\vec{x},w\}}[A_1 \text{ inverts } f_1(\vec{x}) | \forall i, x_i \notin BAD] \cdot Pr_{\vec{x}}[\forall i, x_i \notin BAD] \\
&\leq \frac{2m}{\epsilon_0(m)} \left(\frac{\epsilon_2(m)\epsilon_0(m)}{4m} \right) \cdot 1 + 1 \cdot \left(1 - \frac{\epsilon_0(m)}{2} \right)^{\frac{2m}{\epsilon_0(m)}} \\
&\leq \frac{\epsilon_2(m)}{2} + e^{-m} \\
&< \epsilon_2(m)
\end{aligned}$$

But we assumed that $Pr_{\{\vec{x},w\}}[A_1 \text{ inverts } f_1(\vec{x})] \geq \epsilon_2(m)$ a contradiction. Hence we have shown that $Pr_x[x \text{ is } BAD] < \frac{\epsilon_0(m)}{2}$. We are now ready to estimate the failure probability of A_0 , using the fact that we try procedure Q in case of failure a total of $\frac{4m^2}{\epsilon_2(m)\epsilon_0(m)}$ times:

$$\begin{aligned}
Pr_{\{x,w\}}[A_0 \text{ does not invert } f_0(x)] &= Pr_{\{x,w\}}[A_0 \text{ does not invert } f_0(x) | x \in BAD] \cdot Pr_x[x \in BAD] \\
&+ Pr_{\{x,w\}}[A_0 \text{ does not invert } f_0(x) | x \notin BAD] \cdot Pr_x[x \notin BAD] \\
&\leq 1 \cdot \frac{\epsilon_0(m)}{2} + \left(1 - \frac{\epsilon_2(m)\epsilon_0(m)}{4m} \right)^{\frac{4m^2}{\epsilon_2(m)\epsilon_0(m)}} \cdot 1 \\
&\leq \frac{\epsilon_0(m)}{2} + e^{-m} \\
&< \epsilon_0(m)
\end{aligned}$$

Thus $Pr_{x,w}[A_0 \text{ inverts } f_0(x)] > 1 - \epsilon_0(m)$ contradicting the assumption that f_0 is a weak one-way function.

■

Lecture 2

9 Hard-Core Bits

9.1 Introduction to Hard-Core Bits

In this lecture⁵ our goal is to discuss Hard-Core Bits and draw conclusions from definitions presented along the way. Hard-Core Bits were defined by Blum and Micali in 1982. Informally, a Hard-Core Bit $B(\cdot)$ of a one-way function $f(\cdot)$ is a bit which is almost (i.e., polynomially) as hard to compute as it is to invert f . Blum and Micali showed that a particular number theoretic function (which is believed to be one-way) has a Hard-Core Bit. It was later shown that all (padded) one-way functions have a Hard-Core Bit. We conclude by presenting this proof (due to Goldreich and Levin 1989).

Motivating example (due to Manuel Blum): Consider the problem of gambling on the outcome of a random coin flip with an adversary over a telephone line. If you bet on heads and allow the adversary to flip the coin and then inform you of the outcome, he may cheat and say tails without even bothering to flip the coin. Now suppose that after losing quite a bit of money, you decide to play a more sophisticated game in which both you and the adversary select a random bit and you win if the XOR of the two bits is 1. Unfortunately, it is still unsafe to transmit your random bit in the clear to an un-trustworthy adversary, for your adversary can always cheat by claiming that it selected the same bit.

To keep from being swindled further, you decide on the following commitment protocol to play the game described above fairly. You begin by sending the adversary your bit in a locked safe, then the adversary sends you its bit in the clear, and finally you send the adversary the combination to the safe. Both of you then compute the XOR of the two bits certain that the other party had no unfair advantage playing the game. We use this analogy to motivate the idea that it may be possible to send a commitment of a secret bit to an adversary, without revealing any information as to the value of that bit. Our objective is to develop such a legitimate commitment protocol based on one-way functions.

Assume that we have a one-way function. One (unfair) strategy would be to commit to b by sending $b \oplus x_3$ with $f(x)$, where x_3 is the third bit of x . The flaw with this strategy is that the player can cheat, since $f(x)$ might not have unique inverses. In particular, suppose $f(x)$ has inverses x_1 and x_2 such that the third bit of x_1 and x_2 differ. Then once the adversary presents its random bit in the clear, the player can choose to transmit either x_1

⁵Scribe notes by V.Goyal, O.Pandey, Urwin, Karimi, Cheng; April 2006

or x_2 to the adversary, and clearly will choose to transmit the one which results in a payoff.

What if we assume much more, i.e. that we have a 1-1, length-preserving one-way function? The sender can no longer cheat in the manner described above, but the receiver may still be able to cheat. Just because $f(x)$ is hard to invert does not necessarily mean that any individual bit of $f(x)$ is hard to invert. As an example, suppose we have a one-way function $f(x)$ and another function $g(x) = g(b_1, b_2, b_3, x_4, x_5, \dots, x_n) = b_1 b_2 b_3 f(x_4, x_5, \dots, x_n)$. Now since f is one-way, g is also one-way, yet, given $g(x)$, the three highest-order bits of x are simple to compute.

9.2 Coin Flipping using Bit Commitment Protocol

Details

1. *Alice* flips r_0 and locks the result in a safe deposit box.
2. The locked safe deposit box with r_0 inside is given to *Bob*.
3. *Bob* in turn flips r_2 in the open and sends to result to *Alice*.
4. *Alice* then sends the deposit box combination for *Bob* to open the box containing the outcome of r_0 .
5. *Alice* and *Bob* then exclusive-or the two flips r_0 and r_2 (*ie*) $coin = r_0 \oplus r_2$.

This example illustrates the use of what is called a Bit Commitment (BC) Protocol. It is the idea where *Alice* commits a bit and sends it to *Bob* without revealing the value of the bit. There are two properties that we wish to have in a BC Protocol.

- (1) Given the 'box', *Bob* cannot predict what's 'in it' with probability $\geq \frac{1}{2} +$ a negligible amount.
- (2) After committing, *Alice* cannot change her mind about what is in the 'box'.

This exchange is known as the Commit Phase, where two events take place: hiding and binding. *Alice* places the outcome of r_0 in the safe deposit box is the hiding event. Hiding provides that *Bob* can not predict the outcome of r_0 with probability greater than $\frac{1}{2} + \epsilon$, where ϵ is negligible. Taking the safe deposit box and sending it to *Bob* is the binding event. Once *Alice* commits to the contents of the safe deposit box, she can not change her mind.

The Commit Phase is followed by the De-commit Phase in which *Alice* sends the deposit box combination to *Bob*.

Going further in depth we see that we need to somehow 'commit' the coin flip that *Alice* initially sends to *Bob* (i.e., find an electronic equivalent of the deposit box). As discussed

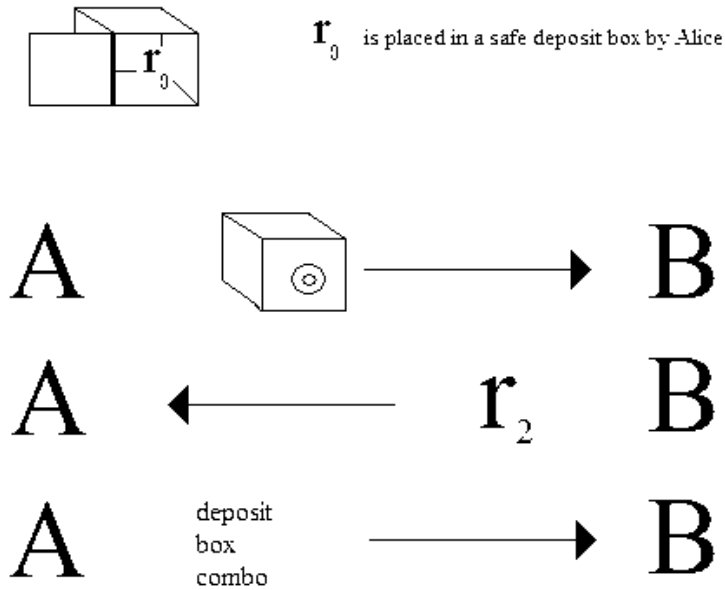


Figure 5: Alice and Bob wish to flip an unbiased coin. Concluding with the $coin = r_0 \oplus r_2$

before, a function that is one-way and 1-1 can still reveal a large part of its input. Instead, we look for some bit of information that is hard to compute. If we can find this bit b , then we can use it to 'commit' Alice's coin flip (i.e. $b \oplus r_0$). This bit b is what we call a Hard-Core Bit and we discuss it thoroughly in the next section.

9.3 Definition of a Hard-Core Bit

These examples motivate the following definition of a Hard-Core Bit due to Blum and Micali. Intuitively, a Hard-Core Bit is a bit associated with a one-way function which is as hard to determine as is inverting the one-way function.

Definition 17 A **Hard-Core Bit** $B(\cdot)$ of a function $f(\cdot)$ is a boolean predicate such that:

- $B(x)$ is easy to compute given x , i.e. in deterministic polytime
- Given $f(x)$, $B(x)$ is hard to guess better than at random:

$$\forall c \forall \text{ probabilistic poly-time } A, \text{ there exists } N_c \text{ such that } \forall n > N_c$$

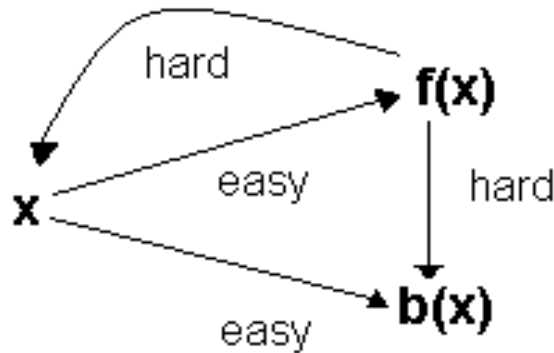


Figure 6: Relationship between a one-way function and a Hard-Core Bit

$$Pr_{\{x,\omega\}}[A(f(x)) = B(x)] < \frac{1}{2} + \frac{1}{n^c}$$

where $|x| = n$, and probability is taken over x and coin-flips ω of A .

9.4 Does the existence of a hard-core bit $B(\cdot)$ for a function $f(\cdot)$ imply that f is a one-way function?

We note first that the existence of a Hard-Core Bit for f does not necessarily imply that the corresponding one-way function is hard. As an example, the almost-identity function $I(b, x) = x$ has a hard-core bit b but is not hard-to-invert in the sense that we have defined in previous lectures. However, if no information is lost by the function f , then the existence of a Hard-Core Bit guarantees the existence of a one-way function. We prove a somewhat weaker theorem below.

Theorem 18 If f is a permutation which has a Hard-Core Bit, then f is a one-way function.

Proof Assume f is not one-way. Then there exists a good inverter for f which correctly computes inverses with probability $q > \epsilon(n)$, where probability is taken over x and coin-flips of A . The predictor for the Hard-Core Bit B first attempts to invert f using this good inversion strategy. If it succeeds in inverting f , it knows x , and can compute $B(x)$ in polynomial time. Otherwise, with probability $1 - q$, it fails to invert f , and flips a coin as its guess for $B(x)$. The predictor predicts B correctly with probability

$$q \cdot 1 + (1 - q) \cdot \frac{1}{2} = \frac{1}{2} + \frac{q}{2} \geq \frac{1}{2} + \frac{\epsilon(n)}{2}$$

Therefore f does not have a Hard-Core Bit, proving the contrapositive. ■

9.5 One-way functions have hard-core bits

The next two lectures are devoted to a proof of the following important theorem, first proved in 1989 by Goldreich and Levin, then simplified (by Venkatesan and Rackoff). It says that if $f_1(x)$ is a strong one-way function, then parity of a random subset of bits of x is a Hard-Core Bit:

Theorem 19 [Goldreich, Levin] Let f_1 be a strong one-way function. Let $f_2(x, p) \equiv (f_1(x), p)$, where $|x| = |p| = n$. Then

$$B(x, p) \equiv \sum_{i=1}^n x_i p_i \pmod{2}$$

is hard-core for f_2 .

Notice that a random subset is chosen by choosing p at random. The Hard-Core Bit of x is simply parity of a subset of bits of x , where the subset corresponds to all bits of x where corresponding bits of p are set to one.

Proof outline

The proof that $B(x, p)$ is a Hard-Core Bit will be by contradiction. We begin by assuming that $B(x, p)$ is not a Hard-Core Bit for f_2 . That is:

$B(\cdot, \cdot)$ is not hard-core: $\exists A_B \in \text{PPT}, \exists c$ such that for infinitely many n ,

$$\Pr_{\{x, p, \omega\}}[A_B(f_2(x, p)) = B(x, p)] > \frac{1}{2} + \frac{1}{n^c} \equiv \frac{1}{2} + \epsilon(n)$$

where A_B is probabilistic poly-time and probability is taken over x, p and coins ω of A_B

We want to show that we can invert f_2 with noticeable probability, proving that f_2 (and likewise f_1) is not a strong one-way function, i.e.:

f_2 is not a strong one-way function: $\exists A_{f_2} \in \text{PPT}, \exists c$ such that for infinitely many n :

$$\Pr_{\{x, p, \omega\}}[A_{f_2} \text{ inverts } f_2(x, p)] > \frac{1}{n^c}$$

where A_{f_2} is probabilistic poly-time and probability is taken over x, p and coins of A_{f_2}

We will show how to construct A_{f_2} using A_B as a subroutine.

Preliminaries

First, let us recall some useful definitions and bounds. Recall that a set of random variables is *pairwise independent* if given the value of a single random variable, the probability distribution of any other random variable from the collection is not affected. That is,

Definition 20 Pairwise independence: A set of random variables X_1, \dots, X_n are pairwise independent if $\forall i \neq j$ and $\forall a, b$:

$$Pr_{\{X_i, X_j\}}[X_i = a \wedge X_j = b] = Pr_{X_i}[X_i = a] \cdot Pr_{X_j}[X_j = b]$$

As an example of pairwise independence, consider distribution of three coins, taken uniformly from: $\{HHH, HTT, THT, TTH\}$. It is easy to check that given an outcome of any one of the three coins, the outcome of any other (of the two remaining) coins is still uniformly distributed. Notice however, that the number of *sample points* is small (only 4). On the other hand, for total (i.e. three-wise) independence we need all 8 combinations.

We are sometimes interested in bounding tail probabilities of large deviations. In particular, recall a Chernoff bound which we already used:

Definition 21 Chernoff bound: Let X_1, \dots, X_m be (totally) independent 0/1 random variables with common probability $0 < p < 1$, and let $S_m = X_1 + X_2 + \dots + X_m$. Then

$$Pr_{\{X_1, \dots, X_m\}}[|S_m - pm| > \delta m] \leq 2e^{-\frac{\delta^2 m}{2}}$$

Notice that as a function of m , the error-probability in Chernoff bound drops exponentially fast. In case of *pairwise independence* we have an analogous, Chebyshev bound. Like the Chernoff bound, the Chebyshev bound states that a sum of identically distributed 0/1 random variables deviates far from its mean with low probability which decreases with the number of trials (i.e. m). Unlike the Chernoff bound, in Chebyshev bound the trials need only be pairwise independent, but the probability drops off only polynomially (as opposed to exponentially) with respect to the number of trials.

Definition 22 Chebyshev bound: Let X_1, \dots, X_m be pairwise independent 0/1 random variables with common probability $0 < p < 1$, and let $S_m = X_1 + X_2 + \dots + X_m$. Then




Figure 7: Computing i^{th} bit.

$$Pr_{\{X_1, \dots, X_m\}}[|S_m - pm| > \delta m] \leq \frac{1}{4\delta^2 m}$$

We also implicitly used before a union bound, which simply states that:

Definition 23 Union Bound: For any two events A and B (which need not be independent), the

$$Pr[A \cup B] \leq Pr[A] + Pr[B]$$

Observation 24 *It is easy to see that if we flip the i^{th} bit of p , take its inner product with x and XOR the result with the original hard core bit, we will recover the i^{th} bit of x . That is, $\langle p \cdot x \rangle \oplus \langle p^i \cdot x \rangle = i^{\text{th}}$ bit of x*

Pictorially:

Observation 25 *Probability Distribution on 3 bits.*

In this observation, Pairwise Independence can be understood through a simple example. Below are three binary variables b_0, b_1, b_2 with equal distribution (i.e. uniformly random).

However, not all three random variables need to be generated. With only two of them the last one can be found using Pairwise Independence.

b_0	b_1	b_2
0	0	0
0	1	1
1	0	1
1	1	0

This can be accomplished by taking any two columns and xor-ing them. All possible choices look random when you apply this property over the variable values.

$b_0b_1 \rightarrow b_0 \oplus b_1$	<i>this looks like b_2</i>
$b_0b_2 \rightarrow b_0 \oplus b_2$	<i>this looks like b_1</i>
$b_1b_2 \rightarrow b_1 \oplus b_2$	<i>this looks like b_0</i>

Above observation demonstrates how to generate three pairwise independent bits from two independent bits using XOR operation. More generally, we can extend this trick over $\log n$ bits. Suppose that we have $\log n$ different bits $b_1, \dots, b_{\log n}$. Generate all possible subsets of these $\log n$ bits. Because each subset will contain one b_i different from any other set, all these sets are pairwise independent. Thus, corresponding to one set, by XORing the bits in that set all together, one obtains a new bit which is pairwise independent from other bits constructed in the same way. As there are approximately n such sets, we end up having n pairwise independent bits.

We use this trick later on in the proof. This is indeed the key idea of the proof. Because we can afford the advantage to go down by at most a polynomial factor, we can correctly guess $\log n$ bits. Using them we can generate n pairwise independent bits which will all be correct. We then use Chebyshev's inequality to bound the probabilities.

Two warmup proofs

To motivate the direction we will be heading for in the full proof, we first consider two scenarios in which the adversary on input $f_1(x)$ and p can guess $B(x, p)$ with probability much greater than a half.

In the following two warmup proofs, we use the following notation to (hopefully) clarify the presentation of the results. Given a string x , we use x^i to denote the string x with the i th bit flipped. We use array notation, $x[j]$, to denote the j th bit of x . Also, when referring to a string in the set of strings P , we use p_k to denote the k th string in the set.

The super-easy proof: Suppose the adversary A_B is able to guess the Hard-Core Bit $B(x, p)$ given $f_2(x, p)$ with probability 1. Then A_B can compute x bit-by-bit in the following manner. To compute $x[i]$, the i th bit of x , choose a random string p , and construct p^i . Since the adversary can compute hard-core bits with certainty, it can compute $b_1 = B(x, p)$ and $b_2 = B(x, p^i)$. By a simple case analysis, $x[i] = b_1 \oplus b_2$. After n iterations of this procedure (i.e. separately for each bit of x), we have the entire string x .

The somewhat easy proof: Now suppose that for every x , the adversary A_B is able to guess the Hard-Core Bit $B(x, p)$ given $f_2(x, p)$ with probability (over p) greater than $\frac{3}{4} + \epsilon(n)$: $Pr_p[A_B(f_2(x, p)) = B(x, p)] > \frac{3}{4} + \epsilon(n)$. Using the same procedure as in the super-easy proof, i.e. for each $x[i]$, we pick a random p and compute p^i , then guess $B(x, p)$ and $B(x, p^i)$ to help us determine $x[i]$. If we let E_1 and E_2 denote the events that the adversary's guesses for $B(x, p)$ and $B(x, p^i)$ are correct. We know that:

$$Pr_p[E_1 \equiv [A_B(f_2(x, p)) = B(x, p)]] > \frac{3}{4} + \epsilon(n)$$

and

$$Pr_p[E_2 \equiv [A_B(f_2(x, p^i)) = B(x, p^i)]] > \frac{3}{4} + \epsilon(n)$$

But these two events are *not* independent. Our guess for $x[i]$ is correct if both E_1 and E_2 occur (we also happen to get lucky if neither E_1 nor E_2 occur, but we ignore this case). We know that $Pr_p[\neg E_1] = \frac{1}{4} - \epsilon(n)$ and $Pr_p[\neg E_2] = \frac{1}{4} - \epsilon(n)$. Hence, by using union bound:

$$Pr_p[E_1 \wedge E_2] = 1 - Pr_p[\neg E_1 \vee \neg E_2] \geq 1 - \left[\left(\frac{1}{4} - \epsilon(n) \right) + \left(\frac{1}{4} - \epsilon(n) \right) \right] \geq \frac{1}{2} + 2\epsilon(n)$$

By employing tricks we have already seen, we can run the procedure for poly-many random p for each $x[i]$ and take the majority answer, which by a Chernoff bound amplifies the probability of success so that all bits of x can be guessed correctly with overwhelming probability.

9.6 One-way function have hard-core bits: The full proof

Now that we have obtained some insight as to how using a predictor of a Hard-Core Bit can help us to invert, we are ready to tackle the full proof. Therefore, we now assume that we are given an algorithm A_B which can compute the Hard-Core Bit with probability $> \frac{1}{2} + \epsilon(n)$ (over x, p and its coin-flips) and show an algorithm A_f (which uses A_B as a black-box) to invert f with noticeable probability.

The main idea of the proof is as follows: from the somewhat easy proof it is clear that we can not use our predictor twice on the same random string p . However, if for a random p we **guess correctly** an answer to $B(x, p) = b_1$, we can get the i 'th bit of x with probability

$\frac{1}{2} + \epsilon(n)$ by asking A_B to compute $B(x, p^i) = b_2$ *only once* for this (p, p^i) pair. So if we guess polynomially many $B(x, p_j) = b_j$ correctly for different random p_j 's we can do it. But we can only guess (with non-negligible probability) logarithmic number of totally independent bits. However, as we will see, we *can* guess (with non-negligible probability) a polynomial number of pairwise independent bits, and hence can do it. Now we go into the details.

Eliminating x from probabilities

In the somewhat easy proof, we assumed that the predictor A_B had $> \frac{3}{4} + \epsilon(n)$ chances for all x . But our A_B does not have such a guarantee. Our A_B guarantees only $\frac{1}{2} + \epsilon(n)$ success probability over all x and p and its coins. From this, we will conclude that there is a sufficiently “large” fraction of x such that we will still have a $> \frac{1}{2} + \frac{\epsilon(n)}{2}$ guarantee (only over the choice of p and *coins*). We will try to invert f only on this fraction of x 's. Thus, we begin by formalizing the notion of a *good* x and restrict our attention to adversaries which have reasonable chance of inverting $f_2(x, p)$ only on good x .

Definition 26 A string x is said to be **good** if $Pr_{p,\omega}[A_B(f(x), p) = B(x, p)] > \frac{1}{2} + \frac{\epsilon(n)}{2}$ where probability is taken over p and coin-flips ω of A_B .

Claim 27 At least an $\frac{\epsilon(n)}{2}$ fraction of x is good.

Proof Suppose not. Then,

$$\begin{aligned} Pr_{x,p,\omega}[A_B(f_2(x, p)) = B(x, p)] &= Pr_{x,p,\omega}[A_B(f_2(x, p)) = B(x, p) | x \text{ is good}] \cdot Pr_x[x \text{ is good}] \\ &\quad + Pr_{x,p,\omega}[A_B(f_2(x, p)) = B(x, p) | x \text{ not good}] \cdot Pr_x[x \text{ not good}] \\ &\leq 1 \cdot \frac{\epsilon(n)}{2} + \left(\frac{1}{2} + \frac{\epsilon(n)}{2}\right) \cdot 1 \\ &= \frac{1}{2} + \epsilon(n) \end{aligned}$$

This yields a contradiction, so the claim holds. ■

Overall strategy

Consider an adversary which attempts to invert $f(x)$ only on the set of good x , and succeeds with probability $> \frac{1}{2}$ on this set. Such an adversary succeeds in inverting $f(x)$ with total probability $\geq \frac{\epsilon(n)}{4}$, which is non-negligible, thereby ensuring that f is not a one-way function. This is exactly what we are going to do.

Our next question is for good x , with what probability does the adversary need to guess each bit $x[j]$ of x correctly in order to ensure that the entire x string is guessed correctly with probability $> \frac{1}{2}$. If the adversary computes each $x[j]$ correctly with probability $1 - \gamma$, then we can upper bound the probability that the adversary's guess for x is incorrect by employing the Union Bound:

$$Pr_{\omega}[A_{f_1}(f_1(x)) \text{ gets some bit of } x \text{ is wrong}] \leq \sum_{i=1}^n Pr_{\omega}[A_{f_1}(f_1(x)) \text{ gets } i\text{th bit wrong}] \leq n\gamma$$

where ω are coin-flips of A_f .

Setting $\gamma < \frac{1}{2n}$ guarantees that the probability that some bit of x is wrong is less than $\frac{1}{2}$, or equivalently, ensures that A_{f_1} guess is correct with probability $> \frac{1}{2}$. That is, if A_f can get each individual bit of x with probability greater than $(1 - \frac{1}{2n})$ then we can use the same procedure to get all bits of x with probability greater than $\frac{1}{2}$ even *if our method of getting different bits of x is not independent!*

Using pairwise independent p 's

Our next goal is to devise a strategy for the adversary to guess each bit $x[j]$ with probability at least $1 - \frac{1}{2n}$. Again, we begin by making an assumption which seems difficult to achieve, prove the result given the far-fetched assumption and then show how to derive the assumption.

Lemma 28 Suppose we are given a collection of $m \equiv \frac{n}{2\epsilon(n)^2}$ pairwise independent p_1, \dots, p_m , where $1 \leq i \leq m$, $|p_i| = n$ and every p_i is uniformly distributed. Moreover, suppose that for every i , we are given a b_i satisfying $x \cdot p_i = b_i$. Then, for good x , we can compute $x[j]$ correctly with probability $\geq 1 - \frac{1}{2n}$ in polynomial time.

Proof The adversary employs the following polytime algorithm.

1. For each $i \in 1, \dots, m$, construct p_i^j by flipping the j th bit of p_i .
2. Compute $b_i^j = x \cdot p_i^j$ by asking A_B .
3. Derive a guess for $x[j]$ as in the “somewhat easy” proof: $g_i = b_i^j \oplus b_i$
4. Take the majority answer of all guesses g_i as the guess for $x[j]$.

We are interested in bounding the probability that the majority of our guesses were wrong, in which case our guess for $x[j]$ is also wrong. Define $y_i = 1$ if g_i was incorrect and $y_i = 0$ otherwise, and let $Y_m = \sum_{i=1}^m y_i$. Using Chebyshev:

$$\begin{aligned}
 \Pr \left[Y_m > \frac{m}{2} \right] &= \Pr \left[Y_m - mp > \frac{m}{2} - mp \right] \\
 &= \Pr \left[Y_m - mp > \left(\frac{1}{2} - p \right) m \right] \\
 &\leq \Pr \left[|Y_m - mp| > \left(\frac{1}{2} - p \right) m \right] \\
 &\leq \frac{1}{4 \left[\left(\frac{1}{2} - p \right)^2 m \right]} \quad [\text{Chebyshev}] \\
 &= \frac{1}{4\epsilon(n)^2 m}
 \end{aligned}$$

Substituting in for $m = \frac{n}{2\epsilon(n)^2}$ ensures that the probability that we misguess $x[j]$ (i.e., that $\Pr[Y_m > \frac{m}{2}]$) is at most $\frac{1}{2n}$, proving the claim. ■

Lemma 29 If we are given uniformly distributed completely independent p_1, \dots, p_l for $l \equiv \lceil \log(m+1) \rceil$ together with b_1, \dots, b_l satisfying $B(x, p_i) = b_i$ then we can construct in polynomial time a pairwise independent uniformly distributed p_1, \dots, p_m , $m \equiv \frac{n}{2\epsilon(n)^2}$ sample of correct equations of the form $B(x, p_i) = b_i$

Proof The proof hinges on the following fact, whose proof we omit because it is a simple case analysis:

Fact 30 Given correct equations $x \cdot p_1 = b_1$ and $x \cdot p_2 = b_2$, then $x \cdot (p_1 \oplus p_2) = (b_1 \oplus b_2)$.

It is easy to see by induction that this fact extends to the case in which there are arbitrarily many b_i and p_i . Therefore, we can generate a large set of new, valid equations by repeatedly

choosing an arbitrary subset of the p_i s, XOR them together; XOR the corresponding b_i s together to form a new equation of the form, for example, $x \cdot p_{1,3,5,7} = b_{1,3,5,7}$. Since there are $(2^l - 1)$ non-empty subsets of a set of size l , by choosing all possible subsets, the new set is of polynomial size $2^{\log l} = m$ and each new equation is polytime constructible. Furthermore, if we look at the symmetric difference of two different subsets, they are pairwise independent, so the entire set of new equations is pairwise independent. ■

Putting it all together

We now have all the machinery to provide a construction for inverting $f_1(x)$ with noticeable probability given a predictor A_B for predicting $B(x, p)$ with probability $> \frac{1}{2} + \epsilon(n)$. Here is the algorithm to invert f_1 .

Algorithm $A_{f_1}(y = f_1(x))$:

- Step 1: Pick a set $P \equiv \{p_1, \dots, p_l\}$ uniformly at random, where $|x| = |p_i| = n$ and $l \equiv \lceil \log(\frac{n}{2\epsilon(n)^2} + 1) \rceil$
- Step 2: Compute pairwise independent $\hat{P} \equiv \{\hat{p}_1, \dots, \hat{p}_m\}$ where $m \equiv 2^l - 1$ and \hat{P} is computed by taking XOR of all possible non-empty subsets of P .
- Step 3: For all $p_i \in P$ choose bits b_1, \dots, b_l randomly.
 - Step 3.1: [Assume that for every $p_i \in P$, $(1 \leq i \leq l)$, $B(x, p_i) = b_i$]
From P , and b_1, \dots, b_l compute for every $\hat{p}_k \in \hat{P}$ bit \hat{b}_k , where where \hat{b}_k are computed by taking XOR of the b_i 's corresponding to $p_i \in P$ used for computing \hat{p}_k , and where $\hat{b}_k \equiv B(x, \hat{p}_k)$ for all $1 \leq k \leq m$.
 - Step 3.2: For j from 1 to n do: [compute all bits $x[j]$ of x]
 - Step 3.2.1: For every $\hat{p}_k \in \hat{P}$, where $1 \leq k \leq m$ ask A_B to predict $c_k \equiv B(x, \hat{p}_k^j)$. Let $b_{\hat{p}_k} \equiv c_k \oplus \hat{b}_k$
 - Step 3.2.2: define $x[j]$ as the majority of $\hat{b}_{\hat{p}_k}$ from step 3.2.1
 - Step 3.3: Check if for $z \equiv (x[1], \dots, x[n])$ after step 3.2 $f_1(z) = y$. If so, **return** z , otherwise output *fail*.

The adversary randomly selects a set of l strings of length n to from a set P . It then iterates through all possible completions $x \cdot p_i = b_i$, of which there are only $2^{\log l} = m$. For each incorrect completion, the adversary will perform a polynomial amount of useless work

which we are not interested in; we focus on the work performed on the correct completion of the set of equations (which we can check in step 3.3). By Lemma 5.10, since the set of l equations is totally independent, we can construct a pairwise independent set of m equations which are also correct. (step 3.1) Now from Lemma 5.9, this set of m equations suffices to invert $f(x)$ if x is good with probability $> \frac{1}{2}$. Early on, we noticed that the existence of a probabilistic poly-time A_{f_1} which succeeds in inverting $f_1(x)$ for good x with probability $> \frac{1}{2}$ proves that we can invert f with probability greater than $\frac{\epsilon(n)}{4}$, since good x occurs with probability greater than $\frac{\epsilon(n)}{2}$. But if we can invert f_1 with probability greater than $\frac{\epsilon(n)}{4}$, f_1 is not a strong one-way function. This completes the proof of the contrapositive, so we have shown that every one-way function has a Hard-Core Bit. ■

Remark: Instead of searching through all possible 2^l bit strings b_1, \dots, b_l in step 3, we can just pick *at random* b_1, \dots, b_l and try it only once. We guessed correctly with probability $\frac{1}{2^l} = \frac{1}{poly}$, hence we will still invert f_1 on good x with $\frac{1}{poly}$ probability.

Lecture 3

10 Private Key Encryption

Consider a game⁶ between three players: Alice, Bob and Eve. In this game, Alice and Bob are communicating over an insecure channel, where Eve can listen to messages, although she cannot modify them. Informally, Bob and Alice win if they are able to communicate without Eve being able to decipher any partial information about the messages. On the other hand, Eve wins if she can discover any information about the messages sent. Notice that Eve does not have to find the whole content of the messages in order to win; for her, it suffices to discover any information hidden in the messages that she did not know when the game started.

In this game, a *plaintext message* m is encrypted by means of a *cipher key*, to produce a *ciphertext* c . Because Eve does not have access to the key used to encrypt the messages, it is said that Bob and Alice communicate by means of a *private key encryption system*. In a private key encryption system, it is necessary that Bob and Alice meet before starting to communicate, so they can agree on a particular key, which should be kept secret, that is, only they must know it. This is demonstrated in Figure 8.

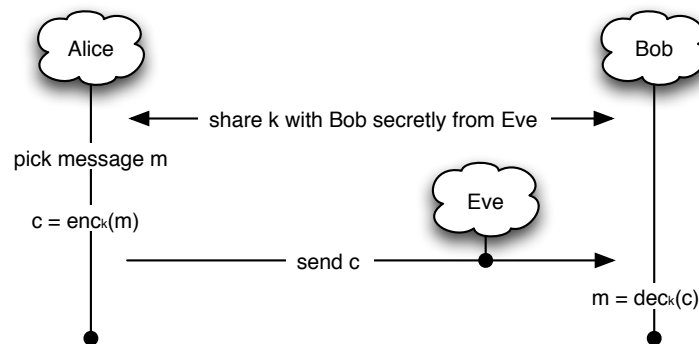


Figure 8: Alice and Bob use a private-key encryption scheme.

Sometimes, in a private key system, an encryption scheme is associated with a message space M and some distribution on M . For example, M may be the set of all strings for a given length. The ciphertext space is the set of all possible plaintext messages $m \in M$ encrypted with each possible key. Note that the ciphertext space does not necessarily equal

⁶Scribe notes by Paul Ray, Mike Welch, Fernando Pereira; January 26th 2006

M , that is, for some combinations of m and key k the same ciphertext can be produced. Some ciphertext strings may not correspond to any plaintext. A formal definition of a Private Key cryptosystem follows below:

Definition 31 We define an encryption scheme as a triple of algorithms (GEN , ENC , DEC):

- *GEN: Key Generation: a randomized polynomial time algorithm that outputs a key K .*
 $K \leftarrow GEN(1^n, R)$
where 1^n is a security parameter and R is a polynomial number of coin-flips.
- *ENC: Encryption: a randomized polynomial time algorithm that takes a plaintext message m and key K and generates a ciphertext C .*
 $C \leftarrow ENC_K(m, R)$
- *DEC: Decryption: a polynomial time algorithm that takes a ciphertext C and key K and outputs a plaintext message m' .*
 $m' \leftarrow DEC_K(C)$
- *Correctness: a cryptographic system is called correct if $DEC_K(ENC_K(m, R)) = m$, provided that m has been generated by the GEN algorithm. Notice that we can allow negligible error probability: $Pr[DEC_K(ENC_K(m, R)) \neq m] < \epsilon$.*

10.1 Examples of Private Key Systems

Cryptosystems that we used in kindergarden are easy to break. Examples include:

Shift cipher (insecure!) For a shift cipher with $m \in \{A...Z\}^n$. $K \leftarrow GEN()$ assigns K a random number in the range $\{0...25\}$. $C \leftarrow ENC_K(m_1m_2...m_l) = c_1c_2...c_l$ such that $c_i = (m_i + K) \text{ modulus } 26$. $m' \leftarrow DEC_K(c_1c_2...c_l) = m'_1m'_2...m'_l$ where $m'_i = (c_i - K) \text{ modulus } 26$.

Example: If $K = 4$, $m = \text{"HELLO"}$ then $C = \text{"LIPPT"}$.

Because there are only 26 different keys, this system is trivial to break.

Caesar cipher (insecure!) For a substitution cipher, we let K be a random permutation of the numbers $\{0...25\}$. Then, $C \leftarrow ENC_K(m_1m_2...m_l) = c_1c_2...c_l$ where $c_i = K(m_i)$. Example: For a limited alphabet $\{A, B, C, D\}$, let K be the mapping $\{A, B, C, D\} \rightarrow \{B, D, A, C\}$. Then for $m = \text{"DAD"}$, $C = \text{"CBC"}$.

Again, this type of cipher is weak, with a very limited key space. For example, if the only allowed symbols are the lower case characters of the English alphabet, then there are 26 possible substitutions! However, in order to break this system, it is not

necessary to try every possible combination. Attacks based on the frequency of each character are very effective against the Caesar cipher.

On the other hand, we will see that the following cryptosystem is secure:

One-time pad For a one-time pad, we let $K \leftarrow \{0, 1\}^n$ with M the set of all binary strings of length l . We let $ENC_K(m) = m \oplus K$ (bitwise). The decryption function $DEC_K(c) = c \oplus K$ (bitwise).

Example: If $m = 11001$ and $K = 01101$, $C = 11001 \oplus 01101 = 10100$.

10.2 Security Criteria

There are several different formal definitions of an information-theoretically *secure cryptographic* system. One is the *Shannon Security Criterion* (Definition 32). According to the Shannon criterion, the probability that a certain message was generated, given the knowledge of a ciphertext, is equal to the *a-priori* probability that any message was produced. That is, the knowledge of the ciphertext does not provide any *a-posteriori* information about the message m , except what was already known from *a-priori* probabilities.

Definition 32 An encryption scheme over message space M is Shannon secure if, for all distributions D over M , for all $m \in M$, and for all ciphers c , the following equation holds:

$$Pr_D[m|c] = Pr_D[m]$$

The equation above means that the a posteriori probability that a message m has been sent, given that the cipher c has been sent, is equal to the a priori probability that message m has been sent.

Claim 33 One-time pad is Shannon secure

Proof For all the possible distributions D of messages over the message space $M = \{0, 1\}^n$, for all keys in the key space $K = \{0, 1\}^n$, for all fixed $m \in M$, and for a fixed $c \in \{0, 1\}^n$:

$$Pr_{D,k}[m|c] = \frac{Pr_{D,k}[m \wedge c]}{Pr_{D,k}[c]} = \frac{Pr_{D,k}[c|m] \cdot Pr_{D,k}[m]}{Pr_{D,k}[c]}$$

$$Pr_{D,k}[c|m] = Pr_{D,k}[m \oplus k = c] = \frac{1}{2^n}$$

Notice that, for a given (m, c) pair, there is only one possible value for k such that $k = m \oplus c$. Also, the probability of a given cipher c been produced is given by:

$$Pr_{D,k}[c] = \sum_{m \in M} Pr_{D,k}[c|m] \cdot Pr_{D,k}[m] = 2^{-n} \cdot \sum_{m \in M} Pr_{D,k}[m] = 2^{-n}$$

Using the latter result in the first equation gives:

$$Pr_{D,k}[m|c] = \frac{2^{-n} \cdot Pr_{D,k}[m]}{2^{-n}}$$

$$Pr_{D,k}[m|c] = Pr_{D,k}[m]$$

$$Pr_D[m|c] = Pr_D[m]$$

■

Unbreakable cryptographic systems can also be characterized by the concept of *Perfect Security* (Definition 34). This definition states that an encryption scheme is perfectly secure if a ciphertext c is an equally likely output for any two messages in the message space.

Definition 34 *An encryption scheme S over message space M is perfectly secure if, for any distribution D over any set $\{m_1, m_2\} \subset M$ of two messages of same length, and for all cipher texts c we have:*

$$Pr_D[m_1|c] = Pr_D[m_2|c] \tag{5}$$

Given Definitions 32 and 34, which is stronger? That is, which gives the user of a cryptosystem the highest level of security? Although they look different, it is possible to show that they are essentially equivalent definitions. Here we prove that Shannon security is a consequence of perfect security.

Claim 35 *If perfect security holds then Shannon security holds.*

Proof

Given arbitrary messages m_1 and m_2 from the message space M , according to perfect security, we have $Pr_D[c|m_1] = Pr_D[c|m_2]$, where D is any distribution over some arbitrary subset of M with just two elements, m_1 and m_2 . Since this equation is true for arbitrary messages m_1 and m_2 , it must hold for any pair of messages in M . Assume that the following probabilities range over any distribution of messages from M :

$$\begin{aligned}
Pr[m|c] &= \frac{Pr[c|m] \cdot Pr[m]}{Pr[c]} \\
&= \frac{Pr[c|m] \cdot Pr[m]}{\sum_{m' \in M} Pr[c|m'] \cdot Pr[m']} \\
&= \frac{Pr[c|m] \cdot Pr[m]}{\sum_{m' \in M} Pr[c|m] \cdot Pr[m']} \\
&= \frac{Pr[c|m] \cdot Pr[m]}{Pr[c|m] \cdot \sum_{m' \in M} Pr[m']} = Pr[m]
\end{aligned}$$

Because $Pr[c|m] = Pr[c|m']$ for any message m' , it is possible to remove the probability $Pr[c|m']$ from the summation. Also, $\sum_{m' \in M} Pr[m'] = 1$, given that a message has been produced. This two remarks justify the last steps of the derivation above.

■

Claim 36 *For any perfectly secure encryption, the number of keys $|K|$ is at least as big as the size of the message space $|M|$.*

Proof Consider an encryption scheme S in which the size of the key space is less than the size of the message space, that is, $|K| < |M|$. Bob and Alice will try to use S to defeat Eve. Given a message $m \in M$, and a key $k \in K$, Bob generates a cipher text $c = ENC_k(m)$, and sends it to Alice. Eve intercepts c and now applies every possible key k' in order to decrypt c . If S is a perfectly secure system, for any message m , and any key k , $Pr_k[m|c] = Pr_k[m]$. However, because $|K| < |M|$, there must be at least one message m such that $Pr_k[m|c] = 0$. Therefore, after knowing c , one of the following situations must hold: $Pr_k[m|c] > Pr_k[m]$, or $Pr_k[m|c] = 0$. In this case, the cryptographic system does not meet the Shannon security criterion, and, consequently, perfect security does not hold.

■

Fact 37 *Shift and Substitution Ciphers do **not** satisfy perfect security.*

Proof This is easy to see using Claim 36. ■

Fact 38 *1-Time Pad is perfectly secure.*

11 Indistinguishability

The problem with 1-Time Pad is that the key must be as long as the message, and we cannot seem to do better if we need information-theoretic security. We will see that using so-called pseudo-random generators and weaker “computational security”, we can make the key of the private-key encryption scheme much shorter. In the context of this discussion, a pseudo-random generator will be defined as a deterministic polynomial-time algorithm that takes as input an initial sequence of truly random bits, called a *seed*, plus a polynomial Q , and outputs a new, apparently “random looking” sequence of size equal to at least $Q(|seed|)$. A simplified scheme is shown in Figure 9.




Figure 9: A pseudo-random generator.

Before moving to the formal definition, it is important to give an intuitive idea about what does it mean for a sequence of symbols to be “apparently random”. A distribution is called pseudo-random if it is not possible to build an algorithm of polynomial complexity that can distinguish that distribution from a truly random distribution of strings of the same length. It is possible to establish an analogy between this concept and the idea underlying the *Turing Test*. This test, proposed by Alan Turing, aims to determine if a computer could be considered intelligent. According to the test, a computer is considered intelligent if it is impossible to distinguish it from a true human being in an interview conducted on-line. In our case, the test aims to distinguish between a truly random and a pseudo-random sequence of bits. Such sequences, according to Definition 39, are said to be from a *Sampleable Distribution*.

Definition 39 *A Sampleable Distribution* – A sampleable distribution is a probabilistic polynomial time algorithm $S(1^n, r)$ that takes a finite sequence of random bits r and 1^n , and produces strings of length n , such that, for some distribution X_n :

$$Pr_r[S(1^n, r) = \alpha] = Pr_{X_n}[X_n = \alpha] \quad (6)$$

The test T , used to determine if two distributions are the same or not, is called a *statistical test*. According to definition 40, given two sampleable distributions X_n and Y_n , it is said that they are indistinguishable with respect to the probabilistic polynomial test T if T cannot determine whether a sample comes from X_n or Y_n . Figure 11 represents a statistical test.

Definition 40 The random distributions X_n and Y_n are said to pass the probabilistic polynomial test T if $\forall c, \exists N$ such that $\forall n > N$:

$$|Pr_{w, X_n}[T_w(X_n) = 1] - Pr_{w, Y_n}[T_w(Y_n) = 1]| < \frac{1}{n^c} \quad (7)$$




Figure 10: Experiment 1: statistical test receiving pseudo-random string.




Figure 11: Experiment 2: statistical test receiving truly random sequence.

According to Definition 41, two distributions are statistically close if the difference in probability that a sample string α comes from one distribution versus the other is negligible.

Definition 41 We say that two distributions X_n and Y_n are statistically close if $\forall c, \exists N$, such that $\forall n > N$:

$$\sum_{\alpha \in \{0,1\}^n} |Pr[X_n = \alpha] - Pr[Y_n = \alpha]| < \frac{1}{n^c} \quad (8)$$

If it is not possible for any polynomial time algorithm to distinguish two distributions, it is said that these distributions are computationally close. Definition 42 formally states this fact.

Definition 42 [Yao] X_n and Y_n are polynomial-time indistinguishable if, and only if, $\forall c$, and $\forall A \in PPT$, $\exists n$ such that $\forall n > N$:

$$|Pr[A(X_n) = 1] - Pr[A(Y_n) = 1]| < \frac{1}{n^c} \quad (9)$$

11.1 Extended Statistical Test

One could argue that, if it was possible to feed the testing algorithm T with several, instead of just one, samples from the random distributions, perhaps it would be easier to distinguish a pseudo-random sequence from a truly random distribution. According to Definition 43, a test that receives more than one sample is called an *extended statistical test*. An extended statistical test T' is a machine that takes a polynomial number of inputs and outputs a single bit in $\{0, 1\}$. Given two distributions X_n and Y_n , T' is given a polynomial number of samples from either X_n or Y_n (not both). Then, if T' determines that the samples come from X_n it outputs bit b , otherwise it outputs bit \bar{b} .

Definition 43 *Extended Statistical Test* – The sampliable distributions X_n , and Y_n pass extended statistical test T' if $\forall c_1, c_2, \exists N$ such that $\forall n > N$:

$$|Pr(T'(X_n^{c_1}, \dots, X_n^{c_2}) = 1) - Pr(T'(Y_n^{c_1}, \dots, Y_n^{c_2}) = 1)| < \frac{1}{n^{c_1}} \quad (10)$$

Claim 44 If X_n and Y_n are sampliable distributions which can be distinguished by a (uniform) extended statistical test T' , then there exist a single sample (uniform) statistical test T which distinguishes X_n and Y_n .

Proof Let $k = \text{poly}(n)$ and $\epsilon(n) = 1/k$. We assume that there exists T' and show how to construct T . Assuming that there exists T' means, w.l.o.g. that

$$Pr_{X_n}(T'(X_1, X_2, X_3, \dots, X_{poly}) = 1) - Pr_{Y_n}(T'(Y_1, Y_2, Y_3, \dots, Y_{poly}) = 1) > \epsilon(n)$$

Consider “hybrids” P_j , for $0 \leq j \leq k$, where in P_j the first j samples come from Y_n and the remaining samples come from X_n :

$$P_0 = x_1 x_2 x_3 x_4 \dots x_k$$

$$P_1 = y_1 x_2 x_3 x_4 \dots x_k$$

$$P_2 = y_1 y_2 x_3 x_4 \dots x_k$$

$$P_3 = y_1 y_2 y_3 x_4 \dots x_k$$

...

$$P_k = y_1 y_2 y_3 y_4 \dots y_k$$

We know that $P_0 - P_k > \epsilon(n)$, and therefore, $\exists j$ such that $P_j - P_{j+1} > \epsilon(n)/k$ (which is another 1/poly fraction!) Consider a distribution:

$$P(\boxed{z}) = y_1 y_2 y_3 \dots y_j \boxed{z} x_{j+2} \dots x_k$$

Notice that if z is a sample from Y_n then $P(z) = P_j$ and if z is a sample from X_n then $P(z) = P_{j+1}$. Hence, if we are given z on which we have to guess which distribution it came from, if we put z in the box above, and somehow fix other locations we could distinguish on a single sample z . Two questions remain: (1) how do we find the correct $j + 1$ position, and (2), how do we fix other values. The answers differ in uniform and non-uniform case:

non-uniform case (i.e. both T' and T are circuits): Since T is a circuit, we can non-uniformly find the correct $j + 1$ value and find values to other variables which maximizes distinguishing probability.

uniform case : Since X_n and Y_n are sampleable, we can fix values different from j to be samples from X_n and Y_n and by guessing correctly j (we guess the position of j correctly with probability 1/poly). The distinguishing probability could be further improved. By experimenting with the distinguisher, we get (again using sampleability of X_n and Y_n !) to check if our choice of j and samples of other positions are good. ■

12 Pseudo-Random Generators

A Pseudo-Random Generator must be able to generate a sequence of symbols which cannot be distinguished, in polynomial time, from a truly random distribution, even though such a sequence is deterministically generated, given a short truly random seed. Definition 45 formally defines the properties of a pseudo-random generator.

Definition 45 A deterministic algorithm $G(\cdot, \cdot)$ is pseudo-random generator if:

- 1. $G(x, Q)$ runs in time polynomial in $|x|, Q(|x|)$ where Q is a polynomial.
- 2. $G(x, Q)$ outputs strings of length $Q(|x|)$ for all x .
- 3. For every polynomial $Q()$, the induced distribution $\{G(x, Q)\}$ is indistinguishable from $\{U_{Q(|x|)}\}$, where $U_{Q(|x|)}$ is a uniform distribution on strings of length $Q(|x|)$.

12.1 Construction

We will now show how we can construct a pseudo-random generator from a one-way permutation. Remember from the previous lecture that for a one-way permutation $f(x)$, it is possible to obtain a hard-core bit b , which is as difficult to predict as inverting f .

Consider the following steps:

1. Pick a one-way permutation function f , random seed X_0 , and a random bit string P , $|X_0| = |P| = n$
2. For i from 1 to $Q(|x|)$ do:
Calculate $X_i = f(X_{i-1})$ and output the hardcore bit $\langle P * X_{i-1} \rangle$
3. Output $X_{Q(|x|)}$ and P
4. Return the sequence of random bits in reverse order.

Remark In the construction above, the bits are returned in “reversed” order. However, since we will show that this is pseudo-random, the order (left-to-right or right-to-left) is unimportant: if it is pseudo-random one way, it must be pseudo-random the other way. We have chosen this particular order of the outputs to make the proof easier.

A graphical representation of the algorithm is shown in Figure 12.

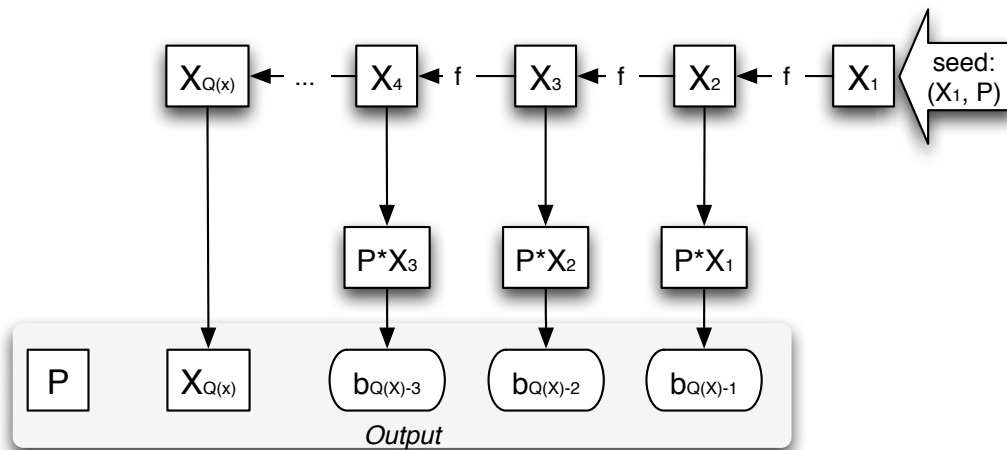


Figure 12: A pseudo-random generator.

The proof that the output of the above algorithm is pseudo-random is shown in two steps: First, we will show that the pseudo-random sequence is unpredictable as defined by Blum

and Micali. Then we show that every unpredictable sequence is pseudo-random (which was shown by Yao).

Informally, a pseudo-random generator G is *unpredictable* if it passes the next-bit-test, defined as follows: given G_1, \dots, G_k , it is hard to predict G_{k+1} , for any k . That is, given the first k bits of the output, it is hard to predict the next bit with probability better than $\frac{1}{2}$

Definition 46 X_n passes the next bit test (is unpredictable) if $\forall \epsilon$ and $\forall A \in PPT$, there $\exists N$ such that $\forall n > N$ and $\forall i, (0 \leq i \leq n)$:

$$Pr_{X_n, \text{coins of } A} [A(\text{first } i \text{ bits } b_1, b_2, \dots, b_i \text{ of } x \in X_n) = b_{i+1}] < \frac{1}{2} + \frac{1}{n^\epsilon} \quad (11)$$

Claim 47 If f is a strong one-way permutation, then G as defined in the previous section is unpredictable.

Proof The proof is by contradiction. We will show that if G does not pass the next bit test, then we can invert a one-way permutation on a random input. The construction is as follows: We are given an adversary $A \in PPT$ which for some prefix $b_1 \dots b_i$ of output from a pseudo-random generator, can compute b_{i+1} with probability $> 1/2 + \epsilon(n) = 1/2 + 1/\text{poly}$. We wish to invert f , that is, find $f^{-1}(y)$. We know that if we can predict a hard-core bit of $f^{-1}(y)$ and p (for a random p), then we can find the inverse of y (with $1/\text{poly}$ probability). We make y the $(n-i)^{\text{th}}$ X_i value of the generator. We can then compute (by applying f to y i times, and computing hard-core bits) the first i bits of the generator in the straightforward fashion. Finally, we feed it to our “next-bit” predictor A . Notice that the next bit is exactly the hard-core bit of y . ■

Remark Notice that we are dealing with f , which is a one-way permutation. Hence, a uniform distribution of the inputs implies a uniform distribution of the outputs, and in the above experiment, the fact that we start with y and compute the first i bits of the output, has the same (uniform) distribution over the seeds.

Claim 48 Ensemble X_n is pseudo-random if and only if X_n is unpredictable.

Proof Proof in one direction is trivial: a next bit test is a type of statistical test. Hence if it passes all polynomial-time statistical tests it passes the next-bit test as well.

In the opposite direction, we must show that passing the next bit test implies passing all polynomial-time statistical tests. We will show a proof by contradiction. That is, if \exists distinguisher D , then \exists a next bit predictor. A distinguisher means that for distributions

$p = X_n$ and $q = Y_n$, $|p - q| > \frac{1}{n^c}$.

By the hybrid argument, there $\exists i$ such that $|P_i - P_{i+1}| > \frac{\epsilon}{l}$. Then, we can construct an $(i + 1)^{st}$ bit predictor as follows:

Give an adversary A an input string $B_1 \dots B_i \hat{b} Y$, where \hat{b} is a random bit. If the adversary outputs 1 on the input, we output \hat{b} , otherwise we output $\bar{\hat{b}}$.

$$\begin{aligned} \Pr[A(B_1 \dots B_i) = b_{i+1}] &= \Pr[\hat{b} = b_{i+1}] * \Pr[A(B_1 \dots B_i b_{i+1} Y) = 1] \\ &\quad + \Pr[b_{i+1} = \bar{\hat{b}}] * \Pr[A(B_1 \dots B_i \bar{b}_{i+1} Y) = 0] \\ &= \frac{1}{2} * P_{i+1} + \frac{1}{2} * q \end{aligned}$$

So, what is q ? To figure it out, let's expand P_i

$$\begin{aligned} P_i = \Pr[A(B_1 \dots B_i Y)] &= \Pr[A(B_1 \dots B_i b_{i+1} Y) = 1] * \frac{1}{2} \\ &\quad + \Pr[A(B_1 \dots B_i \bar{b}_{i+1} Y) = 1] * \frac{1}{2} \\ &= P_{i+1} * \frac{1}{2} + (1 - q) * \frac{1}{2} \end{aligned}$$

or, $q = 1 + P_{i+1} - 2 * P_i$

Substituting for q above, we get:

$$= \frac{1}{2} * P_{i+1} + \frac{1}{2} * q = \frac{1}{2} * P_{i+1} + \frac{1}{2} [1 + P_{i+1} - 2P_i] = \frac{1}{2} [P_{i+1} - P_i] > \frac{1}{2} + \frac{\epsilon}{l} \blacksquare$$

Lecture 4

13 Pseudo-random Generators

First⁷ we note that the existence of pseudo-random generators is equivalent to the existence of one-way functions. This was established by Håstad, Impagliazzo, Levin, and Luby in 1999 [2]. We will omit this proof in one direction due to its complexity.

Theorem 49 *There exists a pseudo-random generator iff there exists a one-way function.*

Proof Here we will only show that any pseudo-random generator is a one-way function.

Let $G : \{0, 1\}^n \rightarrow \{0, 1\}^{\ell(n)}$ be a pseudo-random generator where $\ell(n) \geq n + 1$. This means that G has at most 2^n distinct outputs, in particular, since $\ell(n) \geq n + 1$, at most half of the strings in $\{0, 1\}^{\ell(n)}$ are in the range of G . Now suppose G is not a one-way function, i.e. there exists a machine A that can invert G with non-negligible probability $\epsilon(n)$. We will use this A to construct a machine that can distinguish the output of G from true random. Notice that for strings that are not in the range of G , A cannot find an inverse, since no inverse exists. For an $\ell(n)$ -bit string x , if x is pseudo-random, we can invert x with probability $\epsilon(n)$. If x is truly random, we invert with probability $\epsilon(n)$ multiplied by the probability that x has an inverse, which is less than $\frac{1}{2}$. Thus A inverts a truly random string with probability less than $\frac{\epsilon(n)}{2}$. We construct a distinguisher as follows: given an $\ell(n)$ -bit string x , we try to invert x , if we succeed, we say that x is pseudo-random, otherwise we say that x is random. The probability of each outcome is summarized in the following table.

	Guess $x \in PR$	Guess $x \in R$
$x \in PR$	$\epsilon(n)$	$1 - \epsilon(n)$
$x \in R$	$< \frac{\epsilon(n)}{2}$	$> 1 - \frac{\epsilon(n)}{2}$

Thus the probability that we are correct is greater than $\frac{1}{2} + \frac{\epsilon(n)}{2}$ which is non-negligible since $\epsilon(n)$ is non-negligible. Thus we have created a distinguisher that can distinguish the output of G from true-random with non-negligible probability, which contradicts the assumption that G is a pseudo-random generator. ■

Recall that in Lecture 4, we showed that the existence of a one-way *permutation* implies the existence of a pseudo-random generator. The work of Håstad, Impagliazzo, Levin,

⁷Scribe notes by B. Hemenway, B. Tradewell; 12 Apr 2006

and Luby extends this result to say that any one-way *function* can be used to construct a pseudo-random generator.

We now illustrate some applications of pseudo-random generators.

14 Applications of Pseudo-Random Generators

14.1 P vs. BPP

Suppose we have a machine that can't flip coins, but want to simulate a machine that can.

Say that we have a BPP machine for determining whether a given string x is in the language L . Recall that a BPP machine takes as input the string x , with say $|x| = n$, and a string of random bits, of length polynomial in n (say $Q(n)$), and outputs either a yes or a no. If $x \in L$ then the probability that the output is "yes" is at least $\frac{3}{4}$ (probability taken over all possible strings of $Q(n)$ random bits), and if $x \notin L$ then the probability of a "yes" output is no more than $\frac{1}{4}$.

Now if we have such a BPP machine and an x , there is a simple algorithm which will tell us definitely, not just probabilistically, whether $x \in L$ or not. It goes like this:

- Try all $2^{Q(n)}$ random strings
- Count how many strings give yes and how many give no.
- If yes occurs more frequently $x \in L$; if no occurs more frequently $x \notin L$.

This is an exponential time algorithm. We can improve the speed of this algorithm using a pseudo-random generator. Since the existence of a pseudo-random generator is equivalent to the existence of a one-way function, we have the following theorem:

Theorem 50 [Yao]: *If there exist non-uniform one-way functions then BPP is contained in subexponential time. That is,*

$$BPP \subseteq \bigcap_{\epsilon > 0} DTime(2^{n^\epsilon})$$

The algorithm, call it P' , uses a pseudo-random generator $G : n\text{-bits} \rightarrow Q(n)\text{-bits}$. The algorithm is:

- Cycle through all 2^n possible seeds of G .

- Take the outputs of G as the $Q(n)$ -bit strings to be input into the BPP machine along with the input x .
- Count how many strings give yes and how many give no.
- If yes occurs more frequently $x \in L$; if no occurs more frequently $x \notin L$.

Proof Assume our algorithm P' makes a mistake. We will prove this implies that: We can construct a polynomial-time distinguisher of $\{G(n)\}_{Q(n)}$ and $\{U\}_{Q(n)}$, where $G(n)$ is the output of the pseudo-random generator and U is the uniform, i.e. truly random, distribution.

Suppose there exists a string x , $|x| = n$, on which P' makes a mistake, that is, either $x \in L$ and P' says “no” or $x \notin L$ and P' says “yes”. We will show that in either case, x together with our BPP machine can be used to distinguish truly random strings from the pseudo-random outputs of G , a contradiction. There are two possible errors P' can make, and we examine them individually.

Case 1. $x \in L$ but P' makes a mistake and says no.

This means that on more than half of the pseudo-random strings the output is no but on more than three quarters of random string the output is yes. Hence we have a distinguisher which succeeds with probability $1/4$.

Case 2. $x \notin L$ but P' makes a mistake and says yes.

This means that for more than half of the pseudo-random strings the machine says yes, while for less than a $1/4$ of truly random strings the machine says yes. Hence again we have a distinguisher that succeeds with probability $1/4$. ■

14.2 Bit Commitment Protocol

Many cryptographic protocols use as their basis a bit commitment protocol. Some of these applications are zero knowledge protocols, identification schemes, multi-party protocols, and coin flipping over the phone.

The existence of a good pseudo-random generator allows us to construct a secure bit commitment protocol. The proof was developed by Naor in [3].

Theorem 51 *If there exists any pseudo-random generator $G: \{0,1\}^n \rightarrow \{0,1\}^{3n}$, then there exists a bit commitment protocol.*

Bit Commitment Protocol Formal Definition

The protocol consists of two stages: the commit stage and the reveal stage. The formal definition of the bit commitment protocol follows:

- Before the protocol begins:
 1. A pseudo-random generator, $G(\cdot) : n \mapsto 3n$, known to both Alice and Bob.
 2. Alice is given (as an input), a secret bit b unknown to Bob.
- Commit Stage:
 1. Bob selects bit vector $R = \{0, 1\}^{3n}$ and sends it to Alice.
 2. Alice selects a seed $S = \{0, 1\}^n$ and computes $G(S) = Y$, where $Y = \{0, 1\}^{3n}$.
 3. Alice sends to Bob the vector $Z = \{0, 1\}^{3n}$ where $Z = Y$ if $b = 0$ and $Z = Y \oplus R$ if $b = 1$.
- Reveal Stage:
 1. Alice sends S to Bob.
 2. Bob computes $G(S)$. If $G(S) = Z$, $b = 0$; if $G(S) \oplus R = Z$, $b = 1$; otherwise, repeat the protocol.

Bit Commitment Protocol Proof of Security

In order to prove that this is a secure construction, we must prove what are known as the binding and privacy properties. The binding property requires that Alice cannot change her bit selection and that this is verifiable by Bob. The privacy property requires that Bob cannot determine any information about Alice's bit selection until Alice decommits.

Claim 52 (Binding Property) *Even if Alice has infinite computing power and memory, she cannot cheat with non-negligible probability.*

Proof For Alice to be able to cheat, she needs to find S_0 , S_1 , and Y such that $G(S_0) = Y$, and $G(S_1) = Y \oplus R$. Notice that this is equivalent to finding S_0 , S_1 such that $G(S_0) \oplus G(S_1) = R$. Alice is assumed to have infinite computing resources, so if such a pair S_0, S_1 exists she will find it. To show that Alice cannot cheat with probability greater than $\frac{1}{2^n}$, we only need to show that such a pair exists with probability less than $\frac{1}{2^n}$. There are 2^n n -bit seeds, so there are 2^{2n} pairs of seeds S_0, S_1 . Thus there are at most 2^{2n} $3n$ bit numbers of the form $G(S_0) \oplus G(S_1)$. Since R is $3n$ -bits long, there are 2^{3n} possible choices for R , if R

is chosen at random by Bob, the probability that R is of the form $G(S_0) \oplus G(S_1)$ is less than $\frac{2^{2n}}{2^{3n}} = \frac{1}{2^n}$. ■

Claim 53 (Privacy Property) *Bob cannot predict the bit, b , with probability greater than $\frac{1}{2} + \epsilon(n)$.*

Proof Suppose Bob can predict the bit b with probability $\frac{1}{2} + \epsilon(n)$, i.e. Bob can distinguish $G(S)$ from $G(S) \oplus R$ with probability $\frac{1}{2} + \epsilon(n)$. If Y is a truly random string, so is the string $Y \oplus R$, thus no matter how powerful Bob is, he cannot distinguish Y from $Y \oplus R$ with probability greater than $\frac{1}{2}$. We can use Bob to distinguish pseudo-random strings from random strings as follows: Given a string Z . We generate a random string R and give Bob Z and $Z \oplus R$. If Bob says Z is pseudo-random, we guess that Z is pseudo-random. If Bob says $Z \oplus R$ is pseudo-random, we guess Z is random. The different events with their associated probabilities can be seen in the following table.

	Bob says $Z \in PR$	Bob says $Z \oplus R \in PR$
$Z \in PR$	$\frac{1}{2} + \epsilon(n)$	$\frac{1}{2} - \epsilon(n)$
$Z \in R$	$\frac{1}{2}$	$\frac{1}{2}$

So the probability of success of this distinguisher is $\frac{1}{2} + \frac{\epsilon(n)}{2}$ which is non-negligible since $\epsilon(n)$ is non-negligible. Since G is a pseudo-random generator, no such distinguisher can exist, so we conclude that Bob cannot distinguish $G(S)$ from $G(S) \oplus R$. ■

15 Pseudo-Random Functions

Pseudo-random functions are a useful extension of pseudo-random generators. As the name suggests, pseudo-random functions create strings of seemingly random bits and, for any fixed input, will give the same string of bits every time. We will first define pseudo-random functions with more precision and then show one method of constructing them. Next, we will give a detailed proof of security for pseudo-random functions. Finally, some possible uses of pseudo-random functions will be introduced.

15.1 Definitions

We would like to define a pseudo-random function, loosely as a function that behaves randomly. Since functions are deterministic, a single function cannot be “random”. So we are forced to define pseudo-randomness only for families of functions. Let $\{U\}$ denote the uniform distribution of *all* functions from n -bits to $\ell(n)$ -bits, where $\ell(n)$ is some polynomial.

Now, we can say that a distribution of functions $\{F\}$ is pseudo-random if it “looks like” $\{U\}$. In this context, an individual function in $\{F\}$ is usually written F_S , and S is called the seed.

To make this definition (slightly) more rigorous, we begin by recalling the definition of an oracle Turing machine. An oracle Turing machine can ask a series of questions, one at a time. Each question must wait for an answer before a successive one may be asked. The questions are responded to by an oracle which has access to a single function. When the oracle Turing machine has a response from the oracle, it uses this information to come up with a binary decision. When a decision has been reached, the oracle Turing machine can make another query to the oracle. This can be repeated any polynomial number of times.

A pseudo-random function family $\{F\}$: $\{0,1\}^n \rightarrow \{0,1\}^{l(n)}$ is a collection that an oracle Turing machine cannot distinguish from the uniform distribution of function $\{U\}$ in polynomial-time.

Definition 54 *An family of functions $\{F\}$ is called pseudo-random, if $\{F\}$ is polynomial-time indistinguishable from the uniform function family $\{U\}$, i.e. For all polynomial-time machines A , for all c , there exists an N , such that for all $n > N$:*

$$\left| Pr(A^{\{F\}}(1^n) = 1) - Pr(A^{\{U\}}(1^n) = 1) \right| < \frac{1}{n^c}$$

Where $Pr(A^{\{G\}}(1^n) = 1)$ is the probability that the oracle machine A , given access to some function, $G_S \in \{G\}$, concludes $\{G\} = \{F\}$. The definition means that any *BPP* oracle Turing machine, A , has a negligible probability of distinguishing whether a given function is from $\{F\}$ or $\{U\}$.


15.2 Constructing Pseudo-Random Functions

Having defined pseudo-random functions families, we now move to their construction given a pseudo-random generator as outlined in [1].

Given a pseudo-random generator $G : \{0,1\}^n \rightarrow \{0,1\}^{2n}$ We can construct a collection of 2^n functions indexed by the seed S of G .

Given an n -bit seed S for G , we define the function F_S as follows: First, we compute $G(S) = [S_0, S_1]$, where $|S_0| = |S_1| = n$. Where S_0 and S_1 are the first and second halves of the output string, $G(S)$, respectively. Next, we compute $G(S_0) = [S_{00}, S_{01}]$ and $G(S_1) = [S_{10}, S_{11}]$. Doing this recursively n times will create a tree with height n . Note that this tree cannot be created and stored by any poly-time machine because it is exponential in size. This construction remains practical because we can traverse from the root to any leaf in polynomial time.

The function F_S takes an n -bit input string x , and we consider x as a set of directions for a path down the tree. At level $i : 1 \leq i \leq n$, we take the left branch if $x_i = 0$ and the right branch if $x_i = 1$. For example, if $x = 0100\dots$, we would begin by taking the left branch and compute $G(S_0) = [S_{00}, S_{01}]$. Then we would take the right branch and compute $G(S_{01}) = [S_{010}, S_{011}]$, followed by taking the left twice. Proceed in this manner for all i . The value $F_S(x)$ is the $2n$ -bit leaf obtained by traversing the path defined by x .



This construction gives us a set of 2^n functions from the set of all $(2^{2n})^{2^n}$ functions. We claim that if S is randomly distributed, this distribution of functions $\{F\}$ is indistinguishable from the uniform distribution $\{U\}$.

15.3 Proof of Security

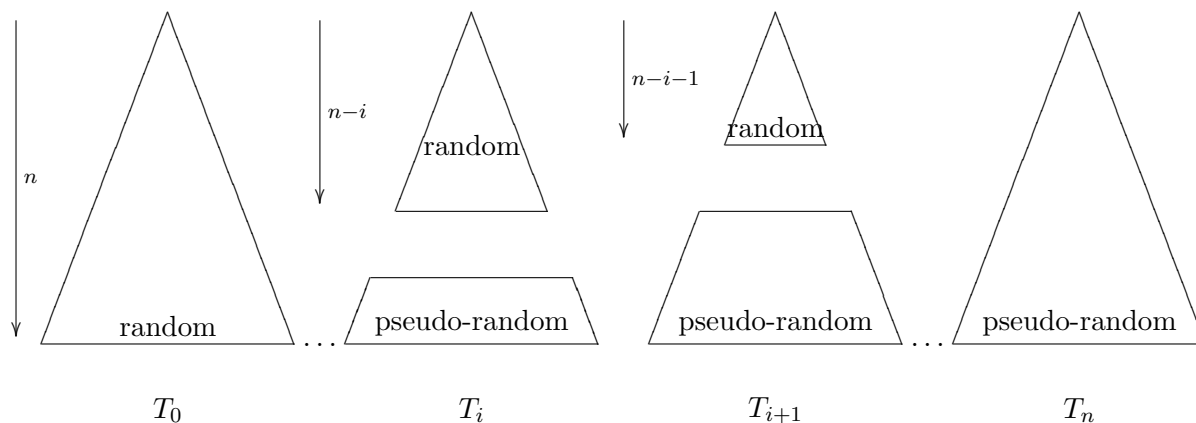
Theorem 55 *The function family $\{F_S\}$, is indistinguishable from the uniform function family $\{U\}$ of all functions taking n -bits to $2n$ -bits.*

Recall from Lecture 4, that if two distributions are indistinguishable in single-sample statistical test, then they are indistinguishable by an extended statistical test. This means that in our proof of indistinguishability of $\{F_S\}$ and $\{U\}$, we can limit our consideration to single-sample tests.

Proof Assume there exists a polynomial-time machine A that given a function F can distinguish with probability significantly greater than $\frac{1}{2}$ whether F came from the $\{F_S\}$ or $\{U\}$. We will use this machine A to distinguish the output of the pseudo-random generator G from true random, a contradiction.

We proceed by hybrid argument. Consider a collection of $n+1$ types of binary trees of depth n , all of whose nodes contain $2n$ -bit strings. The first type of tree T_0 has truly random bits at every node, the second type of tree has truly random bits at every node of depth i , $1 \leq i \leq n-1$, and the leaves of this tree are the output of our pseudo-random generator G applied to the nodes of depth $n-1$. We proceed in this way, where tree T_i has true

randomness through depth $n - i$, and pseudo-randomness thereafter, until we get to the tree T_n , which is completely pseudo-random.



By hypothesis, A can distinguish between T_0 and T_n with some probability p , thus by the triangle inequality, for some i , A can distinguish between T_i and T_{i+1} , with probability at least $p' = \frac{p}{n}$, which remains polynomial in n , if p was polynomial in n . Although we do not know the value i , since there are only n possibilities, we can try them all in polynomial time.

The distribution of the trees of type T_i and T_{i+1} differ only at depth $n - i$, where T_i is truly random, and T_{i+1} is pseudo-random. Given a collection of $2n$ -bit strings $\{Z\}$, we can tell if $\{Z\}$ is random or pseudo-random as follows: construct a tree T where T is random through depth T_{n-i-1} , the strings $\{Z\}$ are placed at depth $n - i$, and the deeper nodes are constructed pseudo-randomly, from the higher nodes. Then ask A to tell you whether this tree is of type T_i or T_{i+1} . If A says it is from T_i then you know your Z are random, and if A says it is from T_{i+1} , then your collection $\{Z\}$ is pseudo-random. There is a problem, which is that the tree T is exponential in size, so you cannot actually construct it. This is easily resolved, by simply pretending you have such a tree T , and when A asks you for a leaf, generating it and storing on the fly using your collection $\{Z\}$. Since A is a polynomial-time machine, it can only ask for polynomially many leaves of your tree T . Thus you only need polynomially many Z to generate the polynomially many paths through the tree that A requests. ■

16 Applications of Pseudo-Random Functions

Knowing how to construct pseudo-random functions, we can look at a few of their applications.

16.1 Message Authentication

One application of pseudo-random functions is in authenticating messages, i.e. authenticating the origin of the message. Alice and Bob first select a pseudo-random function family $\{F_S\}$ and share a seed S to a pseudo-random function, thus they have agreed on a specific function F_S . Alice wishes to send Bob a message m (possibly in plaintext) such that Bob can be sure m was not sent from someone who does not know S .

The protocol can be summarized as follows

1. Alice and Bob share a seed S , this is their secret key which determines a pseudo-random function $F_S : n\text{-bits} \rightarrow n\text{-bits}$.
2. To sign an n -bit message m , Alice sends Bob, $(m, F_S(m))$.
3. To verify that a message (a, b) came from Alice, Bob verifies that $F_S(a) = b$.

A third party, Eve does not know, S , and so she cannot calculate $F_S(m')$ for any message m' .

16.2 Identification Schemes

To prove your identity to a server, it is customary to use a password. This system has the flaw that if you send your password unencrypted (e.g. using telnet) and some is eavesdropping on your communication they will have your password, and will be able to pose as you in the future. If you and the server were able to share a secret S , then using a pseudo-random function F_S , you will be able to send passwords unencrypted with no risk of identity-theft. The scheme works as follows

1. Server and client share a seed S , this is their secret key which determines a pseudo-random function $F_S : n\text{-bits} \rightarrow n\text{-bits}$.
2. For the client to authenticate himself to the server, the server first chooses a random string x , and sends it to the client.
3. The client computes $F_S(x)$, and sends it to the server.

4. Since the server has the seed S , it can easily verify whether the client calculated $F_S(x)$ correctly.

This scheme is secure because F_S is a pseudo-random function, so even if an eavesdropper intercepts a polynomial number of $x, F_S(x)$ pairs, he still cannot calculate $F_S(y)$ for any new value y .

16.3 Private Key Encryption

Using pseudo-random functions, we can design a private key encryption scheme.

The protocol can be summarized as follows

1. Alice and Bob share a seed S , this is their secret key which determines a pseudo-random function $F_S : n\text{-bits} \rightarrow n\text{-bits}$.
2. To encrypt an n -bit message m , Alice chooses a random n -bit string i , and sends $(i, m \oplus F_S(i))$.
3. After receiving (a, b) , to decrypt, Bob calculates $F_S(a) \oplus b$.

If the function F were selected truly randomly from the distribution of all functions, then $F(a)$ would be truly random, and no adversary, no matter how powerful could recover b from $F(a) \oplus b$. If some adversary could recover b from $F_S(a) \oplus b$, this would give us a way of distinguishing random functions from our pseudo-random functions.

Lecture 5

17 Digital Signatures

17.1 Introduction

Say⁸ Bob wants to communicate with Alice over a channel in which Eve can intercept and transmit messages. We have already considered the problem of message security - how Bob can encrypt a message for Alice such that Eve cannot decipher it in polynomial time with non-negligible probability.

Now we consider a different problem: if Bob receives a message purportedly from Alice, how can he be certain it isn't a forgery by Eve? This is particularly true if Bob has published a public key, and thus anyone can send him encrypted messages. One solution is to have Alice "sign" the message in some way that Bob will recognize, and that Eve will be unable to forge. There are a number of properties we would ideally like for such a signature:

- Alice can efficiently sign any message, for some reasonable limit on the message size.
- Given any document D that Alice has not signed, nobody can efficiently forge Alice's signature on D .
- Given a document D and a signature, anyone (not just Bob!) can efficiently tell whether the signature is valid for D .

We introduce digital signatures schemes as a way of accomplishing this.

17.2 Digital Signatures

We now give a definition of a digital signature scheme. A *digital signature scheme* is a triple of poly-time computable algorithms (KeyGen, Sign, Verify) over a message space \mathcal{M} that satisfy the following conditions:

1. KeyGen($1^n, R$) is a probabilistic (with coin flips R) poly-time algorithm that outputs a public key and a secret key pair, (PK, SK)

⁸Scribe notes by Kamil Azhar; April 17th, 2006

Alice has a public key PK

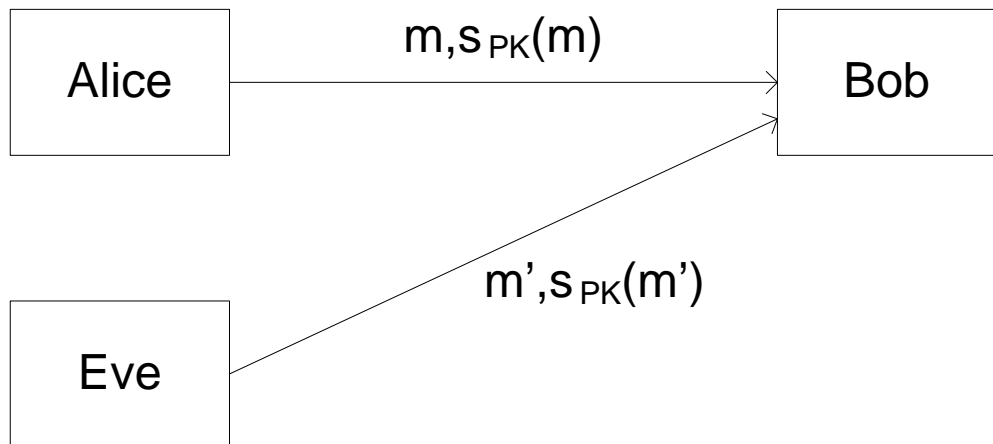


Figure 13: Eve tries to forge a signature.

2. $\text{Sign}(D, PK, SK, R)$ is a probabilistic (with coin flips R) poly-time algorithm that signs a document $D \in \mathcal{M}$ with a signature $\sigma(D)$. Note: $|\sigma(D)|$ should be polynomially related to $|D|$
3. $\text{Verify}(PK, D, s)$ is a (probabilistic) poly-time algorithm that outputs an element of $\{\text{Yes}, \text{No}\}$. It returns **Yes** (with negligible error) if s is a valid signature of D , i.e. $s = \sigma(D)$

Given such a scheme, Alice can set up her document signer. First she generates $(PK, SK) \leftarrow \text{KeyGen}(1^n, R)$ and publishes PK while keeping SK secret. Then when she wants to sign a document, D , she can run the signing algorithm $\sigma(D) \leftarrow \text{Sign}(D, PK, SK, R)$ and sends the pair $(D, \sigma(D))$ to Bob. Bob can then verify the signature by running $\text{Verify}(PK, D, \sigma(D))$. Two important properties of a digital signature scheme are correctness and privacy. By correctness, we mean that if PK and SK are generated according to KeyGen , then for all $D \in \mathcal{M}$, $\text{Verify}(PK, D, \sigma(D)) = \text{Yes}$. Hence, a legally signed message can never be rejected. Privacy is discussed in detail in the next section on security. We also say that Eve forges a signature if she can produce a D and a $\sigma(D)$ (that was not signed by Alice) such that $\text{Verify}(PK, D, \sigma(D)) = \text{Yes}$ with non-negligible probability.

17.3 Security of a Digital Signature Scheme

When we talk about security for a digital signature scheme, we consider an adversary, Eve, who attempts to send a message to Bob and try to forge Alice's signature. What possible information does Eve have access to before attacking the system? Here are some reasonable assumptions that have been proposed:

- Eve only knows the public key. (**Key-only Attack**)
- Eve has seen a set of messages $\{m_1, \dots, m_k\}$ with their corresponding signatures. The set of messages is given to her but not chosen by her. (**Known Message Attack**)
- Eve chooses a fixed set of messages $\{m_1, \dots, m_k\}$ (there are two cases, where the messages are chosen independently of the public key or not) and gets to see the signatures of those messages. (**Chosen Message Attack**)

Once this information is given to her, what does it mean for the signature scheme to be broken by Eve?

- Eve computes the secret key. This is as bad as it gets because now Eve can sign *any* message she wants. (**Total Break**)
- Eve computes a poly-time algorithm that can forge a signature for any message. (**Universal Forgery**)
- Eve can forge a signature for a particular message of her choice. (**Selective Forgery**)

The previous examples will ultimately motivate our definition of security for signature schemes.

17.4 Security over multiple messages

Consider the following scheme proposed by Diffie and Hellman: let (Gen, f, f^{-1}) be a one-way permutation, where Gen outputs a (public) key k and a (private) trapdoor td . Given a document D , Sign gives the signature $\sigma(D) = f_k^{-1}(D)$. To verify a signature, any party can compute whether $f_k(\sigma) = D$. Therefore this indeed defines a legitimate signature scheme. But how secure is this scheme?

It turns out that the scheme is not secure. An adversary can pick a random σ and define D to be $f_k(\sigma)$. Note that D could be completely meaningless as a document, but nevertheless the adversary has given a new document and a valid signature. Message independence requires that no one should be able to forge random documents. Also a forger cannot sign a new document regardless of the distribution of messages.

This motivates the following definition of security over multiple messages, due to Goldwasser, Micali and Rivest:

Definition 56 *A digital signature scheme is existentially unforgeable under an adaptive chosen-message attack if for all $A \in PPT$ who is allowed to query $Sign$ polynomially many times (such messages may be dependent on both previously chosen messages and the public key) cannot forge any new message.*

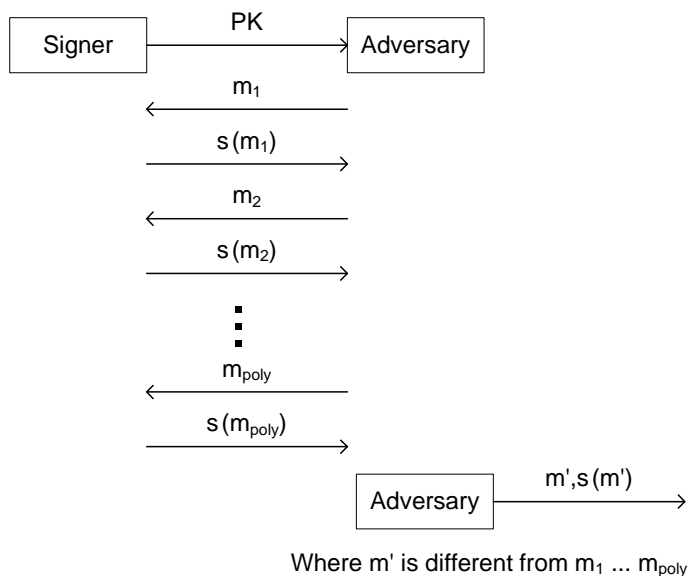


Figure 14: An Adaptive Chosen Message Attack

So we cant use documents as input of 1-way functions, since the distribution will have certain outputs which are easy to invert. Also as seen above, random, meaningless documents can be signed. Thus the above scheme is not secure according to our defintion of security, which is very robust.

17.5 Lamport's 1-time signature scheme

The following is a construction of a one-time signature scheme (for messages of length n) out of a one-way permutation, $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$.

1. $\text{KeyGen}(1^n, R)$ will randomly select $2n$ elements from $\{0, 1\}^n$. We will label them $x_1^0, x_1^1, x_2^0, x_2^1, \dots, x_n^0, x_n^1$. Then we compute $y_i^b = f(x_i^b)$ for all $1 \leq i \leq n$ and $b \in \{0, 1\}$. The secret key SK and the public key PK will be

$$SK = \begin{pmatrix} x_1^0 & x_2^0 & \dots & x_n^0 \\ x_1^1 & x_2^1 & \dots & x_n^1 \end{pmatrix} \quad PK = \begin{pmatrix} y_1^0 & y_2^0 & \dots & y_n^0 \\ y_1^1 & y_2^1 & \dots & y_n^1 \end{pmatrix}$$

2. $\text{Sign}(m, PK, SK, R)$ will use $m = m_1 m_2 \dots m_n$ to index SK meaning it will return $(x_1^{m_1}, x_2^{m_2}, \dots, x_n^{m_n})$. For example if $m = 100 \dots 1$ then we will return the selected entries:

$$\sigma(m) = \begin{pmatrix} x_1^0 & \boxed{x_2^0} & \boxed{x_3^0} & \dots & x_n^0 \\ \boxed{x_1^1} & x_2^1 & x_3^1 & \dots & \boxed{x_n^1} \end{pmatrix}$$

3. $\text{Verify}(PK, m, s)$ where $m = m_1 m_2 \dots m_n$ and $s = (s_1, s_2, \dots, s_n)$ checks that $f(s_i) = y_i^{m_i}$ and returns **Yes** if they are equal for all $1 \leq i \leq n$. Continuing our previous example where $m = 100 \dots 1$ we check the equalities in the selected entries:

$$s = (s_1, s_2, \dots, s_n) \quad \begin{pmatrix} y_1^0 & \boxed{f(s_2) = y_2^0} & \boxed{f(s_3) = y_3^0} & \dots & y_n^0 \\ \boxed{f(s_1) = y_1^1} & y_2^1 & y_3^1 & \dots & \boxed{f(s_n) = y_n^1} \end{pmatrix}$$

Claim Assuming f is a one-way permutation, this scheme is secure against an adversary that is allowed only *one* query for the signature of a message m of his choice, then has to come up with $m' \neq m$ and a forgery $\sigma(m')$.

Proof We shall prove the contrapositive of the claim. We start by assuming there is an adversary A that can forge signatures with non-negligible probability, then we show that A can be used to invert f with non-negligible probability. Suppose A can forge a signature with probability $> \epsilon$ conditioned over all m and PK (because f is a permutation, PK simply has a uniform distribution). Then we construct an algorithm to invert y as follows:

$(PK, SK) \leftarrow \text{KeyGen}(1^n, R)$
 $i' \leftarrow \{1, \dots, n\}; b' \leftarrow \{0, 1\}$
 Replace $y_{i'}^{b'}$ in PK by y .
 Give A this new PK , and he will request a signature for $m = m_1 m_2 \dots m_n$
if $m_{i'} = b'$ **then** FAIL; **else** send A the correct signature
 A will then output a forged signature (s_1, s_2, \dots, s_n) for a different message m'
Output if $m'_{i'} = b'$ **then** FAIL; **else** return $s_{i'}$

Notice y is uniformly distributed because f is a permutation, so the modified PK s look like they are also uniformly distributed, which means that the adversary will invert with the same probability $> \varepsilon$. The first place our algorithm can fail is if the adversary asks us to sign a message that has b' as its i' -th bit because we do not know $f^{-1}(y)$. This occurs with probability $\frac{1}{2}$. The second place our algorithm can fail is if the message generated by our adversary did *not* pick out b' as its i' -th bit, which means that we did not get the inverse to y from him. Because $m' \neq m$, it must differ by at least one bit, which means the chance that it differs on the i' -th bit is $\frac{1}{n}$.

Provided that our algorithm did not fail, then the answer it returns will be $x = s_{i'}$. Because s is a valid signature for m' , it must be the case that $f(x) = f(s_{i'}) = y_{i'}^{m_{i'}} = y_{i'}^{b'} = y$. Thus the algorithm has succeeded in inverting y . Combining all the probabilities, we have a probability $\frac{1}{2} \cdot \frac{1}{n} \cdot \varepsilon$ of inverting f . Thus, f cannot be one-way, which proves the contrapositive. ■

Remark This scheme is only secure for signing one message, because the signature reveals part of your secret key. For example, if you signed $00 \dots 0$ and $11 \dots 1$, then your *entire* secret key has been revealed.

Some drawbacks of this scheme are that the public key has size $2n^2$ and that it can only be used to securely sign *one* message. In the next section we will construct a scheme that can sign many messages.

17.6 Signing multiple messages

We saw a secure scheme that could sign one message, and from this we can build a scheme that can sign many messages. All the schemes we will present in this section will require the signer to save a “state” based on how many messages have been already signed. Under the formal definition of a digital signature scheme this is not allowed, but we will relax this condition. Secure, stateless signature schemes have been first achieved by Goldreich in 1986.

One way we can sign N messages is to generate N secret key and public key pairs $\{(PK_0, SK_0), \dots, (PK_N, SK_N)\}$ under the Lamport 1-time signature scheme. Then to

sign the i -th message, one can simply sign it using (PK_i, SK_i) under the Lamport scheme. This way of signing multiple messages is highly impractical because our public key is unreasonably huge and we need to know a priori how many messages to sign.

By introducing hash functions into our schemes, we can make a great deal of improvement on the lengths of messages we can sign, and how many we can sign. For example, one may first hash a message using a collision resistant hash function and then sign the hashed document. Merkle in 1989 constructed a multi-time signature scheme using trees by “hashing-down” public keys to a single root which will be the master public key.

17.7 Introduction to Hash Functions

We want to define the notion of a hash function. Intuitively, it should be a function that is easy to compute, and the output should be shorter in length than the input.

Definition 57 A hash function $h : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is a deterministic, poly-time computable function such that $n > m$.

Because h is length decreasing, the function is many-to-one, i.e. there exists a pair (x, x') such that $h(x) = h(x')$. In cryptographic constructions using a family of hash functions we want to have control over how these collisions occur. We consider a family of hash functions defined by a key-gen algorithm $\text{Gen}(1^n, R) \rightarrow h$ where R is some randomness. We define the following three notions of collision resistance on the family, \mathcal{H} , of hash functions generated by $\text{Gen}(1^n, R)$:

Definition 58 (Collision resistance) $(\forall A \in PPT)(\forall c > 0)(\exists N_c)(\forall n \geq N_c)$

$$\Pr_{H,R}[h \leftarrow H_n; (x_n, y_n) \leftarrow A(h) : h(x_n) = h(y_n)] \leq \frac{1}{n^c}.$$

Definition 59 (Universal one-way [Naor-Yung]) $(\forall A \in PPT)(\forall c > 0)(\exists N_c)(\forall n \geq N_c)$

$$\Pr_{H,R}[(x_n, \alpha) \leftarrow A(1^n); h \leftarrow H_n; y_n \leftarrow A(\alpha, h, x_n) : h(x_n) = h(y_n)] \leq \frac{1}{n^c}.$$

Definition 60 (Two-wise independence) $(\forall A \in PPT)(\forall c > 0)(\exists N_c)(\forall n \geq N_c)$

$$\Pr_{H,R}[(x_n, y_n) \leftarrow A(1^n); h \leftarrow H_n : h(x_n) = h(y_n)] \leq \frac{1}{n^c}.$$

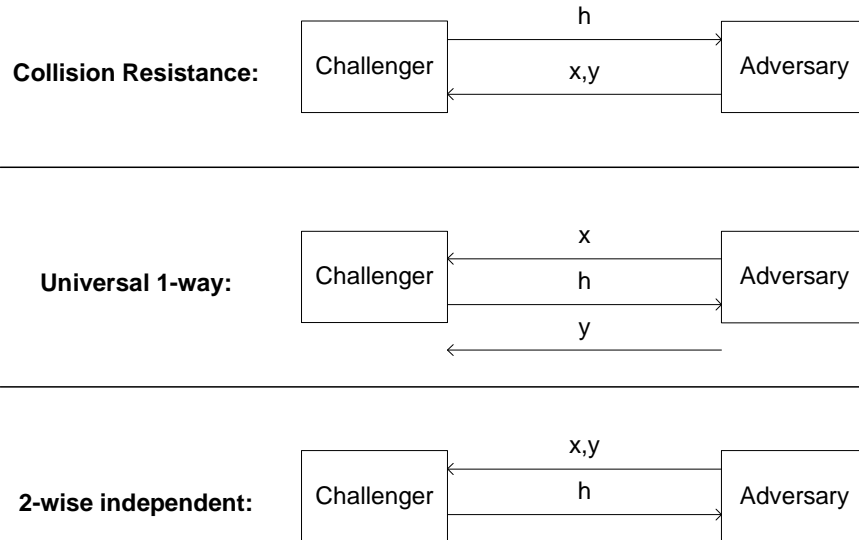


Figure 15: Three games of collision resistance.

In collision resistant hash functions, h is chosen at random, and the adversary must come up with x and y , such that $h(x)=h(y)$. In Universal one-way hash functions, the adversary chooses an x , and the challenger selects h , and the adversary must select y , such that $h(x)=h(y)$.

17.8 Naor-Yung trees

We will present an improvement of Merkle's signature scheme based on the work of Naor-Yung. The construction assumes the one-time security of Lamport's 1-time signature, $(\text{KeyGen}, \text{Sign}, \text{Verify})$ as well as the existence of a family of universal 1-way hash functions $\{h : \{0, 1\}^{4n^2} \rightarrow \{0, 1\}^n\}$ which will be used in the construction of our signature scheme which signs $n^{\log(n)}$ messages of length n . The main concept for signing is to first pretend we have a complete tree of height, say, $k = \log^2(n)$ which will have a secret key and a public key at each node. Only the public key PK of the root of this tree will be published, which means our public key size is independent of the number of messages we need to sign. To sign the i -th message, we sign it on the i -th leaf using Lamport's 1-time signature, then include the public keys of the nodes in the path from the leaf to the root and their siblings. To make sure this path is authentic, we also need to have each parent sign the public keys of its two children. This is accomplished by concatenating the public keys of the two children then applying a hash to it, then signing the result. All of this information is to be included

in the signature, but the good news is that the size of the signature does not grow as the number of signed messages increases. Notice that because we only pretended to have such a tree, some of these values may need to be computed and stored on the fly, but still this only takes poly-time to accomplish. Also notice that because our tree has more than polynomially many leaves, no polynomially bounded adversary can exhaust all of the leaves, so that we can always sign a message when an adversary asks for one.

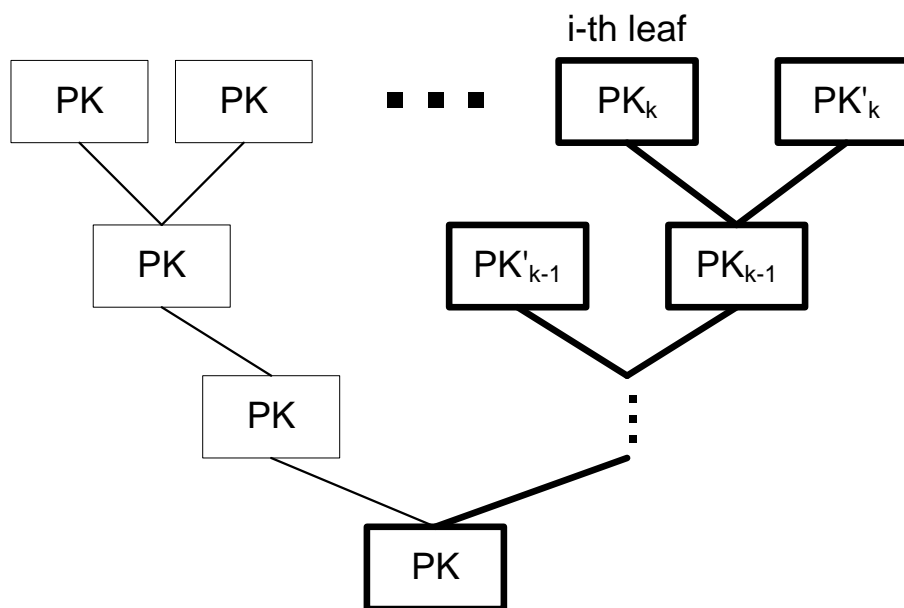


Figure 16: Signing the i -th message.

More formally, we can define a triple $(\text{MKeyGen}, \text{MSign}, \text{MVerify})$ with state information as follows:

- $\text{MKeyGen}(1^n, R)$ will generate the keys for the root of our tree from the 1-time generator $(PK, SK) \leftarrow \text{KeyGen}(1^n, R)$
- $\text{MSign}(m, PK, SK, R)$ will use state information to sign a message. To sign the i -th message, m , we first set up some notation. Let n_k^j denote the j -th node (reading the tree from left to right) of depth $k = \log^2(n)$. Then by this notation $n_{k-1}^{j/2}$ denotes its parent, $n_{k-2}^{j/4}$ denotes its grandparent, and so on. Let $n_k = n_k^i, n_{k-1} = n_{k-1}^{i/2}, \dots$ and let $r = n_1$ denote the root of our tree. Let (PK_ℓ, SK_ℓ) be the keys corresponding

to node n_ℓ for $1 \leq \ell \leq k$ with $(PK_1, SK_1) = (PK, SK)$. Also, let (PK'_ℓ, SK'_ℓ) denote the keys for the sibling of n_ℓ for $2 \leq \ell \leq k$. Finally, we compute $\sigma \leftarrow \text{Sign}(m, PK_k, SK_k, R)$ and $\sigma_\ell \leftarrow \text{Sign}(h(PK_{\ell+1}PK'_{\ell+1}), PK_\ell, SK_\ell, R)$ for $1 \leq \ell \leq k-1$, and output $(PK_\ell, PK'_\ell, \sigma_\ell, \sigma)$ as our signature for m .

- $\text{MVerify}(PK, m, s)$ will first check that s is of the form $s = (PK_\ell, PK'_\ell, \sigma_\ell, \sigma)$. Then it will run $\text{Verify}(PK_k, m, \sigma)$ and $\text{Verify}(PK_\ell, h(PK_{\ell+1}PK'_{\ell+1}))$ and return **Yes** if they both pass.

We mention as a side note that the space requirements can be reduced to a constant if one uses pseudorandom functions to generate the public keys. Because pseudorandom outputs are indistinguishable from a uniform distribution, such a construction is equally secure.

This (stateful) digital signature scheme we constructed is existentially unforgeable adaptively secure if the Lamport 1-time scheme is secure and universal one-way hash functions exist. The sketch of the proof is as follows:

Sketch of Proof Assume for the contrapositive that there exists a poly-time adversary A that can forge a signature with non-negligible probability, ε after $p = \text{poly}(n)$ steps. Because there is a path of signatures from each leaf down to the root, two cases can occur in a forgery (1) A found a collision to h or (2) A can sign a message on an existing leaf or a message different from $h(PK_\ell PK'_\ell)$ on an existing node. One of the two cases occur with probability at least $\varepsilon/2$ so either we can show h is not universal 1-way, or we can show f is not one-way, which proves the contrapositive. ■

We summarize this result as showing secure digital signatures exist if one-way permutations exist (for the Lamport scheme to work), and universal 1-way hash functions exist. In the next section we will show how to construct a universal 1-way hash function from a one-way permutation.

17.9 One-way Permutations Imply Universal 1-way Hash Functions

We will construct a family of universal 1-way hash functions from a one-way permutation f . The hash functions we construct will take n bits to $n-1$ bits and they will be indexed by $h = (a, b)$ where $a, b \in GF(2^n)$. The algorithm for hashing a string x of length n is to apply $y \leftarrow f(x)$ then $z \leftarrow \text{chop}(ay + b)$ to $n-1$ bits (operations are taken over $GF(2^n)$). By the fact that f and the linear map $ay + b$ are both 1-1 and chop is 2-1, our hash function is a 2-1 mapping from $\{0, 1\}^n \rightarrow \{0, 1\}^{n-1}$. We claim that this is a family of universal 1-way hash functions if f is a one-way permutation.

Proof Assume for the contrapositive that this family is not universal 1-way. Let A be a poly-time adversary that chooses x and is given h chosen from a uniform distribution can

find a x' such that $h(x) = h(x')$ with probability $> \varepsilon$. Then to invert $y' = f(\gamma)$, we first look at x and we solve for (a, b) to satisfy the equation $chop(af(x)+b) = chop(af(\gamma)+b)$. Because f is a permutation, and the fact that this linear equation does not skew the distribution of the (a, b) returned, the hash function $h = (a, b)$ looks as if it were chosen truly from a uniform distribution. Then A will return x' such that $chop(af(x) + b) = chop(af(x') + b)$, but $f(x)$ and $f(\gamma)$ are the only two solutions to that linear equation, so $f(x') = f(\gamma) = y'$. Thus we can invert f with probability $> \varepsilon$, proving that it is not one-way. ■




Figure 17: Setting a trap.

Assuming that there exists a family of one-way permutations: $f_{poly(n)}, f_{poly(n)-1}, \dots, f_{n+1}$, such that $f_i : \{0, 1\}^i \rightarrow \{0, 1\}^i$, we may construct a family of universal 1-way hash functions $h : \{0, 1\}^{poly(n)} \rightarrow \{0, 1\}^n$ as follows:

```

 $x_0 \leftarrow \{0, 1\}^{poly(n)}$ 
for  $k = 0$  to  $poly(n) - n - 1$ 
     $a_k \leftarrow GF(2^{poly(n)-k}); b_k \leftarrow GF(2^{poly(n)-k})$ 
     $x_{k+1} \leftarrow chop(a_k f_{poly(n)-k}(x_k) + b_k)$ 
Output  $x_{poly(n)-n}$ 

```

Proof Idea

Assume that there was a poly-time adversary A who could break this construction with probability $> \varepsilon$. Then we can pick a random location in our chain to set a trap by solving

for (a_k, b_k) at that level as before. If the adversary finds a collision, then at some level in our construction there will be a collision. The probability of that level being equal to the one on which we set the trap is $\frac{1}{poly}$. This will allow for us to invert the one-way permutation at that location with probability $> \frac{\epsilon}{poly}$, thus contradicting the one-way property of the function.

■

Thus we have shown the following theorem:

Theorem 61 *Hash functions exist if one-way permutations exist.*

17.10 Application to Signatures

Assuming that there exists a family of one-way permutations: $f_{4n^2}, f_{4n^2-1}, \dots, f_{n+1}$, such that $f_i : \{0, 1\}^i \rightarrow \{0, 1\}^i$, we may construct a family of universal 1-way hash functions $h : \{0, 1\}^{4n^2} \rightarrow \{0, 1\}^n$ as seen in the previous section. This gives the following result:

Theorem 62 (Naor-Yung) *Secure digital signatures exist if one-way permutations exist.*

To conclude on this topic, we mention that it has also been shown:

Theorem 63 (Rompel) *Secure digital signatures exist iff any one-way function exists.*

17.11 Hash functions for Commitment Schemes

Consider the scheme where a sender is given a hash function $h : \{0, 1\}^{2n} \rightarrow \{0, 1\}^n$ by a receiver. Let h be applied on every $2n$ bits x , and then iteratively on every $2n$ bits of the result producing the final output a .

Claim 64 *The output a is a good commitment for x .*

Sketch of Proof The receiver gets no information on x which follows from the way we constructed our hash function. On the other hand, the sender cannot cheat because in order to cheat, he would be able to find a collision for the hash function h . ■

Lecture 6


18 Oneway Trapdoor Permutations

Recall⁹ that a oneway function, f , is easy to compute, but hard to invert. Formally, for all PPT adversaries A , there is a c such that for eventually all n ,

$$\Pr [A(f(x)) \in f^{-1}f(x)] < \frac{1}{n^c}$$

with the probability taken over $|x| = n$ and coin flips of A .

A oneway, trapdoor function is a oneway function f , which becomes easy to invert when given some extra information, t , called a trapdoor.



We formalize this as follows.

Definition 65 A oneway trapdoor function is a parameterized family of functions $\{f_k : D_k \rightarrow R_k\}_{k \in K}$ with K , D_k , and $R_k \subseteq \{0, 1\}^*$.

1. Key, trapdoor pairs are PPT sampleable: there is a polynomial p and PPT algorithm **GEN** such that $\mathbf{GEN}(1^n) = (k, t_k)$, with $k \in K \cap \{0, 1\}^n$, and $|t_k| \leq p(n)$. Call k a key, and t_k the trapdoor for f_k .
2. Given k , the domain D_k is PPT sampleable.
3. f_k^{-1} is computable, given a trapdoor t_k : there is an algorithm I , such that $I(k, t_k, f_k(x)) = x$, for $x \in D_k$.
4. For all PPT A , the following is negligible:

$$\Pr [A(k, f_k(x)) \in f_k^{-1}f_k(x)]$$

where k is sampled by **GEN**, and the asymptotics are relative to the security parameter.

⁹Scribe notes by V.Goyal, M.Chov, K.Leung, J.Salomone; April 2006

In this definition, (1) is saying that we can randomly generate a function from the family, and its trapdoor. The size of the trapdoor information must be polynomial in the size of the key. (3) says that an instance f_k is invertible, given its description k , and its trapdoor t_k . (4) says that $\{f_k\}$ is a oneway family. For clarity, we will often let the k be implied, and write (f, f^{-1}) , instead of (k, t_k) .

Note that it is important to use a family of functions. If we try to make the above definition for a single function, (4) will fail. There is always some adversary A , with a description of the trapdoor t , and the inverter I , “hard-wired” into its description. This adversary will always be able to invert.

19 Public Key Encryption

A public key encryption scheme (say, for entity A) consists of three algorithms, **KEYGEN** for key generation, and **ENC** and **DEC**, for encryption and decryption, respectively. Given a security parameter, 1^n , **KEYGEN** should return two keys, PK and SK . The idea is that PK is made public, and is used by any other entity B , as input to **ENC**, to encrypt a message for A . SK is kept secret by A , and is used in **DEC** to decrypt a ciphertext, and recover the original message. We will define the semantic security of this system so that no adversary E can recover the message, even with knowledge of the public key PK .

$$(PK, SK) \leftarrow \mathbf{KEYGEN}(1^n, r)$$

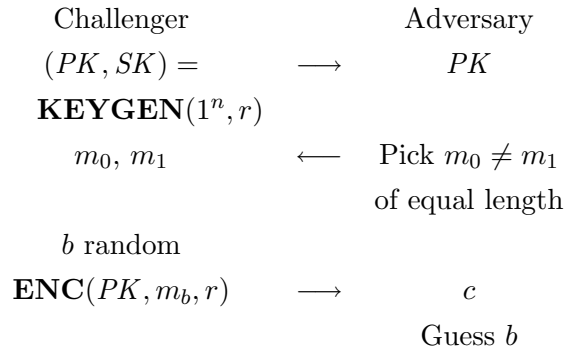
$$c \leftarrow \mathbf{ENC}(PK, m, r)$$

$$m' \leftarrow \mathbf{DEC}(PK, SK, c)$$

Of course, in the above procedure, we want $m' = m$, so that we recover the original message m . We demand that the scheme be *correct*: if (PK, SK) is generated by **KEYGEN**, then for all messages m ,

$$\mathbf{DEC}(PK, SK, \mathbf{ENC}(PK, m, r)) = m. \tag{12}$$

To define semantic security, consider the following game. Challenger uses **KEYGEN** to generate a key pair (PK, SK) and publishes PK . Adversary, given PK , picks distinct messages m_0 and m_1 , of equal length, and sends them to Challenger. Challenger picks a random bit b , and then sends to Adversary the ciphertext $c = \mathbf{ENC}(PK, m_b)$.



We say that the cryptosystem is secure if Adversary can then guess b with probability which deviates only negligibly from $\frac{1}{2}$.

Remark For this definition to work, we needed **ENC** to be probabilistic. Otherwise, Adversary could simply compute **ENC**(m_0) and **ENC**(m_1), and compare them to c , thus determining b .

19.1 Example: PK Cryptosystem from Oneway Trapdoor Permutations

A semantically secure, public key cryptosystem can be constructed from a oneway, trapdoor permutation. The algorithms are as follows¹⁰.

KEYGEN($1^n, r$):

1. **compute** $(f, f^{-1}) := \mathbf{GEN}(1^n)$.
2. Pick a string p , uniformly at random, for computing hard-core bits.
3. **return** $PK = (f, p), SK = f^{-1}$.

Encryption and decryption are done bit-wise on the plaintext and ciphertext.

ENC($((f, p), m, r)$):

1. Pick x at random from the domain of f .
2. **compute** $c := (p \cdot x) \oplus m$.
3. **compute** $d := f(x)$.
4. **return** ciphertext (c, d) .

¹⁰Recall that $p \cdot x = \bigoplus_{1 \leq i \leq n} p[i]x[i]$, where $|p| = |x| = n$.

DEC $((f, p), f^{-1}, (c, d))$:

1. **compute** $x := f^{-1}(d)$.
2. **compute** $m := (p \cdot x) \oplus c$.
3. **return** m .

Clearly this cryptosystem is correct; it is also semantically secure. If an adversary could distinguish two messages m_0 and m_1 , then by a hybrid argument, it could distinguish two messages m'_0 and m'_1 , which differ in only one bit. We could then use this adversary to compute hard-core bit, $p \cdot x$, knowing only $f_s(x)$.

19.2 Stronger Definitions of Security for Public Key Encryption: Security against Chosen Ciphertext attacks

CCA-1 Security

This is also called security against Lunch time attacks. CCA-1 security is defined by the following game.

KenGen The challenger generates a key pair $(PK, SK) \leftarrow \text{KenGen}(1^k)$ and declares the public key PK .

Decryption The adversary queries the challenger with a ciphertext and gets back its decryption with the key SK . The adversary is allowed to do this for polynomially many adaptively chosen ciphertexts.

Challenge The adversary submits two equal length messages m_0 and m_1 with $m_0 \neq m_1$. The challenger flips a random coin b , and encrypts m_b with PK . The ciphertext is passed to the adversary.

Guess The adversary outputs a guess b' of b .

Adversary wins the above game if its guess b' of the challenger bit b is correct. The advantage of an adversary \mathcal{A} in this game is defined as $\Pr[b' = b] - \frac{1}{2}$.

Definition 66 *A public key encryption scheme is secure in the CCA-1 model of security if all polynomial time adversaries have at most a negligible advantage in the above CCA-1 game.*

CCA-2 Security

CCA-2 security game is similar to the CCA-1 security game except for an additional decryption phase after the Challenge phase. The game is as follows.

KeyGen The challenger generates a key pair $(PK, SK) \leftarrow \text{KeyGen}(1^k)$ and declares the public key PK .

Decryption Phase 1 The adversary queries the challenger with a ciphertext and gets back its decryption with the key SK . The adversary is allowed to do this for polynomially many adaptively chosen ciphertexts.

Challenge The adversary submits two equal length messages m_0 and m_1 with $m_0 \neq m_1$. The challenger flips a random coin b , and encrypts m_b with PK . The ciphertext is passed to the adversary.

Decryption Phase 2 Same as decryption phase 1.

Guess The adversary outputs a guess b' of b .

Adversary wins the above game if its guess b' of the challenger bit b is correct. The advantage of an adversary \mathcal{A} in this game is defined as $\Pr[b' = b] - \frac{1}{2}$.

Definition 67 A public key encryption scheme is secure in the CCA-2 model of security if all polynomial time adversaries have at most a negligible advantage in the above CCA-2 game.

20 Some Cryptographic Assumptions

20.1 Finite, Abelian Groups

Recall that an abelian group is a collection of elements G , with a binary operation \star on G , satisfying:

$$(\forall a, b, c \in G) (a \star b) \star c = a \star (b \star c) \text{ (Associativity)}$$

$$(\forall a, b \in G) a \star b = b \star a \text{ (Commutativity)}$$

$$(\exists 1 \in G) (\forall a \in G) 1 \star a = a$$

$$(\forall a \in G) (\exists a^{-1} \in G) a \star a^{-1} = e$$

Call 1 the identity element of G , and a^{-1} the inverse of a . The order of a finite group G is the number of elements in the group, denoted $|G|$. A useful fact is that if $|G| = n$ then for any element a , $a^n = 1$.

We will usually be concerned with a specific type of abelian group: Call $g \in G$ a *generator* iff $G = \{g^n \mid 0 \leq n < |G|\}$. In case G has a generator, say that G is *cyclic*, and write $G = \langle g \rangle$.

We wish to generate finite, cyclic groups randomly. Fix a PPT algorithm **GROUP**, which samples a finite, cyclic group, given a security parameter 1^n . In other words, if

$$(G, p, g) \leftarrow \mathbf{GROUP}(1^n),$$

then G is a (binary description of a) finite group, $p = |G|$, and g is a generator.

20.2 Discrete Logarithm Problem

Suppose we are given a cyclic group G , of order p , with generator g , and a group element $a \in G$. The Discrete Logarithm Problem is to find an integer k , such that $g^k = a$. In other words, to compute $k = \log_g(a)$. The Discrete Logarithm Assumption says that this is computationally hard.

Assumption 68 (DLA) For any PPT algorithm A

$$\Pr \left[g^k = a : (G, p, g) \leftarrow \mathbf{GROUP}(1^n); a \xleftarrow{R} G; k \leftarrow A(G, p, g, a) \right]$$

is negligible in n .

Many financial transactions are done using a **GROUP** which returns $G = \mathbb{Z}_p$ for p a prime.

20.3 Decisional Diffie-Hellman Problem

The Decisional Diffie-Hellman Problem is similar to the Discrete Log Problem, except that one tries to distinguish powers of a generator, rather than trying to compute a log. Suppose we are given a group G , of order p , with generator g . Then integers $x, y, z \in \mathbb{Z}_p^*$ are selected randomly. From this, two sequences are computed:

$$\langle G, p, g, g^x, g^y, g^z \rangle (\text{Random sequence})$$

$$\langle G, p, g, g^x, g^y, g^{xy} \rangle (\text{DDH sequence})$$

The DDH problem is to determine which sequence, Random or DDH, we have been given. The DDH Assumption is that the DDH Problem is hard.

Assumption 69 (Decisional Diffie-Hellman) Let G be a sampled group of order p , with generator g . Pick $x, y, z \in \mathbb{Z}_p^*$ uniformly at random. Then it is asymptotically difficult (with respect to the security parameter), for a PPT adversary A to distinguish $(G, p, g, g^x, g^y, g^{xy})$ from (G, p, g, g^x, g^y, g^z) .

Remark The DDH assumption is stronger than the DLP assumption. Computing discrete logarithms would allow one to trivially distinguish g^{xy} from g^z , for a random z .

21 The ElGamal Public Key Cryptosystem

The security of the ElGamal cryptosystem is based on the difficulty of the DDH problem. The algorithms are:

KEY(1^n):

1. **compute** $(G, p, g) := \mathbf{GROUP}(1^n)$.
2. Sample $x \in \mathbb{Z}_p^*$, uniformly at random.
3. **compute** $w := g^x$.
4. **return** $PK = (G, p, g, w)$, $SK = x$.

ENC((G, p, g, w), m) (for $m \in G$):

1. Sample $r \xleftarrow{R} \mathbb{Z}_p^*$.
2. **compute** $c := w^r m$, $d := g^r$.
3. **return** ciphertext (c, d) .

DEC((G, p, g, w), x , (c, d)):

1. **compute** $m := \frac{c}{d^x}$.
2. **return** m .

To see that the cryptosystem is correct, compute $cd^{-x} = w^r mg^{-rx} = g^{xr} mg^{-rx} = m$. It is also secure, assuming the DDH assumption holds.

Theorem 70 *ElGamal is semantically secure, if the DDH assumption holds.*

Proof Suppose we have a PPT adversary A , which breaks ElGamal's semantic security. We can use it to construct an algorithm A' , which solves the DDH problem. A' is given a sequence $\langle G, p, g, g_1, g_2, g_3 \rangle$ and must decide whether this is a Random Sequence or a DDH Sequence. A' will play the semantic security "game", using A 's responses to identify the sequence, thus solving the DDH problem.

$A'(G, p, g, g_1, g_2, g_3) :$

1. **compute** messages $(m_0, m_1) := A(G, p, g, g_1)$.
2. Pick $b \in \{0, 1\}$ uniformly at random.
3. **compute** A 's guess $b' := A(g_2, g_3 m_b)$.
4. **if** $b' = b$ **then return** 1 **else return** 0.

A' takes an input (G, p, g, g_1, g_2, g_3) (with G, p, g sampled). (G, p, g, g_1) is used as an El-Gamal public-key, which is given to A . The adversary returns a pair of messages m_0, m_1 , which it can distinguish. After selecting a random bit b , $(g_2, g_3 m_b)$ is returned to A , as a potential ciphertext. Then A returns b' , its guess for b . If $b' = b$ we return 1, which we interpret as identifying the DDH sequence. Otherwise, we return 0, identifying the Random sequence.

Note that if we give A' the input $(G, p, g, g^x, g^y, g^{xy})$, then $(g_2, g_3 m_b) = (g^y, (g^x)^y m_b)$. This is a valid ciphertext encryption of m_b , with public key (G, p, g, g^x) , and secret key x . Since A can distinguish m_0 from m_1 , it will guess $b' = b$ correctly. In this case A' will output 1 with as high a probability as A can distinguish the messages.

On the other hand, if we give input (G, p, g, g^x, g^y, g^z) for independently chosen z , $g_3 m_b = g^z m_b$ will just be a random element of G . Thus $g^z m_0$ and $g^z m_1$ will have equal probability of appearing in the ciphertext. So A will not be able to guess b , and A' will output 0 with high probability (and output 1 with low probability).

Thus A' can solve the DDH problem with non-negligible probability, assuming that A can break the semantic security of ElGamal. ■

22 The Cramer-Shoup Cryptosystem

The Cramer-Shoup cryptosystem in its fully strong form satisfies the strongest security definition, i.e., CCA-2. It is also quite efficient. The security of Cramer-Shoup relies on difficulty of the DDH problem.

Before introducing the full Cramer-Shoup cryptosystem, we introduce some simplified versions of the same to explain the key ideas. The full version will build upon these ideas later on.

22.1 Modifying the El-Gamal Cryptosystem

We describe a modification to El-Gamal and the corresponding security proof. Though at first, it seems like the scheme is CCA-1 secure, we explain why the security proof fails in that case. Thus, the scheme is just semantically secure and achieves nothing more in terms of security compared to El-Gamal. However, it serves our purpose of introducing the key ideas of Cramer-Shoup cryptosystem.

We assume we have a group G of prime order q , where q is large. Let g_1 and g_2 be two randomly chosen generators of the group. The cryptosystem is as follows.

Key Generation (1^k) Choose x, y at random from \mathbb{Z}_q . Define $h = g_1^x g_2^y$. Return $PK = \langle g_1, g_2, h \rangle$ and $SK = \langle x, y \rangle$.

Encryption (PK, m) Choose r at random from \mathbb{Z}_q . Return $E_{PK}(m) = (g_1^r, g_2^r, h^r \cdot m)$.

Decryption ($SK, E_{PK}(m) = (u, v, e)$) Decryption goes as follows.

$$D_{SK}(u, v, e) = \frac{e}{u^x \cdot v^y} = \frac{h^r \cdot m}{(g_1^r)^x \cdot (g_2^r)^y} = \frac{h^r \cdot m}{(g_1^x \cdot g_2^y)^r} = m.$$

Theorem 71 *The modified El-Gamal encryption scheme is semantically secure, if the DDH assumption holds in group G .*

Proof

We prove the above by contradiction. Suppose there exists a PPT adversary \mathcal{A} that breaks the semantic security of the modified El-gamal scheme with probability $\frac{1}{2} + \epsilon$ where ϵ is significant. We show how to build a PPT simulator \mathcal{B} that can play the DDH game with advantage $\frac{\epsilon}{2}$.

The input to the algorithm \mathcal{B} is (g_1, g_2, g_3, g_4) which is either a DDH tuple or a random tuple. Algorithm \mathcal{B} runs as follows. Create an instance of the modified El-Gamal scheme by randomly selecting x, y from \mathbb{Z}_q and setting $PK = \langle g_1, g_2, h = g_1^x g_2^y \rangle$, $SK = \langle x, y \rangle$. It then passes PK to \mathcal{A} and gets as return the set (m_0, m_1) . It flips a bit b and selects one of the messages m_b . Encrypt m_b as follows $C = E_{PK}(m_b) = (g_3, g_4, g_3^x g_4^y \cdot m)$. Algorithm \mathcal{B} passes C onto \mathcal{A} and obtains its guess of the bit flipped b' . If $b = b'$, \mathcal{B} outputs that the tuple is a "DDH tuple". Else it outputs that it is a "random tuple". Now we have the following claims.

Claim 1 *If the given tuple (g_1, g_2, g_3, g_4) is a DDH tuple, then the above experiment is identical to the real semantic security game.*

This is easy to see. Let's represent the tuple as $(g_1, g_2 = g_1^a, g_3 = g_1^b, g_4 = g_1^{ab})$. The

encryption in this case will be

$$C = E_{PK}(m) = (g_3, g_4, g_3^x g_4^y \cdot m) = (g_1^b, g_2^b, (g_1^x)^b (g_2^y)^b \cdot m) = (g_1^b, g_2^b, (g_1^x g_2^y)^b \cdot m)$$

which is a perfectly valid encryption of message m with randomness b using the key pair PK, SK .

Claim 2 *If the given tuple (g_1, g_2, g_3, g_4) is a random tuple, then the above experiment reveals no information about the bit b . That is, the bit b is information theoretically hidden from the adversary \mathcal{A} .*

Given the ciphertext $C = (g_3, g_4, g_3^x g_4^y \cdot m)$ and the public key $PK = (g_1, g_2, h = g_1^x g_2^y)$, computing the bit b is equivalent to computing the group element $\gamma = g_3^x g_4^y$. Taking the discrete log w.r.t. g_1 ,

$$\log_{g_1}(\gamma) = x \cdot \log_{g_1}(g_3) + y \cdot \log_{g_1}(g_4)$$

Let's consider what other information does the adversary have regarding the numbers x and y . During the whole experiment, the only other quantity adversary gets to see about these randomly chosen numbers is $h = g_1^x g_2^y$. Taking the discrete log w.r.t. g_1 again,

$$\log_{g_1}(h) = x + y \cdot \log_{g_1}(g_2)$$

The above two equations are different (i.e., not a multiple of each other) in case the given tuple is NOT a DDH tuple. Consider x and y as variables, the above two equations can be satisfied for any arbitrary choice of γ . This means, given any group element γ , there always exists a choice of the pair (x, y) which satisfies the above equations. Hence, γ is information theoretically hidden from the adversary. This implies that the adversary cannot guess the bit b with a probability better than $\frac{1}{2}$.

The above two claims imply the following. In case of DDH tuple, the simulator can give a correct answer with probability $\frac{1}{2} + \epsilon$. In case of a random tuple, the corresponding probability is just $\frac{1}{2}$. Thus, overall the simulator guesses correctly with probability $\frac{1}{2} + \frac{\epsilon}{2}$ giving it a non-negligible advantage of $\frac{\epsilon}{2}$ in the DDH game. This is a contradiction. Hence, modified El-gamal is secure in the semantic security game.

■

Lets consider the security of the above system for CCA-1. Since the simulator has the secret keys x and y , it decrypt challenge ciphertexts asked by the adversary in the CCA-1 game. However does this reveal any additional constraint on x and y ? Not if the encryptions are prepared in the honest way and are valid. However, it is possible to choose the ciphertexts maliciously such that decryption reveals additional information about x and y . For example, consider the ciphertext (g_1^r, g_2^{r-1}, h^r) . Although this is an invalid ciphertext, the simulator has no way of detecting that. Decrypting it will reveal g_2^y which fixes y completely. Thus,

the above proof fails for the CCA-1 case. We do not know whether this scheme is CCA-1 secure or not. If it is, we surely need a different proof than the above to prove security.

In the following sections, we will introduce additional techniques to overcome this problem and construct cryptosystems which satisfy stronger security definitions.

22.2 The Cramer-Shoup Lite Cryptosystem

We assume we have a group G of prime order q , where q is large. Let g_1 and g_2 be two randomly chosen generators of the group. The cryptosystem is as follows.

Key Generation (1^k) Choose x, y, a and b at random from \mathbb{Z}_q . Define $h = g_1^x g_2^y$ and $c = g_1^a g_2^b$. Return $PK = \langle g_1, g_2, h, c \rangle$ and $SK = \langle x, y, a, b \rangle$.

Encryption (PK, m) Choose r at random from \mathbb{Z}_q . Return $E_{PK}(m) = (g_1^r, g_2^r, h^r \cdot m, c^r)$.

Decryption ($SK, E_{PK}(m) = (u, v, e, w)$) First check whether $w = u^a \cdot v^b$. If this is not the case, output fail, the decryption has failed. Else, rest of the decryption goes as follows.

$$D_{SK}(u, v, e, w) = \frac{e}{u^x \cdot v^y} = \frac{h^r \cdot m}{(g_1^r)^x \cdot (g_2^r)^y} = \frac{h^r \cdot m}{(g_1^x \cdot g_2^y)^r} = m.$$

Theorem 72 *The Cramer-Shoup Lite cryptosystem is secure against non-adaptive chosen ciphertext attack if the DDH assumption holds in the group G .*

Proof

We prove the above by contradiction. The proof mostly goes similar to the previous proof. Suppose there exists a PPT adversary \mathcal{A} that breaks the non-adaptive CCA-1 security of the Cramer-Shoup Lite with probability $\frac{1}{2} + \epsilon$ where ϵ is significant. We show how to build a PPT simulator \mathcal{B} that can play the DDH game with advantage $\frac{\epsilon}{2}$.

The input to the algorithm \mathcal{B} is (g_1, g_2, g_3, g_4) which is either a DDH tuple or a random tuple. Algorithm \mathcal{B} runs as follows. Create an instance of the Cramer-Shoup Lite scheme by randomly selecting x, y, a, b from \mathbb{Z}_q and setting $PK = \langle g_1, g_2, h = g_1^x g_2^y, c = g_1^a g_2^b \rangle$, $SK = \langle x, y, a, b \rangle$. It then passes PK to \mathcal{A} and gets as return the set of messages to decrypted. It decrypts the set using the secret key SK and gets the challenge set (m_0, m_1) . It flips a bit b and selects one of the messages m_b . Encrypt m_b as follows $C = E_{PK}(m_b) = (g_3, g_4, g_3^x g_4^y \cdot m, g_3^a g_4^b)$. Algorithm \mathcal{B} passes C onto \mathcal{A} and obtains its guess of the bit flipped b' . If $b = b'$, \mathcal{B} outputs that the tuple is a "DDH tuple". Else it outputs that it is a "random tuple". Now we have the following claims.

Claim 1 *If the given tuple (g_1, g_2, g_3, g_4) is a DDH tuple, then the above experiment is identical to the real semantic security game.*

This claim is easy to verify similar to the previous proof for Modified El-Gamal.

Claim 2 *If the given tuple (g_1, g_2, g_3, g_4) is a random tuple, then the above experiment reveals no information about the bit b . That is, the bit b is information theoretically hidden from the adversary \mathcal{A} .*

The proof of this claim is similar to the previous proof. The adversary does not have enough constraints on x and y enabling him to compute the group element $\gamma = g_3^x g_4^y$. All γ 's are possible and equally likely.

Here, the adversary also gets to see the decryption of a set of non-adaptively chosen messages. In the following subclaim, we prove that this power doesn't give him any additional constraint on x and y . Rest of the proof for claim 2 is same as before.

Subclaim 2.1 *The adversary \mathcal{A} , except with a negligible probability, does not get any additional constraints on x and y as a result of the non-adaptive decryption queries which it is allowed to ask.*

Consider a ciphertext $C = (u, v, e, w)$ submitted by the adversary for decryption. Now there are two possible cases:

Case 1 $\log_{g_1}(u) = \log_{g_2}(v) = r$. In this case, the adversary gets no additional information. To see this, consider what he learns as a result of the decryption. He learns the following:

$$m = \frac{e}{u^x v^y}$$

Lets denote $\log_{g_1}(g_2)$ by α as a shorthand. Taking log of both the sides of the above equation:

$$\log_{g_1}(m) = \log_{g_1}(e) - x \cdot \log_{g_1}(u) - y \cdot \alpha \cdot \log_{g_2}(v) = \log_{g_1}(e) - r(x + \alpha \cdot y)$$

But he already has the constraint $h = g_1^x g_2^y$, or:

$$\log_{g_1}(h) = x + \alpha \cdot y$$

The above two equations are linearly dependent. Hence, he doesn't obtain any additional constraint on x and y .

Case 2 $\log_{g_1}(u) \neq \log_{g_2}(v)$. In this case, except with negligible probability, the decryption oracle outputs "fail". We show that the adversary gains only a negligible amount of information about a and b which will not help him to obtain any further constraints on x and y .

To see this, consider the constraints the adversary have on a and b . We have $c = g_1^a g_2^b$, or, taking the discrete log:

$$\log_{g_1}(c) = a + \alpha \cdot b$$

There is no other constraint known to the adversary at the time of preparing the set of messages to be decrypted. Now, let $r = \log_{g_1}(u)$ and $r' = \log_{g_2}(v)$. For the check $w = u^a v^b$ to succeed,

$$\log_{g_1}(w) = a.r + \alpha.b.r'$$

Since $r \neq r'$, this equation is linearly independent of the constraint equation known to the adversary. Hence, the probability of this equation being satisfied is the same as the probability of adversary guessing a and b correctly with only one constraint on them. This probability is exponentially small. Thus, the decryption oracle outputs "fail" except with negligible probability.

The only information that the adversary learns about a and b from the above is an inequality about them. That is, he can rule out one possibility of the value of the pair (a, b) . Since he is allowed to ask only polynomially many queries, he can rule out only polynomially many possibilities out of an exponential space. Thus, the number of possibilities ruled out is still negligible.

This proves the subclaim 2.1 and hence claim 2.

The above two claims imply the following. In case of DDH tuple, the simulator can give a correct answer with probability $\frac{1}{2} + \epsilon$. In case of a random tuple, the corresponding probability is just $\frac{1}{2}$ (a negligible amount). Thus, overall the simulator guesses correctly with probability $\frac{1}{2} + \frac{\epsilon}{2}$ giving it a non-negligible advantage of $\frac{\epsilon}{2}$ in the DDH game. This is a contradiction. Hence, Cramer-Shoup Lite is secure in the non-adaptive CCA-1 game.

■

Now we consider the above scheme for CCA-2 security. Its definitely not CCA-2 secure because, given the challenge ciphertext (u, v, e, w) , the adversary can re-randomize it as (u^r, v^r, e^r, w^r) . It will still remain a valid ciphertext of the same message. Thus in the decryption phase 2, adversary can ask for the decryption of the re-randomized ciphertext and output the result as the decryption of the challenge ciphertext.

22.3 The Full Cramer-Shoup Cryptosystem

We are now ready to present the Full Cramer-Shoup Cryptosystem. It builds upon the ideas presented earlier to achieve CCA-2 security. We assume we have a group G of prime order q , where q is large. Let g_1 and g_2 be two randomly chosen generators of the group. We also assume the availability of a collision resistant hash function¹¹ H . The cryptosystem is as follows.

¹¹Collision resistant hash functions can be build assuming discrete log (which is implied by our assumption DDH). A simple construction is $H(a, b) = g^a h^b$, where g and h are two different generators of the group and the discrete log of one is unknown w.r.t. to the other.

Key Generation (1^k) Choose x, y, a, b, a', b' at random from \mathbb{Z}_q . Define $h = g_1^x g_2^y$, $c = g_1^a g_2^b$ and $d = g_1^{a'} g_2^{b'}$. Return $PK = \langle g_1, g_2, h, c, d, H \rangle$ and $SK = \langle x, y, a, b, a', b' \rangle$.

Encryption (PK, m) Choose r at random from \mathbb{Z}_q . Return $E_{PK}(m) = (u = g_1^r, v = g_2^r, e = h^r \cdot m, w = (c \cdot d^\alpha)^r)$ where $\alpha = H(u, v, e)$.

Decryption ($SK, E_{PK}(m) = (u, v, e, w)$) First check whether $w = u^{a+\alpha a'} \cdot v^{b+\alpha b'}$. If this is not the case, output fail, the decryption has failed. Else, rest of the decryption goes as follows.

$$D_{SK}(u, v, e, w) = \frac{e}{u^x \cdot v^y} = \frac{h^r \cdot m}{(g_1^r)^x \cdot (g_2^r)^y} = \frac{h^r \cdot m}{(g_1^x \cdot g_2^y)^r} = m.$$

Theorem 73 *The Full Cramer-Shoup cryptosystem is secure against CCA-2 attacks if the DDH assumption holds in the group G and the function H used is collision resistant.*

Proof

We prove the above by contradiction. The proof mostly goes similar to the proof of Cramer-Shoup Lite with one exception.

Suppose there exists a PPT adversary \mathcal{A} that breaks the CCA-2 security of the full Cramer-Shoup cryptosystem with probability $\frac{1}{2} + \epsilon$ where ϵ is significant. We show how to build a PPT simulator \mathcal{B} that can play the DDH game with advantage $\frac{\epsilon}{2}$.

The input to the algorithm \mathcal{B} is (g_1, g_2, g_3, g_4) which is either a DDH tuple or a random tuple. Algorithm \mathcal{B} runs as follows. Create an instance of the full Cramer-Shoup scheme by randomly selecting x, y, a, b, a', b' from \mathbb{Z}_q and setting $PK = \langle g_1, g_2, h = g_1^x g_2^y, c = g_1^a g_2^b, d = g_1^{a'} g_2^{b'}, H \rangle$, $SK = \langle x, y, a, b, a', b' \rangle$. It then passes PK to \mathcal{A} and gets the challenge set (m_0, m_1) . It flips a bit b and selects one of the messages m_b . Encrypt m_b as follows $C = E_{PK}(m_b) = (g_3, g_4, g_3^x g_4^y \cdot m, g_3^{a+\alpha a'} g_4^{b+\alpha b'})$, where $\alpha = H(g_3, g_4, g_3^x g_4^y \cdot m)$. Algorithm \mathcal{B} passes C onto \mathcal{A} and obtains its guess of the bit flipped b' . Throughout the above, whenever \mathcal{A} asks for the decryption of a ciphertext, \mathcal{B} answers using its secret key SK . Now, if $b = b'$, \mathcal{B} outputs that the tuple is a "DDH tuple". Else it outputs that it is a "random tuple". Now we have the following claims.

Claim 1 *If the given tuple (g_1, g_2, g_3, g_4) is a DDH tuple, then the above experiment is identical to the real semantic security game.*

This claim is easy to verify similar to the Cramer-Shoup Lite proof.

Claim 2 *If the given tuple (g_1, g_2, g_3, g_4) is a random tuple, then the above experiment reveals no information about the bit b . That is, the bit b is information theoretically hidden from the adversary \mathcal{A} .*

The proof of this claim is similar to the proof of Cramer-Shoup Lite. The adversary does not

have enough constraints on x and y enabling him to compute the group element $\gamma = g_3^x g_4^y$. All γ 's are possible and equally likely.

Here, the adversary gets to see the decryptions of a number of messages which can be adaptively chosen even after seeing the challenge ciphertext. Let a ciphertext chosen by the adversary for decryption be (u, v, e, w) . To prove that learning its decryption doesn't give him any additional constraint on x and y , we consider two cases as before.

Case 1 $\log_{g_1}(u) = \log_{g_2}(v)$. In this case, the proof goes exactly as before. We prove that adversary gets a constraint which is linearly dependent upon the constraints already known to him.

Case 2 $\log_{g_1}(u) \neq \log_{g_2}(v)$. In this case, except with negligible probability, the decryption oracle outputs "fail" and hence the adversary gains none or negligible information.

To see this, consider the constraints the adversary have on a, b, a' and b' . We have $c = g_1^a g_2^b$ and $d = g_1^{a'} g_2^{b'}$ or, taking the discrete log:

$$\log_{g_1}(c) = a + \log_{g_1}(g_2).b$$

$$\log_{g_1}(d) = a' + \log_{g_1}(g_2).b'$$

Suppose the adversary is given the challenge ciphertext $(u^* = g_3, v^* = g_4, e^* = g_3^x g_4^y . m, w^* = g_3^{a+\alpha a'} . g_4^{b+\alpha b'})$. Let $\alpha = H(u^*, v^*, e^*)$, $r = \log_{g_1}(g_3)$ and $r' = \log_{g_2}(g_4)$ with $r \neq r'$ because the tuple given is random. As a result of this knowledge, he also learns the following (third) constraint:

$$\log_{g_1}(w^*) = r(a + \alpha a') + \log_{g_1}(g_2).r'(b + \alpha b')$$

Now we look at the following three possible subcases:

Subcase 1 $(u, v, e) = (u^*, v^*, e^*)$ **but** $w \neq w^*$. It is easy to see that the decryption will return "fail" in this case and hence the adversary gains negligible information.

Subcase 2 $(u, v, e) \neq (u^*, v^*, e^*)$ **but** $H(u^*, v^*, e^*) = H(u, v, e)$. The simulator has found a collision in the hash function H . This leads to a contradiction.

Subcase 3 $H(u^*, v^*, e^*) \neq H(u, v, e)$. Let $\alpha' = H(u, v, e)$, $\hat{r} = \log_{g_1}(u)$ and $\hat{r}' = \log_{g_2}(v)$ with $\hat{r} \neq \hat{r}'$ (recall that we are in case 2). For this decryption query to not return "fail", we should have the following check to succeed:

$$\log_{g_1}(w) = \hat{r}(a + \alpha' a') + \log_{g_1}(g_2).\hat{r}'(b + \alpha' b')$$

Since $r \neq r'$, $\alpha \neq \alpha'$ and $\hat{r} \neq \hat{r}'$, we can show that the above constraint is linearly independent of the previous three constraints. Lets denote $\log_{g_1}(g_2)$ by ℓ as a shorthand.

To verify linear independence, we take the four equations, write them in the matrix form and perform Gaussian Elimination. Following is the coefficient matrix corresponding to the above four equations:

$$\mathbf{M} = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ r & lr' & r\alpha & lr'\alpha \\ \hat{r} & l\hat{r}' & \hat{r}\alpha' & l\hat{r}'\alpha' \end{pmatrix}$$

Gaussian Elimination performs elementary row operations to put the coefficient matrix into the upper triangular form. Following is the matrix M converted to the upper triangular form using Gaussian Elimination. To verify linear independence, we can just see that none of the rows or the columns of the converted matrix is all zero.

$$\mathbf{M} = \begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & \alpha \\ 0 & 0 & 0 & \frac{\ell^2 r^2 r' \alpha' \hat{r}' - \ell^2 r r'^2 \alpha \hat{r} - \ell^2 r' \alpha r \hat{r}' + \ell^2 r'^2 \alpha \hat{r} - \ell r^2 \alpha' r \hat{r}' - \ell r^2 \alpha r' \hat{r}}{r(\ell \alpha' r' \hat{r} - \alpha' r \hat{r}' + \alpha r \hat{r} - \ell \alpha r \hat{r}')} \end{pmatrix}$$

Hence, the probability of fourth equation being satisfied is the same as the probability of adversary guessing a, b, a' and b' correctly with only three constraints on them. This probability is exponentially small. Thus, the decryption oracle outputs “fail” except with negligible probability.

This proves the claim 2.

The above two claims imply the following. In case of DDH tuple, the simulator can give a correct answer with probability $\frac{1}{2} + \epsilon$. In case of a random tuple, the corresponding probability is just $\frac{1}{2}$. Thus, overall the simulator guesses correctly with probability $\frac{1}{2} + \frac{\epsilon}{2}$ giving it a non-negligible advantage of $\frac{\epsilon}{2}$ in the DDH game. This is a contradiction. Hence, the Full Cramer-Shoup cryptosystem is secure in the CCA-2 game.

■

References

- [1] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM*, 33(4):792–807, 1986.
- [2] Johan Hastad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.

- [3] M. Naor. Bit commitment using pseudorandomness. *Journal of Cryptology: the journal of the International Association for Cryptologic Research*, 4(2):151–158, 1991.

Lecture 7

1 Introduction to Interactive Proofs

1.1 Introduction

A traditional, Euclidean-style proof for an assertion consists of a prover who simply outputs a proof¹. Someone reading the proof, a verifier, then decides whether or not the proof is correct. The observation has been made that there could be an advantage to letting the verifier *interact* with the prover. This may allow the assertion to be proven faster or with the release of less information than would be the case if the verifier were passive.

Our general framework consists of a prover P who is allowed an arbitrary exponential amount of time, and a verifier V who is allowed only poly-time. Both P and V are allowed to flip coins and they communicate to each other by sending messages. Note that since V is a poly-time machine, only a poly-number of messages will be sent between P and V .

The programs for P and V define the protocol PV . The input, typically an assertion of the form $x \in L$, is presented to both P and V , and P tries to convince V that the assertion is true. If V is convinced, then V accepts.

1.2 Definition of IP [Goldwasser, Micali, Rackoff]

Definition 10.1 An Interactive Proof for a language L is a protocol PV for a Prover and a Verifier such that:

- **Completeness:** If $x \in L$ then P has a good chance of convincing V that $x \in L$

$$\forall c > 0 \exists N \text{ s.t. } \forall x \in L \text{ where } |x| > N$$

$$\Pr_{\text{coins of } V, P} [PV(x) \text{ makes } V \text{ accept}] > 1 - \frac{1}{|x|^c}$$

- **Soundness:** If $x \notin L$ then every P' has little chance of convincing V that $x \in L$

$$\forall P' \forall c > 0 \exists N \text{ s.t. } \forall x \notin L \text{ where } |x| > N$$

$$\Pr_{\text{coins of } P'} [P'V(x) \text{ makes } V \text{ accept}] < \frac{1}{|x|^c}$$

IP is defined to be the class of languages which have Interactive Proofs.

¹Scribe notes taken by: Vassilis Papavassiliou, October 4, 11 1994

1.2.1 IP for Graph Nonisomorphism

Disclaimer: In these notes we will write interactive proofs for two languages: graph-isomorphism (GI) and graph-non-isomorphism (GNI). These languages are chosen because, besides the belief that V (a PPT machine) cannot recognize them, they provide a convenient framework in which to study the notions of interactive proofs and zero-knowledge. In addition, protocols for GI and GNI can be translated into protocols for certain hard number-theory problems. One would not really base a real system on GI or GNI.

Many of these interactive proofs will rely on the ability to produce random permutations of graphs. We do this by creating a random permutation and applying it to the graph.

$x \in GNI$ iff x is a pair of graphs (G_0, G_1) and $G_0 \not\sim G_1$

We abbreviate this as $x = \{G_0 \not\sim G_1\}$

The following protocol is an interactive proof for GNI . V picks either G_0 or G_1 at random and generates a random permutation of that graph. V sends this new graph to P who responds by telling V which graph V originally picked. Repeat this $k = |x|$ times. V accepts if P is right every time. In tabular form:

	P	$x = \{G_0 \not\sim G_1\}$	V
1			Generate a random bit b
2		$\leftarrow G' \leftarrow$	Generate a random permutation Π . Let $G' = \Pi(G_b)$
3	Determine b' s.t. $G' \sim G_{b'}$	$\rightarrow b' \rightarrow$	Reject if $b' \neq b$
4			Repeat steps 1-3 k times. Accept if $b' = b$ every time.

This protocol is an interactive proof because:

- **Completeness:** If the graphs are not isomorphic then only one of G_0 or G_1 will be isomorphic to G' , so P will always be able to determine b .
- **Soundness:** If the graphs are isomorphic then G' could have come from either G_0 or G_1 with equal probability, so any prover P' can only guess b . P' would have to guess correctly $|x|$ times. This has probability $\frac{1}{2^{|x|}}$.

1.2.2 Protocol (P1): Interactive Proof for Graph Isomorphism

$x \in GI$ iff x is a pair of graphs (G_0, G_1) and $G_0 \sim G_1$

We abbreviate this as $x = \{G_0 \sim G_1\}$

The following protocol is an interactive proof for GI . P generates a random permutation of G_0 . P sends this new graph to V who responds with a bit b . P then responds to V 's request by sending the permutation which maps the new graph that P generated to G_b . V checks that this

permutation is actually an isomorphism between the graphs. Repeat this k times one after the other (i.e. sequentially). V accepts if P was able to send a correct permutation every time. In tabular form:

		$x = \{G_0 \sim G_1\}$	
	P	communication	V
1	Generate a random permutation Π_1 . Let $G' = \Pi_1(G_0)$	$\rightarrow G' \rightarrow$	
2		$\leftarrow b \leftarrow$	Generate a random bit b .
3	Determine Π_2 s.t. $G' = \Pi_2(G_b)$	$\rightarrow \Pi_2 \rightarrow$	Reject if $G' \neq \Pi_2(G_b)$
4			Repeat steps 1-3 k times sequentially. Accept if $G' = \Pi_2(G_b)$ every time.

This protocol is an interactive proof because:

- **Completeness:** If the graphs are isomorphic then G' is isomorphic to both G_0 and G_1 , so P can always send an isomorphism between G' and G_b .
- **Soundness:** If the graphs are not isomorphic then G' is isomorphic to at most one of G_0 and G_1 . Any prover P' would be able to send an isomorphism between G' and G_b only if G' was originally created as a permutation of G_b . Thus P' would have to guess which b V will send. The probability that P' can do this k times is $\frac{1}{2^k}$.

2 Introduction to Zero Knowledge

The obvious interactive proof protocol for GI is for P to simply send V the isomorphism between G_0 and G_1 . This corresponds to the traditional, non-interactive way of proving things². However this protocol has the undesirable feature of revealing to V much information. In particular, V now knows an isomorphism between G_0 and G_1 . We desire an interactive proof which still convinces V that the graphs are isomorphic without revealing so much information to V . In fact, we don't want P to reveal *anything* to V beyond that the graphs are isomorphic. Such proofs are called zero-knowledge, introduced by Goldwasser, Micali and Rackoff (GMR-85). (P1) from the previous lecture is one such protocol for GI.

2.1 Motivating story

(This is a story also due to [GMR]) One night in a small town there is a murder. Verry Fire, the local reporter, hears about the murder and wants to write a story for her newspaper. She goes to a pay phone and calls the detective to get the facts of the murder case, but the detective simply tells her "There was a murder" and he hangs up. She calls back several times, but every time

²Scribe notes taken by: Vassilis Papavassiliou, October 11 1994

she just hears the phrase “There was a murder.” Verry already knew that there was a murder, so she certainly didn’t need to call the detective to obtain this information. She could have just saved her money and generated this phrase herself. Feeling a little frustrated, she decides to call the police chief. The chief, who enjoys playing games with reporters, flips a coin, and if the coin is heads, the chief says “There was a murder” and hangs up. If the coin is tails, he says “No comment” and hangs up. Verry calls back several times and sometimes she hears the first phrase while other times she hears the second phrase. Her conversation with the chief, however, still hasn’t given her any new information for her column. She could have just flipped a coin herself and generated the chief’s phrases with the same probability distribution. Verry proceeds to write her column, but she could have written the column without talking to the detective or the chief, because her conversations with them were *zero-knowledge*.

2.1.1 Definition of ZK

Definition 11.2 For a protocol PV , let $PV(x)$ represent *view* of the conversation from verifiers point of view on input x . Specifically, $PV(x)$ consists of:

- The messages sent back and forth between P and V
- The coins of V

In a later section we will justify including the coins of V in this definition. Let $[PV(x)]$ represent the distribution of points $PV(x)$ taken over the coins of P and V . In general for a machine S , let $[S(x)]$ be the distribution of outputs of S on input x taken over the coins of S .

Definition 11.3 An Interactive Proof PV for a language L is Zero Knowledge if

$$\forall V' \quad \exists S_{V'} \in PPT \quad \text{s.t.} \quad \forall x \in L$$

$$[S_{V'}(x)] \simeq [PV'(x)]$$

$S_{V'}$ is a PPT machine which knows V' and which on input $x \in L$ outputs points of the form $PV(x)$ defined above. Intuitively, the existence of $S_{V'}$ shows that V' does not *need* P to generate the output distribution $[PV'(x)]$. V' , a PPT machine, could have generated the distribution itself. Therefore P does not transfer any knowledge to V' (beyond the fact that $x \in L$) which V' could not have generated itself.

The quantifier is over all verifiers V' , even cheating verifiers. That is, PV' doesn’t *have* to be an interactive proof for L . The only goal of V' may be to extract information from P . Even for such cheating verifiers, there must exist a simulator with \simeq output distribution.

The existence of a PPT simulator for a verifier V' means that the interaction of P and V' is zero-knowledge. If the interaction of P with every verifier V' is zero-knowledge, then the protocol PV is zero-knowledge. In this case, note that the protocol PV' for any verifier V' is

also zero-knowledge, but remember that this protocol is not necessarily an interactive proof that $x \in L$.

In the definition of ZK , we use a simulator whose distribution on $x \in L$ is \simeq to the distribution of $PV(x)$. There are actually three different definitions of ZK corresponding to the three different definitions of \simeq of distributions:

Definition 11.4 The three variants of ZK are:

- **Perfect ZK =**

The distributions are exactly equal. This is the strictest definition of ZK .

- **Statistical $ZK \stackrel{s}{\simeq}$**

The distributions are statistically close. Recall that:

Distributions $\{X_n\}$ and $\{Y_n\}$ are **statistically close iff**

$$\forall c \exists N \text{ s.t. } \forall n > N$$

$$\sum_{\alpha \in \{0,1\}^n} \left| \Pr_{\{X_n\}} [X_n = \alpha] - \Pr_{\{Y_n\}} [Y_n = \alpha] \right| < \frac{1}{n^c}$$

In other words $\forall c \exists N \text{ s.t. } \forall n > N$ even a machine which is allowed exponential time must take at least n^c samples before it can distinguish $\{X_n\}$ and $\{Y_n\}$. Thus more than a polynomial number of samples are required to distinguish the distributions.

- **Computational $ZK \stackrel{c}{\simeq}$**

The distributions are computationally indistinguishable to poly-time machines. Recall that:

Streams $\{X_n\}$ and $\{Y_n\}$ are **poly-time indistinguishable iff**

$$\forall c \forall A \in PPT \exists N \text{ s.t. } \forall n > N$$

$$\left| \Pr_{\{\{X_n\}, A's \text{ coins}\}} [A(X_n) = 1] - \Pr_{\{\{Y_n\}, A's \text{ coins}\}} [A(Y_n) = 1] \right| < \frac{1}{n^c}$$

In other words $\forall c \exists N \text{ s.t. } \forall n > N$ any polynomial time machine must have running time at least n^c before it can distinguish $\{X_n\}$ and $\{Y_n\}$. Thus no poly-time machine can distinguish the streams.

distributions are equal \Rightarrow statistically close \Rightarrow computationally indistinguishable

The converses may not be true. In particular, if 1-way functions exist, then *computational*

indistinguishability $\not\Rightarrow$ statistical closeness:

$\{X_{2n}\}$ is the output distribution of a pseudo-random number generator $G : \{0, 1\}^n \rightarrow \{0, 1\}^{2n}$.
 $\{Y_{2n}\}$ is the output distribution of a true random number generator for $2n$ -bit strings.

$\{X_{2n}\}$ contains at most 2^n different strings.
 $\{Y_{2n}\}$ contains 2^{2n} different strings.

Therefore these distributions are statistically distinguishable, however since G is a pseudo-random number generator, they are computationally indistinguishable.

2.1.2 Requiring the Simulator to Output the Coins of V

The output of the simulator consists of the messages passed back and forth between P and V' as well as the random bits used by V' . We will now show why it is necessary for the simulator to output the random bits of V' .

Consider the following protocol for graph-isomorphism. This is clearly not a very good way to prove graph-isomorphism, but it illustrates the need for the simulator to output the coins of V' .

First, a description with words: The input is G_0 and G_1 . V randomly selects one of the two graphs, say G_0 , and sends to P a random permutation of that graph. P randomly selects one of the two graphs, say G_1 , and sends back the permutation which maps the graph it received from V to G_1 . If P and V happened to choose different graphs (in the above example V chose G_0 and P chose G_1) then V will be able to determine an isomorphism between G_0 and G_1 , and so V will accept the graphs as being isomorphic. Repeat everything k times, and if V is never able to determine an isomorphism between G_0 and G_1 , then V rejects. In table form, the protocol is:

		$x = \{G_0 \sim G_1\}$	
	P	communication	V
1			Generate a random bit b
2		$\leftarrow G' \leftarrow$	Generate a random permutation Π_1 . Let $G' = \Pi_1(G_b)$
3	Generate a random bit c		
4	Determine Π_2 s.t. $G_c = \Pi_2(G')$	$\rightarrow \Pi_2 \rightarrow$	Accept if $G_b = \Pi_2(\Pi_1(G_b))$. That is, accept if $\Pi_1 \circ \Pi_2$ is an isomorphism between $G_0 \sim G_1$
5			Repeat steps 1-4 k times. Reject if $G_{b'} \neq \Pi_2(\Pi_1(G_b))$ every time.

This protocol is an interactive proof because:

- **Completeness:** If the graphs are isomorphic then V will determine an isomorphism iff $b \neq c$. The probability that all k trials have $b = c$ is $\frac{1}{2^k}$, so the probability that P fails to

convince V that the graphs are isomorphic is negligible.

- **Soundness:** If the graphs are not isomorphic then V will never accept since it only accepts if it can determine an isomorphism between G_0 and G_1 .

The protocol is not zero-knowledge because V learns an isomorphism between G_0 and G_1 . Thus we cannot construct a simulator for V with the appropriate output distribution where the output consists of the messages passed back and forth between P and V as well as the random bits of V . However, if we only required the simulator to output the messages passed back and forth between P and V , and in the definition of zero-knowledge interactions we only required the distributions to be \simeq on points consisting only of the messages passed back and forth, then the interaction of P and V is zero knowledge because, as shown below, we can construct an appropriate simulator for V . Remember that to prove zero-knowledge for a protocol, we would have to show how to simulate any verifier V' , not just V . All we will show is that in the particular case of P talking to V , if we don't require the simulator to output the random bits of V , then we will incorrectly conclude that the interaction between P and V is zero-knowledge.

Code for (Pseudo)Simulator S_V

1. FOR $i:=1$ TO k DO

(a) pick a random bit b

(b) pick a random permutation Π . Let $G' = \Pi(G_b)$

(c) output messages:



(d) END FOR LOOP

2. END PROGRAM /* The simulator has output k points */

If we consider only the messages passed back and forth between P and V , then the above simulator produces a distribution of points identical to the one generated by the real P talking to V : for $x \in L$, G' is uniformly distributed over all the graphs isomorphic to G_0 (or G_1 since $G_0 \sim G_1$), and Π is a map between G' and either G_0 or G_1 with equal probability for each. This would lead us to believe that the interaction of P and V is zero-knowledge, even though V may learn an isomorphism between G_0 and G_1 .

Now consider the messages passed back and forth between P and V *and* the random bits of V . If we have the above simulator also output the bit b from which G' was created (b is suppose to be the random bit of V), then $S_V(x)$ no longer has output distribution identical to $[PV(x)]$ for $x \in L$. In particular, the points in $[PV(x)]$ have the graph G_c to which Π maps G' independent from bit b . The points in $[S_V(x)]$ always have Π mapping G' to G_b . This justifies including the random bits of V in the view of what the verifier sees in the definition of a zero-knowledge interaction.

2.2 Proving a Protocol is Zero Knowledge

Let PV be the protocol (P1) for graph isomorphism which was defined earlier. It was proved previously that PV obeys Completeness and Soundness, so PV is an interactive proof. In this section we will prove that PV is zero-knowledge. We will do this by constructing, for any verifier V' , a simulator whose distribution on $x \in GI$ is the same as that for $PV(x)$. The simulator will depend on the notions of saving the state and restarting a Turing Machine.

The verifier is a special kind of Turing Machine. It consists of a finite state control, an input tape, output tape, work tape, random tape, input communication tape and output communication tape. The information on these tapes and the state of the control completely defines the state of the verifier. Thus, we can save the state of the verifier by saving this information. Say we save the state of the verifier at time t . We then put something on the verifier's input communication tape, let the verifier continue running and observe the behavior of the verifier. We can now restore the verifier to the state we saved at time t . If we now put the same thing on the verifier's input communication tape and let the verifier run again, we will observe exactly the same behavior we did before; the Turing Machine has no way of remembering what we had it do before we restored its state.

2.2.1 Story-time

Jay Lentil, host of a popular late-night television show, convinces the great, world-famous magician Flippo to be on his show. To the amazement of all the viewers, Flippo proceeds to flip a normal coin and have it come up heads 100 times in a row. Not to be outdone, Dave Numberman, host of a competing show, tries to find a magician to match this incredible feat. Dave however doesn't find a suitable magician, so he disguises his assistant Paul in a magician's costume. Dave's show is pretaped, so in the studio, he has Paul flip a coin again and again until it has come up heads 100 times. Then the tape is edited to remove all the coin flips that came up tails. When the tape is shown that night, the viewers are amazed to see a man flipping a coin and having it come up heads 100 times in a row!

2.2.2 Construction of a Simulator for (P1)

The previous story showed how it is possible to run an experiment many times and pick out only the successful experiments. If the chance of success is high enough ($\frac{1}{2}$ in the story) then we can get the required number of successful experiments quickly. We will use this for our simulator $S_{V'}$; we will save the state of the verifier V' , run an experiment on V' , output the results if they are successful, restore the state of V' , run another experiment, etc. We continue until we have the required number of successful experiments. For our simulator, a successful experiment corresponds to a point the simulator can output. At the end, the k points that $S_{V'}(x)$ has output have \simeq distribution to $[PV'(x)]$.

Code for Simulator $S_{V'}$

1. pick at random a random tape R for V'
 2. FOR $i:=1$ TO k DO: /* simulate 3-round atomic protocol k times */
 - (a) record state of S'
 - /* record configuration of FSM control of V' */
 - /* record work tape configuration of V' */
 - /* record head positions of work and random tapes readers */
 - (b) set $DONE:=FALSE$
 - (c) WHILE (not $DONE$) DO
 - i. pick bit c at random
 - ii. pick permutation Π at random
 - iii. compute $H_i = \Pi_i(G_c)$
 - iv. send H_i to V' and get bit b from V'
 - v. if $b=c$
 - then
 - $DONE:=TRUE$
 - output random tape R and messages:

$\rightarrow H_i \rightarrow$
$\leftarrow b \leftarrow$
$\rightarrow \Pi_i \rightarrow$
 - else
 - reset V' to previously saved state
 - vi. END WHILE LOOP
 - (d) END FOR LOOP
3. END PROGRAM /* The simulator has output k points */

Each experiment is successful half the time. In particular, it is successful when $b = c$. Therefore in $2k$ expected time, $S_{V'}(x)$ outputs k points. If $S_{V'}(x)$ outputs k points, then the distribution of these points is identical to $[PV'(x)]$. This is true because $S_{V'}$ chooses H_i exactly the same way as P does, and V responds with the same b because it has no way of knowing if it is talking to P or $S_{V'}$.

The only complication comes from the fact that there is an exponentially small chance the simulator will fail to terminate and so fail to produce the desired output distribution. Therefore $[S_{V'}]$ is statistically close to $[PV'(x)]$, not exactly equal. To get perfect zero-knowledge, we can run a brute-force, exponential time algorithm for graph isomorphism in parallel to the simulator. If the exhaustive algorithm finishes before the simulator, we use its results instead (i.e. if it gives an isomorphism between G_0 and G_1 , use this isomorphism to create the k points). The expected running time of this method is polynomial since the chance that $S_{V'}$ takes a long time is negligible, but there is a small chance it will run in exponential time. Therefore, we have to modify our definition of ZK slightly to allow simulators which are *expected PPT* rather than

PPT.

2.2.3 Parallel Version of (P1)

Protocol (P1) runs a 3-round protocol k times, so there are $3k$ total messages sent. Communication is expensive in reality, so we would like a way to minimize the number of rounds required by an interactive proof. This would save us the overhead on each separate message. One thing we could think of for (P1) is to do all k runs in parallel so that we have only 3 (albeit larger) messages sent. The Parallel (P1) protocol is:

		$x = \{G_0 \sim G_1\}$	
	P	communication	V
1	For $i=1$ to k : Generate a random permutation Π_i^1 . Let $G'_i = \Pi_i^1(G_0)$	$\rightarrow G'_1, \dots, G'_k \rightarrow$	
2		$\leftarrow b_1, \dots, b_k \leftarrow$	Generate random bits b_1, \dots, b_k .
3	For $i=1$ to k : Determine Π_i^2 s.t. $G'_i = \Pi_i^2(G_{b_i})$	$\rightarrow \Pi_1^2, \dots, \Pi_k^2 \rightarrow$	Accept iff $\forall i G'_i = \Pi_i^2(G_{b_i})$

The above protocol is still an interactive proof for GI because:

- **Completeness:** If the graphs are isomorphic then P will be able to provide an isomorphism between G'_i and G_{b_i} for all i .
- **Soundness:** If the graphs are not isomorphic then V will only accept if some prover P' can guess right all k times. The probability that this happens and P' fools V is $\frac{1}{2^k}$.

Unfortunately, the above protocol is believed not to be zero-knowledge. In particular, the simulator we created for (P1) will not work for Parallel (P1). This is true because a simulator for a parallel verifier produces a successful experiment only if it guess all k bits b_1, \dots, b_k correctly *simultaneously*. It can do this with probability only $\frac{1}{2^k}$. In fact, if the above protocol is zero-knowledge, that would imply that $GI \in BPP$, as was shown by [Goldreich, Krawczyk]. However, later on we will see that with appropriate modifications to the protocols we can make a constant-round ZK for GI.

2.2.4 Application: Interactive Passwords

Imagine that you are in Berkeley and that you want to login to a computer in New York. The computer in New York requires a password, so you send your password over the network to New York, however anyone can tap the network cable and learn your password by monitoring the line. This is clearly not acceptable. You need a system for proving to the computer in New York that you really are who you say you are, without sending any information over the network line that can be intercepted:

1. P and V get together and generate “hard” $G_0 \sim G_1$ and Π , the isomorphism between them (warning: in practice, we will not use GI, since we do not know which graphs are hard to find isomorphism for, but rather some algebraic problem, but for now lets assume that we can somehow find a pair of graphs for which is it hard to find an isomorphism, to make our example simpler)
2. Whenever P wants to login, P and V run the ZK proof for GI. If P is able to convince V that $G_0 \sim G_1$, then V allows P to login.

Why is this secure? The answer is that ZK protocols can not be repeated by listeners (i.e. it is not transferable!). Why? Because whatever listener heard on the wire during login, he could have generated all by himself, since there is a simulator for ZK!

3 Number of rounds for ZK

The topic of today’s lecture³ is the number of rounds of interaction needed for perfect and statistical ZK proofs. We will also consider the special subclass of Interactive proofs (called Arthur-Merlin proofs), in which verifier can not keep secrets from the prover.

3.1 Arthur-Merlin Protocols

Today’s topic explores special type of Interactive Proof, called Arthur-Merlin proofs, introduced by Babai and Moran. The name is drawn from the Arthurian legend, where Merlin is the all powerful (exponential time) prover and Arthur is the polynomial time verifier. Arthur is restricted to generating public random coin tosses (i.e. which Merlin can see) as opposed to Interactive Proofs of [GMR], where the verifier can secretly (from the prover) generate the coin tosses. Moreover, the only messages Arthur is allowed to send to Merlin is results of his coin-tosses. (Recall that in interactive proofs, Verifier can make arbitrary polynomial computations in order to generate questions for the Prover. Notice in the setting where coins are public, there is no need to send anything but the coin-flips, since whatever questions verifier can compute based on his public coin-flips, prover can compute just as well).

Complexity remarks : Goldwasser and Sipser have shown that if there exists an interactive protocol for a language L , then it can be transformed into an AM protocol for a language L . The transformation they present does not (as far as we know) preserve Zero-Knowledge. Also, in the [Shamir,LFKN] proof that $IP = PSAPCE$, it is in fact shown that $AM = PSAPCE$ as far as languages membership is concerned. This does not tell us anything about ZK, though. Fortunately, Impagliazzo and Yung (and [BGGHKMR]) have shown that everything in IP is in computational ZK. As far as perfect and statistical ZK, it was shown by Fortnow, and Aiello and Hastad that only languages in the second level of the polynomial-time hierarchy (in single round AM intersect co-AM) can be proven in statistical ZK.

³Scribe notes taken by: Wendy Heffner, October 13, 1994. These notes are based on the previous notes by Sivan Toledo of the same lecture given at MIT in 1992.

3.2 Public/Private coins and The number of rounds of interaction

As stated above, the AM protocol is different from IP in that it restricts the verifier to generating public coin tosses, and it restricts verifier's messages to random strings only. Let us revisit the proof for graph non-isomorphism (GNI) to see if this is either or both ZK and AM.

P	input: G_0, G_1	V
		Privately generate a coin flip b
	$\leftarrow \Pi(G_b) = G' \leftarrow$	Generate random permutation Π of graph G_b
Calculate to which graph G' is isomorphic, send back that subscript.	$\rightarrow c \rightarrow$	
		If $c = b$ OK.

In that protocol, the verifier generates a secret bit b and sends to the prover a random permutation $\Pi(G_b) = G'$ isomorphic to one of the two input graphs G_0 or G_1 . Then the prover sends back a bit c indicating the graph G_c to which G' is isomorphic. The verifier then checks the value of c to see if matches b . If the graphs G_0 and G_1 are isomorphic, the cheating prover will be caught with probability $1/2$. If the graphs are not isomorphic, the verifier will be always convinced.

Clearly, this protocol is not AM. The verifier cannot publicly generate the coin flips in this protocol. In addition, this protocol is not ZK. Suppose the verifier V had a third graph G_{new} that (s)he knew to be isomorphic to either G_0 or G_1 . The cheating V could substitute the graph G_{new} for G' in the first step and have the prover P show to which graph G_{new} was isomorphic. Using this technique V , therefore, can gain additional knowledge.

In the future lectures, we will see how to design a perfect ZK protocol for GNI. Could we construct a perfect ZK protocol for GNI which is also AM? The answer depends on the number of rounds of interaction between prover and verifier: Goldreich and Krawczyk have shown that *any* (perfect or statistical) ZK protocol with constant number of rounds which is AM for language L implies that $L \in BPP$. However, if we allow *private* coins, we can design a constant number of rounds perfect ZK protocol for GNI. This we will see next time, this time, we will look at graph-isomorphism problem, and show a constant-round (*private coins perfect ZK protocol for it, due to [Bellare, Micali, Ostrovsky 1990]*).

3.3 Private coins, Constant rounds protocol for Graph-Isomorphism

Let's look again at the Graph Isomorphism protocol we devised in the last lecture. In that protocol, the prover sends a random graph C isomorphic to the two input graphs G_0 and G_1 , obtained by choosing one of them at random and randomly permuting it. Then the verifier sends a random bit b , and the prover has to show the isomorphism between C and G_b .

P	communication	V
Privately generate a coin flip x and randomly permute graph G_x .	$\rightarrow \Pi(G_x) = C \rightarrow$	
	$\leftarrow b \leftarrow$	Generate a public coin flip b .
Show isomorphism between G_b and C	$\rightarrow G_b \sim C \rightarrow$	

If the graphs are indeed isomorphic, the verifier always will be convinced. If they are not, the cheating prover will be caught with probability $1/2$. We want to amplify the probability of catching a cheating prover. As we have seen in previous lecture, doing so by repeating the above atomic protocol k times sequentially will amplify the probability to $1 - 2^{-k}$, and will preserve also the ZK property. But, this requires repeating the atomic protocol sequentially, so the number of rounds raises to $3k$.

Can we squeeze all k repetitions of the atomic protocol into fewer steps? Let us consider what happens when we send all k graphs isomorphic to the input graphs at once, then send all k query bits in one round, and then send all the answers to the queries in one round.

P	communication	V
Privately generate a coin flips x_1, \dots, x_k and randomly permute graph G_{x_1}, \dots, G_{x_k} .	$\rightarrow \Pi(G_{x_1}) = C_1, \dots, \Pi(G_{x_k}) = C_k \rightarrow$	
	$\leftarrow b_1, \dots, b_k \leftarrow$	Generate a public coin flip b_1, \dots, b_k .
Show isomorphism between each pair G_{b_i} and C_i	$\rightarrow (G_{b_1} \sim C_1), \dots, (G_{b_k} \sim C_k) \rightarrow$	

This protocol is still an interactive proof, meaning that if the graphs are isomorphic the verifier will be convinced, and if they are not he will detect this with probability $1 - 2^{-k}$. Additionally, the protocol is an AM proof, however, the ZK property cannot be established any more.

The problem is that when the simulator sends the k graphs H_1, \dots, H_k to a verifier, the verifier may send in return query bits that depend on those graphs. This means that if the simulator tries to reset the verifier and send k new graphs that were generated from G_0 and G_1 according to the query bits, the verifier might ask different queries. Depending on luck is also a bad strategy here. While the verifier has no way to know from which of the input graph was each H_i generated, the chances that the query bits will match the generation pattern of the H_i 's is only 2^{-k} . So it will take the simulator exponential (in k) expected time to generate a valid conversation. Note that even knowing the code of the verifier is not enough. The verifier might choose the query bits according to some hard to invert hash functions of the H_i 's.

What should we do? Intuitively, we would like to modify the protocol such that the verifier will have to commit to its query bits before seeing the H_i 's. Of course the commitment should not reveal the bits to the verifier in any way, otherwise he might generate the graphs H_i accordingly,

and the protocol will no longer be an interactive proof. Having such a bit commitment mechanism, the idea is that a simulator will be able to get the committed bits, then send some H_1, \dots, H_k . Then the verifier de-commits, or reveals its bits. Now the simulator rewinds the verifier to the state it was in just before receiving the H_i 's, and sends another sequence of k graphs, which were generated according to the bits (which must be the same, since the verifier committed to them), and generate a conversation.

We face however a major problem. When we used encryption, we used it to hide information from the verifier. This worked because the verifier has only polynomial time to try to decipher the messages. But our prover has infinite computational power, so the verifier cannot hide anything by encryption. The solution is to modify a bit our requirements from the bit commitment protocol. We will devise a mechanism that will ensure that the prover has no way at all to know the query bits, since this is essential for maintaining the IP properties. We will however let the verifier cheat sometimes, that is to ask a different query than the one he committed to. This has no relevance to the IP properties, but it might affect the ZK property. However, we will make sure that if the verifier changes his bits, then he already knows the isomorphism between G_0 and G_1 , so he doesn't gain any knowledge from the protocol.

P	communication	V
Generate 2 random graphs A_0, A_1 by permuting G_0 twice.	$\rightarrow A_0, A_1 \rightarrow$	
	$\leftarrow \Pi_1(A_{b_1}) \dots \Pi_k(A_{b_k}) \leftarrow$	Generate k random bits $b_1 \dots b_k$ and k random permutation $\Pi_1 \dots \Pi_k$.
Generate k random graphs $H_1 \dots H_k$ by permuting G_0	$\rightarrow H_1 \dots H_k \rightarrow$	
P checks that all (b_i, Π_i) for all $i = 1, \dots, k$ are valid. If this is <i>not</i> the case, P stops the conversation.	$\leftarrow (b_1, \Pi_1) \dots (b_k, \Pi_k) \leftarrow$	(Send the query bits, with "proofs" that these were indeed committed.)
(Send a proof that P could not decipher the committed bits.)	$\rightarrow G_0 \approx A_0 \approx A_1 \rightarrow$	
(Send the actual proofs.)	$\rightarrow H_1 \approx G_{b_1} \dots H_k \approx G_{b_k} \rightarrow$	

It is easy to see that the protocol is indeed an IP protocol. Since V accepts only if in the step before last P proves that A_0 and A_1 are isomorphic, V can be sure that P could not know what are b_1, \dots, b_k . So in fact in this sense the protocol is similar to the previous (non ZK) one, since the query bits are sent all after P sends H_1, \dots, H_k . To prove that the protocol is indeed a ZK protocol we have to show a simulator for any verifier. The simulator works in phases, where each phase is divided into two sub-phases. If a phase succeeds, the simulator generates a conversation, and is done. Since the probability of success will be shown to be constant, the expected number of phases of the simulator will be constant. In the first sub-phase, the simulator tries to simulate a conversation under the assumption that the verifier he has is honest, hence the

simulator is in the honest mode. He generates A_0 and A_1 from G_0 , sends them to the verifier, then gets the k permuted copies of them (the commitments), and sends the k random graphs, generated arbitrarily from G_0 or G_1 . Then the verifier de-commits and reveals the bits. The simulator then rewinds the verifier to the state just before the k random graphs were sent, and now sends k permuted graphs, but that were generated from G_0 or G_1 according to the query bits. If the verifier is an honest one, and actually de-commits the same bits, we are done, since the simulator can run the protocol to completion. If, however, the verifier sends different bits, we are out of luck. Note that the simulator cannot actually declare the verifier a cheater, since this is not something that a prover could find out in a real conversation. In that case the simulator moves to the second subphase, the cheating mode. Now the simulation starts by generating A_0 from G_0 and A_1 from G_1 . Note that the verifier cannot distinguish between these modes, since he always sees two random isomorphic graphs. The simulation proceeds as before, the verifier is rewinded, and given new H_i 's. If he is honest now, and de-commits to the same bits, we are again out of luck, since the simulator will not be able to demonstrate the isomorphism between A_1 and G_0 . But if the verifier cheats again, we can complete the simulation. Consider the bit b_i that the verifier changed. In the first try, he demonstrated that some graph is isomorphic to A_0 say. In the second, he demonstrated that the same graph is isomorphic to A_1 . By doing so, he gave the simulator the isomorphism between A_0 and A_1 , and therefore between G_0 and G_1 . So now there is no problem to answer all the queries, and to show that A_1 is isomorphic to G_0 .

The crucial point is that because the verifier cannot distinguish between the honest and cheating mode, the probability that he will cheat in the honest mode and be honest in the cheating mode is at most $1/4$. So the expected number of phases is constant.

The protocol which appears in [Bellare,Micali,Ostrovsky 1990], is a perfect ZK protocol. The way we presented the correctness proof however, only shows that the protocol is a statistical ZK protocol, since a verifier might behave in some deterministic way (that is, be honest or cheat) for some A_0 and A_1 's, so it is not clear that the distribution will be completely identical. The protocol is a perfect ZK protocol, however, but the proof of this fact is somewhat more complicated. A similar protocol can be devised for Quadratic Residuosity (and any other random self-reducible problem).

4 Number-theoretic ZK protocols and Compound ZK protocols

4.1 Some Number Theory

In this lecture⁴, we consider Perfect ZK proofs for a number theoretic language. Before doing that, we need some facts from number theory. Let's start with some definitions.

Definition 13.5 $Z_N^* = \{x | 1 \leq x \leq N, \gcd(N, x) = 1\}$.

Definition 13.6 $x \in QR(Z_N^*)$ if

- $\exists w \in Z_N^*$ such that $w^2 = x \pmod N$, and

⁴Scribe notes taken by: Shuzo Takahashi, October 18, 1994

- $\left(\frac{x}{N}\right) = 1$ where $\left(\frac{x}{N}\right)$ is the Jacobi symbol.

For the definition of the Jacobi symbol, see p. 30 of Dana Angluin's lecture note. We make the following assumption:

If N is a product of two large primes and $\left(\frac{x}{N}\right) = 1$, then it is hard to decide if x is a QR or not. Assumption :

Fact 13.7 $\left(\frac{x}{N}\right)$ can be computed in polynomial time in given N even if you do not know the factorization of N .

Fact 13.8 Suppose $N = P_1P_2$ (a product of two different primes). Given $x \in QR(Z_N^*)$, there are four different square roots of $x \pmod N$, say $y, -y, z, -z$. Then we have $\gcd(y + z, N) = P_1$ or P_2 . In particular, if you know these four square roots, it is easy to factor N .

Proof: Since $y^2 \equiv x \pmod N$ and $z^2 \equiv x \pmod N$, $y^2 \equiv z^2 \pmod N$. So there is K such that $y^2 - z^2 = KN$. Since $y^2 - z^2 = (y + z)(y - z)$, we have $P_1|y + z$ or $P_1|y - z$. Also, $P_2|y + z$ or $P_2|y - z$. From these, it is easy to see $P_1|y + z$ xor (exclusive or) $P_2|y + z$. Thus, we have $\gcd(y + z, N) = P_1$ or P_2 . ■

On the other hand,

Fact 13.9 if you know the factorization of N , then it is easy to check whether x is a QR or not.

Fact 13.10 If both q and x are squares, then qx is a square. If q is a square and x is not a square, then qx is not a square.

4.2 A Perfect ZK Proof for QR in $3k$ rounds [GMR]

Now we describe an IP protocol for QR. In this protocol, an input is a pair (N, x) , and the prover needs to convince the verifier that $x \in QR(Z_N^*)$.

Protocol 1

P	communication	V
Generate a square q at random.	$\rightarrow q \rightarrow$	
	$\leftarrow b \leftarrow$	Generate a random bit b .
Send \sqrt{q} (if $b = 0$) and send \sqrt{qx} (if $b = 1$).	$\rightarrow \sqrt{q}$ or $\sqrt{qx} \rightarrow$	

Repeat the above k -times sequentially.

It is easy to see that this is an IP proof. Notice the similarity between the protocol for the graph isomorphism and Protocol 1. Actually, this protocol can be translated into the graph isomorphism protocol. Now we show that Protocol 1 is a statistical ZK proof (it can also be shown that the above protocol is perfect ZK by running, in parallel to the simulator below, an "exponential search" simulator). The following is a statistical ZK simulator:

Description of a Simulator

- (1) Set the state of the verifier as usual.
- (2) Pick a bit b' and $l \in Z_N^*$ at random. Then if $b' = 0$, set $q = l^2 \pmod N$ and if $b' = 1$ then $q = l^2 x^{-1} \pmod N$.
- (3) Then send q to the verifier.
- (4) If $b = b'$ (where b is a random bit generated by the verifier), then we can supply \sqrt{q} or \sqrt{qx} (depending on $b = 0$ or $b = 1$) to the verifier. Otherwise, reset the state to (1) and repeat (2)-(4).

The distribution created by taking only the successful repetitions is equal to the distribution of the prover-verifier conversation in Protocol 1.

Remark: the above protocol is also perfect ZK, where the above simulator is augmented by adding low-probability exponential-time search, as before.

4.3 A Perfect ZK proof for QR in 5 rounds

In this section, we give a perfect ZK proof for QR which takes only 5 rounds. (it is a simplification of a [BMO-90] protocol for this problem.) Again, the input is (N, x) and the prover needs to convince the verifier that $x \in QR(Z_N^*)$.

Protocol 2

P	communication	V
Randomly generate $s \in Z_N^*$, and calculate $z = s^2$.	$\rightarrow z \rightarrow$	
	$\leftarrow y_1, \dots, y_k \text{ (commitment)} \leftarrow$	Randomly generate bits b_1, \dots, b_k and $r_1, \dots, r_k \in Z_N^*$, and calculate $y_i = z^{b_i} r_i^2$.
Randomly generate squares q_1, \dots, q_k .	$\rightarrow q_1, \dots, q_k \rightarrow$	
P checks that $y_i = z^{b_i} r_i^2$ for all $i = 1, \dots, k$. If this is <i>not</i> the case, P stops the conversation.	$\leftarrow r_1, \dots, r_k, b_1, \dots, b_k \leftarrow$	De-commit y_1, \dots, y_k , i.e., show r_1, \dots, r_k and b_1, \dots, b_k to the prover.
Send a proof that the prover couldn't decipher the committed bits.	$\rightarrow \sqrt{z} \rightarrow$	
Send $\sqrt{q_i}$ or $\sqrt{q_i x}$ for each i depending on $b_i = 0$ or $b_i = 1$.	$\rightarrow \sqrt{q_i}$ or $\sqrt{q_i x} \rightarrow$	

We can show that Protocol 2 is a perfect ZK proof. The proof is similar to the [BMO-90] five-round graph isomorphism simulator..

4.4 Another Example of Perfect ZK Proof: Proving an "OR" of GI

Let $L = \{ \langle (G_0, G_1), (C_0, C_1) \rangle \mid \text{either } G_0 \sim G_1 \text{ or } C_0 \sim C_1 \}$. In the following protocol, the prover needs to convince the verifier that at least one of two pairs of graphs is isomorphic.

Protocol 3

P	communication	V
Randomly generate bits b_1 and b_2 , and graphs G' and C' such that $G' \sim G_{b_1}$ and $C' \sim C_{b_2}$.	$\rightarrow G', C' \rightarrow$	
	$\leftarrow b \leftarrow$	Randomly generate a bit b .
Choose b_1' and b_2' such that $b = b_1' \oplus b_2'$, and such that $G' \sim G_{b_1'}$ or $C' \sim C_{b_2'}$.	$\rightarrow b_1', b_2', G' \sim G_{b_1'}, C' \sim C_{b_2'} \rightarrow$	

Repeat the above k -times sequentially

Now we claim:

Claim 13.11 The above protocol is an IP proof for the language L .

Proof: If $x \notin L$, that is, $G_0 \not\sim G_1$ and $C_0 \not\sim C_1$, then the prover cannot find b_1', b_2' described in Protocol 3 at least half the time. So the verifier will reject with probability $\geq \frac{1}{2}$. On the other hand, if $x \in L$, say $G_0 \sim G_1$, then the prover can change b_1 to b_1' (if it is necessary) and take $b_2' = b_2$ so that $b = b_1' \oplus b_2'$ for b which is sent by the verifier. So the prover can always convince the verifier. ■

Next we claim:

Claim 13.12 The protocol is a statistical ZK proof

Proof: The following is a simulator:

Simulator

- (1) Record the state of the verifier.
- (2) Pick b_1, b_2 at random and generate G', C' such that $G' \sim G_{b_1}$ and $C' \sim C_{b_2}$.
- (3) Send G' and C' to the verifier.
- (4) If $b = b_1 \oplus b_2$ (where b is a random bit generated by the verifier), then we can supply b_1, b_2, G_{b_1} , and C_{b_2} to the verifier. Otherwise reset the state of the verifier to (1) and repeat (2) - (4).

The distribution created by taking only the successful repetition is statistically close (i.e. with exponentially small probability the protocol is always guessing wrong, and then we can not proceed) to the distribution of the original prover-verifier conversation in the protocol. Thus, Protocol 3 is a statistical ZK proof. ■

Remark: the above protocol is also perfect ZK, where the above simulator is augmented by adding low-probability exponential-time search, as before.

5 Perfect Zero-Knowledge Graph Non-Isomorphism

In this lecture we consider ZK protocol for GNI, due to [GMW]⁵.

Let $L = \{(G_0, G_1) | G_0 \not\sim G_1\}$. Now, the prover's task is to convince the verifier that two graphs G_0 and G_1 are not isomorphic. The following is an IP protocol for L .

Protocol 4

P	communication	V
	$\leftarrow G' \leftarrow$	Randomly generate a bit b and G' such that $G' \sim G_b$.
Find b such that $G' \sim G_b$.	$\rightarrow b \rightarrow$	

Repeat the above k -times.

Now we claim:

Claim 14.13 Protocol 4 is an IP proof for the language L .

Proof: If $G_0 \not\sim G_1$, then the prover can always tell b which is sent by the verifier. On the other hand, if $G_0 \sim G_1$, then the prover can send b at most half the time. ■

However, Protocol 4 is **not** a ZK proof (of any type)! Suppose the verifier has a graph C such that $G_0 \sim C$ or $G_1 \sim C$; but the verifier does not know which is the case. If the verifier sends C to the prover and the prover answers faithfully, then the verifier knows which one of G_0 and G_1 is isomorphic to C . This is extra information. Also, notice that an obvious simulation does not work because a verifier V' might have a biased coin. A ZK proof for the graph non-isomorphism exists, to design it, let us first consider a modification of the above protocol:

Here's another interactive protocol for graph non-isomorphism (GNI) – P wants to prove to V that $G_0 \not\cong G_1$.

Prover	Communication	Verifier
		Flip a coin $b \in \{0, 1\}$
	$\leftarrow C_0, C_1 \leftarrow$	Generate random permutations $C_0 \cong G_b$ and $C_1 \cong G_{\bar{b}}$
Guess bit b	$\rightarrow b \rightarrow$	

(1) This is an interactive protocol:

- If $G_0 \not\cong G_1$, then $C_0 \not\cong C_1$, so P can always distinguish between them and thus guess the bit.

⁵Scribe notes taken by: Shuzo Takahashi, October 18, 1994; Sanjoy Dasgupta October 27, 1994

- If $G_0 \cong G_1$, then $C_0 \cong C_1$, so no prover can distinguish them; thus any prover has at most a $1/2$ chance of fooling V . This can be reduced to $1/2^k$ by repeating the process k times.

(2) However, the protocol is not clearly zero-knowledge: we run into difficulties constructing a simulator for it.

The basic problem is that V may not know the bit b before the prover tells him what it is, so that he actually gains information. To overcome this, we force V to start by proving to P that he knows b (so that he can't possibly have gained anything from P 's answer). Here is the expanded protocol:

Prover	Communication	Verifier
		Flip a coin $b \in \{0, 1\}$
	$\leftarrow C_0, C_1 \leftarrow$	Generate random permutations $C_0 \cong G_b$ and $C_1 \cong G_{\bar{b}}$
		V proves that he knows b :
	$\leftarrow D_1, \dots, D_{2k} \leftarrow$	Generate graphs D_1, \dots, D_{2k} , $\{D_{2i-1}, D_{2i}\} \cong \{C_0, C_1\}$
Flip coins q_1, \dots, q_k	$\rightarrow q_1, \dots, q_k \rightarrow$	Find k pairs of isomorphisms π_1, \dots, π_k , such that $\pi_i : \{D_{2i-1}, D_{2i}\} \rightarrow \{G_0, G_1\}$ if $q_i = 0$, and $\pi_i : \{D_{2i-1}, D_{2i}\} \rightarrow \{C_0, C_1\}$ if $q_i = 1$
Check the π_i	$\leftarrow \pi_1, \dots, \pi_k \leftarrow$	
Find bit b	$\rightarrow b \rightarrow$	

Notation:

$\{A, B\} \cong \{C, D\}$ means that $(A \cong C \wedge B \cong D)$ or $(A \cong D \wedge B \cong C)$.

$\pi : \{A, B\} \rightarrow \{C, D\}$ means that $\{A, B\} \cong \{C, D\}$ and π consists of the two relevant mappings.

To get an idea of what is going on in the “proof of knowledge” phase, consider the pair (D_1, D_2) . P will ask V to either demonstrate a mapping $\{D_1, D_2\} \rightarrow \{G_0, G_1\}$ or $\{D_1, D_2\} \rightarrow \{C_0, C_1\}$. Since V doesn't know beforehand which mapping he will be asked for, he must know both of them (otherwise, he'll fail with probability $1/2$). If he knows both of them, then he automatically knows a mapping $\{G_0, G_1\} \rightarrow \{C_0, C_1\}$ and he therefore knows b . So the chance that he doesn't know b but manages to give the right answers (for all k pairs) anyway is $1/2^k$. Here's a statistical ZK simulator $S_{\hat{V}}$ for this protocol:

1. Get (C_0, C_1) and $(D_1, D_2), \dots, (D_{2k-1}, D_{2k})$ from \hat{V} .
2. Record \hat{V} 's state (the usual).
3. Send random bits q_1, \dots, q_k , and get responses π_1, \dots, π_k .

4. Reset \hat{V} to its old state (in 2).

5. Send new random bits q'_1, \dots, q'_k , and get responses π'_1, \dots, π'_k . With high probability ($1 - 1/2^k$), $q_i \neq q'_i$ for some i . For this i , we have $\pi_i : \{D_{2i-1}, D_{2i}\} \rightarrow \{G_0, G_1\}$ and $\pi'_i : \{D_{2i-1}, D_{2i}\} \rightarrow \{C_0, C_1\}$. Therefore, we know b .


The distribution created by successful runs of this simulator is statistically equivalent to that created by the prover-verifier conversation. It is not necessarily perfect ZK because of the tiny chance of failure. To make the simulator perfect ZK, we can augment it with a low-probability exponential-time search.

6 ZK for ALL of NP

6.1 ZK Proofs for all of NP under physical assumptions

In our protocolgraph non-isomorphism, the verifier had to convince the prover that he knew a certain bit. Now we'll look at a situation where the prover must commit to a bit and then reveal it later to the verifier. For the time being, let us implement this using a physical assumption – safes. Later, we'll show to how simulate this using one-way permutations. Figure 1 shows the protocol.

Figure 1: Bit commitment using a safe:



We'll use safes to exhibit a ZK protocol for Graph 3-Colourability (G3C), an NP-complete problem – see Figure 2. It is statistical zero-knowledge, since the conversations can be simulated as shown in Figure 3.

Figure 2: ZK proof for G3C:

Figure 3: ZK simulator for G3C:

6.2 ZK proofs for NP using one-way permutations

6.2.1 Implementing safes

We'll show how to simulate a safe using a one-way permutation f :

Comittal:

- P wants to commit to bit b . He randomly chooses p, x such that $|p| = |x|$.
- P sends $p, f(x), (x \cdot p) \oplus b$ to V .

Decommittal:

- P sends x, b to V .

Why does this system work?

(1) It is perfectly binding – because f is a permutation, P is fully committed to x . He is thus committed to $(x \cdot p)$ and therefore to b .

(2) If the verifier (who operates in probabilistic polynomial time) can deduce any information about b , then he can invert f , as we showed in our proofs about hard-core bits.

It makes a big difference whether f is a uniform or non-uniform one-way permutation, as we'll see below.

6.2.2 ZK protocol for Graph 3-Colourability

Here is a more formal statement of the ZK protocol for G3C:

Prover	Communication	Verifier
$C = 3$ -colouring of graph G		
Repeat k times:		
Permute(C)	→ Commit to C →	
	← u, v ←	Pick an edge $e = (u, v) \in G$
	→ Decommit colouring of u, v →	Check colouring of u, v

Intuition for proof: If graph G has $|E|$ edges, then the probability that P doesn't know a coloring but manages to fool V is at most $(1 - \frac{1}{|E|})^k$. So for $k \sim \ell|E|$, this probability is less than $1/e^\ell$. The simulator $S_{\hat{V}}$ follows the usual pattern:

Do k times:

1. Pick an edge $e = (u, v)$ in G and color u, v differently. Assign all other vertices the same color.
2. Commit to the coloring of G .
3. Record the state of \hat{V} .
4. If \hat{V} asks for an edge other than e , reset its state, and repeat step (4). Otherwise, decommit vertices u, v .

As before, this is a statistical ZK simulator, because it has a (small) chance of failure, but it can be made into a perfect ZK simulator by adding a low-probability exponential search.

The proof of the protocol is rather involved, and we'll touch on some of the important issues below.

6.2.3 Uniform vs. Non-Uniform Bit Commitment

One possible problem with the protocol is that we've assumed that our bit commitment scheme is secure in this context. Consider, for instance, what happens if it is based on a permutation f which is uniform one-way but not non-uniform one-way. That is, f (and thus the bit commitment) can be cracked with non-negligible probability by a family of poly-sized circuits (one circuit for each input size) or equivalently, by a polynomial time Turing machine which receives some additional "advice" as input (this advice being the description of the relevant circuit). Well, it might just happen that the advice that the verifier needs to break the bit commitment scheme is precisely the statement " G is 3-colourable"! Who knows? – this statement might, by some strange mapping, be exactly equivalent to the description of a circuit that can invert the function f . Thus it is crucial that f be non-uniform one-way.

6.2.4 ZK subroutines

Here's another difficulty – our interactive protocol for $G3C$ consists of a loop which is repeated many times. We can show that a single iteration of this loop is zero-knowledge, but it is somewhat more complicated to show that the entire protocol is zero-knowledge, since with each successive iteration, the verifier has more information (specifically, the conversation from previous rounds). To handle this, we introduce a new term:

Definition An interactive protocol (P, V) for a language L is **auxiliary-input zero-knowledge** if \forall verifiers $\hat{V} \exists$ simulator $S_{\hat{V}} \in PPT$ such that:

$$\forall h, \forall x \in L, \quad P\hat{V}(x, h) \simeq S_{\hat{V}}(x, h)$$

Here, h corresponds to the past history between the prover and verifier. If a protocol is auxiliary-input zero-knowledge, then we'll see in a later lecture that it can be used as a subroutine without any problems. That is, it can be used in another program which calls it a polynomial number of times, and these multiple calls taken together will be auxiliary-input zero-knowledge.

6.2.5 Main proof

The core of the proof consists of the following:

Claim Say the bit-commitment scheme is based on a permutation f . If for infinitely many graphs, the verifier can often (with non-negligible probability) decide if a committed (hidden) graph coloring is correct, then f is not non-uniform one-way.

Proof Intuition. Let's first restate the claim. Without loss of generality, assume there's a infinite sequence of graphs G_1, G_2, \dots where the size of G_i is i , on which the verifier V can often

distinguish a hidden correct coloring from a hidden incorrect coloring. For each graph G_i , fix (1) a correct 3-coloring C_i and (2) an incorrect coloring C'_i in which the vertices of one edge are colored the same as in C_i and the remaining vertices all have the same color. Let our bit-commitment relation (a hard-core predicate of f) be $g(\cdot)$. With a slight abuse of notation, let $g^*(C_i)$ denote a committal of coloring C_i . So we know that $\forall i, V$ can often distinguish between $g^*(C_i)$ and $g^*(C'_i)$. We'll use this fact to construct a family of poly-sized circuits P_i , each of which can invert $g(\cdot)$ on outputs of length i , with non-negligible probability.

Look at a specific pair (C_i, C'_i) . Consider a sequence of colorings $C'_i = D_1, D_2, D_3, \dots, D_k = C_i$ such that D_{j+1} is the same as D_j except for one vertex v_j , whose coloring is changed to that in C_i . Using our old hybrid argument, we can show that there is some j for which V can tell apart D_j and D_{j+1} with non-negligible probability. So here's our circuit P_i : it has G_i, D_j, D_{j+1} , and v_j hard-wired into it, and on input $I = g(b) = (f(x), (x \cdot p) \oplus b, p)$, it asks V if there's a difference between (1) D_{j+1} and (2) D_j with I substituted for $g(\text{coloring of } v_j)$. The circuit is right whenever V is.

The proof is slightly more complicated than this (for instance, the coloring of v_k actually uses two bits since there are three possible colors), but these are the main ideas behind it.

7 ZK: a fine print

Herein⁶ we will cover the the following:

- Uniform versus non-uniform zero-knowledge: a danger
- Auxiliary-input zero-knowledge
- A compiler for honest zero-knowledge \Rightarrow general zero-knowledge

7.1 Uniform versus non-uniform zero-knowledge

Thus far, we have implemented a secure bit-commitment mechanism for zero-knowledge proofs using one-way functions. A question remains as to whether we need uniform or non-uniform one-way functions for the proofs to actually be zero-knowledge.

It turns out we must assume our one-way function f is non-uniformly secure, i.e., that it is secure not only against polynomial-time adversaries, but also a family of circuits, one per input length.

Recall that a uniform one-way function cannot be broken by a polynomial time adversary, but can be broken by a family of circuits or a polynomial-time adversary with advice.

A non-uniform one-way function can't be broken by either.

As an example, in a zero-knowledge proof for graph-coloring, the description of the circuit that breaks the particular uniform f , or the advice necessary to break it, might be embedded in the (potentially very bizarre) graph itself.

Remark: Note that for all G of a certain input length, no advice will help.

⁶Scribe notes taken by: Todd Hodes, October 27, 1994

7.2 Auxiliary-input zero-knowledge

Usually, a zero-knowledge proof is a subroutine in some larger application. Thus, much additional information might be available to the verifier. For example, the verifier might know how to color half of a certain graph. We now discuss how a verifier, V , gets no additional information, even given some auxiliary input.

Definition 16.14 An interactive proof PV for a language L is auxiliary-input zero-knowledge if:

$$\forall \hat{V} \exists S_{\hat{V}} \in \text{PPT s.t. } \forall h \forall x \in L \\ [P\hat{V}(x, h)] \cong [S_{\hat{V}}(x, h)]$$

i.e. the distributions are indistinguishable.

Here, the simulator has two inputs, where h is an arbitrary string: the auxiliary input.

The benefit of auxiliary-input zero-knowledge proofs is that they can be combined, and they stay zero-knowledge. For example, let's look at the special case where a single protocol is repeated multiple times. We first prove that the protocol is "good," i.e. that the probability of error is negligible, and then prove it is auxiliary-input zero-knowledge. Thus, assume we have a protocol PV that shows $x \in L$ with probability of $\frac{1}{2}$. Running this protocol Q times gives us a new protocol, P_QV_Q .

Claim 16.15 If PV is auxiliary-input zero-knowledge, then P_QV_Q is also auxiliary-input zero-knowledge.

Proof:

Suppose \exists a distinguisher that can distinguish distributions $P_1V_1, P_2V_2, P_3V_3, \dots$ from the distributions S_1, S_2, S_3, \dots , where the simulator S_n has the previous partial history as auxiliary input.

Consider "hybrids" P_j , for $0 \leq j \leq k$, where in P_j the first j samples come from P_nV_n and remaining samples come from S_n :

$$\begin{aligned} P_0 &= (P_1V_1)(P_2V_2)(P_3V_3) \dots (P_kV_k) \\ P_1 &= S_1(P_2V_2)(P_3V_3) \dots (P_kV_k) \\ P_2 &= S_1S_2(P_3V_3) \dots (P_kV_k) \\ &\vdots \\ P_k &= S_1S_2S_3 \dots S_k \end{aligned}$$

We know $P_0 - P_k > \frac{1}{2^n}$, and therefore $\exists j$ such that $P_j - P_{j+1} > \frac{1}{k2^n}$, another $\frac{1}{\text{poly}}$ fraction. Consider a distribution:

$$P(\boxed{z}) = S_1S_2S_3 \dots S_j\boxed{z}(P_{j+1}V_{j+1}) \dots (P_kV_k)$$

If z is a sample from S_n then $P(z) = P_j$, and if z is a sample from P_nV_n then $P(z) = P_{j+1}$. Assuming we find $j+1$ and fix the other values correctly, such a distinguisher could be used to distinguish this single sample, \boxed{z} , which is a contradiction. ■

7.3 A compiler for honest zero-knowledge \Rightarrow general zero-knowledge

It would be nice if we knew that a verifier following a zero-knowledge interactive proof protocol wouldn't "cheat" and try to obtain additional information from the prover. Since this is unlikely, we will now look at a procedure to take a honest ZK protocol and convert it into one where it is assured that no information is revealed. In other words, we will show how to convert an interactive proof which is zero-knowledge for an honest verifier into an interactive proof which is zero-knowledge for any verifier.

The first step in the procedure is to design a protocol where V must follow the instructions exactly, an honest zero-knowledge protocol. The second step is to "compile" this into a new protocol which is guaranteed to divulge no information, regardless of the actions of the verifier.

First, we describe a procedure developed by [Blum and Micali]: **flipping coins into a well**:

A person P_1 stands far enough away from a well so as not to be able to look into it. Another person, P_2 , stands at the edge of the well. P_1 then flips coins into the well. P_2 can discern the result of the coin tosses by looking into the well, and also prevent P_1 from seeing them.

This can be formalized as follows:

		Tossing Coins Into a Well	
	P_1	communication	P_2
1		$\leftarrow \text{commit}(b_1) \leftarrow$	Generate a random bit b_1
2	Generate a random bit b_2	$\rightarrow b_2 \rightarrow$	$r = b_1 \oplus b_2$

If multiple bits are needed, they can all be done in parallel, still using only two communications.

This procedure has the following properties:

1. P_2 cannot control the outcome of the coin flip
2. P_1 doesn't know the outcome of the coin flip

Note that V will never have to decommit her bit if they use a "good" commitment scheme.

Given this procedure, if we let $P_1 = P$ and $P_2 = V$, Prover can pick the coins for V . We must now enforce that V acts honestly. One way of doing this is by demanding that each time V sends a message m to P , V also sends a zero-knowledge proof of the following NP statement:

"According to my committed secret random tape and previous random history, m is the message I was supposed to send."

However, how does a verifier commit to a prover, if the prover has arbitrary computational power? This is the topic of our next lecture.

Lecture 8

1 Bit-Commitment (BC): two variants

1.1 Introduction

These notes¹ explain two versions of Bit Commitment and the construction of Bit Commitment protocols based on cryptographic protocols.

Let's recall the problem of Bit Commitment (BC). There are two communicating parties, a sender S , and a receiver R . BC takes place in two stages. First, in the *commit* stage, a bit b is committed to, then in the *reveal* stage the bit is revealed. In order to make this protocol effective we want it to possess the following two properties:

- R has no knowledge of the value of b until S wishes him to know it.
- S cannot change the value of his commitment, i.e., *decommit* to a different value after committing to it.

There are two properties of a BC protocol:

1. Security: The complexity of R knowing the value of b , i.e., how well b is “hidden”.
2. Binding: The complexity of S being able to “cheat” (change the value of his commitment without R detecting it), i.e., how “binding” is the commitment to the sender.

We will see two versions:

1. Computationally Secure/Perfectly Binding.
2. Perfectly Secure/Computationally Binding.

1.2 Computationally Secure/Perfectly Binding BC (CS/PB)

Computationally Secure/Perfectly Binding BC has the following properties:

1. After commitment b is well defined, i.e., sender will *never* be able to cheat and decommit both a 0 and a 1.
2. b is hidden only computationally.

Note: S can have arbitrary complexity but R must only be of polynomial-time complexity.

The following is an example of a CS/PB BC protocol. Let f be a one-way permutation. Let $HCB(x)$ be a Hard Core Bit of a string x generated using f , then the following is a CS/PB BC protocol of bit b :

¹Scribe notes taken by: Elan Amir, October 25, 27 1994

S	communication	R
Generate a random string x . Let $c = b \oplus HCB(x)$.		
Commitment	$\rightarrow f(x), c \rightarrow$	
Decommitment	$\rightarrow x, b \rightarrow$	Verify $c = b \oplus HCB(x)$.

Let's examine the protocol in detail and see why it is in fact CS/PB. In order to cheat the sender has to be able to find a value of x with the property that

$$HCB(f^{-1}(f(x))) = \{0, 1\}$$

But f is a permutation so this is impossible. Therefore the protocol is Perfectly Binding. On the receiver's side, in order for the receiver to determine b from the information he is given, he needs to determine the value of $HCB(x)$ which is computationally difficult for a polynomial-time receiver. Therefore, the protocol is Computationally Secure.

1.2.1 Extending CS/PB BC Protocol Construction to all One-Way Functions

We have seen a CS/PB BC protocol that requires the use of one-way permutations. We now show that such a protocol can be devised using *any* one-way function f .

The first step of the construction uses the fact established in [ill] that any one-way function can be used to build a Pseudo-Random Number Generator (PRG). [naor] completes the construction with the result that any PRG can be used to construct a CS/PB BC protocol. We now prove this result.

Let $G : \{0, 1\}^n \rightarrow \{0, 1\}^{3n}$ be a PRG. Let $C(g, r) = c$, $g, r, c \in \{0, 1\}^{3n}$ where the bits of c are defined as follows:

$$c_i = \begin{cases} g_i \oplus b & \text{if } r_i = 1 \\ g_i & \text{if } r_i = 0 \end{cases}$$

Now consider the following BC protocol of bit b .

S	communication	R
	$\leftarrow r \leftarrow$	Choose a random $3n$ -bit string r .
Choose a random n -bit seed s .		
Let $g = G(s), c = C(g, r)$		
Commitment	$\rightarrow c \rightarrow$	
Decommitment	$\rightarrow b, s \rightarrow$	Verify c .

Let's examine the properties of this protocol. First, we claim that the sequence c is still pseudo-random. To see this, observe that if $b = 0$, then c is just the output of a PRG. If $b = 1$, then c is the output of a PRG with a random set of flipped bits. The latter case is pseudo-random since if it was not, i.e., we could construct a distinguisher to distinguish between c and a truly random value, we could use this distinguisher to distinguish all pseudo-random numbers, which contradicts the assumption of PRG existence. Since c is pseudo-random, b is computationally hidden, and the protocol is Computationally Secure.

Next, for a given r , consider what the sender needs to do in order to cheat. He must find two seeds s_1 and s_2 such that $G(s_1)$ agrees with $G(s_2)$ on all bit positions where $r_i = 0$ and disagrees with $G(s_2)$ on all bit positions where $r_i = 1$. A simple counting argument will show that given a random r the probability of the existence of such a pair is exponentially small. We observe that each pair of seeds corresponds to a single choice of r (since the bits of r are defined according to the corresponding bits of s_1 and s_2 as described above). Next, notice that there exist 2^{2n} pairs of n -bit seeds. The correspondence between seed pairs and r implies that there exist at most 2^{2n} values of r for which there exist a pair of seeds s_1 and s_2 that can be used to cheat. However, there exist 2^{3n} $3n$ -bit strings for r so that the probability that given a random r there exist a pair of seeds s_1 and s_2 that can be used to cheat is $2^{2n}/2^{3n} = 2^{-n}$.

Therefore we have shown that any one-way function can be used to construct a CS/PB BC protocol.

1.3 Perfectly Secure/Computationally Binding BC (PS/CB)

Perfectly Secure/Computationally Binding BC has the following properties:

1. After commitment, b is perfectly hidden, i.e.,

$$\forall \text{conversations } c \text{ of commitment, } P(c|b = 0) = P(c|b = 1)$$

where the probabilities are over the coin-flips of S and R .

2. A polynomial-time sender can not cheat, i.e., decommit to both a 0 and a 1.

Here we note that S must be of polynomial-time complexity while R can have arbitrary complexity.

Before moving on to an example of a PS/CB BC protocol we introduce the notion of claw-free functions. Informally, two one way permutations f_0, f_1 are **claw-free** if it is computationally intractable to find in polynomial-time two values x_0, x_1 such that $f_0(x_0) = f_1(x_1)$.

We are now ready for a PS/CB BC protocol. Let f_0, f_1 be two claw-free one-way permutations. Then the following is a PS/CB BC protocol of bit b :

S	communication	R
Generate a random string x . Let $z = f_b(x)$.		
Commitment	$\rightarrow z \rightarrow$	
Decommitment	$\rightarrow x, b \rightarrow$	Verify $z = f_b(x)$.

Why is this PS/CB? In order for the sender to cheat, he has to find two values x_0, x_1 such that $f_0(x_0) = f_1(x_1)$, that way when he decommits he can send x_b . But this property contradicts the fact that $f_{0,1}$ are claw-free for a polynomial-time sender, so this protocol is Computationally Binding. Note that the alternative of finding a value of z such that $f_{0,1}(x) = z$ is impossible since $f_{0,1}$ are permutations. On the receiver's side, since $f_{0,1}$ are permutations, there exist x_0, x_1 such that $f_0(x_0) = f_1(x_1) = z$, therefore only the knowledge of z does not enable the receiver to

obtain any information on b since the sender could have computed z with $b = 0$ or $b = 1$ with equal probability. Therefore the protocol is Perfectly Secure.

So we see that the above protocol is PS/CB providing that we can find claw-free functions. We now show a Number Theoretic construction of a PS/CB BC protocol which uses a pair of such functions.

Let b be a bit to be committed. Then the following protocol is PS/CB.

S	communication	R
		Generate 2 large primes p_0, p_1 . Let $n = p_0 p_1$ Generate a random string w . Let $s = w^2 \pmod n$ Generate k random values v_i , $1 \leq i \leq k$. Let $t_i = v_i^2 \pmod n$
	$\leftarrow n, s, t_1, t_2, \dots, t_k \leftarrow$	Proof the s is a square
Choose k bit values $c_i, 1 \leq i \leq k$	$\rightarrow c_1, c_2, \dots, c_k \rightarrow$	
	$\leftarrow \sqrt{t_1 s^{c_1}}, \dots, \sqrt{t_k s^{c_k}} \pmod n \leftarrow$	End proof that s is a square.
Commitment. Choose a large number r . Define $M_b(r) = s^b r^2 \pmod n$	$\rightarrow M_b(r) \rightarrow$	
Decommitment	$\rightarrow (b, r) \rightarrow$	Verify $M_b(r) = s^b r^2$.

The proof rests assumes the fact that factoring cannot be done in polynomial time. First, let's examine the proof of the fact that s is a square. Note that if s is not a square then if $b_i = 1$, $\sqrt{t_i s^{c_i}}$ will not be defined unless t_i is not a square and the product $t_i s^{c_i}$ is a square. In that case, when $b_i = 0$, $\sqrt{t_i}$ will not be a square. So we see that the probability of R being able to cheat on the fact that s is a square is 2^{-k} .

Next, notice that since s is square, both values that the sender could send are squares, i.e., in the quadratic residue of n , and therefore have the same distribution. As a result, the receiver learns no information on the value of b from the commitment which implies that the protocol is Perfectly Secure. From the sender's point of view, in order for him to cheat, he must be able to factor s , since then he would be able to send the pair $(b, r\sqrt{s})$ or the pair (b, r) , depending on what he chose to decommit to. However, since we assume that factoring s is computationally impossible for a polynomial-time machine, the fact that the sender has only polynomial complexity assures that the protocol is Computationally Binding. Note that the above discussion shows that M_b is a pair of claw-free functions.

1.3.1 Construction of a PS/CB BC using any One-Way Permutation

We now show a PS/CB BC protocol that removes the restriction of claw-free functions. The following protocol is presented in [novy]. The sender wishes to commit to a bit b . Let $B(x, y) = x \cdot y \pmod 2$, $x, y \in \{0, 1\}^n$. Let f be a one-way permutation on $\{0, 1\}^n$.

The following is the commitment stage:

1. S generates a random n -bit string x and computes $y = f(x)$.
2. R selects a linearly independent set of n -bit vectors $\{h_1, h_2, \dots, h_{n-1}\}$.
3. For j from 1 to $n - 1$
 - R sends h_j to S .
 - S sends $c_j = B(h_j, y)$ to R .
4. The $n - 1$ iterations of the above step define $n - 1$ linearly independent equations of the form $c_j = B(h_j, y)$. Therefore there are exactly two n -bit vectors y_0, y_1 that are solutions to these equations. Define y_0 to be the lexicographically smaller of the two vectors. Both S and R compute y_0 and y_1 (this does not place any restrictions on the complexity of S or R since a set of linear equations can be solved in polynomial time). Observe now that y is equal to either y_0 or y_1 . Let

$$c = \begin{cases} 0 & \text{if } y = y_0 \\ 1 & \text{if } y = y_1 \end{cases}$$

5. S computes c and sends it to R .

To decommit S sends b and x to R who verifies the computation of c and that in fact $y = f(x)$ solves the set of equations.

The above protocol is clearly perfectly secure. How do we show that it is computationally binding? Observe that in order to cheat the sender must be able to know $f^{-1}(y_0)$ and $f^{-1}(y_1)$ since then he will be able to send the appropriate inverse corresponding to the bit value he wishes to decommit to according to the definition of c . However, we now show that this is impossible. In particular, we show that if the sender can invert both solutions to the equations, we can use that sender as a subroutine to an algorithm which can invert a one-way function. The construction of such an algorithm is as follows.

Define the I_f to be the following algorithm which “communicates” with a cheating sender \hat{S} . The input to I_f is an n -bit string w . $B(x, y)$ is defined as before.

The core of the algorithm is the following loop:

1. Record the entire state of \hat{S} .
2. Pick h_1 at random and send it to \hat{S} .
3. Upon receipt of c_1 from \hat{S} check if $B(h_1, w) = c_1$. If so, proceed to h_2 (linearly independent of h_1), otherwise reset \hat{S} to the state recorded in (1) and goto (2), i.e., choose a new h_1 .

I_f continues this loop until either it succeeds $n - 1$ times, i.e., it accepts h_1, h_2, \dots, h_{n-1} , or it fails n times, i.e., it needs to choose a new h_i in step (3) n times over the runtime of the algorithm (not necessarily consecutively).

Here is an idea of what happens (the actual proof is more complicated): We now examine the state of I_f when it terminates. In the first case, we have obtained $n - 1$ linear equations to which w is a solution and which we know that \hat{S} holds the inverse to, since it computed its c_i 's using the same check we performed in step (3). Therefore, at the decommittal stage \hat{S} will send


us $f^{-1}(w)$. In a sense we have “forced” \hat{S} into inverting our w by constraining the equations to contain w as a solution, thereby obtaining the inverse by the assumption that \hat{S} could invert both solutions. In the latter case, we observe that if \hat{S} has avoided w n times this means that we already have n equations which are sufficient to *solve for* w . This means that \hat{S} has already guessed w , which ((can be shown to happen over the choice of w) with negligible probability. We see therefore that if S can cheat, we can invert a one-way function. Therefore, the protocol is Computationally Binding.

2 Two-prover Model

2.0.2 Setup

All of our previous ² discussions of zero-knowledge proofs have been based on the notion of “secure envelopes” or “locked safes,” where we can place a bit into a container in a *commitment* stage and then only open the envelope in a *decommitment* stage. We will now examine a **two-prover model** for zero-knowledge proofs, first described by [Goldwasser, Benor, Wigderson, and Kilian].

The basic framework for the two-prover model consists of two (or more) provers (P_1, P_2) who are unable to communicate with each other, but who can communicate with a Verifier (V) using a prescribed protocol. Graphically, the model looks like the following:



We could imagine many ways of implementing this. For example:

- The two provers could be devices (such as ATM cards) that are inserted into the verifying device (such as an ATM machine). By having the verifier physically isolate the two provers, it can be sure no communication occurs.
- The two provers could simply have a wall between them. We illustrated this last version for the sake of simplicity.

2.0.3 Implementation of Bit-Commitment (BC)

So far, we have shown that there exist zero-knowledge proofs for all languages in NP if we have bit-commitment. So, in order to implement zero-knowledge proofs in the two-prover model, we

²Scribe notes taken by: Todd Hodes, November 1 1994

can simply show how to implement bit-commitment, and all the other machinery follows from there.

The nice properties of the two-prover model is that it model is “perfect:” no cheating is possible, and everything is hidden information-theoretically. Also note that there are no cryptographic assumptions here, only physical ones.

First we’ll describe a slightly broken protocol, *Weak Two-Prover Bit-Commitment*, where we allow the provers to cheat with a probability of $\frac{1}{2}$. Then we’ll describe how to fix it to obtain *Strong Two-Prover Bit-Commitment*, where the provers chances of cheating can be made arbitrarily small (negligible).

2.0.4 Weak Bit-Commitment

The following is a protocol for *Weak Two-Party Bit-Commitment*.

The provers P_1 and P_2 have a bit b they want to commit. They choose two bits b_0 and b_1 such that $b_0 \oplus b_1 = b$. V then picks a bit i , sends it to P_1 , and asks P_1 to return b_i . This is the *commitment* stage. V then asks for b_0 and b_1 from P_2 , and verifies the reply. This is the *decommitment* stage.

In tabular form:

Assume P_1 and P_2 have chosen b , a bit to commit, and a set of bits b_0 and b_1 such that $b_0 \oplus b_1 = b$.

Weak Two-Party Bit-Commitment					
	P_1	communication	V	communication	P_2
1		$\leftarrow i \leftarrow$	Generate a random bit i		
2		$\rightarrow b_i \rightarrow$	(commitment ends)		
3			de-commit phase	$\leftarrow b_0$ and $b_1 \leftarrow$	
4			$b = b_0 \oplus b_1$		

We see that the provers can cheat by changing either b_0 or b_1 , but not both. The cheating will go unnoticed with probability $\frac{1}{2}$.

2.0.5 Strong Bit-Commitment

We now show how to implement *Strong Two-Party Bit-Commitment* using the weak protocol as a subroutine. We simply repeat the procedure for weak bit-commitment k times, and the probability of tricking the verifier decreases to $\frac{1}{2^k}$. This probability can be made arbitrarily small by varying k .

The protocol is illustrated as follows:

Assume P_1 and P_2 have chosen b , the bit to commit, and k pairs of bits (b_0, b_1) , (c_0, c_1) , $(d_0, d_1), \dots$ such that $b = b_0 \oplus b_1 = c_0 \oplus c_1 = d_0 \oplus d_1 = \dots$

		Strong Two-Party Bit-Commitment			
	P_1	communication	V	communication	P_2
1		$\leftarrow i_0, i_1, i_2, \dots \leftarrow$	Generate a random vector of k bits i_0, i_1, i_2, \dots		
2		$\rightarrow b_{i_0}, c_{i_1}, d_{i_2}, \dots \rightarrow$			
3				$\leftarrow (b_0, b_1), (c_0, c_1), (d_0, d_1), \dots \leftarrow$	
4			$b = b_0 \oplus b_1 = c_0 \oplus c_1 = d_0 \oplus d_1 = \dots$		

Again, the provers can cheat by changing either b_0 or b_1 , and either c_0 or c_1 , and either d_0 or d_1 , etc. But this time, the cheating will go unnoticed with a probability of only $\frac{1}{2^k}$.

2.0.6 Some extensions

It was mentioned earlier that $IP = PSPACE$. Incidentally, $MIP = NEXP$. That is, the set of languages with multiple-prover interactive proofs is the same as the set of languages where membership can be determined in non-deterministic exponential time (!). This is clearly a *huge* class of languages.

Additionally, multiple-prover interactive proofs with more than two provers can be simulated by two-prover interactive proofs. Showing how this simulation is done is left as an exercise.

Finally, a quick note of the importance of the class MIP. Multiple-prover interactive proofs are quite useful for two important reasons:

- The model gives a method to obtain strong complexity results, like PCP, but this is outside the scope of this course
- They provide information-theoretical security
- They translate into identification schemes

One example of the usefulness of MIP is a mentioned above two-card ATM scheme, where no information is revealed if the cards are not able to communicate.

3 Introduction to Oblivious Transfer

3.1 Oblivious Transfer

In this lecture we introduce the oblivious transfer primitive³. The concept of **oblivious transfer** was developed by [Rabin, 79] for digital signatures.

Here is a story to explain the concept: Imagine I have a very smart parrot that is old and about to die. Also imagine that I have a bit b I want to commit to my friend. I whisper the

³Scribe notes taken by: Todd Hodes, November 1, 1994

value of the bit to the parrot. The parrot flies over to my friend, and with probability $\frac{1}{2}$ tells my friend the value of b , or with probability $\frac{1}{2}$ dies and tells my friend nothing.

The sender (me) has no idea if the output was either the bit b or nothing. The receiver (my friend) does know if she gets the bit or not.

Thus, if we indicate the case where the receiver gets the bit as b , and the case where she gets nothing as $\#$, then

$$\Pr(b) = \Pr(\#) = \frac{1}{2}.$$


In a sense, one can simply think of OT as a noisy channel.

This game is useful in that it can be used to play many other games; it can be thought of to be “complete” in this sense. We will now describe some other such games.

3.1.1 One-out-of-two string oblivious transfer

This game comes from [Even, Goldreich, Lempel] and can be abbreviated 1-2-String OT.


Imagine a traitor, T , who is trying to sell one of two secrets to some enemy of our country. The enemy, E , doesn’t want T to know which secret he or she has purchased. If we call the two secrets S_0 and S_1 , and the selection bit c , then the situation looks like this:



More formally, a sender inputs two strings, S_0 and S_1 , while a receiver inputs a bit c , and gets as output S_c . At the end of the game, the sender doesn’t know anything about the selection bit c , and the receiver doesn’t know anything about message S_{1-c} .

3.1.2 One-out-of-two oblivious transfer

This game is nearly identical to 1-2-String OT, but instead of having T input strings S_0 and S_1 , he or she inputs single bits b_0 and b_1 . The output in this case is the single bit b_c . 1-2 OT looks like so:



Again, the sender doesn’t know anything about the selection bit c , and the receiver doesn’t know anything about message S_{1-c} .

It has been shown by [Crepeau] that $OT \Leftrightarrow 1-2 OT$. The constructions for both directions are as follows:

Claim 19.1 $1-2 OT \Rightarrow OT$

Proof: We prove by construction.

1. sender picks bit i at random and sets $S_i = b$ and $S_{1-i} = 0$.
2. Run the 1-2 OT protocol with inputs S_0, S_1 .
3. Have the sender tell the receiver which S_i had the actual message.
4. Thus, with probability $\frac{1}{2}$, R gets the bit b , and with probability $\frac{1}{2}$, R gets nothing. This is exactly the outcome of OT.

Now for the other direction... ■

Claim 19.2 OT \Rightarrow 1-2 OT

Proof: Again, we prove by construction.

1. S generates a random string of bits $c_1 \dots c_{poly}$, and sends it to the receiver over the OT channel.
2. R gets something of the form $c_1 c_2 \# \# c_5 \# \dots c_{poly}$ at the other end of the OT channel
3. R selects two disjoint subsets of $\frac{1}{3}$ of the indices from $1 \dots poly$; one of these sets should only contain indices of bits that were successfully received.
4. R sends both sets of indices to S in the clear.
5. S then sends $S_0 \oplus$ (all the bits in the first subset), and $S_1 \oplus$ (all the bits in the second subset.)
6. Using the Chernoff Bound, the probability of successfully getting either more than $\frac{2}{3}$ or less than $\frac{1}{3}$ bits can be made negligible:

$$Pr[|\sum X_i - E(X)| \geq \frac{E(X)}{3}] < e^{-\frac{E(X)}{18}}$$

Thus, with overwhelming probability, R can construct one subset with all of the bits known, but can't construct a second subset with all bits known. In this case, just a single message will get through, which is exactly the outcome of 1-2 OT. ■

4 More on Oblivious Transfer, Connection to ZK

4.1 OT \implies BC

Recall⁴ that in Oblivious Transfer (OT), S has a bit that is sent to R . R either receives that bit, or receives the $\#$ sign, which is interpreted as knowledge that a bit was sent, but no knowledge of what that bit was. Each possible outcome occurs with probability $\frac{1}{2}$, and S has no knowledge of what R actually received. Recall also that we showed OT to be equivalent to 1-2-OT, where S has two bits b_0 and b_1 , and R has a bit c . R receives b_c , but S has no knowledge of which bit R received.

We will show that given a black box for performing OT, we can perform Bit Commitment (BC) [kilian]. Say that S wishes to commit one bit b to R . S chooses $b_1, b_2, b_3 \dots b_n$ such that $b = b_1 \oplus b_2 \oplus b_3 \oplus \dots b_n$. Then, S sends $b_1, b_2, b_3 \dots b_n$ to R through the OT channel. With probability $1 - \frac{1}{2^n}$, at least one of the bits b_i does not go through, and thus the bit b will be information theoretically hidden.

Decommitment is performed by S sending all the bits $b_1, b_2, b_3 \dots b_n$ in the clear (i.e. through a channel such that they all get through). Sending an incorrect value for b requires changing at least one of the bits $b_1 \dots b_n$. But, since S does not know which bits actually were correctly received by R , changing any bit b_i in the decommitment stage will result in being caught with probability $\frac{1}{2}$. This can be amplified by performing the commitment many times, where the value of the committed bit has to be the same every time.

4.2 OT + PRG \implies Communication efficient ZK

We first define what communication efficient ZK is (due to [kilian,micali,ostrovsky]). This is motivated by the story of the traveling Theoretical Computer Scientist who proves theorems during his travels. He is only able to send postcards home. In order to ensure that he gets credit for his work, he sends Zero Knowledge proofs on these postcards, but since he does not receive replies, the Zero Knowledge proofs cannot be interactive, and hence the term communication efficient.

So, say there exists a Prover P , and a Verifier V . They are allowed k rounds of communication in a pre-processing stage. During this stage, neither the prover nor the verifier knows what the prover will be proving, (which corresponds to the traveler not knowing what he will prove before he leaves for his trip). After the pre-processing stage, all further communications from P to V will not be interactive. We will show an implementation where P is able to send a polynomial number of Theorems, of the form (Theorem 1, Proof 1).

To help explain this technique, we first give a ZK proof for the existence of a Hamiltonian Cycle on graph G , due to Manuel Blum:

Blum's protocol:

⁴Scribe notes taken by: Micah Adler, November 3, 1994

		$x = \{G \text{ is Hamiltonian}\}$	
	P	communication	V
1	Generate Π , a random permutation of G	\rightarrow Commitment of $\Pi \rightarrow$	
2	Find $C(\Pi(E))$, a cycle on the edges of $\Pi(G)$	\rightarrow Commitment of $C(\Pi(E)) \rightarrow$	
3		$\leftarrow b \leftarrow$	Generate b , a random bit
4	if $b = 0$	\rightarrow Decommitment of $\Pi \rightarrow$	
5	if $b = 1$	\rightarrow Decommitment of $C(\Pi(G)) \rightarrow$	

Each iteration of this technique can only succeed with probability $\frac{1}{2}$ if the graph is not Hamiltonian, and thus the probability of being able to cheat can be made very small by repeated iterations.

In order to give a communication efficient ZK proof for this problem, we will use 1-2-String OT, and Pseudo Random Generation (PRG). Recall that in PRG, function f takes a string s and returns a string $f(s)$ that has length polynomial in s , such that $f(s)$ is difficult to distinguish from a truly random string.

The pre-processing phase proceeds as follows: P selects pairs of strings $(a_0, a_1), (b_0, b_1), \dots, (z_0, z_1)$, and V selects bits b_1, b_2, \dots, b_{26} . P and V then use 1-2-String OT to transmit $a_{b_1}, b_{b_2}, \dots, z_{b_{26}}$ to V . Thus, V will know exactly one string of each pair, but P will not know which one. These strings can then be used as the seeds to a PRG, f . Then, for each pair (i_0, i_1) , P can send two messages M_0 and M_1 , as $M_0 \oplus f(i_0)$ and $M_1 \oplus f(i_1)$, where $M_0 \oplus f(i_0)$ represents the strings M_0 and $f(i_0)$ bitwise XORed together. V , who also has access to the function f will be able to decode exactly one of M_0 and M_1 , but P does not know which one.

A communication efficient ZK proof for Hamiltonian Cycle then works as follows. P sends, in the clear, a commitment of Π , a permutation on G and a commitment of $C(\Pi(G))$, a Hamiltonian cycle on the edges $\Pi(G)$. Then, P sends Π XORed with $f_{a_0}(\text{theorem} - \text{number})$, and $C(\Pi(G))$ XORed with $f_{a_1}(\text{theorem} - \text{number})$. Thus, since V will be able to decode exactly one of them, but P does not know which one, P has only a $\frac{1}{2}$ probability of being able to cheat if the graph is not Hamiltonian. But, since we had several such pairs of strings, this can be boosted by repeating this process for the other pairs of strings. Since the pseudo random functions can be used polynomially many times in the length of the original string, we can perform this same procedure for a polynomial number of proofs.

Lecture 9

1 Non Interactive Zero Knowledge (NIZK) Proofs

1.1 Definitions

The NIZK¹ model includes two entities: a prover P and a verifier V . These two entities share a random sequence of bits which is denoted by R . We assume that R is chosen by a trusted third party at random and not under the control of either P or V , though they both agree on R . The interaction between the prover and the verifier consists of a single message that the prover sends to the verifier, where the prover's goal is to write down in this message a proof that an input x ($|x| = n$), belongs to a language L . The verifier which is probabilistic polynomial time machine check the proof against the common random string R . Only a polynomial size (in n) R is used. The prover has an auxiliary input w , that is a witness to the statement $x \in L$. The prover can use this witness to compute the noninteractive proof.

Definition 1 A pair of algorithms (P, V) , where V is PPT, is called a noninteractive proof system for a language L if it satisfies:

- *Completeness: The verifier accepts a honest prover message almost always. For every $x \in L$ where $|x| = n$ and all witnesses w for x ,*

$$\Pr[V(x, R, P(x, w, R)) \text{ accepts}] > 1 - \delta(n)$$

- *Soundness: Any cheating prover P^* can not convince V to falsely accept $x \notin L$, except with negligible probability. For every algorithm P^* and every $x \notin L$,*

$$\Pr[V(x, R, P^*(x, R)) \text{ accepts}] < \delta(n)$$

where $\delta(n)$ is a negligible function in n and random variable R is uniformly distributed in $\{0, 1\}^{\text{poly}(n)}$. The string R is called the common reference string or common random string (CRS).

Definition 2 A non interactive proof system (P, V) for a polynomial time relation R is zero-knowledge if there exists a PPT simulator M such that for any $(x, w) \in R$ the ensembles $M(x)$ and $(x, R, P(x, w, R))$ are computationally indistinguishable.

¹Scribe notes by Ratzin, Chandran, Moriarty; April 4th 2005

1.2 The Hidden Bits Model

The hidden-bits model is a helpful model. It is used for the design of NIZK proof systems. In this model the CRS is uniformly selected as before, however in this model the CRS is given only to the prover but hidden from the verifier. In order to prove that $x \in L$, the prover sends to the verifier a proof that consists of two parts: some bit positions specification in the CRS that the prover has chosen to reveal and a "certificate". The verifier can check only the bits of the CRS that the prover has chosen to reveal, but can never learn the bits of CRS the prover has not revealed. The verifier also checks the common input and the "certificate".

More formally, the prover is given a CRS R where $|R| = n$. The prover sends to the verifier a set of revealed bits $I \subseteq 1, 2, \dots, n$. In addition the prover sends $R_I = \{R_i\}_{i \in I}$ a sub string of R at position I and the certificate π .

Definition 3 A pair of PPT algorithms (P, V) is a NIZK proof system in the hidden-bits model if the following conditions are satisfied:

- *Completeness: For every $x \in L$ where $|x| = k$ and all witness w for x ,*

$$\Pr[V(x, R_I, I, \pi) \text{ accepts}] > 1 - \delta(n)$$

- *Soundness: For every $x \notin L$ and every algorithm P^* (no matter how P^* chooses I),*

$$\Pr[V(x, R_I, I, \pi) \text{ accepts}] < \delta(n)$$

- *Zero-knowledge: There exists a PPT simulator M such that the ensembles $M(x)$ and (x, R_I, I, π) are computationally indistinguishable.*

Construction of NIZK system from Hidden-Bits system

The hidden-bits model is not realistic model, However it is an important model due to two reasons:

1. It is a clean model that can help the design of proof systems.
2. It is very easy to transform hidden-bits model into non-interactive system.

Given (P, V) , a hidden-bits proof system for L , we would like to convert it to (P', V') a NIZK proof system for L in the CRS model. The idea is to use the CRS to "simulate" the hidden-bits string. This is done by treating the CRS as a sequence of images of a one-way trapdoor permutation, and setting the hidden-bits string to be the hard core bits of the respective pre-images. By letting the prover have access to the trapdoor, he is able to "see" the hidden-bits and also to reveal bits in position of his choosing.

In order to generate the trapdoor function we use an algorithm $Gen(1^k)$ which is a randomized algorithm that output a pair of functions (f, f^{-1}) where f^{-1} is called the "trapdoor" for f . Furthermore, f is always a permutation over $\{0, 1\}^k$, and $f^{-1}(f(x)) = x$ for all $x \in \{0, 1\}^k$. Finally, we let h denote a hard-core bit for this trapdoor permutation family. Formally, this means that for all PPT algorithms A the following is negligible:

$$|Pr[(f, f^{-1}) \leftarrow Gen(1^k); x \leftarrow \{0, 1\}^k; y = f(x) : A(1^k, f, y) = h(x)] - \frac{1}{2}|$$

We define (P', V') as follows:

Common input: $x \in \{0, 1\}^n$

Common reference string: $r = (r_1, \dots, r_n)$, where each $r_i \in \{0, 1\}^k, |r| = n \times k$.

$P'(r, x, w)$

$(f, f^{-1}) \leftarrow Gen(1^k); \setminus *P'$ runs Gen to obtain $(f, f^{-1}) * \setminus$

For $i = 1$ to $p(k)$ do

$b_i = h(f^{-1}(r_i)); \setminus *P'$ computes an n -bits string b that are the hard core bits for $f * \setminus$

$(\pi, I) \leftarrow P(b_1, \dots, b_{p(k)}, x, w); \setminus *P'$ runs P to obtain π and $I * \setminus$

Output $(\pi, I, \{f^{-1}(r_i)\}_{i \in I}, f)$;

$V'(r, x, (f, \pi, I, \{z_i\}_{i \in I}))$

$\setminus *P'$ checks the validity of the values $\{z_i\}_{i \in I}$ and calculates the HC bits $\{b_i\}_{i \in I} * \setminus$

For all $i \in I$

If $f(z_i) = r_i$ then

let $b_i = h(z_i)$;

else stop and output 0;

Output $V(\{b_i\}_{i \in I}, I, x, \pi)$;

Claim 4 (P', V') is a NIZK proof system for L in the CRS model.

Proof Sketch : The completeness of the transformed proof system is easy to see, as the prescribed P' runs P as a subroutine. For soundness, consider first a fixed trapdoor permutation (f, f^{-1}) . As argued above, this results in a uniformly-random string R (if R is chosen after (f, f^{-1})) as seen by the cheating prover. So, soundness of the original proof system implies that a prover can only cheat, using this (f, f^{-1}) , with probability at most 2^{-2k} . But a cheating prover can choose (f, f^{-1}) as he pleases! However, summing over all 2^k possible choices of (f, f^{-1}) (we assume here that legitimate output of Gen are easily decidable and that Gen uses at most k random bits on security parameter k) shows that the probability of cheating is at most 2^{-k} over the choice of r , which is negligible.

For zero-knowledge, we can convert a simulator for P to a simulator for P' . For the unrevealed bits, we choose uniformly a string $s \in \{0, 1\}^n$ and put it in the CRS in the corresponding position. For the revealed bits, for each bit $= x$ we choose a string $s \in \{0, 1\}^n$ such that $b(r) = x$, and put $f(r)$ in the CRS in the corresponding position. The fact that b is a hard-core bit leads to an output of P' that is computationally indistinguishable from the verifier's view. ■

1.3 A basic protocol: NIZK for any $L \in NP$ in the Hidden-Bits Model

We now construct a NIZK proof system for the Hamiltonicity of directed graph languages in the hidden-bits model. Note that since this is an NP-complete property, this protocol implies similar results for any $L \in NP$.

$$L_0 = \{G \mid G \text{ is a directed graph with a Hamiltonian cycle}\}$$

Recall that a Hamiltonian cycle in a graph is a sequence of edges that forms a cycle such that the cycle passes through every vertex in the graph exactly once. In our construction, a graph with n vertices will be represented by an $n \times n$ adjacency matrix, such that entry (i, j) in the matrix is 1 iff there is an edge from vertex i to vertex j in the graph. In such representation, an n -vertex graph can be identified with a string of length n^2 . We assume that this string is drawn uniformly from the set of cycle graphs.

Given H a random Hamiltonian cycle on n nodes (We assume here that H comes from the hidden bit model and is guaranteed to be a cycle graph). H is represented by an adjacency matrix that is a permutation matrix such that there exists a single '1' in each row and column. Let H' be the encoding matrix (the "commitment" of the adjacency matrix of H) that has some hidden bits such that the prover can read these hidden bits but the verifier can not. Given such H' , the prover wants to use it in order to prove to the verifier that some graph G includes a Hamiltonian cycle. We define (P, V) as follows:

P(H', G, w)

Calculate π , a permutation that maps H onto the Hamiltonian cycle of G .

Calculate I to be the set of positions in H' that correspond (under π) to non-edges in G .

Output π, I and the values in $\pi(H')$ that correspond to non-edges in G , $\{H'_i\}_{i \in I}$.

V($\{H'_i\}_{i \in I}, I, G, \pi$)

Verify that π is a permutation.

Verify that I contains all the positions that correspond to non-edges in G .

Accept the proof iff the values of all the non-edges in G are 0.

Claim 5 (P, V) is a NIZK proof system for L_0 in the hidden-bits model.

Proof Sketch :

Completeness: It is easy to see that the prover which can be either infinitely powerful or a polynomial time machine with knowledge of w can perform the protocol.

Soundness: We know that the hidden bits string H' is guaranteed to be a cycle graph, by the assumption of the distribution on H' . If the verifier accepts, then the n 1's that remain unopened in $\pi(H')$ correspond to edges in G which implies that G contains a cycle.

Zero Knowledge: Define simulator M as follows:

M(G)

Pick a random permutation π on the vertices of G .

Let I be the set of positions that correspond to non-edges in G .

Set the values of all the revealed bits H'_I to 0.

Output $\pi, I, \{H'_i\}_{i \in I}$.

■

1.4 Modified Construction

In the previous section we had an assumption that the hidden-bits string was drawn uniformly from the set of cycle graphs. However in the actual hidden-bits model, the string is chosen uniformly at random (this property was used in the conversion from the hidden-bits model to the model in which common random string is available).

In this section we modify the previous construction; we do so by showing how to generate a random directed cycle graph from a uniformly-random string, with noticeable probability. In order to reach this goal we will work with a CRS that is viewed as a $n^3 \times n^3$ Boolean matrix such that an entry has the value "1" with probability n^{-5} and "0" otherwise, where n denotes the number of vertices in the graph. It is easy to obtain such biased bits from a uniform random string by simply parsing the original string in blocks of size $5 \cdot \log_2 n$, and setting a matrix entry to 1 only if all bits in the block are 1, and 0 otherwise.

Definitions A *permutation matrix* is a binary matrix in which each row and each column contain only a single value of "1". A *Hamiltonian matrix* is a permutation matrix that corresponds to a cycle; i.e., viewed as an adjacency matrix, it corresponds to a directed cycle graph. An $n^3 \times n^3$ matrix is *useful* if it contains an $n \times n$ Hamiltonian sub-matrix and all other $n^6 - n^2$ entries are 0.

Claim 6 $Pr[M \text{ is useful}] = \Theta(\frac{1}{n^2})$

Proof

First we will prove the following

$$Pr[M \text{ contains exactly } n \text{ 1's}] = \Omega\left(\frac{1}{n}\right)$$

The proof is as follows. Let X be a random variable denoting the number of 1's in M , and let $p = n^{-5}$. Then X follows a binomial distribution with expectation $p \cdot n^6 = n$ and variance $p(1-p)n^6 < n$. Applying Chebyshev's inequality ($Pr[|X - \mu| \geq k] \leq \frac{\sigma^2}{k^2}$), where X is a random variable with mean μ , variance σ^2 and $k > 0$, we see that

$$Pr[|X - n| > n] \leq \frac{n}{n^2} = \frac{1}{n}$$

This implies

$$\sum_{i=1}^{2n} Pr[X = i] = 1 - Pr[|X - n| > n] > 1 - \frac{1}{n}$$

Since $Pr[X = i]$ is maximum at $i = n$, we have

$$Pr[M \text{ contains exactly } n \text{ 1's}] = Pr[X = n] > \frac{\sum_{i=1}^{2n} Pr[X = i]}{2n} > \frac{(1 - \frac{1}{n})}{2n} > \frac{1}{3n}$$

This proves that $Pr[M \text{ contains exactly } n \text{ 1's}] = \Omega\left(\frac{1}{n}\right)$.

Given that M contains exactly n 1's, a "birthday paradox" argument shows that, with probability $\Omega(1)$, no row or column of M contains more than a single 1. This means that with probability $\Omega\left(\frac{1}{n}\right)$, the matrix M contains a permutation sub-matrix. Now, there are $n!$ permutation matrices, and $(n-1)!$ Hamiltonian matrices. Thus, the probability that a random permutation matrix is a Hamiltonian matrix is $\frac{1}{n}$.

Let $n - ones$ denote the event that M has exactly n 1's, let $perm$ denote the event that M contains a permutation sub-matrix, and let Ham denote the event that M has a Hamiltonian sub-matrix. Putting everything together, the probability that M is *useful* is at least

$$Pr[n - ones] \times Pr[perm|n - ones] \times Pr[Ham|perm] = \Omega\left(\frac{1}{n}\right) \times \Omega(1) \times \Omega\left(\frac{1}{n}\right) = \Omega\left(\frac{1}{n^2}\right)$$

■

Construction of a proof system in the Hidden Bits Model for Hamiltonian Cycle

- Common input:
 - A directed graph $G = (V, E)$ with $|V| = n$, where n is also the security parameter.

- CRS: A uniformly-distributed string viewed as $n^3 \times n^3$ Boolean matrix M with probability $= \frac{1}{n^5}$ of each entry to be "1".
- Prover:
 - If M is not useful, reveal all the entries of M .
 - Otherwise if M is useful, let H denote the $n \times n$ Hamiltonian sub-matrix of M . Reveal all $n^6 - n^2$ entries of M that are not in H .
 - Find a function π that maps H to G .
 - Reveal the $(n^2 - |E|)$ entries in H corresponding to the non-edges of G .
- Verifier (when M is useful):
 - Verify that π is a permutation.
 - Verify that the prover has revealed all the entries that correspond to non-edges in G .
 - Accept the proof iff the values of all the non-edges in G are 0.

We argue that the above is an NIZK (single theorem) proof system for Hamiltonian Cycle in the hidden-bits model:

Completeness: If M is not useful, the prover can easily convince the verifier by simply revealing all entries. When M is useful, the argument for the basic problem holds.

Soundness: The soundness is no longer perfect, but instead holds with all but negligible probability. Following the argument for the basic problem, soundness holds with probability 1 whenever there exists at least one useful matrix M . The probability that there is no such matrix is at most $(1 - \Omega(\frac{1}{n^2}))^{n^3} \leq e^{-\Omega(n)}$, which is negligible.

Zero-Knowledge: Here we simply need to modify the simulator that was given in the basic case. The simulator now proceeds in n^3 sequential iterations as follows: in the i^{th} iteration, it generates $n^6 \times 5 \log_2 n$ uniformly-random bits. If this defines a matrix which is not useful, the simulator simply outputs these bits and moves to the next iteration. If this defines a useful matrix M with Hamiltonian sub-matrix H , the simulator outputs all $n^6 - n^2$ entries of M that are not in H , and then runs the basic problem's simulator. Note that this simulation will ignore H , and will instead just output a permutation π and reveal a 0 for every non-edge in G .

2 Multiple NIZK Proofs Based on a Single Random String

In the NIZK system that we presented in the previous sections, we do not know if in general the zero knowledge property is preserved when the same CRS is used to prove more than one statement. We only consider NIZK proof systems which the real prover does not need to be stronger than polynomial time. We call a NIZK system which allows proving polynomially-many proofs, a "general" NIZK proof system.

2.1 The Transformation

This model is based on the concept of witness indistinguishability. Informally witness indistinguishability means that it is impossible (except with negligible probability) to distinguish which of two possible witnesses $w_1, w_2 \in w(x)$ is used in a proof for an NP-statement.

Using the assumption that one-way functions exist, one can modify any NIZK proof system for any NP complete language to a new non-interactive proof system for which witness indistinguishability holds. Based on the assumption on the existence of one-way functions, we can build a pseudo-random generator that extends n -bits seed to $3n$ -bit pseudo-random strings (which are polynomial indistinguishable from truly random $3n$ -bit string).

Given (P, V) a NIZK proof system with polynomial time prover for language L , we would like to convert it to (P'', V'') , a general NIZK system for L .

- The common random string R is composed from two segments:
 - The reference statement y which is the first $3n$ bits. This $3n$ -bit string is interpreted as a pseudo-random string. It is an NP-statement and its witness (if exists such a witness) is the n bit seed that generates y .
 - The rest of the CRS which will be used as a reference string R' for the protocol (P, V) .

- $P''(\mathbf{R}, \mathbf{x}, \mathbf{w})$
 - Construct a new NP-statement $x\#y$;
 - Calculate a witness w' to this statement
(w' is either a witness for x or w' is a seed for y);
 - Reduce $x\#y$ to an instance X of the NP language L ;
 - Reduce w' to a witness W for the language L ;
 - Run $P(X, W, R)$;

- $V''(\mathbf{R}, \mathbf{x})$
 - Construct a new NP-statement $x\#y$;
 - Reduce $x\#y$ to an instance X of the NP language L ;
 - Run $V(X, R, P(X, W, R))$;

Claim 7 (P'', V'') is a NIZK proof system under the assumptions that (P, V) is a NIZK proof system and one-way functions exist.

Proof Sketch :

Completeness: Prover P can derive a witness to X from the witness to x , and therefore executes the NIZK proof system (P, V) .

Soundness: y is chosen as a truly random string. Therefore only negligible number of strings can be pseudo-random strings.

Zero Knowledge: Define simulator M for (P'', V'') as follows:

- Replace the reference statement y by a pseudo-random string y' , this replacement is indistinguishable to polynomial time observers.
- For any statement (x, w) that is transformed into a statement X , use the seed y' instead of w in order to prove X .

■

Feige and Shamir proved that "Any non-interactive zero knowledge proof system with polynomial time prover is also a non-interactive witness indistinguishable proof system". Using this proof, we can see that since (P, V) is a zero knowledge it is witness indistinguishable. Feige, Lapidot and Shamir proved that "Any noninteractive indistinguishable proof system is also in general witness indistinguishable". By this proof we can see that even polynomial many executions of (P, V) using the same CRS are witness indistinguishable.

Lecture 12

*Lecture date: Wednesday May 10th, 2006**Scribe: Paul Bunn*

1 A CCA-1 Secure Encryption Scheme

The main purpose of this lecture is to introduce two encryption schemes, the first being secure under CCA-1 (lunchtime attack), and the second being secure under CCA-2.

1.1 Description of the CCA-1 Scheme

The following encryption scheme was described by Naor and Yung in “Public-key Cryptosystems Provably Secure Against Chosen Ciphertext Attacks.” It will be shown below that the Naor-Yung encryption scheme is secure under a CCA-1 attack, provided the following assumptions hold:

- Semantically Secure Encryption Schemes exist
- Non-Interactive, Zero-Knowledge (NIZK) Proofs exist

We describe the Naor-Yung Encryption Scheme by detailing the methods **KEYGEN**, **ENCRYPT**, and **DECRYPT**. Assume that **GEN** takes in a security parameter 1^n and returns PK_1, PK_2, SK_1 , and SK_2 , corresponding to two Semantically Secure encryption schemes.

KEYGEN:

1. **GEN**($1^n, R$) is run to obtain PK_1, PK_2, SK_1 , and SK_2 .
2. The public encryption key is PK_1 and PK_2 together with a Common Reference String (CRS) to be used for a NIZK proof.

ENCRYPT: Given a message m , the encryption of m will be a 3-tuple $(E_{PK_1}(m_1), E_{PK_2}(m_2), NIZK_{CRS}(m_1 = m_2))$, where $NIZK_{CRS}$ is a non-interactive zero-knowledge proof that $m_1 = m_2$ (for an encryption, $m = m_1 = m_2$).

DECRYPT: To decrypt a ciphertext $(E_{PK_1}(m_1), E_{PK_2}(m_2), NIZK_{CRS}(m_1 = m_2))$:

1. Verify the NIZK proof that $m_1 = m_2$.
2. Use *either* of the secret keys SK_1 or SK_2 to decrypt the corresponding $E_{PK_i}(m_i)$ (for $i = 1$ or 2) to find the original message m .

1.2 Discussion on the Security of NY Encryption

Correctness of the NY Encryption Scheme is clear. In this section we outline the proof of its security:

Theorem 1 *The above described NY Encryption Scheme is secure under a CCA-1 attack provided Semantically Secure Encryptions schemes exist and NIZK proofs exist.*

Proof (Sketch) Suppose that an Adversary \mathcal{A} exists that can win a CCA-1 attack with non-negligible advantage. We construct an Adversary \mathcal{A}' that has a non-negligible advantage in a Semantic Security attack. See Figure 1 for a schematic picture of the following procedure: Denote by Ch' and \mathcal{A}' the players in the Semantic Security game, and by Ch and \mathcal{A} the players in the CCA-1 game for which \mathcal{A} has a non-negligible advantage. We will actually denote the challenger in the CCA-1 game as Ch/\mathcal{A}' rather than simply Ch , because \mathcal{A}' will play the role of the challenger in the CCA-1 game.

1. Ch' sends PK to \mathcal{A}' .
2. Ch/\mathcal{A}' picks any PK' for which he knows SK' . He then “flips a bit” c to decide whether to set $PK_1 = PK$ and $PK_2 = PK'$ or vice-versa (to make the discussion notationally more convenient, we assume the first case). Ch/\mathcal{A}' then makes public PK_1, PK_2 , and CRS .
3. **Phase I:** \mathcal{A} queries Ch/\mathcal{A}' to decrypt chosen ciphertexts:

$$E(m) = (E_{PK_1}(m_1), E_{PK_2}(m_2), NIZK_{CRS}(m_1 = m_2))$$

4. Ch/\mathcal{A}' verifies that $m_1 = m_2$ via $NIZK_{CRS}$. Next Ch/\mathcal{A}' responds by decrypting $E(m)$ with respect to his known private key SK_2 , i.e. Ch/\mathcal{A}' responds with $m = SK_2(E_{PK_2}(m_2))$.
5. **Challenge Phase:** \mathcal{A} sends two messages m_0 and m_1 to Ch/\mathcal{A}' .
6. Ch/\mathcal{A}' relays m_0 and m_1 to Ch' .
7. Ch' picks a random $b \in \{0, 1\}$ and returns $E_{PK_1}(m_b)$ to Ch/\mathcal{A}' .
8. Ch/\mathcal{A}' responds as follows: Ch/\mathcal{A}' “flips a bit” d , and returns to \mathcal{A} :

$$(E_{PK_1}(m_b), E_{PK_2}(m_d), SIM-NIZK_{CRS}(m_b = m_d)).$$

Note that it is possible that $m_d \neq m_b$.

- If they *are* equal, then \mathcal{A} can distinguish SIM-NIZK proof from the real NIZK proof with only negligible probability.
- If they are *not* equal, then \mathcal{A} may recognize the simulated proof from the real thing. In this case, \mathcal{A} knows Ch/\mathcal{A}' “cheated,” and did not send two encryptions of the same message. However, \mathcal{A} has no idea which of the two encryptions that he receives corresponds to the encryption from Ch versus the encryption guessed by Ch/\mathcal{A}' . Therefore, \mathcal{A} cannot deliberately sabotage his game (e.g. by guessing b intentionally wrong) with Ch/\mathcal{A}' , and has a 50% chance of guessing b correctly, regardless of how he guesses (\mathcal{A} has lost his advantage in guessing b correctly). Thus, if \mathcal{A} aborts the game when he recognizes the false SIM-NIZK proof, then Ch/\mathcal{A}' knows b with 100% certainty (it is the opposite of his guess d). Otherwise, if \mathcal{A} continues to play, he has a 50% chance of guessing b correctly.

We use this analysis below.

9. \mathcal{A} guesses b' .
10. Ch/\mathcal{A}' relays b' as his guess to Ch' .

Analysis of the Above Game

We separate the analysis into the two following 2 cases:

1. $\mathbf{d} = \mathbf{b}$. Then \mathcal{A} can distinguish SIM-NIZK from NIZK with negligible probability δ_1 . We assume Ch/\mathcal{A}' loses if this happens. Otherwise, \mathcal{A} maintains his ϵ advantage, and guesses b' correctly with probability $1/2 + \epsilon$.
2. $\mathbf{d} \neq \mathbf{b}$.
 - (a) If \mathcal{A} can distinguish SIM-NIZK from NIZK and decides to abort, then Ch/\mathcal{A}' wins with 100% probability. Let's say this case happens with probability δ_2 .
 - (b) Otherwise, if \mathcal{A} keeps playing, then he has a 50% of guessing $b' = b$ correctly, and this probability is thus passed on to Ch/\mathcal{A}' . This case happens with probability $1/2 * (1 - \delta_2)$




Figure 1:

Since $Pr[\text{case 1}] = Pr[\text{case 2}] = 1/2$, we add everything up:

$$\begin{aligned}
 Pr[Ch/\mathcal{A}' \text{ guesses correctly}] &= Pr[Ch/\mathcal{A}' \text{ guesses correctly}|\text{case 1}] * Pr[\text{case 1}] \\
 &\quad + Pr[Ch/\mathcal{A}' \text{ guesses correctly}|\text{case 2a}] * Pr[\text{case 2a}] \\
 &\quad + Pr[Ch/\mathcal{A}' \text{ guesses correctly}|\text{case 2b}] * Pr[\text{case 2b}] \\
 &= [(1 - \delta_1) * (1/2 + \epsilon) * (1/2)] + [1 * \delta_2] + [(1/2 * (1 - \delta_2)) * 1/2] \\
 &= 1/2 + 3/4 * \delta_2 + \epsilon/2 - \delta_1 * (1/4 + \epsilon/2) \\
 &\geq 1/2 + \epsilon/2 - \delta,
 \end{aligned}$$

where $\delta = \delta_1 * (1/4 + \epsilon/2)$ is negligible.

■

Although the Naor-Yung encryption scheme is secure under a CCA-1 attack, it is NOT secure under a CCA-2 attack. For example, if the $NIZK_{CRS}$ proof is such that the last bit is extraneous, then an adversary could break this encryption scheme with 100% probability in a CCA-2 attack (In Phase II, Adversary simply changes the challenge encryption in the trivial way of flipping the last bit of the NIZK proof, and then the response from challenger reveals m_b). The next section introduces a new encryption scheme that will be (provably) secure under CCA-2.

2 A CCA-2 Secure Encryption Scheme

The following is the encryption scheme developed by Dolev, Dwork, and Naor (DDN) as described in “Non-Malleable Cryptography” in 1991. It will be shown below that this encryption scheme is secure under a CCA-2 attack if the following 3 assumptions hold:

- Semantically Secure Encryption Schemes exist
- Strong Digital Signature Schemes exists
- Non-Interactive, Zero-Knowledge (NIZK) Proofs exist

where a Strong Digital Signature Scheme means that not only is it “hard” to sign a given message (without knowledge of the signature key), but it is also “hard” to find a *second* valid signature of some message, even if you are given one signature of the message. Note that it is sufficient to assume the existence of a one-way trapdoor permutation, since this implies the existence of the 3 statements above. Note that the converse of this statement is not known, i.e. whether the existence of a Semantically Secure Encryption scheme implies the existence of one-way trapdoor permutations. What *is* known is that both the top

assumption and the third assumption imply the existence of one-way functions, which in turn imply the existence of Strong Signature Schemes. Thus, only the first and the third assumptions above are necessary.

2.1 Description of the DDN Encryption Scheme

We describe the DDN Encryption Scheme by detailing the methods **KEYGEN**, **ENCRYPT**, and **DECRYPT**. Assume that **GEN₁** takes in a security parameter 1^n and returns (\vec{PK}, \vec{SK}) , where \vec{PK} is a $2 \times n$ matrix of public keys (coming from some Semantically Secure encryption scheme), and \vec{SK} is the corresponding $2 \times n$ matrix of private keys. We view \vec{PK} as:

PK_1^0	PK_2^0	...	PK_n^0
PK_1^1	PK_2^1	...	PK_n^1

Table 1: \vec{PK}

Also assume **GEN₂** takes in security parameter 1^k and returns a pair (SIG-PK, SIG-SK) corresponding to a Digital Signature Scheme (e.g. Lamport's one-time scheme).

KEYGEN:

1. **GEN₁**($1^n, R$) is run to obtain (\vec{PK}, \vec{SK}) .
2. The public encryption key is \vec{PK} together with a Common Reference String (CRS) to be used for a NIZK proof.

ENCRYPT: Given a message m , the encryption of m will be a 3-tuple (SIG-PK, α , $\sigma(\alpha)$), computed as follows:

1. **GEN₂**($1^k, R$) is run prior to *each* encryption to obtain a new (SIG-PK, SIG-SK). SIG-PK is the first part of the 3-tuple encryption.
2. The middle portion of the encryption, α , consists of $n + 1$ parts. The first n parts come from n different encryptions of m : $E_{PK_1}(m_1), \dots, E_{PK_n}(m_n)$ (described below in step 3), and the last part is a NIZK proof (using CRS) that $m_1 = \dots = m_n$. Note, here $m = m_1 = \dots = m_n$.
3. SIG-PK is used as a *selector* from the $2 \times n$ matrix of public keys \vec{PK} . In particular, if we represent SIG-PK in binary as $\text{SIG-PK} = v_1 v_2 \dots v_n$, then:

$$\alpha = E_{PK_1^{v_1}}(m_1) \dots E_{PK_n^{v_n}}(m_n), \text{NIZK}_{\text{CRS}}(m_1 = \dots = m_n)$$

4. $\sigma(\alpha)$ is the signature of α using SIG-SK.

DECRYPT: To decrypt a ciphertext (SIG-PK, α , $\sigma(\alpha)$):

1. Using SIG-PK, verify that $\sigma(\alpha)$ is a valid signature of α .
2. Verify the NIZK proof that $m_1 = \dots = m_n$ (this involves using SIG-PK as a selector).
3. Use *any* of the secret keys SK_i^{vi} (corresponding to the *selection* by SIG-PK) to decrypt the corresponding $E_{PK_i^{vi}}(m_i)$ to find the original message m .

2.2 Discussion on the Security of DDN Encryption

Correctness of the DDN Encryption Scheme is clear. In this section we outline the proof of its security:

Theorem 2 *The above described DDN Encryption Scheme is secure under a CCA-2 attack provided Semantically Secure Encryptions schemes exist, Digital Signature schemes exist, and NIZK proofs exist.*

Proof (Sketch) Suppose that an Adversary \mathcal{A} exists that can win a CCA-2 attack with non-negligible advantage. We construct an Adversary \mathcal{A}' that has a non-negligible advantage in an Extended Semantic Security attack (see def. of Extended Semantic Security, and its equivalence to the ordinary Semantic Security attack, in the Appendix below). See Figure 2 for a schematic picture of the following procedure: Denote by Ch' and \mathcal{A}' the players in the Extended Semantic Security game, and by Ch/\mathcal{A}' and \mathcal{A} the players in the CCA-2 game for which \mathcal{A} has a non-negligible advantage.

0. Ch/\mathcal{A}' runs $\mathbf{GEN}_2(1^k, R)$ to obtain (SIG-PK $_{ch}$, SIG-SK $_{ch}$). He will use these to encrypt during the challenge phase. Let SIG-PK $_{ch}$ be denoted in binary: SIG-PK $_{ch} = w_1w_2 \dots w_n$.
1. Ch' sends $\{PK_1, \dots, PK_n\}$ to \mathcal{A}' .
2. \mathcal{A}' creates a $2 \times n$ matrix, and places $\{PK_1, \dots, PK_n\}$ into the slots according to SIG-PK $_{ch} = w_1w_2 \dots w_n$. For instance, \mathcal{A}' sets $PK_i^{w_i} = PK_i$ (see chart below).
 \mathcal{A}' then picks an additional n public keys (for which he knows the corresponding secret keys) and fills in the rest of the matrix. This is the \vec{PK} that is made public in the CCA-2 game, along with CRS. Note that Ch/\mathcal{A}' only knows half of the values of \vec{SK} .

3. **Phase I:**

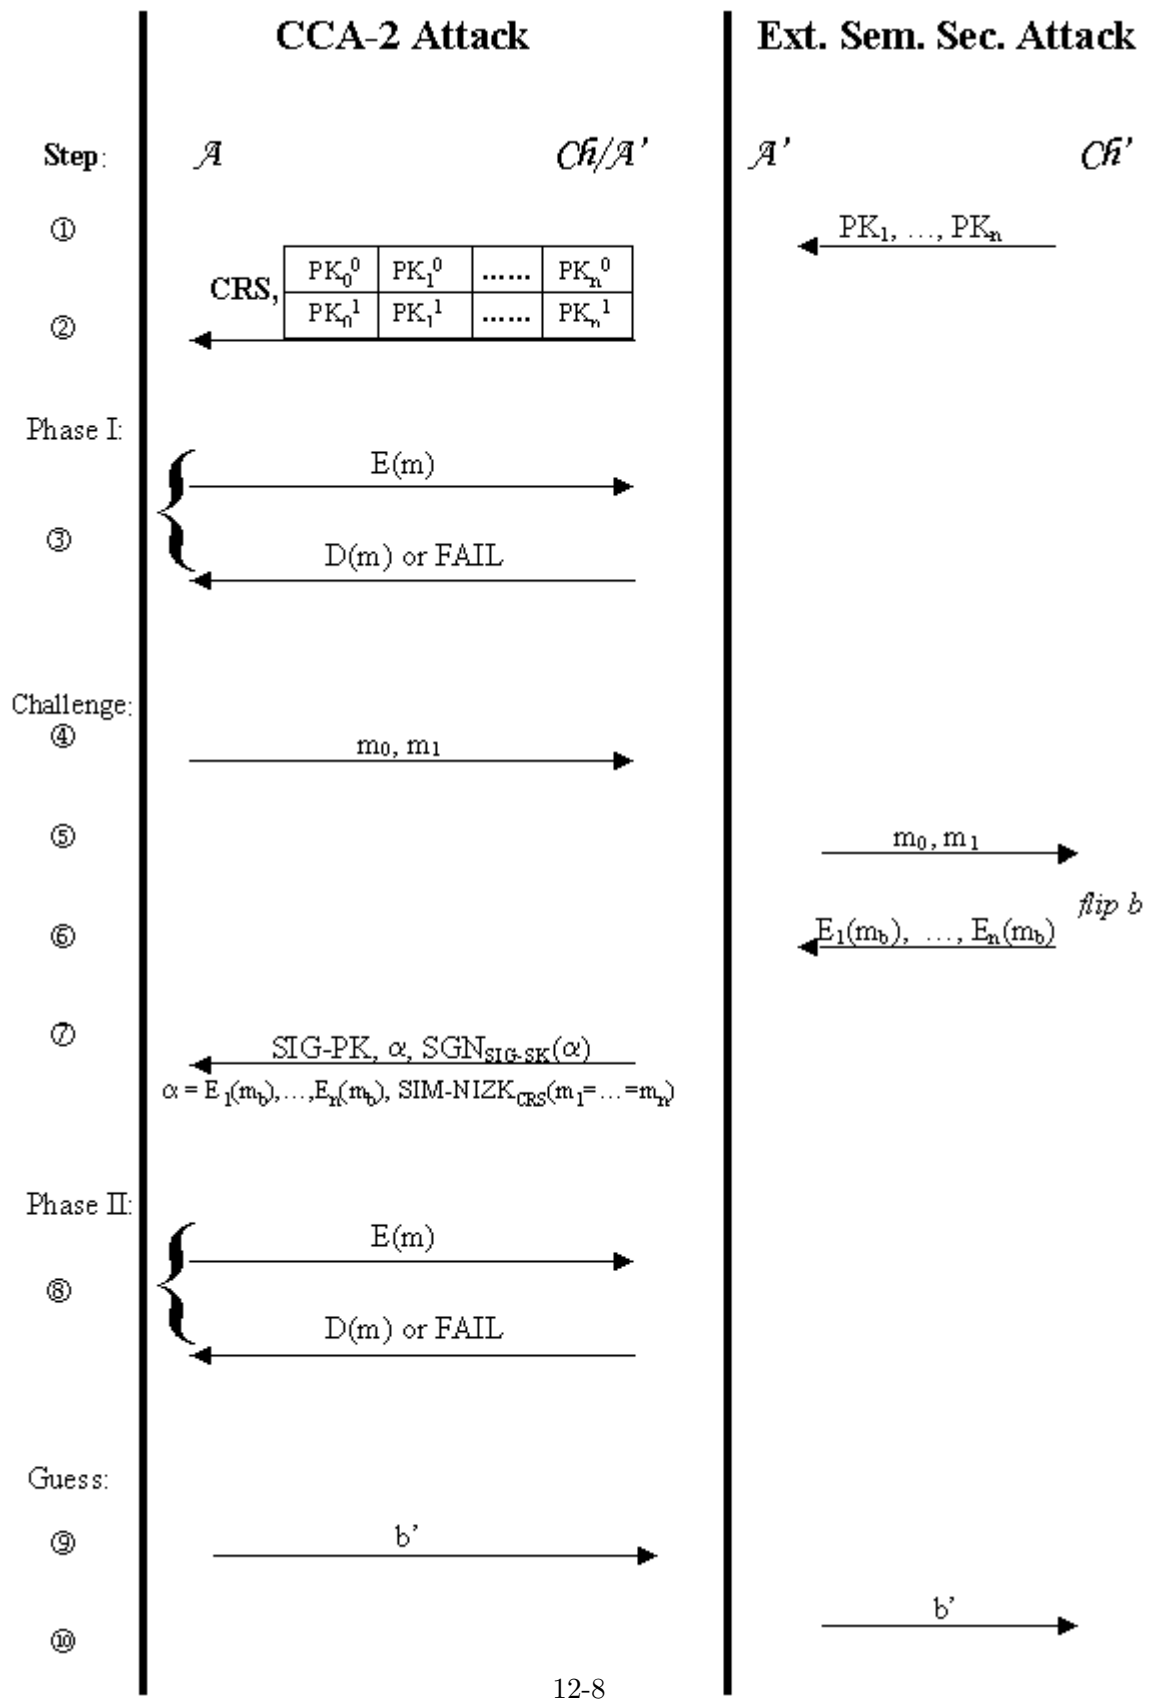


Figure 2:

PK_1			\dots		PK_n
	PK_2	PK_3	\dots	PK_{n-1}	

Table 2: Example: $SIG-PK_{ch} = 011\dots10$

- \mathcal{A} queries Ch/\mathcal{A}' to decrypt chosen ciphertexts $E(m) = (SIG-PK, \alpha, \sigma(\alpha))$.
 - Ch/\mathcal{A}' verifies the signature $\sigma(\alpha)$ using $SIG-PK$, then verifies that $m_1 = \dots = m_n$ via $NIZK_{CRS}$. Next Ch/\mathcal{A}' responds as follows:
 - If $SIG-PK \neq SIG-PK_{ch}$ (for example bit i differs), then Ch/\mathcal{A}' can decrypt since he knows $SK_i^{-w_i}$. In particular, Ch/\mathcal{A}' returns $m = D_{SK_i^{-w_i}}(E_{PK_i^{-w_i}}(m))$.
 - If $SIG-PK = SIG-PK_{ch}$, then Ch/\mathcal{A}' cannot decrypt. In this case, \mathcal{A}' must proceed in his Extended Semantic Security game with no additional information, and CCA-2 game is terminated. Note that if this happens, not only has \mathcal{A} stumbled upon $SIG-PK_{ch}$, but he has also figured out how to sign with this. The probability of this happening is discussed below.
4. **Challenge Phase:** \mathcal{A} sends two messages m_0 and m_1 to Ch/\mathcal{A}' .
 5. Ch/\mathcal{A}' relays m_0 and m_1 to Ch' .
 6. Ch' picks a random $b \in \{0, 1\}$ and returns $E_{PK_1}(m_b) \dots E_{PK_n}(m_b)$ to Ch/\mathcal{A}' .
 7. Ch/\mathcal{A}' responds to \mathcal{A} with $E(m_b) = (SIG-PK_{ch}, \alpha, \sigma(\alpha))$ by setting:
$$\alpha = E_{PK_1}(m_b) \dots E_{PK_n}(m_b), \text{SIM-NIZK}_{CRS}(m'_b \text{ s are equal}),$$
where $SIM-NIZK_{CRS}$ is a simulated proof that the m'_b s are equal. (Since Ch/\mathcal{A}' does not know what m_b is, he cannot provide a “real” proof. However, it is a proof of a valid statement since indeed the m'_b s are all equal). $\sigma(\alpha)$ is then the signature (using $SIG-SK_{ch}$) of α .
 8. **Phase II:**
 - \mathcal{A} queries for more $E(m) = (SIG-PK, \alpha, \sigma(\alpha))$ with the one restriction that $E(m)$ differs in some way from $E(m_b)$.
 - Ch/\mathcal{A}' responds as in Phase I.
 9. \mathcal{A} guesses b' .
 10. Ch/\mathcal{A}' relays b' as his guess to Ch' .

Analysis of Above Algorithm

As long as the above game can be played without early termination, the ε -gap that \mathcal{A} had in the CCA-2 attack is transferred to \mathcal{A}' in the Extended Semantic Security attack:

$$\Pr[b' = b | \text{CCA-2 attack did not terminate early}] = \Pr[\mathcal{A} \text{ guesses correctly}] = 1/2 + \varepsilon$$

where ε is some non-negligible function. Let's let γ be the probability that the above game is terminated early (γ will be calculated shortly). Then because the probability of \mathcal{A}' guessing correctly in the case that the above game terminates early is $1/2$, we have:

$$\begin{aligned} \Pr[\mathcal{A}' \text{ guesses correctly}] &= \Pr[b' = b | \text{CCA-2 attack didn't fail}] \Pr[\text{CCA-2 attack didn't fail}] \\ &\quad + \Pr[b' = b | \text{CCA-2 attack **did** fail}] \Pr[\text{CCA-2 attack **did** fail}] \\ &= (1/2 + \varepsilon) * (1 - \gamma) + (1/2) * (\gamma) \\ &= 1/2 + (\varepsilon * (1 - \gamma)) \end{aligned} \tag{1}$$

It remains to show that $(\varepsilon * (1 - \gamma))$ is non-negligible, i.e. that γ is negligible. To understand γ , we must understand all the ways the CCA-2 game could terminate early. This could only happen for one of the following reasons:

1. In **Phase I**, \mathcal{A} queries an encryption that Ch/\mathcal{A}' cannot decrypt. Note that this happens only if \mathcal{A} happens to query an encryption whose SIG-PK **EXACTLY** matches SIG-PK_{ch}. Since \mathcal{A} has no way of knowing SIG-PK_{ch}, this has probability $1/2^n$. Furthermore, this also means that \mathcal{A} figured out how to sign using SIG-PK_{ch}, which by our assumption of signature schemes, can only happen with negligible probability. Thus, the probability of terminating early due to this reason is a negligible function, call it δ_0 .
2. In **Challenge Phase**, \mathcal{A} can distinguish SIM-NIZK_{CRS} from a genuine NIZK_{CRS}. By the third assumption, this happens with negligible probability δ_1 .
3. Since \mathcal{A} must alter his queries in **Phase II** from the response of the challenge, we investigate the ways he can do this in a case analysis:
 - \mathcal{A} **modifies $E(m_b)$ by changing only $\sigma(\alpha)$** . But by our assumption that the signature scheme is *Strong*, \mathcal{A} has a negligible probability of successfully coming up with a new signature for α , even though he has already seen one valid signature.
 - \mathcal{A} **modifies $E(m_b)$ by changing only SIG-PK, or by changing only α , or changing both (but not $\sigma(\alpha)$)**. In this case, Ch/\mathcal{A}' easily rejects this encryption as invalid, on the grounds that the signature is no longer valid (with overwhelming probability).

- \mathcal{A} modifies $\mathbf{E}(m_b)$ by changing α and $\sigma(\alpha)$. But this requires \mathcal{A} to forge a signature for his new α , which by our assumption on signature schemes can happen with only negligible probability.
- \mathcal{A} modifies $\mathbf{E}(m_b)$ by changing **SIG-PK** and $\sigma(\alpha)$. Because \mathcal{A} did not change α , the SIM-NIZK proof is no longer valid, since it relied on SIG-PK_{ch} as a selector. Thus, Ch/\mathcal{A}' rejects this as an invalid encryption.

Thus, the total probability that \mathcal{A} can utilize Phase II in a different way than Phase I is negligible, call it δ_2 . Otherwise, \mathcal{A} simply queries new encryptions, in which case Ch/\mathcal{A}' can again decrypt these with probability $(1 - \delta_0)$.

Thus, we have:

$$\begin{aligned}
 \gamma = \Pr[\text{CCA-2 attack failed}] &\leq \Pr[Ch/\mathcal{A}' \text{ cannot decrypt in Phase I or Phase II}] \\
 &\quad + \Pr[\mathcal{A} \text{ can distinguish SIM-NIZK}] \\
 &\quad + \Pr[\mathcal{A} \text{ gets extra info. in Phase II}] \\
 &= (\# \text{ of queries } \mathcal{A} \text{ makes in the two phases}) * (\delta_0) + \delta_1 + \delta_2 \\
 &= (\text{polynomial in } n) * (\delta_0) + \delta_1 + \delta_2 \\
 &= \delta_3
 \end{aligned}$$

where δ_3 is negligible since any polynomial times a negligible function is still negligible, and the sum of any three negligible functions is still negligible. ■

3 An Introduction to Commitment Schemes

We conclude this lecture with a new topic: Commitment Schemes. To motivate the development of commitment schemes, we consider the following example. Suppose Alice and Bob are talking on the phone and want to “flip a coin.” How can they do this in a fair way? Suppose Bob is going to flip the coin. If he reveals the result of the flip before Alice communicates her guess, then Alice can alter her guess to guarantee herself victory. On the other hand, if Alice communicates her guess first, then Bob can lie about the outcome of the flip.

We use the above situation as motivation for the following definition (see Figure X):

Definition A *Commitment Scheme* is a protocol for communication between two people (Alice and Bob) involving:

- The Commitment. This is a message $com(b)$ sent from Alice to Bob

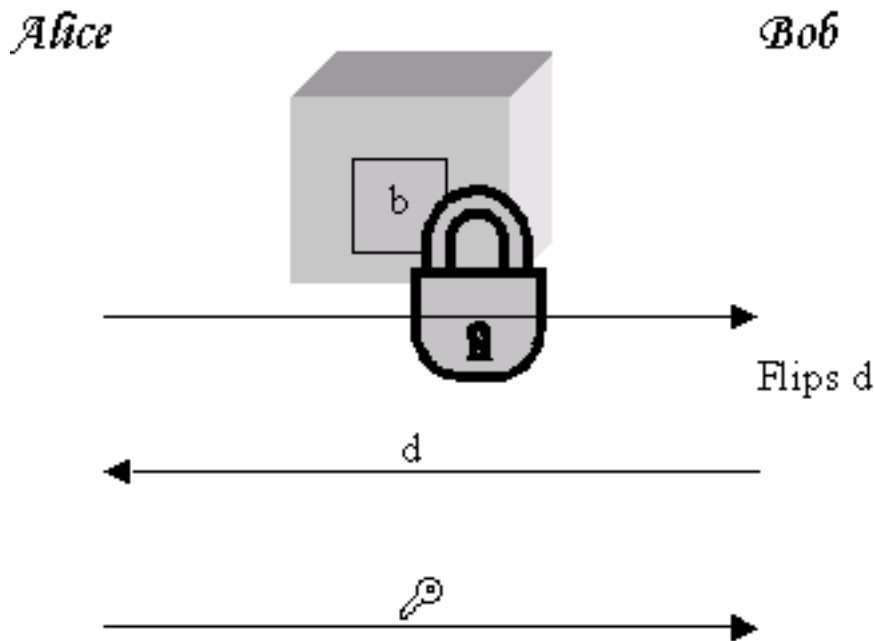


Figure 3:

- The Decommitment. This is a message $dec(b)$ sent from Alice to Bob at some later time

and satisfying:

- **Correctness.** Given $com(b)$, Alice should only be able to decommit in 1 way.
- **Privacy.** Bob can only determine b based on $com(b)$ with probability $\leq 1/2 + \epsilon$, for some negligible function ϵ .

See Figure 3 for a schematic representation of a Commitment Scheme, and Figure 4 for a sketch of a commitment scheme.

3.1 Naor's Commitment Scheme

Using Pseudo Random Generators (PRG), Naor developed a commitment scheme satisfying the above definition. Recall that a Pseudo Random Generator is a deterministic function




Figure 4:

from n bits to m bits ($m > n$) whose output distribution is indistinguishable by any poly-time machine from a random distribution. Here is a brief description of this scheme (see Figure 5):

1. Let G be a PRG:

$$G : \{0, 1\}^n \rightarrow \{0, 1\}^{3n}$$

2. Bob picks a random string R of $3n$ bits and sends this to Alice.
3. (**Commitment**) Alice picks a bit b and a random “seed” s (a string of bits of length n), and does the following:
 - If $b = 0$, then Alice sends to Bob $G(s)$
 - If $b = 1$, then Alice sends to Bob $G(s) \oplus R$ (bitwise XOR)
4. (**Decommitment**) Alice sends s to Bob.

Analysis of Naor Commitment

Claim The above protocol satisfies the Correctness and Privacy requirements of a Commitment Scheme.




Figure 5:

Proof

Privacy: Suppose that Bob can determine b with probability $\geq 1/2 + \epsilon$, for some non-negligible function ϵ . Notice that if $G(s)$ were truly random, then Bob would have no way of distinguishing $G(s)$ from $G(s) \oplus R$. Thus, his non-negligible advantage must come from an ability to distinguish the distribution from PRG from a truly random distribution, which contradicts our assumption concerning the PRG.

Correctness: We claim that Alice can “cheat” with negligible probability (specifically, with probability $\leq 1/2^n$). For what does it mean if Alice can decommit in 2 different ways? It means that after Alice has sent $com(b)$ to Bob, she can either open it (decommit) as $b = 0$ OR as $b = 1$. But $com(b)$ is a string of length $3n$, and to be able to decommit in either of the two possible ways, Alice must have found an s_1 and some other s_2 such that:

$$com(b) = G(s_1) \quad \text{AND}$$

$$com(b) = G(s_2) \oplus R$$

Or equivalently:

$$G(s_1) \oplus G(s_2) = R$$

Thus, Alice must find/select 2 seeds s_1 and s_2 that satisfy the above equation. Lets approach this problem from a different perspective: Given the 2^{3n} possible values of R , how many possible values can $G(s_1) \oplus G(s_2)$ achieve? Since there are 2^n possible seeds of length n , from which Alice must select 2 of them, there are:

$$\binom{2^n}{2} \leq (2^n)^2 = 2^{2n}$$

possible choices for s_1 and s_2 , thus yielding 2^{2n} possible values of $G(s_1) \oplus G(s_2)$. Thus, Alice has a probability less than $2^{2n}/2^{3n} = 1/2^n$ chance of finding strings s_1 and s_2 that will allow her to cheat.

4 Appendix

Here we define Extended Semantic Security, and prove its equivalence to Semantic Security.

Definition. An Extended Semantic Security Attack is identical to a Semantic Security Attack, with the exception that there are k public keys $\{PK_1, \dots, PK_n\}$ (where n can be anything), and the challenger responds in the challenge phase with:

$$E(m) = E_{PK_1}(m) \dots E_{PK_n}(m).$$

Theorem 3 *An encryption scheme defined by $E(m) = E_{PK_1}(m) \dots E_{PK_n}(m)$ is secure under an Extended Semantic Security attack \Leftrightarrow the encryption scheme $E(m) = E_{PK_i}(m)$ is secure under a Semantic Security attack for every i .*

Proof Note that the forward direction is trivial via a contrapositive argument: If an adversary has a non-negligible advantage in breaking $E(m) = E_{PK_i}(m)$ for some i in a Sem. Sec. attack, then it is easy to construct an adversary who has the same advantage in breaking $E(m) = E_{PK_1}(m) \dots E_{PK_n}(m)$ in an Extended Sem. Sec. attack.

We prove the backward direction via a hybrid argument. Suppose that an adversary (lets call him \mathcal{A}) exists with a non-negligible ϵ advantage in winning an Extended Sem. Sec. attack. We modify the Extended Semantic security game as follows. The challenger still makes public $\{PK_1, \dots, PK_n\}$. \mathcal{A} then sends his challenge m_0, m_1 to the challenger as before. However, now the challenger flips n bits, $\{b_1, \dots, b_n\}$, and returns $E_{PK_1}(m_{b_1}), \dots, E_{PK_n}(m_{b_n})$ to \mathcal{A} . As a guess, \mathcal{A} sends c_1 , simply trying to guess b_1 . Note that if $b_1 = \dots = b_n$, then \mathcal{A} maintains his ϵ advantage.

We measure the advantage of this adversary \mathcal{A} in the following $n + 1$ games:

- **Game 0:** In the response to m_0, m_1 , the challenger flips a bit b and responds with: $E_{PK_1}(m_{b_1}), \dots, E_{PK_n}(m_{b_n})$, where $b = b_1 = \dots = b_n$.
- **Game 1:** In the response to m_0, m_1 , the challenger flips a bit b and responds with: $E_{PK_1}(m_{b_1}), \dots, E_{PK_n}(m_{b_n})$, where $b = b_1 = \dots = b_{n-1}$ AND $\neg b = b_n$.
- **Game i:** In the response to m_0, m_1 , the challenger flips a bit b and responds with: $E_{PK_1}(m_{b_1}), \dots, E_{PK_n}(m_{b_n})$, where $b = b_1 = \dots = b_{n-i}$ AND $\neg b = b_{n-i+1} \dots = b_n$.

We say \mathcal{A} “wins” if he guesses b correctly. Note that in Game 0, \mathcal{A} has a non-negligible advantage in winning, whereas in Game n, \mathcal{A} has a non-negligible chance of NOT winning. Thus, in one of the intermediate games, \mathcal{A} ’s probability of winning must change from non-negligible to negligible. If this happened in Game i, we say that \mathcal{A} is “sensitive” to PK_{n-i+1} . It is clear to see how to break an encryption scheme that uses PK_{n-i+1} in a Sem. Sec. attack, hence completing the contrapositive. ■

Lecture 11

1 Commitment Schemes: a Review

We recall that a commitment scheme ¹ is a protocol by which a sender A can send a piece of information $\text{Com}(m)$ to a receiver B which effectively “commits” A to having chosen the value m , while revealing nothing about m to B . Then at a later time the sender can send $\text{Dec}(m)$ which allows B to efficiently calculate m from $\text{Com}(m)$. There are two separate properties which are important here, the first is the *binding* property, this is the idea that the commitment $\text{Com}(m)$ binds A to the message m . Commitment schemes can be computationally binding, meaning that a computationally bounded sender A cannot decommit $\text{Com}(m)$ to a different message m' , or perfectly binding meaning that each commitment $\text{Com}(m)$ has only one valid decommitment. In this lecture we examine only computationally binding commitment schemes. The second property of a commitment scheme is the *hiding* property, this is the idea that give $\text{Com}(m)$ the receiver B cannot recover m . Again, we can make a distinction between computationally hiding commitment schemes, where no computationally bounded receiver B can recover m from $\text{Com}(m)$, and perfectly hiding commitment schemes, where even an infinitely powerful receiver cannot recover m from $\text{Com}(m)$. The protocols examined in this lecture will be perfectly hiding. In this lecture we will assume that the sender A and the receiver B have access to a Common Reference String, CRS . It will be important, as we shall see, that the Common Reference String be generated by a trusted third party, for having “extra” information about the CRS can be used to great advantage. It is exactly this extra information which will enable us to create Equivocable and Non-Malleable Commitment schemes.

For a more in depth examination of Commitment Protocols see Lecture 8.

2 Non-Malleable Commitment

While every commitment scheme must have the binding and hiding properties, in practice sometimes you would like a commitment scheme to do more. Suppose you are committing to bids at an auction. If you commit to a bid b , if given your commitment $\text{Com}(b)$ another bidder can compute $\text{Com}(b + 1)$ you will lose the auction, even if they cannot recover your bid b . This type of breach of security is not covered in the hiding and binding properties of a commitment scheme. To express this property a new definition called *Non-Malleability* was introduced by Rackoff and Simon. Loosely we would like to have a commitment protocol

¹Scribe notes by B. Hemenway, 15 May 2006

such that given any relation R on messages it is just as hard to find $\text{Com}(y)$ such that $R(x, y) = 1$ given $\text{Com}(x)$ as it is to find such a y without $\text{Com}(x)$.

If we denote our message space \mathcal{M} , $\text{Com}(\cdot) : \mathcal{M} \rightarrow \mathcal{C}$, and R is a polynomial-time computable relation $R : \mathcal{M} \times \mathcal{M} \rightarrow \{0, 1\}$, then we define

Definition 1 A Commitment Protocol (Com, Dec) with security-parameter k is said to be Non-Malleable if for any message m_1 and any relation R , for all polynomial-time machines $A : \mathcal{C} \rightarrow \mathcal{M}$ there exists a simulator S such that

$$\left| \Pr(R(A(\text{Com}(m_1)), m_1) = 1) - \Pr(R(S(1^k), m_1) = 1) \right| < \frac{1}{\text{poly}(k)}$$

2.1 Equivocable Commitment

An Equivocal commitment protocol is one where a carefully chosen CRS will allow a sender to break the binding property or “equivocate” their commitment. Informally, a commitment protocol is equivocable if *given* the CRS if you calculate $\text{Com}(x)$ you can only decommit in one way, but there exists a Equivocator E which can generate CRS' which is indistinguishable from a “good” CRS but for which E can decommit $\text{Com}(x)$ as y for any y in the message space \mathcal{M} .

For ease of notation, we break our equivocator into two parts E_1 and E_2 where E_1 generates CRS' along with some secret information s , and E_2 which, given s can equivocate any commitment made under CRS' .

Definition 2 A commitment protocol is called Equivocable if there exists an Equivocator E_1, E_2 such that

1. E_1 outputs CRS' , s .
2. Given a commitment $\text{Com}(m_1)$, for any message m_2 in the message space \mathcal{M} , $E_2(s, m_2) = \text{Dec}(m_1)$.

Where the two distributions CRS and CRS' are identical.

2.2 Pedersen Commitment

We now introduce an Equivocable Commitment Scheme created by Torben Pedersen [?]. Let q be a prime, and p a prime with $p \equiv 1 \pmod{q}$. We know $(\mathbb{Z}/p\mathbb{Z})^*$ is cyclic of order

$p - 1$, so we let G be the unique subgroup of order q in $(\mathbb{Z}/p\mathbb{Z})^*$. Then G is cyclic, and we let g, h be distinct generators of G . We define the Pedersen Commitment scheme as follows,

$$\mathcal{CRS} = G, p, q, g, h$$

and given $x \in G$, choose $r \in G$ at random and compute

$$\text{Com}(x) = g^x h^r$$

Then

$$\text{Dec}(x) = (x, r)$$

Claim 3 *Assuming that it is hard to calculate discrete logs in G no polynomial-time adversary can open any commitment in two ways.*

Proof We will show that opening a commitment in two ways is as hard as calculating a discrete log in G . If we can find x, r and x', r' such that

$$g^x h^r = g^{x'} h^{r'}$$

then we can calculate $\text{DLog}_g(h)$, for if $h = g^a$, then

$$\begin{aligned} g^x h^r = g^{x'} h^{r'} &\Leftrightarrow x + ra = x' + r'a \\ &\Leftrightarrow a = \frac{x - x'}{r - r'} \end{aligned}$$

Thus decommitting in two different ways is as hard as calculating a discrete log in G . ■

Claim 4 *For any $x, r \in G$ and any x' there is an r' such that*

$$g^x h^r = g^{x'} h^{r'}$$

Proof This is clear as g and h both generate G . If $h = g^a$, then

$$r' = (x - x')a^{-1} + r$$

■

While Pedersen Commitment is secure, it is malleable. Given a commitment C to a message m , gC is a commitment to the message $m + 1$.

$$\begin{array}{ccc}
 C & Adv & B \\
 \xrightarrow{C=\text{Com}(m)} & & \\
 & \xrightarrow{g \cdot C} & \\
 \xrightarrow{m,r} & & \\
 & \xrightarrow{m+1,r} &
 \end{array}$$

The Pedersen Commitment scheme is also Equivocable since anyone who knows $\text{DLog}_g(h)$ can decommit $\text{Com}(x)$ to any $y \in G$, the equivocator E_1 need only generate

$$\mathcal{CRS} = G, p, q, (g = h^a), h$$

Where the secret information is the exponent a . Then given a commitment $\text{Com}(x) = g^x h^r$, since $g^x h^r = g^y h^{a(x-y)+r}$ the equivocator E_2 can decommit this message as $(y, a(x-y) + r)$ for any $y \in G$.

2.3 Extended Pedersen Commitment

We now examine a simple extension of the Pedersen Commitment scheme that allows us to commit multiple two messages at once. As before, let q be a prime, $p \equiv 1 \pmod q$ a prime, and G the cyclic subgroup of elements of order q in $(\mathbb{Z}/p\mathbb{Z})^*$. Now let g_1, g_2, h be distinct generators of G . Then we can commit two messages as

$$\text{Com}(m_1, m_2) = g_1^{m_1} g_2^{m_2} h^r$$

for a randomly generated r . To decommit, we send

$$\text{Dec}(m_1, m_2) = (m_1, m_2, r)$$

A similar argument shows that the extended Pedersen Commitment scheme is secure, malleable and equivocable.

2.4 An Efficient Non-Malleable Commitment Scheme

We now present a Non-Malleable Commitment Scheme originally described in [?]. The idea will be to use the equivocable nature of the Extended Pedersen Commitment scheme to

allow us to construct a simulator which can perform as well as any adversary. As in the Pedersen Commitment scheme let q be a prime, p a prime with $p \equiv 1 \pmod q$, and G the subgroup of $(\mathbb{Z}/p\mathbb{Z})^*$ of elements of order q . Let g_1, g_2, g_3 be three distinct generators of the cyclic group G , let $H : G \rightarrow \mathbb{Z}/q\mathbb{Z}$ a hash function and MAC a Message Authentication Code.

A few notes about the components. While $G \simeq \mathbb{Z}/q\mathbb{Z}$ as groups, we represent elements of G as elements in $\mathbb{Z}/p\mathbb{Z}$, so the hash function H does reduce length in some way. If r_1, r_2, B are n -bit strings, for concreteness we may assume the $\text{MAC}_{r_1, r_2}(B) = r_1 B + r_2$. where r_1, r_2, B are viewed as elements in $GF(2^n)$.

The commitment protocol is as follows:

We set our Common Reference String as

$$\text{CRS} = G, p, q, g_1, g_2, g_3, H$$

To commit to a message m , committer chooses $r_1, r_2, r_3, r_4 \in G$ at random, and sends

$$\text{Com}(m) = \left(\underbrace{g_1^{r_1} g_2^{r_2} g_3^{r_3}}_A, \underbrace{(g_1^{H(A)} g_2)^m g_3^{r_4}}_B, \underbrace{\text{MAC}_{r_1, r_2}(B)}_C \right)$$

Here you can view A as a commitment to r_1, r_2 using the Extended Pedersen Protocol.

Claim 5 *This Commitment scheme is non-malleable, i.e. given a commitment (A, B, C) to a message m , creating a commitment (A', B', C') to a related message m' that a simulator could not have found without (A, B, C) is as hard as calculating DLog in G , finding a collision in H or breaking the MAC.*

Proof Given a commitment (A, B, C) to a message m . Suppose there exists a polynomial time adversary \mathcal{A} who can find a commitment $(A', B', C') \neq (A, B, C)$. And given (m, r_1, r_2, r_3) , \mathcal{A} decommits (A', B', C') as (m', r'_1, r'_2, r'_3) where $R(m, m') = 1$.

To reach a contradiction, we consider three cases

Case 1 Suppose $A = A'$.

Then we consider two cases if $(r_1, r_2, r_3) = (r'_1, r'_2, r'_3)$, then if $B = B'$, we have

$$C = \text{MAC}_{r_1, r_2}(B) = \text{MAC}_{r'_1, r'_2}(B') = C'$$

Thus $(A, B, C) = (A', B', C')$, a contradiction. On the other hand, we showed that r_1, r_2 are information-theoretically hidden from the \mathcal{A} until \mathcal{A} receives the decommitment (m, r_1, r_2, r_3) , so finding $C' = \text{MAC}_{r'_1, r'_2}(B') = \text{MAC}_{r_1, r_2}(B')$ violates the security

of the MAC.

If $(r_1, r_2, r_3) \neq (r'_1, r'_2, r'_3)$, then since A is a commitment to both r_1, r_2 and r'_1, r'_2 using the Extended Pedersen Protocol, we can use \mathcal{A} as a subroutine to violate the binding property of the Extended Pedersen Protocol.

Case 2 Suppose $A \neq A'$ but $H(A) = H(A')$.

This represents a collision in our hash function H which was assumed to be collision resistant.

Case 3 Suppose $A \neq A'$ and $H(A) \neq H(A')$.

In this case we use the equivocability of the Extended Pedersen Protocol construct a simulator \mathcal{A}' which, after interacting with the adversary \mathcal{A} , can generate a commitment to a related message m' with essentially the same probability of success as \mathcal{A} , but *without* access to (A, B, C) .

We begin by constructing an equivocable commitment generator as follows, let Equiv and set

$$\begin{array}{l} \text{Equiv}(1^k) \\ \hline p, q, G \text{ are selected as normal} \\ r, s, t, r_2, u \leftarrow \mathbb{Z}/q\mathbb{Z} \\ A = g_1^r, g_3^s \\ g_2 = g_1^{-H(A)} g_3^t \\ r_1 = r + H(A)r_2 \\ r_3 = s - tr_2 \\ B = g_3^u, C = \text{MAC}_{r_1, r_2}(B) \end{array}$$

Then given a message m , the equivocator can decommit (A, B, C) as a commitment to m by setting $r_4 = u - tm$, and sending (m, r_1, r_2, r_3, r_4) . Thus the Equivocator can create commitments that can be decommitted to any message. It is important to notice that \mathcal{CRS}' generated by Equiv has exactly the same distribution as the valid \mathcal{CRS} , and hence is computationally indistinguishable from a “real” Common Reference String.

We now show how to construct the simulator \mathcal{A}' .

First, consider the interaction of \mathcal{A} with the real committer. The committer commits to a message m_1 , gives the commitment to the adversary, and the adversary must come up with a commitment to a related message m_2 . Since the commitment scheme is equivocable, the possible distribution of m_1 given $\text{Com}(m_1)$ is exactly the same as the original distribution of messages \mathcal{M} . Thus we can ignore this step. Now the adversary’s interaction with the committer looks like

1. \mathcal{A} commits to a message m_2 .

2. m_1 is chosen at random from the message space \mathcal{M} and revealed to \mathcal{A} .
3. The adversary \mathcal{A} decides whether to decommit.
4. \mathcal{A} succeeds if it decommits m_2 and $R(m_1, m_2) = 1$

The simulator \mathcal{A}' works as follows,

1. Using the equivocal commitment generator **Equiv**, \mathcal{A} repeats the above procedure until the adversary reveals a message m_2 .
2. A message m_1 is chosen at random from the message space \mathcal{M} .
3. The simulator succeeds $R(m_1, m_2) = 1$.

Note, that if the adversary \mathcal{A} always decommits its commitment, then it is easy to see that the simulator will succeed with the same probability as the adversary. It is possible that the adversary may refuse to decommit $\text{Com}_{\mathcal{A}}$. While this will decrease the adversary's probability of success when interacting with the real committer, it is possible that it may decrease the simulator's probability of success even further.

We must show that the probability of success of the adversary (in the first game) is no larger than the probability of success of the simulator (in the second game). Before we begin, we make a few definitions to simplify notation

- Let $R'(m', m) = R(m', m)$ and \mathcal{A} decommits to m when given m' .
- Let $M_{\mathcal{A}}$ be the random variable denoting the message the adversary *commits* to.
- Let $M'_{\mathcal{A}}$ the random variable denoting the message the adversary decides to *decommit*. Since the adversary does not decommit to every message it commits to, the random variables $M_{\mathcal{A}}$ and $M'_{\mathcal{A}}$ do not necessarily have the same distribution.
- Let $q_m = \Pr[D_{\mathcal{A}} = m]$.
- Let $p_m = \Pr_{m' \in \mathcal{M}}[R'(m', m) = 1]$.
- Let $P_{\mathcal{A}}$ be the probability that the adversary succeeds.
- Let $P_{\mathcal{A}'}$ be the probability that the simulator succeeds.

In this notation we are trying to prove $P_{\mathcal{A}'} \geq P_{\mathcal{A}}$. First, we have

$$\begin{aligned}
 \Pr [D'_{\mathcal{A}} = m_2] &= \Pr_{m' \in \mathcal{M}} [D_{\mathcal{A}} = m | R'(m', m) = 1] \\
 &= \Pr[q_m | p_m] \\
 &= \frac{q_m p_m}{\sum_m q_m p_m}
 \end{aligned}$$

Now, we also have

$$\begin{aligned} P_{\mathcal{A}'} &= E_{m \in D_{\mathcal{A}'}}[p_m] \\ &= \frac{\sum_m q_m p_m^2}{\sum_m q_m p_m} \\ &= \frac{E_{m \in D_{\mathcal{A}'}}[p_m^2]}{E_{m \in D_{\mathcal{A}'}}[p_m]} \end{aligned}$$

Noticing that $P_{\mathcal{A}} = E_{m \in D_{\mathcal{A}}}[p_m]$, we have

$$P_{\mathcal{A}'} = \frac{E[p_m^2]}{E[p_m]}$$

Now applying the moment inequality $E[X^2] \geq E[X]^2$, we have

$$P_{\mathcal{A}'} = \frac{E[p_m^2]}{E[p_m]} \geq E[p_m] = P_{\mathcal{A}}$$

which is what we were trying to prove.

■

3 A General Non-Malleable Commitment Scheme

In the previous section we showed how to build an efficient non-malleable commitment scheme out of a specific number theoretic problem. We sketch a similar construction of a non-malleable commitment scheme from any one-way function. This gives us a general existence result. This construction was originally described in [?] and predates the explicit construction given in the last section.

We note that the existence of a pseudo-random function is equivalent to the existence of a one-way function [?], and the existence of a one-way function is equivalent to the existence of a digital signature scheme [?]. We will use both in the subsequent constructions.

3.1 Equivocable Commitment in the General Setting

Recall that in Lecture 5, we showed how to create bit-commitment scheme using any pseudo-random generator as originally proposed in [?]. Here, we quickly review that protocol

modifying it slightly to put it into the Common Reference String framework, then show that it is equivocal.

If G is a pseudo-random generator taking n -bits to $3n$ -bits, and we take the string R as the first $3n$ bits of the \mathcal{CRS} , then

- Commit Stage:
 1. Alice selects a seed $S = \{0, 1\}^n$ and computes $G(S) = Y$, where $Y = \{0, 1\}^{3n}$.
 2. Alice sends to Bob the vector $Z = \{0, 1\}^{3n}$ where $Z = Y$ if $b = 0$ and $Z = Y \oplus R$ if $b = 1$.
- Reveal Stage:
 1. Alice sends S to Bob.
 2. Bob computes $G(S)$. If $G(S) = Z$, $b = 0$; if $G(S) \oplus R = Z$, $b = 1$; otherwise, repeat the protocol.

We will not repeat the proofs of privacy and binding here.

The new observation is that this scheme is equivocal, and we will show how to create an equivocator $E = (E_1, E_2)$. The equivocator is constructed as follows.

- E_1 picks seeds S_0 and S_1 for G and computes $G(S_0)$, $G(S_1)$ and sets the first $3n$ bits of \mathcal{CRS}' to be $G(S_0) \oplus G(S_1) = R'$. Where \oplus denotes bitwise exclusive-or (addition in \mathbb{F}_2^n). The important property of R' is that $G(S_0) = G(S_1) \oplus R'$.
- E_1 sends the commitment $\text{Com}' = G(S_0)$.
- To decommit Com' as $b \in \{0, 1\}$ E_2 sends $\text{Dec}' = S_b$.

3.2 Non-Malleable Commitment

We now use the equivocal commitment scheme outlined in the last section to build Non-Malleable Commitment Scheme. The construction will be similar to the Non-Malleable Commitment scheme we have already seen. Let G be a pseudo-random generator, and Com the bit-commitment scheme based on G . We can break our \mathcal{CRS} into blocks of $3n$ -bits, so that each block is suitable for encrypting one bit using Com .

$$\mathcal{CRS} = \underbrace{\begin{array}{|c|c|c|} \hline \alpha_1 & \dots & \alpha_n \\ \hline \end{array}}_{\alpha} \underbrace{\begin{array}{|c|c|c|} \hline \beta_1^0 & \dots & \beta_m^0 \\ \hline \beta_1^1 & \dots & \beta_m^1 \\ \hline \end{array}}_{\beta}$$

So each α_i and β_i^j is $3n$ -bits long. The commitment works as follows, to commit the bit b

1. The committer picks an n -bit seed S to G , and uses the α portion of the \mathcal{CRS} to commit S .
We use $\alpha_1, \dots, \alpha_n$ to commit each bit of S
Set $A = (\text{Com}(\alpha_1, S_1), \dots, \text{Com}(\alpha_n, S_n))$
So A is just a commitment to S .
2. Consider A as an m -bit string, A_0, \dots, A_m . We use the bits of A as selector for β and we commit the same bit b with m different keys.
Set $B = (\text{Com}(\beta_1^{A_1}, b), \dots, \text{Com}(\beta_m^{A_m}, b))$.
3. Finally we compute $\text{MAC}_S(B)$.
This can be done by taking the $3n$ -bit string $G(S)$ breaking it in two as $r_1, r_2 = G(S)$ and computing
 $C = \text{MAC}_{r_1, r_2}(B)$ as before.
4. Send the commitment (A, B, C) .

The privacy and binding properties of this scheme are clear, and the proof of non-malleability goes through much as before.

4 Non-Malleable NIZK

4.1 Review

Recall that a Zero Knowledge proof, by which, for any NP -complete language L a prover P can convince a verifier V that $x \in L$ which satisfies two properties

- Completeness:
The prover can prove any true statement.
- Soundness:
The prover cannot prove any false statement with non-negligible probability.

In [?] it was shown that a NIZK protocol can be constructed out of any one-way trapdoor permutation.

4.2 Definitions

We now extend our discussion of non-malleability to the subject of Non-Interactive Zero Knowledge proofs.

We proceed as in [?]. Intuitively, what we would like from a NIZK proof is that verifier, having seen a proof p can prove no more than he could prove before having received the proof p . We can imagine a slightly stronger requirement, that after asking for proof of any statement of its choosing, no adversary can create a proof of a new statement that they could not have proven before. This is called *adaptive non-malleability*.

Definition 6 Let (P, V) be a NIZK scheme, then (P, V) is called *Adaptively Non-Malleable* if for any probabilistic polynomial time adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ there exists a probabilistic polynomial time algorithm \mathcal{A}' such that if we define the following two experiments E_1, E_2

E_1	E_2
$(x', w') \leftarrow \mathcal{A}_1(\mathcal{CRS})$ $p \leftarrow P(x, w, \mathcal{CRS})$ $(x', w', p') \leftarrow \mathcal{A}_2(x, w, p, \mathcal{CRS})$ return 1 if $p \neq p'$ and p is a valid proof for $x' \in L$	$(x', p', w') \leftarrow \mathcal{A}'(\mathcal{CRS})$ return 1 if p' is a valid proof that $x' \in L$

we have

$$|\Pr(E_1 = 1) - \Pr(E_2 = 1)| < \frac{1}{\text{poly}(k)}$$

An alternative, and sometimes desirable, property is that even after seeing simulated proofs of false statements, the adversary can still only prove true statements. This is called *Simulation Soundness* and is due to Sahai [?].

Definition 7 Let (P, V) be a NIZK scheme with simulator S_1, S_2 , then (P, V) is called *Simulation-Sound* if for any probabilistic polynomial time adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, if we define the following experiment

$$\begin{array}{l}
E \\
\hline
\mathcal{CRS} \leftarrow S_1(1^k) \\
x \leftarrow \mathcal{A}_1(\mathcal{CRS}) \\
p \leftarrow S_2(x, \mathcal{CRS}) \\
(x', p') \leftarrow \mathcal{A}_2(x, p, \mathcal{CRS}) \\
\text{return true if} \\
p \neq p' \text{ and} \\
x' \notin L \text{ and} \\
V(x', p') = 1.
\end{array}$$

we have

$$\Pr(E = 1) < \frac{1}{\text{poly}(k)}$$

4.3 Construction

We now show how to construct a Non-Interactive Non-Malleable Zero Knowledge proof for a *single* statement, using as building blocks a NIZK protocol (P, V) and a signature scheme $(\text{Gen}, \text{Sign}, \text{Ver})$. Set the Common Reference String

$$\mathcal{CRS} = \begin{array}{|c|c|c|} \hline \beta_1^0 & \dots & \beta_m^0 \\ \hline \beta_1^1 & \dots & \beta_m^1 \\ \hline \end{array}$$

Where each block β_i^b is a Common Reference String for the NIZK protocol (P, V) . Then for some NP language L , to prove a statement $x \in L$

1. Pick SK, VK keys for the signature scheme.
2. Using VK as a selector, for each bit VK_i of VK , let p_i be the proof that $x \in L$ using $\beta_i^{VK_i}$ as reference string.
3. Output

$$(\underbrace{VK}_A, \underbrace{p_1, \dots, p_m}_B, \underbrace{\text{Sign}_{SK}(B)}_C)$$

The completeness, and soundness properties of this scheme follow from the completeness and soundness properties of (P, V) . We prove the Non-Malleability of this scheme. The proof proceeds in the same manner as the proof of Claim 5.

Claim 8 *This NIZK protocol is non-malleable.*

Proof Suppose the adversary \mathcal{A} has received a proof (A, B, C) that $x \in L$. Suppose the adversary comes up with a proof (A', B', C') that $x' \in L$. We examine two cases

- case 1: If $A = A'$, then if $B = B'$, we must have $C = C'$. Since copying the same message is forbidden, we must have $B \neq B'$. Since $VK = A = A' = VK'$, the adversary \mathcal{A} must have forged a signature $C' = \text{Sign}_{VK'}(B')$ which contradicts the security hypothesis of the signature scheme.
- case 2: If $A \neq A'$, then VK and VK' must be different in at least one bit i . Thus the adversary must have come up with a proof that $x' \in L$ under (P, V) , with Common Reference String $\beta_i^{VK'}$. Thus either \mathcal{A} could already prove $x' \in L$, in which case this could be simulated, or \mathcal{A} can break the security of (P, V) , a contradiction.

■

Notice that the Non-Malleability property is lost if we prove a second statement, since then the adversary will have seen both β_i^0 and β_i^1 for some i .

4.4 Repeated Proofs

The protocol described in the last section allows Non-Malleable proofs of only one statement. We now study a protocol that allows repeated proofs, this scheme was originally presented in [?], and improved in [?].

Let (P, V) denote a (possibly malleable) NIZK protocol. Let $(\text{Sign}, \text{Ver})$ be a public key signature scheme, and view the Common Reference String as

$$\mathcal{CRS} = \boxed{VK_1 \mid \mathcal{CRS} \text{ For } (P, V)}$$

Where VK_1 is a verification key for $(\text{Sign}, \text{Ver})$. Note that if \mathcal{CRS} is truly random, then polynomial-time algorithm can compute SK_1 corresponding to VK_1 .

The prover behaves as follows, to prove $x \in L$ the prover

1. Pick (SK_2, VK_2) a key pair for $(\text{Sign}, \text{Ver})$.
Set $A = VK_2$.
2. Let B be a proof of the statement
“ $x \in L$ or I can compute $\text{Sign}_{SK_1}(VK_2)$.”
using (P, V) , and VK_2 as the selector.

3. Let $C = \text{Sign}_{SK_2}(B)$.
4. Send (A, B, C) .

Note that anyone who knows SK_1 can prove anything, thus we can always create a simulator which can simulate any adversary.

5 Open Questions

5.1 Ring Homomorphic Encryption

We have seen malleable encryption schemes, for example in El-Gamal where p is a prime, g is a generator for the group $(\mathbb{Z}/p\mathbb{Z})^*$, and $h = g^a$ then we encrypt a message m as

$$E(m) = (g^r, m \cdot h^r)$$

For some random $r \in (\mathbb{Z}/p\mathbb{Z})^*$ This gives

$$E(m_1) \cdot E(m_2) = E(m_1 \cdot m_2)$$

Thus the encryption scheme preserves the group operation in $(\mathbb{Z}/p\mathbb{Z})^*$. On the other hand

$$E(m_1) + E(m_2) \neq E(m_1 + m_2)$$

So this encryption scheme does not preserve both ring operations in $\mathbb{Z}/p\mathbb{Z}$.

The open question is whether we can construct a semantically secure encryption scheme over some ring that preserves *both* ring operations. Such highly-malleable encryption schemes are sought after as they have many applications in Private Information Retrieval and Secure Multiparty Computations.

5.2 “Homomorphic” NIZK

A similar question can be posed for NIZK proofs, but instead of whether they respect underlying ring operations we would like them to respect logical operations. The question becomes, can we construct a NIZK protocol, such that given $NIZK(A)$, and $NIZK(A \Rightarrow B)$, we can compute $NIZK(B)$.

Lecture 12

1 Private Information Retrieval

Here's the game: There¹ is a database, DB , which contains an n -bit vector, X . The user, U , has an index $1 \leq i \leq n$ and U wants to learn X_i without revealing i to X . (see Figure 1) For example, if X is the database of all of the stocks on NASDAQ and U wants to know how the stock for IBM is doing without anyone knowing that he is interested in IBM, where we assume that the user knows the address of each stock.




Figure 1

A naive solution could be 1-way communication: DB tells U the contents of the entire database, X , without U asking for anything in particular.


The problem with this is that it has communication complexity of n (communication complexity is defined to be the total number of bits that get sent back and forth). We want a solution with the total communication complexity *strictly* $< n$.

The history of this problem: [CGKS] proved two theorems: it is impossible information-theoretically for a single database to solve this problem with communication complexity $< n$, but if there are two or more identical copies of the database which are not allowed to communicate with each other during the protocol execution, then it is solvable with communication complexity $< n$. (see Figure 2)

For the solution using two databases: the best complexity that we know of is $O(n^{1/3})$. It is unknown if $O(n^{1/3})$ can be reduced, but some weak lower bounds are known.

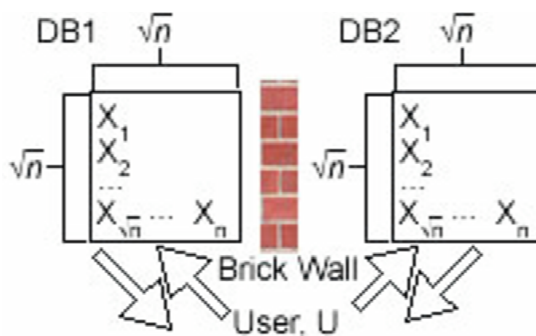
¹Scribe notes by Susan Chebotariov; May 1, 2006

Figure 2



First, we will explore an easier solution that is $O(n^{1/2})$. In this solution, we have $DB1$ and $DB2$ which are identical copies of our n -bit vector, X . We can express the databases in $\sqrt{n} \times \sqrt{n}$ matrices as follows. (see Figure 3)

Figure 3



U wants to know some bit i . Since both $DB1$ and $DB2$ are organized as $\sqrt{n} \times \sqrt{n}$ matrices, this bit, i , is located in some column j which is the same for both $DB1$ and $DB2$. So, U picks at random a string Z such that $|Z| = \sqrt{n}$. U sends this string to $DB1$. It is completely random, and therefore reveals nothing about the index. The database uses Z to generate an answer, A , to send back to U . A is calculated one bit at a time in the following manner: for all Z_i , bitwise sum up mod 2 all of the columns in $DB1$. That is, line Z up with the columns in $DB1$; if the first bit of Z is on (i.e., $Z_1 = 1$), then use the first column otherwise ignore the first column; if $Z_2 = 1$, then use the second column otherwise ignore the second column; and so on until finally, if $Z_{\sqrt{n}} = 1$, then use the last column otherwise ignore the last column. Then, consider the values in each of the non-ignored columns of $DB1$: A_1 = the values of the first row in the non-ignored columns bitwise-XORed together. A_2 = the values of the second row in the non-ignored columns bitwise-XORed together. $A_{\sqrt{n}}$ = the values of the last row in the non-ignored columns bitwise-XORed together.

Then, U flips the j th bit of Z , where j is the column that holds the i th bit of the database,

and sends that message, Z' , to $DB2$. Since Z was completely random, flipping a single bit in a completely random string is still a completely random string, Z' . Therefore, sending Z' to $DB2$ reveals absolutely nothing about i to $DB2$, as long as $DB1$ and $DB2$ do not communicate during the protocol execution. $DB2$'s answer, A_2 , is calculated in the same way that A_1 was calculated, and so on.

Then, U calculates bitwise $A_1 \oplus A_2$. Since Z is identical to Z' except for one bit, think about what happens to all columns that were summed up both in $DB1$ and $DB2$. Each column appears both in $DB1$ and $DB2$ except for the j th column which appears only in one of the two databases (chosen at random). Thus, all columns are added twice and cancel each other (recall that if you XOR any bit twice it is equal to zero) except for the j th column, which is added only once and hence is preserved. In this way, U can learn the value of the j th column that contains the i th bit of the string without revealing anything to $DB1$ and $DB2$ as long as they do not communicate. This solution has communication complexity = $4\sqrt{n}$: two messages of \sqrt{n} bits from U to $DB1$ and $DB2$ and two answers of size \sqrt{n} from each database.

Why is this important? This game relates to a problem in databases called “Locally Decodable Codes.” Given a message, m , one can compute the code-word $C(m)$ such that even if a constant fraction (say 1/10) of all the bits are changed by the adversary, there exists an efficient procedure to recover m back. The drawback of error-correction codes is that the entire code-word, $C(m)$, must be read to recover any bit of m . In databases, people may wish to have locally decodable codes so that, again, after the adversary corrupts a constant fraction of the entire code-word, $C(m)$, you only have to read a constant number of bits to recover bit m_i .

Claim: Private Information Retrieval (PIR) is equivalent to Locally Decodable Codes (LDC).

Proof Idea (*we will only show $PIR \Rightarrow LDC$*) Assume there is a solution to Private Information Retrieval (PIR) which only required U to send each database $O(\log n)$ bits and each database responded with a single bit, which allowed U to recover the i th bit of the database. If such a solution exists, then you can construct a locally decodable code. The user, U , will ask a question of length q where $q = O(\log n)$. There are 2^q possible questions: Q_1, Q_2, \dots, Q_{2^q} and U asks question Q_j . Each question has a single bit answer from each database. To form LDC from PIR, the answers from each database can be concatenated together to form a string of length = 2^{q+1} , which would be the code-word, $C(m)$. The answer to question Q_j , which U wants to learn would correspond in two locations in $C(m)$. (see Figure 4)

Even if the adversary has corrupted a constant fraction of the database (say 1/10), then there is a high probability that Q_j is uncorrupted in both $DB1$ and $DB2$. (see Figure 5)

Figure 4

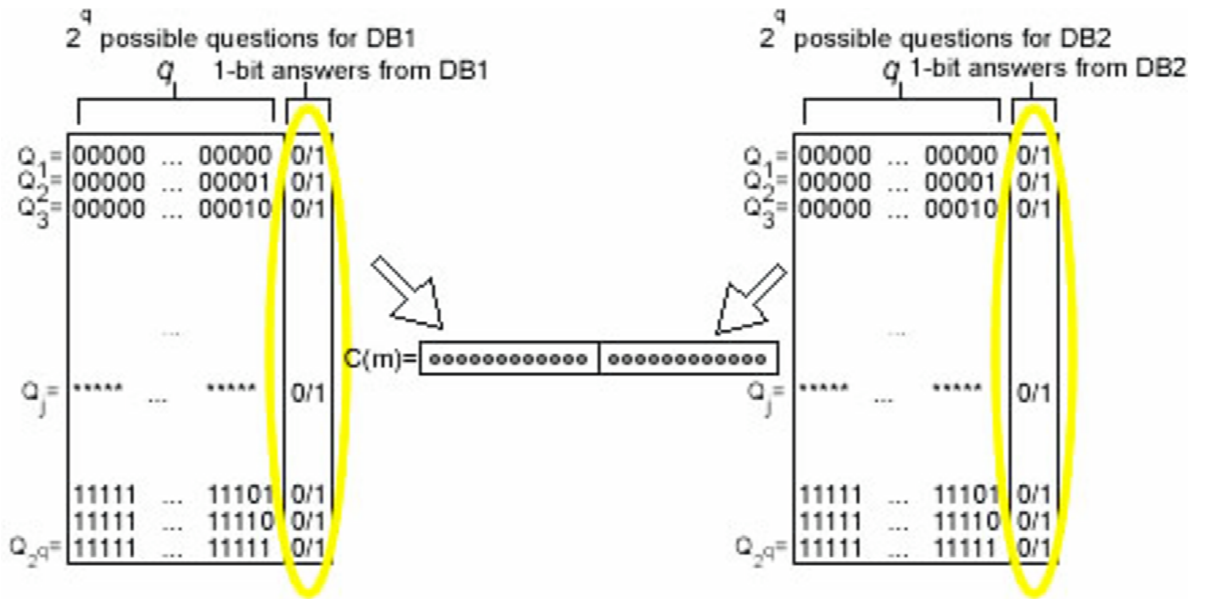
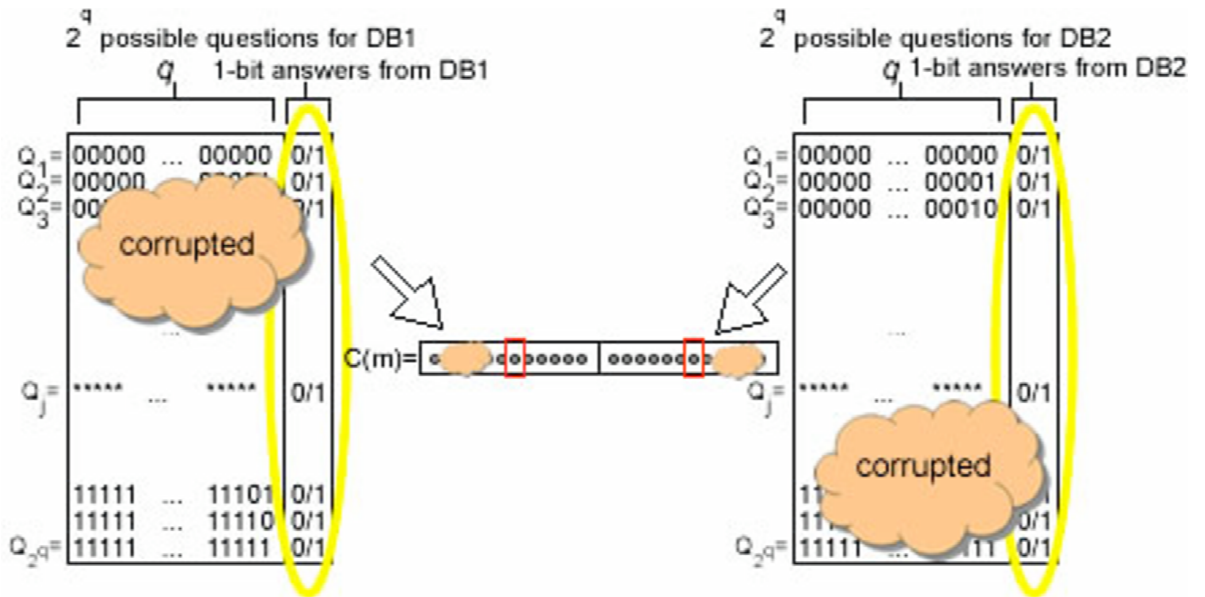


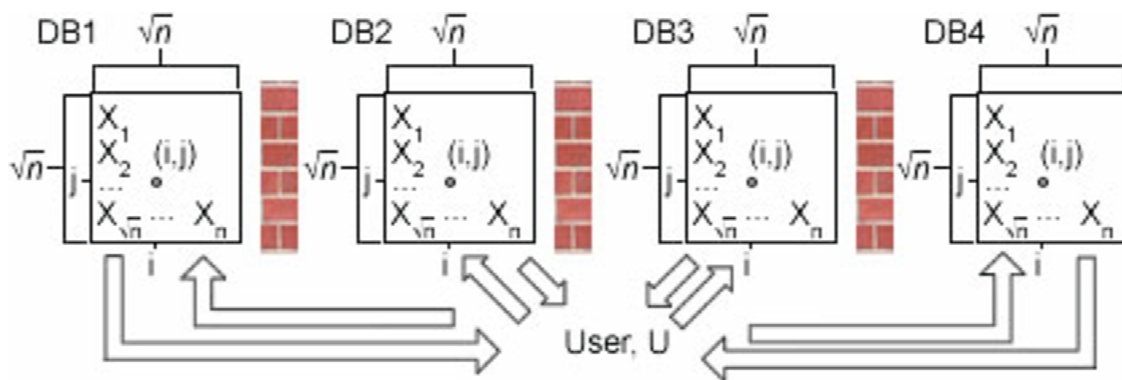
Figure 5



So, the bits of interest can be recovered. Therefore, PIR \Rightarrow LDC.

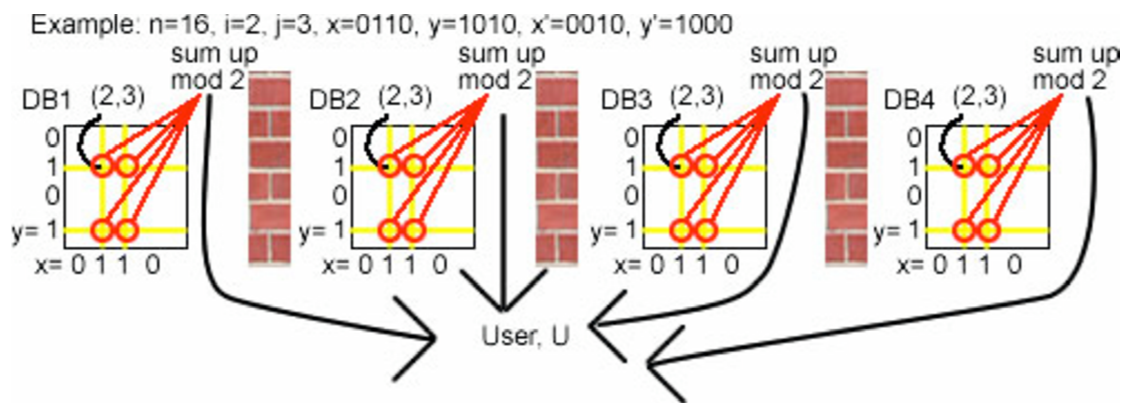
Now we want a solution that is $O(n^{1/3})$. For this solution, we start with four databases: $DB1, DB2, DB3,$ and $DB4$. Consider the bit of interest to be at the (i, j) position in each database. (see Figure 6)

Figure 6



U asks $DB1$ questions x and y where $|x| = |y| = \sqrt{n}$ and gets a 1-bit answer. Then, U asks $DB2$ x' and y , where $x' = \{x \text{ with the } i\text{th bit toggled}\}$ and gets a 1-bit answer. Next, U asks $DB3$ questions x and y' , where $y' = \{y \text{ with the } j\text{th bit toggled}\}$ and gets a 1-bit answer. Finally, U asks $DB4$ questions x' and y' and gets a 1-bit answer. Each database calculates its answer by taking the induced sub matrix of x and y , summing it up mod 2, and sending that single bit back to the user. (see Figure 7)


Figure 7



Therefore, from all four databases, U will end up getting a total of four bits back. If U calculates the bitwise-XOR for those bits, the answer is exactly (i, j) .

Correctness $DB1$ sums up mod 2 all of the bits that are to be included, $DB2$ sums all of the bits except those where $x = i$, $DB3$ sums all of the bits except where $y = j$, and $DB4$ sums all of the bits except those where $x = i$ or $y = j$. Bits that are not located in the i th column or the j th row are all used 4 times. Bits that are located in the i th column but not the j th row are used two times (in the answers from $DB1$ and $DB3$). Bits that are located in the j th row but not the i th column are used two times (in the answers from $DB1$ and $DB2$). The bit that is located in the i th column and the j th row (i.e., the bit of interest) is used three times (in the answers from $DB1, DB2$, and $DB3$). Therefore, all of the bits in the database are used an even number of times except for the bit at (i, j) which is only used three times, so, after a bitwise-XOR, (i, j) is the only thing left. In this manner, U is able to determine the value of (i, j) . Since the four databases cannot communicate with each other during the protocol execution, x, y, x' , and y' appear to be random strings to each of the databases, so none of the databases can determine the value of i or j . (see Figure 8)

Figure 8



Complexity This is still $O(n^{1/2})$, but note that each database sends only a one bit answer.

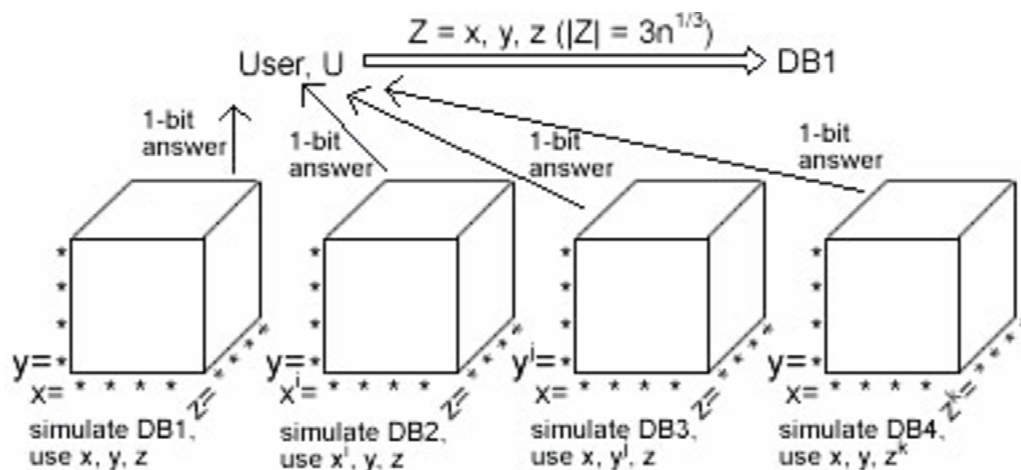
Now we can show a solution that is $O(n^{1/3})$ To get $O(n^{1/3})$, use eight databases which each store the data in 3-dimensional cubes, with the length of each side $= \sqrt[3]{n}$. U asks each database a variation of (x, y, z) and gets back one bit. In order, these are the questions U asks: (x, y, z) (x^i, y, z) (x, y^j, z) $(x, y, z^k) \mid (x, y^j, z^k)$ (x^i, y, z^k) (x^i, y^j, z) (x^i, y^j, z^k)

However, we do not want to deal with eight databases, so let one database, $DB1$, simulate the first four databases and a second database, $DB2$, simulate the last four databases.

U asks $DB1$ $Z = (x, y, z)$ where $|x| = |y| = |z| = n^{1/3}$, so $|Z| = 3n^{1/3}$. $DB1$ simulates the first question by leaving Z alone and calculating the 1-bit answer using the same method

as before. *DB1* simulates the second question by toggling all positions of x one at a time sending the 1-bit answer back to U and simulates the third question by toggling all positions of y one at a time sending the 1-bit answer back to U and simulates the fourth question by toggling all positions of z one at a time sending the 1-bit answer back to U . (see Figure 9)

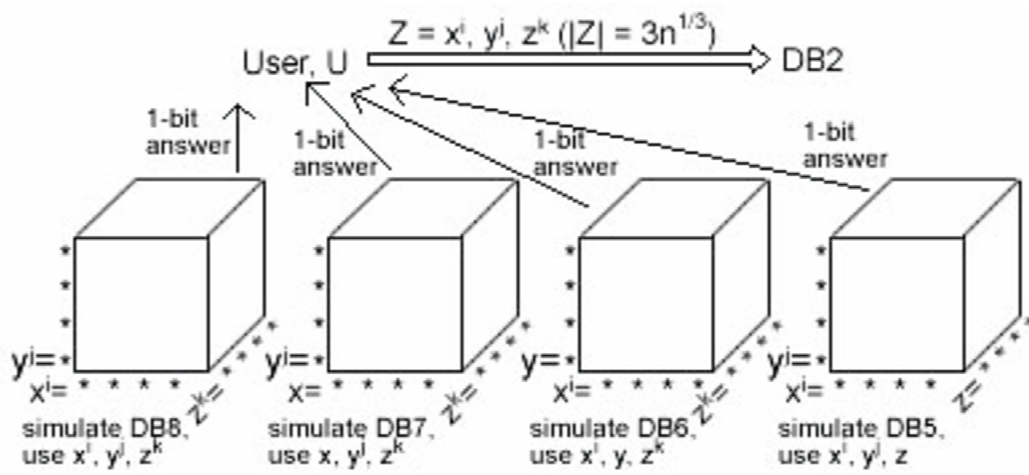
Figure 9



Similarly, U asks $DB2$ $Z' = (x^i, y^j, z^k)$. $DB2$ simulates the last question by leaving Z' alone. $DB2$ simulates the second to the last question by untoggling all positions of z one at a time and simulates the third to the last question by untoggling all positions of y one at a time and simulates the fourth to the last question by untoggling all positions of x one at a time. For each simulation, $DB2$ computes a single bit answer and sends that back to U . (see Figure 10)

In this way, the complexity for two databases is $O(n^{1/3})$ and this is the known bound.

Figure 10



Lecture 13

1 Oblivious Transfer

1.1 Rabin Oblivious Transfer

Rabin oblivious transfer¹ is a kind of formalization of “noisy wire” communication. The objective is to simulate a random loss of information. Formally, a Rabin OT machine models the following behavior. The sender S sends a bit b into the OT machine. The machine then flips a coin, and with probability $\frac{1}{2}$ sends b to the receiver R , and with probability $\frac{1}{2}$ sends ‘#’ to R to signify that a bit was sent, but the information was lost in the transfer. S does not know which output R received.




Figure 1: Rabin oblivious transfer

Remark Note that this may be simulated by a sufficiently noisy wire, provided that the wire transmits faithfully a good proportion of bits and at the same time loses a good proportion of bits, replacing them with noise that is distinguishable from information.

1.2 One-Out-of-Two Oblivious Transfer (1-2-OT)

Even, Goldreich and Lempel formulated a notion of oblivious transfer that has proven useful in various applications. In this situation, S sends an ordered pair of bits (b_0, b_1) into the 1-2-OT machine. R then gives the machine a bit i , indicating which input he would like to receive. The machine outputs b_i and discards b_{1-i} . R knows that R has one of the bits, but not which one.

¹Ruzan Shahinian, Tim Hu; 14,16 March 2006




Figure 2: One-out-of-two oblivious transfer

1.3 Implementing Oblivious Transfer Using 1-2-OT

The two games described above are information theoretically equivalent, as we will see in the following two sections.

Given a 1-2-OT machine as a black box, the protocol for implementing Rabin oblivious transfer is as follows. Here we only show the reduction for honest players. Square brackets indicate where an exchange takes place.

1-2-OT \Rightarrow Rabin OT

1. S has a bit b which he wants to transmit with probability $\frac{1}{2}$ to R .
2. S flips random bits r and l .
3. If $l = 0$, S inputs (b, r) into the 1-2-OT machine, and if $l = 1$, S inputs (r, b) .
4. [1-2-OT] R specifies a bit i to the 1-2-OT machine. Note that the 1-2-OT machine outputs b if and only if $i = l$.
5. [$S \rightarrow R$] S sends the value of l to R in the clear.

After this transfer, R will compare the value of i he picked and the value of l that was sent to him. If $i = l$, then he knows that the bit he received from the 1-2-OT machine was b . If $i \neq l$, then he knows that he was passed the random bit r ; in other words, he received no information about b . So R has exactly $1/2$ probability of receiving the intended bit.

1.4 Implementing 1-2-OT Using Oblivious Transfer

Given a Rabin OT primitive as a black box, the protocol for implementing a one-out-of-two oblivious transfer is as follows. Again, we only show the reduction for honest players.

Rabin OT \Rightarrow 1-2-OT

1. [OT] S inputs a large (say, length $3n$) string of random bits \vec{s} into the OT machine, which relays the bits to R , replacing approximately half of them with '#'.
2. [$R \rightarrow S$] R sends to S two sets of disjoint indices $I_0, I_1 \subset \text{dom}(\vec{s})$ chosen at random satisfying:
 - (a) I_0 and I_1 are of size n .

- (b) One of the sets corresponds to a random subset of places in \vec{s} where R received perfect information, i.e. no ‘#’s. The index of this set (either I_0 or I_1) acts as the i in the description of 1-2-OT transfer.
 - (c) The other set is chosen at random.
3. [$S \rightarrow R$] S chooses the two bits (b_0, b_1) that he would like to send by 1-2-OT, and sends to R $(b_0 \oplus_{i \in I_0} s_i, b_1 \oplus_{i \in I_1} s_i)$.

Note that in step 1, by the Chernoff bound, the probability that R received less than n or more than $2n$ many ‘#’s is exponentially small. Therefore, except for an exponentially small number of trials, in step 2 it is possible for R to find an index set satisfying (b), and the set chosen in (c) must contain at least one ‘#’. Thus he knows exactly one of $\bigoplus_{i \in I_0} s_i$ and $\bigoplus_{i \in I_1} s_i$, and he can calculate exactly one of (b_0, b_1) .

1.5 Implementation of 1-2 OT

Alternatively we may implement oblivious transfers from cryptographic assumptions. First we will need the notion of an enhanced function.

Definition 1 *A function $f : X \rightarrow Y$ is **enhanced** if there is a polynomial time sampling algorithm that samples Y with the same distribution as f , but for this sampling algorithm, it is hard to invert f .*

Note that for f a one-way permutation an enhancing algorithm is immediate, by picking y at random. Now, the protocol.

1. S fixes an enhanced trapdoor permutation f for which he knows the inverse, and a hard-core predicate P for f .
2. [$S \rightarrow R$] S sends f and P to R .
3. [$R \rightarrow S$] Depending on the selection bit i , R takes a random x_i and computes $y_i = f(x_i)$. It also randomly generates $y_{1-i} \in \text{ran } f$ (this is possible since f is an enhanced trapdoor permutation). Then R sends back y_0 and y_1 .
4. [$S \rightarrow R$] S then computes $x_0 = f^{-1}(y_0)$ and $x_1 = f^{-1}(y_1)$, and sends $b_0 \oplus P(x_0)$ and $b_1 \oplus P(x_1)$ to R .

Now, since R knows x_i , it can compute $P(x_i)$, and then compute b_i . Assuming that R was honest in step 3, and really chose y_{1-i} without knowing the inverse, it cannot compute the

inverse x_{1-i} and hence its hard-core bit. Thus R cannot compute b_{1-i} . Because S does not know the selection bit i , he has no idea whether R got b_0 or b_1 .

Now suppose R has $i = 0$, but cheats by not obtaining y_1 randomly, but instead chose an x_1 at random then computing $y_1 := f(x_1)$. Since f is a uniform distribution, S would have no way of detecting such behavior. R would then in the end know both b_0 and b_1 . The protocol works if R is “honest but curious,” but what if he is malicious? The solution to this problem will be explained later in this lecture. For now let us work in an honest-but-curious model.

2 Two Party Secure Computation

Consider the plight of two millionaires who want to find out who is richer, without letting each other know how much money they have. This is the problem of two party secure computation. In a two party secure computation, two parties A and B want to cooperatively run an algorithm where neither party has a complete set of parameters.

We are given A, B , and a polynomial size circuit $f(\vec{a}, \vec{b})$ consisting of AND and XOR gates which they would like to compute. The problem is that only A has access to the first half of the input (say, \vec{a}) and only B has access to the second half of the input (\vec{b}). How will they compute $f(\vec{a}, \vec{b})$ without sharing knowledge of \vec{a} and \vec{b} ? We will build a protocol such that at every intermediate stage of the computation of f , A and B will have a “share” of the output of that stage. The value of each wire is represented as two bits, one bit held by A and the other by B , such that their XOR is the value of the wire. Initially, inputs held by A will be split into two such bits for each input bit, where A gives to B a share, and B does the same with its inputs. Now they have to compute the circuit consisting of \oplus and \cdot gates, maintaining the secrecy.

Formally, A will have a record of bits T^A used in the computation and B will have a similar record T^B such that the actual bits used in the computation of $f(\vec{a}, \vec{b})$ are $\{x^A \oplus x^B : x^A \in T^A; x^B \in T^B\}$. Moreover A and B will never be required to reveal information about their shares T . This is done as follows.

Initialization

1. A generates a random string $a^{\vec{B}}$ and computes $a^{\vec{A}} := \vec{a} \oplus a^{\vec{B}}$.
2. [$A \rightarrow B$] A sends $a^{\vec{B}}$ to B .
3. B generates a random string $b^{\vec{A}}$ and computes $b^{\vec{B}} := \vec{b} \oplus b^{\vec{A}}$.
4. [$B \rightarrow A$] B sends $b^{\vec{A}}$ to A .

XOR gates: A and B have some $x = x^A \oplus x^B$ and $y = y^A \oplus y^B$, where A knows x^A and y^A , and B knows x^B and y^B , and they wish to compute $x \oplus y$. But since $x \oplus y = (x^A \oplus x^B) \oplus (y^A \oplus y^B) = (x^A \oplus y^A) \oplus (x^B \oplus y^B)$, each party can compute their own share of the sum without cooperation from the other party.

1. A computes $(x \oplus y)^A := x^A \oplus y^A$.
2. B computes $(x \oplus y)^B := x^B \oplus y^B$.

AND gates: A and B have some $x = x^A \oplus x^B$ and $y = y^A \oplus y^B$ and they wish to compute $x \cdot y$.

$$\begin{aligned} x \cdot y &= (x^A \oplus x^B) \cdot (y^A \oplus y^B) \\ &= (x^A \cdot y^A) \oplus (x^B \cdot y^A) \oplus (x^A \cdot y^B) \oplus (x^B \cdot y^B) \end{aligned}$$

Now A can compute $x^A \cdot y^A$ and B can compute $x^B \cdot y^B$, but without revealing their share of x and y , they must compute $x^B \cdot y^A$ and $x^A \cdot y^B$. This is done by oblivious transfer. Let M be a 1-2-OT machine. We first handle the case of $x^A \cdot y^B$.

1. A generates a random bit r^A .
2. A inputs the pair $((x^A \cdot 0) \oplus r^A, (x^A \cdot 1) \oplus r^A)$ to M .
3. B inputs y^B to M .
4. [1-2-OT] M outputs $(x^A \cdot y^B) \oplus r^A$ to B , who stores this as w^B .

Note that $x^A \cdot y^B = r^A \oplus w^B$. Also, B does not get any information from A about x^A , and A does not get any information from M about y^B . The case $x^B \cdot y^A$ is done similarly.

1. B generates a random bit r^B .
2. B inputs the pair $((x^B \cdot 0) \oplus r^B, (x^B \cdot 1) \oplus r^B)$ to M .
3. A inputs y^A to M .
4. [1-2-OT] M outputs $(x^B \cdot y^A) \oplus r^B$ to A , who stores this as w^A .

Finally, A and B can assemble their shares.

1. A computes $(x \cdot y)^A := (x^A \cdot y^A) \oplus r^A \oplus w^A$.
2. B computes $(x \cdot y)^B := (x^B \cdot y^B) \oplus r^B \oplus w^B$.

Lastly, when they compute the output of the “output” wire of the circuit, they can combine their shares and learn the output of the function.

3 Coin Flip Into the Well

We wish to have Alice assign a random bit to Bob over which he has no control; yet, Alice should have no knowledge of the bit. In real life, we might ask Bob to stand beside a deep well, deep enough to be inaccessible to Bob, but shallow enough that the bottom is still visible. Alice will stand from some distance away and toss a coin into the well for Bob, but she is not allowed near the well. Now, it is true that, unless the coin has some magical power of its own, Bob may simply lie about the outcome of the coin toss, as Alice would (and should) never know. Let us suspend this concern until later; first we model this game in practical terms. We will give the coin magical powers later.

Let c be a predetermined commitment scheme.

1. B flips a random bit r_B .
2. [$B \rightarrow A$] B sends $c(r_B)$ to A .
3. [$A \rightarrow B$] A sends a random bit r_A to B .
4. B computes the result of the coin flip $r := r_B \oplus r_A$.

Of course, B can still assign r arbitrarily, as A has no way of decommitting $c(r_B)$ by herself.

3.1 Malicious Players

Let us return to the problem of protocols which call for B to make secret, but honest, coin flips. Suppose we have reached such a point in a hypothetical exchange. It is B 's turn to talk, and the protocol requires him to send some message $f(T, \vec{r})$, where f is a deterministic function of T , the transcript of the conversation recorded so far, and a secret random \vec{r} . We will modify this protocol to prohibit B from fixing \vec{r} in his favor.

1. B generates a random string \vec{r}_1 .
2. B sends $c(\vec{r}_1)$ to A .
3. A sends a random bit \vec{r}_2 to B .
4. B computes the result of the coin flip $\vec{r} := \vec{r}_1 \oplus \vec{r}_2$.
5. B sends the message $\alpha := f(T, \vec{r})$.

Now, how can B prove that his message α was according to protocol? That is, he must convince A that $\alpha = f(T, \vec{r}_1 \oplus \vec{r}_2)$, for some \vec{r}_1 that is the decommitment of $c(\vec{r}_1)$. Now, f must be a polynomial time algorithm, since B has only polynomially many computing resources. So the statement

$$(\exists \text{ a decommitment scheme } d)[\alpha = f(T, d(c(\vec{r}_1)) \oplus \vec{r}_2)]$$

is an NP-statement. Thus, by NP-completeness it can be reduced to graph 3-colorability, that is, “ G is 3-colorable” for some graph G which can be computed in polynomial time and hence both A and B can agree upon. B ’s proof to A consists of a zero-knowledge proof that G is indeed 3-colorable.

Though this protocol is polynomial-time, it is inefficient. For each message sent by B , A and B exchange extra messages to ensure that B follows the honest protocol. This can be improved on various counts, as we will see next quarter.

3.2 Example: “Poker Over the Phone”

This technique, combined with two-party secure computation, can also be generalized to simulate an objective third party. Suppose a protocol calls for a third party M to output some message $f(T, \vec{r})$ to A but not to B , where f, T, \vec{r} are as before. There are two obstacles here; A and B must jointly generate a random \vec{r} which neither has knowledge of, and they must compute f without letting B know the result. Both problems are easily solved given previous constructions.

1. A generates a random string \vec{r}^B into the well for B .
2. B generates a random string \vec{r}^A into the well for A . \vec{r}^A and \vec{r}^B are A ’s and B ’s shares, respectively, to the input \vec{r} .
3. A and B compute $(f(T, \vec{r}))^A, (f(T, \vec{r}))^B$ by two-party secure computation.
4. B sends $(f(T, \vec{r}))^B$ to A .
5. A computes M ’s output $f(T, \vec{r}) = (f(T, \vec{r}))^A \oplus (f(T, \vec{r}))^B$.

4 SINGLE DATABASE Private Information Retrieval

Private Information Retrieval, or PIR, is the following game. Bob has a database of n bits, b_1, \dots, b_n . Alice has an index i , and wants to know the value of b_i , but does not want to reveal i (or any information about i) to Bob. This can be trivially accomplished by having Bob send the entire database to Alice. Here, we will discuss single-database PIR.

Kushilevitz and Ostrovsky showed that against a polynomial-time adversary, there is a PIR scheme with a communication complexity of n^ϵ for $\epsilon > 0$.

We will use a public-key encryption scheme where, if m_1 and m_2 are bits, $E(m_1) \cdot E(m_2) = E(m_1 \oplus m_2)$ (this is a *homomorphic encryption scheme*).

A simplified version of the PIR scheme, with communication complexity $O(\sqrt{n})$, is as follows. The database is logically partitioned into \sqrt{n} blocks each of length \sqrt{n} . Alice generates a public and private key, and sends the public key to Bob, along with \sqrt{n} pairs of encrypted values, $(x_1^0, x_1^1), \dots, (x_n^0, x_n^1)$, which Bob will use as representations of 0 and 1 for each block. Each x_j^0 is an encryption of 0. Each x_j^1 for $j \neq i$ is also an encryption of 0; x_i^1 is an encryption of 1. Using the respective x_j^0 and x_j^1 for each block, Bob computes the encrypted sum of the first bit in each block, the second bit in each block, and so on. Since for all blocks except the one containing b_i , the contribution to the encrypted result is 0, the result becomes an encryption of the block containing b_i . Bob sends the result back to Alice, who decrypts the value corresponding to b_i .

Now consider what happens if the database is partitioned into $n^{\frac{1}{3}}$ blocks each of size $n^{\frac{2}{3}}$. Alice and Bob go through the same steps as above, producing a result of length $kn^{\frac{2}{3}}$, where k is the length of the encryption of one bit. Alice only really wants k of those bits (the ones representing b_i), though, so Alice and Bob can use PIR again with a “database” that is the result of the initial PIR. This will produce a result of length $O(n^{\frac{1}{3}})$, yielding an overall communication complexity of $O(n^{\frac{1}{3}} + \epsilon)$. Observe that this can still be a single-round protocol, since Alice can just send all of the necessary encryption pairs for both iterations in her first message. Kushilevitz and Ostrovsky showed that by applying PIR recursively, communication complexity can be reduced to $O(n^\epsilon)$ for $\epsilon > 0$.

4.1 1-2-OT and PIR Imply 1- n -OT

One-out-of- n oblivious transfer is similar to 1-2-OT save that the sender has n bits instead of just two, and the receiver has a $\lg n$ -bit selector i instead of a single bit b . 1- n -OT also bears some similarity to PIR, but it is a different beast; the difficulty with PIR is to provide privacy for the receiver and keep communication complexity less than n , while for 1- n -OT, the difficulty is to provide privacy for both the sender and the receiver.

Naor and Pinkus showed how 1- n -OT can be built using 1-2-OT (more precisely, 1-2-OT for strings) and PIR. The sender chooses $2 \lg n$ keys for a private-key encryption system, labeling them $K_1^0, K_1^1, K_2^0, K_2^1, \dots, K_{\lg n}^0, K_{\lg n}^1$. He then encrypts each of his n bits $\lg n$ times with keys selected by the binary representation of the address of each bit; for instance, b_0 is encrypted with $K_1^0, K_2^0, \dots, K_{\lg n}^0$, and b_{n-1} is encrypted with $K_1^1, K_2^1, \dots, K_{\lg n}^1$. The receiver then uses 1-2-OT $\lg n$ times to get one of K_j^0 or K_j^1 for each j , choosing based on the binary representation of his selector i . Finally, the receiver uses PIR to get the

encryption of b_i and uses the keys to decrypt it.

4.2 PIR Implies Collision-Resistant Hashing

Recall that a collision-resistant hash function is a function h from n bits to m bits, where $m < n$, such that no polynomial-time adversary can find x and y such that $x \neq y$ and $h(x) = h(y)$. Any nontrivial single-round PIR scheme (that is, one with sublinear communication complexity) can be used as a collision-resistant hash function. To build such a gadget, select some i and fix a PIR query for i . Let $h(x)$ be the PIR response for the fixed query on a database x .

Claim 2 *The above construction is a collision-resistant hash function.*

Proof Suppose some adversary could find x and y such that $x \neq y$ and $h(x) = h(y)$. Since the output from h is the same for both inputs, bit i must be the same in x and y . Because $x \neq y$, they differ in at least one bit, which can't be b_i ; therefore, the adversary has learned something about i . This is a contradiction because PIR reveals no information about i . ■


Lecture 14

1 Two Party Secure Computation

1.1 Definitions

Suppose¹ two players each have secret inputs and wish to compute the value of a function of those inputs without sharing the secrets. For example, two millionaires might wish to know which one is richer without revealing their net worths.

In general, we have two players, A and B , with secret inputs, x and y , respectively. They will exchange messages to try to compute $f(x, y)$ for some f in such a way as to minimize what B can discover about x and what A can discover about y .



Definition 1 *The view of a player A after a particular run of the communication protocol is the triple (T, x, R) , where T is the transcript of the conversation, x is A 's input, and R is A 's randomness. The view of B is defined in a similar way.*

Definition 2 *For $f(x, y)$ any poly-time computable function and honest-but-curious players A and B , a two party communication protocol is secure for player B if the following holds: $\forall x, y, y'$ if $f(x, y) = f(x, y')$ then for all randomness r_A and r_B and every transcript, T ,*

$$Pr_{\{r_A, r_B\}}(\mathcal{T}(x, y, r_A, r_B) = T) = Pr_{\{r_A, r_B\}}(\mathcal{T}(x, y', r_A, r_B) = T)$$

Where $\mathcal{T}(x, y, r_A, r_B)$ denotes the transcript generated on inputs x and y and randomness r_A and r_B . Security for A is defined in a similar manner.

¹Scribe notes by Ben Tradewell, May 22nd 2006

Intuitively, this means that all inputs that give a certain output are equally likely to produce any given transcript. Thus, once the players know the output, the transcript gives no extra information about what the inputs are.


Definition 3 For $f(x, y)$ any poly-time computable function and a potentially dishonest player A , a two party communication protocol is secure for player B if the following holds: $\forall y, \forall A' \in PPT, \exists$ a simulator $S_{A'} \in PPT$ with access to A' such that $\forall z \in \text{range}(f)$, then $[S_{A'}(z)] = [T(A', B, x, z)]$. Here, T denotes the transcript of the interaction of A' and B that produces output z with y as B 's input and $[R]$ denotes the distribution of the random variable R based on the randomness of the machines involved. Security for A is defined in a similar way.

This definition is based on that of zero-knowledge. For any behavior of a dishonest A , there is a simulator that can produce a distribution of transcripts that is the same as that of the real transcripts produced. The simulator has no knowledge of B 's input, so no knowledge of y can be revealed by running the simulator. Since the fake transcripts that reveal no knowledge cannot be distinguished from the real transcripts, the real transcripts must reveal no information on y . Thus, the protocol is secure from B 's point of view.


1.2 Set-up and Examples

In this section, we consider specifically the case where $f(x, y)$ is a polynomial size circuit of AND (\cdot) and XOR (\oplus) gates. Circuits of AND and XOR gates can reproduce the behavior of any logical circuit (one of AND, OR, and NOT gates) so this is not an overly strong assumption. Also, we consider A and B to be honest-but-curious players. Assuming [1-2-OT] we can implement two party secure communication.

An example circuit with A having bit inputs a_1, a_2, a_3 and B having bit inputs b_1, b_2, b_3 :




Example 1 XOR. Suppose x and y are bits and $f(x, y) = x \oplus y$. Then once B knows $f(x, y)$, he can find x by computing $x = f(x, y) \oplus y$. Thus x is determined by the result, so there is no harm in sending it in the transcripts. Thus the protocol can just be that A sends x to B and B sends y to A .



Example 2 AND. Again x and y are bits, but this time $f(x, y) = x \cdot y$. If $f(x, y) = 0$ and $y = 0$ then B is given no information about x , so we cannot send x as part of the transcript. To get around this, B uses [1-2-OT] as follows: He prepares a [1-2-OT] box that will transmit $0 \cdot y$ if $i = 0$ and $1 \cdot y$ if $i = 1$. Then A simply sends $i = x$ to the [1-2-OT] box. The resulting value is $f(x, y)$. B does not know which value A has chosen, so he gets no information about x and A does not get to see the other value, so she only finds out $f(x, y)$,

which she must know at the end of the protocol anyway.



1.3 Secret Sharing

Consider the following game on parameters (t, n) between a "dealer" and n players, P_1, \dots, P_n : The dealer knows a secret s and wishes to provide "shares," s_1, \dots, s_n , to the players such that if any t players compare their shares it is not possible to discover s , but if any $t + 1$ players compare their shares they can determine s .


For $s \in GF(2^n)$, we can implement this using polynomials in $GF(2^n)$. The dealer first produces distinct, non-zero $\alpha_1, \dots, \alpha_n \in GF(2^n)$ and gives this list to all the players. Next, he picks a random degree t polynomial, f , that goes through the point $(0, s)$. Then he sends secret $s_i = f(\alpha_i)$ to the player i . If any $t + 1$ players compare their shares they will be able to reconstruct the degree t polynomial f and can then compute $s = f(0)$. If any t players compare their shares, they will not be able to reconstruct f and so will not be able to determine s . In fact, as there will be one degree t polynomial corresponding to their t shares and any particular value of s , the players do not even gain any information about what s .

1.4 A Solution for AND/XOR Circuits Using 1-2-OT for Honest-but-Curious Players

This section is due to [Yao] and [GMW], adapted from [KKMO].


Two honest-but-curious players, A and B , each have a portion of the inputs to a circuit and wish to determine the output of the circuit without revealing their inputs. They can do this using secret sharing. Each player will have a share of the inputs to a given gate and will use them to compute shares of the value of the gate. In this manner they can move up the tree until they eventually have computed shares of the output. They then compare their

shares to find the final value. In this particular scheme, the secret at any wire is always determined by the adding the two shares (mod 2). If z and x are inputs with shares z_0, z_1 and x_0, x_1 to a Δ gate the output is $c = z\Delta x$ with shares c_0, c_1 :




First, the players need to generate shares of their inputs. Each bit of the input corresponds to one wire in the circuit. If A knows the input to a particular wire, she picks a bit a_0 at random and computes $a_1 = a \oplus a_0$ and sends a_1 to B and keeps a_0 for herself. Then the share that B has is random and the shares add to the value of the wire. B does a similar thing and sends b_0 to A for his inputs. For the rest of this section, "wire α has value $z = z_0 \oplus z_1$ " means that z is the value of wire α and that A has share z_0 and B has share z_1 .

For XOR gates, the players simply add their shares. If wire 1 has value $z = z_0 \oplus z_1$ and wire 2 has value $x = x_0 \oplus x_1$ then the output is $z \oplus x = (z_0 \oplus x_0) \oplus (z_1 \oplus x_1)$. So no communication is necessary for XOR gates and hence this is private.



AND gates are more complicated. If wire 1 has value $z = z_0 \oplus z_1$ and wire 2 has value $x = x_0 \oplus x_1$ then the output is $xz = (z_0 \oplus z_1)(x_0 \oplus x_1) = z_0x_0 \oplus z_0x_1 \oplus z_1x_0 \oplus z_1x_1$. A can compute z_0x_0 and B can compute z_1x_1 , but neither can compute the other values without communicating. In fact, if B does compute z_0x_1 , this might tell him what z_0 is, violating secrecy. To get around this, the players randomize the values of the mixed terms. Player A picks a random bit r and prepares a [1-2-OT] box that will transmit $r \oplus z_0 \cdot 0$ if $x_1 = 0$ and $r \oplus z_0 \cdot 1$ if $x_1 = 1$. B then chooses the appropriate value and thus knows $r \oplus z_0x_1$. As r is random and B can only see one of the prepared values, the value he chooses will look random to him. As A cannot see which value B chooses, she cannot determine his share. B then does a similar procedure, choosing random bit s and using [1-2-OT] to send $s \oplus z_1x_0$

to A . Then $xz = (z_0x_0 \oplus r \oplus s \oplus z_1x_0) \oplus (z_1x_1 \oplus s \oplus r \oplus z_0x_1)$. A knows z_0x_0 and r and gets $s \oplus z_1x_0$ from B . B knows z_1x_1 and s and gets $r \oplus z_0x_1$ from A . Thus each player knows the appropriate share. The randomization ensures that the players do not learn the value of the other share.



In this manner, they compute shares to the entire circuit, gate by gate. Once they have computed the last gate, they simply reveal their shares and add them to determine the value of the circuit.

1.5 Aborting

While the players in this game are assumed to be honest-but-curious and therefore can't actually lie or abort, this is not the case for dishonest players. For dishonest players, we can try to compel honest-but-curious behavior using "coin-flips into the well" and ZK proofs. However, the issue of aborting must also be considered. This is particularly a problem in the last step where the dishonest players exchange shares. If A sends her share to B , then B will have both shares and will thus know the output. If he does not reveal his share to A , then A cannot determine the output. There are a number of strategies to deal with this problem.

One method is for A and B to exchange encryptions of their shares and then begin revealing the decryption keys one bit at a time. That way, if B quits at the last bit, then A only has one unknown bit to check and so can quickly guess the decryption key. However, B might have much greater computing power than A , so B might quit with, say, 100 bits to go and still be confident that he can break A 's encryption, yet A might have no hope of breaking a 101 bit key in any reasonable period of time. Also, it might be the case that the first m bits are more useful to B than they are to A , so if B quits after m bits, he will have a huge advantage. For instance, if it is known that A 's decryption key is of the form $a_1a_2 \dots a_m 0 \dots 0$ and B 's is of the form $0 \dots 0b_1b_2 \dots b_m$ then after m bits, B will know enough to decrypt A 's message, but A will know nothing.

Another method uses the idea of "splitting a bit" [Cleve]. Player A defines a probability distribution and random variables X_1, \dots, X_{poly} each of which has an expectation of $\frac{1}{2} + \frac{1}{poly}$ if her bit is 0 and $\frac{1}{2} - \frac{1}{poly}$ if her bit is 1. Similarly, B generates Y_1, \dots, Y_{poly} . Then they

give NIZK proofs that their random variables are biased appropriately. They then begin revealing their variables one at a time. Once the players get enough variables, they can determine the bias, so this will transmit the correct bits. Also, if B stops early, then he only has one more random variable than A , so his guess of what A 's bit is will not be appreciably better than A 's guess of what B 's bit is.

2 Multiparty Secure Computation Using Secret Channels

The method in the previous section can be expanded to any number of players, so long as every pair of players has a private [1-2-OT] channel for communication. In this section we discuss a protocol due to [BGW] for a similar game that is secure in an information theoretic sense (i.e. no cryptology will be used), assuming the players have secure channels. Also, now we assume that we are working with inputs in a finite field of size strictly greater than the number of players and that the function in question is a circuit of addition and multiplication gates.

Suppose n players, P_1, \dots, P_n have secret inputs x_1, \dots, x_n and they wish to compute $f(x_1, \dots, x_n)$ in secret, where f is a polynomial circuit of $+$ and \cdot gates. They are allowed no cryptography assumptions, but every pair of players has a secret channel over which they may exchange messages without being overheard. Also, in this game the players are assumed to have infinite computing power. A privacy level, t , is determined beforehand as follows:

Definition 4 *For $t < n$, we say the protocol has "t-privacy" if any group of t players that compare their views cannot determine the inputs of any other player better than the a priori distribution based on the output of f and the conspirators' inputs.*

Also, in this game, we allow for the possibility that a certain number of players might not be honest. These malicious players, sometimes called "Byzantine", do not have the restrictions on their behavior that honest players do. Furthermore, the malicious players are assumed to be chosen by some "behind the scenes" adversary, but this adversary might have the option of choosing the players to corrupt before the game begins ("static") or at runtime ("dynamic"). A protocol that can still run in the presence of t malicious players is called "t-resilient."

Example 3 *A protocol that is statically $\frac{n}{4}$ -resilient, but not dynamically $\frac{n}{4}$ -resilient. Choose some set of $k = O(\log n^2)$ players at random to be the "leaders". Each player then creates a $(\frac{k}{2}, k)$ share of his information and gives one share to each leader. The leaders now have shares of all of the information, so they can use some statically $\frac{n}{3}$ -resilient protocol to compute the result and share it with all the players. In the static case, the expected*

number of corrupted leaders is $\frac{k}{4}$, assuming the leaders are indeed chosen at random. Thus the chance that more than $\frac{k}{3}$ of the leaders are corrupted is negligible and the protocol will almost certainly run correctly. In the dynamic case, however, all leaders can be corrupted so the protocol has no hope of running correctly.

Definition 5 *Byzantine Agreement.* A group of soldiers are trying to decide if they should attack a heavily fortified position. If they all attack at once, the attack will be successful, but if only part of the army attacks, they will be defeated. Using secret communication the soldiers must agree to all attack or all wait. Furthermore, each soldier has their own opinion about whether to attack now or later and if all the loyal soldiers think the same, then they must agree to that course of action. To complicate matters, a group of soldiers are traitors and will lie to try to confuse the loyal soldiers.

Theorem 6 [PSL] *The loyal soldiers can succeed so long as no more than one third of the soldiers are traitorous.*

Theorem 7 [BGW] *If $t < \frac{n}{3}$ then for any f , a poly size circuit of addition and multiplication gates, there is a protocol that is t -private and dynamically t -resilient.*

Theorem 8 [RB] *$t \leq \frac{n}{2}$ can be achieved if broadcast channels are assumed, and if you permit an ϵ chance of the protocol failing to produce a result, for any $0 < \epsilon$.*

Proof We will show only the honest-but-curious case, with no broadcast channels. In this case, $t < \frac{n}{2}$ can be achieved.

This proof is similar to the two party proof, but uses the secret sharing scheme outlined in section 1.3. Each player constructs a random degree t polynomial with constant term equal to their secret input and distributes the shares to the other players. As was discussed above, no group of t players can get any information about the other players' inputs from this scheme. At the end of the computation, the players will just exchange shares so that all know the output.

At each gate, the players will have a share of the value of each input to the gate and they will compute shares of the output. Let the two inputs be represented by two degree t polynomials p and q , so player i has shares, $p(\alpha_i)$ and $q(\alpha_i)$ and the value of the wires are $p(0)$ and $q(0)$.

For addition gates, again the players just add their shares. Then the output value is $p(0) + q(0) = (p + q)(0)$. Note that $p + q$ is a degree t polynomial with the appropriate constant term and each player's share is $p(\alpha_i) + q(\alpha_i) = (p + q)(\alpha_i)$. No communication is necessary, so this computation is private.

For multiplication gates, we would like to do a similar operation. It is the case that pq does have the desired constant term and that each player can compute $(pq)(\alpha_i) = p(\alpha_i)q(\alpha_i)$, but pq could have degree up to $2t$. Also, pq is not random (we know it can't be irreducible, for example) so we will need to randomize it.

To randomize, each player produces a random degree $2t$ polynomial, g_i , with constant term zero and distributes the shares of g_i to the other players. The new polynomial to be reduced is $h = pq + \sum_{i=1}^n q_i$. This polynomial is random and $h(0) = pq(0) + \sum_{i=1}^n q_i(0) = pq(0) + \sum_{i=1}^n 0 = pq(0)$, so h is also a degree $2t$ polynomial that encodes the output of the gate.

To reduce the degree of h , we need the following lemma:

Lemma 9 *Let S be an n -vector of secret inputs and A a known $n \times n$ matrix, then the players can t -privately compute $S \cdot A$ in such a way that only player i knows the i th component of the result.*

Proof The i th player only needs to compute a linear combination of the values, where the coefficients come from the i th column of the matrix. We have proved that they can secretly compute addition, so we just need to show it for scalar multiplication. However, for any c , $c \cdot p(0) = (cp)(0)$ and cp has degree t so each player can just use $c \cdot p(\alpha_i)$ as their new secret share. To ensure that only the correct player gets the result of the linear combination, the other players reveal their shares to that player only. Do this for each player to get the result. ■

For $h(x) = h_0 + h_1x + \dots + h_{2t}x^{2t}$ we will compute the shares for the truncated polynomial $r(x) = h_0 + h_1x + \dots + h_tx^t$. What we want to do is find a known A so that if S is the n -vector of shares for h , then SA is the n -vector of shares for r .

Recall that the Vandermonde matrix on x_1, \dots, x_n is

$$\begin{pmatrix} 1 & 1 & \dots & 1 \\ x_1 & x_2 & \dots & x_n \\ \vdots & \vdots & \ddots & \vdots \\ x_1^{n-1} & x_2^{n-1} & \dots & x_n^{n-1} \end{pmatrix}$$

and has determinant $\prod_{i < j} (x_i - x_j)$.

Let $H = (h_0, \dots, h_{2t}, 0, \dots, 0)$ and $K = (h_0, \dots, h_t, 0, \dots, 0)$ be n -vectors. Let B be the Vandermonde matrix on $\alpha_1, \alpha_2, \dots, \alpha_n$. Let $P(x_1, \dots, x_n) = (x_0, \dots, x_t, 0, \dots, 0)$ be the linear projection onto the first t coordinates. It is easy to see by multiplying out that $H \cdot B = S$, $H \cdot P = K$, and $K \cdot B = R$. $\det B \neq 0$, as the α_i are distinct and non-zero, so we may let $A = B^{-1}PB$. Then $SA = SB^{-1}PB = HPB = KB = R$, as desired.

Now we apply the lemma to obtain that each player can compute the appropriate share of r . Note that r is a random degree t polynomial (as h was random) and $r(0) = h_0 = h(0) = p(0)q(0)$. Thus r is an appropriate polynomial for the output of the gate.

As in the two player version, the players compute each gate in sequence and then compare their shares to find the final output. ■

3 References

- [BGW] M. Ben-Or, S. Goldwasser, A. Wigderson. Completeness Theorems for Non-Cryptographic Fault-Tolerant Distributed Computation. STOC88.
- [Cleve] R. Cleve. Controlled Gradual Disclosure Schemes for Random Bits and their Applications. Proc. Advances in Cryptology, July 1989.
- [GMW] O. Goldreich, S. Micali, and A. Wigderson. How to Play any Mental Game. STOC87.
- [KKMO] J. Kilian, E. Kushilevitz, S. Micali, R. Ostrovsky. Reducibility and Completeness in Private Computations. SIAM Journal on Computing, 29(4): 1189-1208, 2000.
- [PSL] M. Pease, R. Shostak, and L. Lamport. Reaching Agreement in the Presence of Faults. J. ACM, 27(2), 1980.
- [RB] T. Rabin, M. Ben-Or. Verifiable Secret Sharing and Multiparty Protocols with Honest Majority. STOC89.
- [Yao] A. C. Yao. Protocols for Secure Computations. Proc. of 23rd FOCS, 1982.

Lecture 15

In this lecture¹, we will study three different topics: Yao's garbled circuits, efficient ZK-proofs using PCP reductions and Barak's non-black-box techniques for constructing ZK-proofs.

1 Yao's Garbled Circuits

In previous lectures we have seen how two parties with private inputs x and y respectively, can compute any function $f(x, y)$ such that no party learns any information about the input of the other party. It also ensures that parties do not learn any value on intermediate wires during the circuit evaluation. At the end of the protocol, the parties learn the values corresponding to the output wires of the circuit. Evaluating f in such a way that parties learn only what is implied by the values corresponding to the output wires and their own input, is known as secure function evaluation (SFE).

In previous lectures, we have seen information theoretically secure protocols for doing this task for all players when $t < \frac{n}{3}$ (or $t < \frac{n}{2}$ if the broadcast is given for free). The presented protocols had the round complexity linear in the depth of the circuit. No protocols are known to evaluate f in the multi-party (more than two players) setting in above mentioned manner with constant round complexity and information theoretic security. We remark that information theoretically secure SFE for only two players where inputs of both the players remain information theoretically hidden, is impossible.

Yao's garbled circuit technique shows how to evaluate any polynomial time computable function securely in constant number of rounds based on computational assumptions. Let us assume that the parties are A and B with private inputs x and y respectively. Every polynomial-time computable function f can be represented as a polynomial-size circuit. So let f have a well known public circuit C . One of the parties, say B , *garbles* this circuit and her own inputs (garbling the input means that it selects two independent random strings for each input value one string representing 0 and another representing 1). The garbled circuit is nothing but a set of tables, one for each gate, explaining how to do the computation for that gate (explained later).

Once the tables are prepared for each gate, B sends to A all these tables and the garbled values corresponding to her own inputs. A then executes oblivious transfer with B to obtain the garbled values for its inputs (i.e., one of the two strings). Now A has both the garbled circuit and the garbled inputs. The function f can now be computed gate by gate. For each

¹Scribe notes by Omkant Pandey, 24th May 2006

2 Efficient Zero Knowledge Proofs [Kilian 91]

Consider an interactive proof system consisting of a prover P and a verifier V . Assume that both the prover and the verifier are polynomially bounded. Let $x \in L$ be the theorem input to both the prover and the verifier for some language L in \mathbf{NP} . We wish to construct zero knowledge proofs for all languages in \mathbf{NP} such that the total communication complexity of the protocol will be $O(\text{poly}(k) + \text{poly}(\log(|x|)))$, where k is some security parameter.

These efficient proofs are based on an important result called the PCP theorem. The PCP theorem is an important developments in theoretical computer science. We shall review the basics related to the PCP theorem and use them for our end goal: construction of efficient zero knowledge proofs for all languages in \mathbf{NP} . We will also review the construction of Merkle trees and Merkle-tree based commitments.

2.1 Probabilistically Checkable Proofs (PCP)

Consider an instance X of 3-SAT. We say that $X \in 3\text{-SAT}$ iff there exists an assignment to the variables of X such that all of its clauses are satisfied. The PCP theorem can be stated as follows:

Theorem 1 (PCP Theorem) *Any 3-SAT instance X can be converted into another 3-SAT instance X' such that:*

- If $X \in 3\text{-SAT}$ then $X' \in 3\text{-SAT}$.
- If $X \notin 3\text{-SAT}$, any assignment to the variables of X' violates at least α fraction of clauses of X' , for some constant α (say $\alpha = \frac{3}{4}$).
- $|X'| = \text{poly}(|X|)$

It follows from the PCP theorem that by reading only a constant number of bits in the proof, the verifier can determine the membership of X in 3-SAT with some constant probability of being correct². If we repeat the process $O(\log |X'|) = O(\log |X|)$ times, the error probability will reduce to $\frac{1}{\text{poly}(|X'|)} = \frac{1}{\text{poly}(|X|)}$.

We remark that the proof of PCP theorem provides a polynomial time algorithm to convert X into X' . If $X \in 3\text{-SAT}$, then let the string w' denote a satisfying assignment to the variables of X' . Let \circ denote concatenation and let $\text{PCP}(X) = w' \circ X'$. Later on, we will denote $w' \circ X'$ as y .

²Simply pick a random clause in X' and check that it is true. Since $\alpha = \frac{3}{4}$, all clauses in X are not true and a random clause will be true only with probability $\frac{1}{4}$

Figure 1: Merkle tree on string y .

2.2 Merkle Trees

In this section we briefly review the construction of Merkle trees and commitment schemes using them.

Constructing Merkle Trees

Let $h : \{0, 1\}^{2k} \rightarrow \{0, 1\}^k$ be a collision resistant hash function. Let y be a string of length n . Merkle tree of a string y , using h , is computed as follows. String y is divided into blocks of length k . Using these blocks as leaf nodes, nodes at upper level are computed by hashing two fresh consecutive nodes from the lower level. This process is carried on until we are left with only one string of length k , called the root of the tree. This process is shown in figure 1.

We shall refer to the whole tree by $\text{Merkle}(y)$.

Commitments using Merkle Trees

Merkle trees can be used to commit to long strings (e.g., $n = O(2^k)$ in above construction of Merkle tree). These commitments can then be selectively opened to expose only selected parts of y . Details follow.

Let A be the committer and B be the receiver in a commitment protocol. B sends to A , a collision resistant function h as above. Now, following is the commitment protocol.

- **Commit.** To commit to a string y , A computes the Merkle tree on y as above and sends the root of this tree to B .
- **Decommit.** For the decommit phase, suppose A wants to open only the first block y_1 of y . Then, A sends to B the siblings of all the nodes in the path from y_1 to the root node.

Example: In the figure of previous subsection, the path from y_1 to the root node $X^{(\ell)}$ has following nodes: $X_1^{(1)}, X_1^{(2)}, \dots, X_1^{(\ell-1)}$. Thus, to decommit, A will send to B , the following nodes: $X_2^{(1)}, X_2^{(2)}, \dots, X_2^{(\ell-1)}$. A also provides y_1 and y_2 .³

Now B reconstructs the root node from these nodes and verifies that it is as was committed during the commit phase. An important property of Merkle commitments is that it is binding: if committer finds two different strings leading to the same root node, it means it has found collisions on the hash function h .

Notice that Merkle commitments as described above are binding but not hiding. If the receiver B selects the same hash function for each commitment, it can always identify whether A is committing to same y or not. Thus, to make Merkle commitments hiding, instead of sending the root of $\text{Merkle}(y)$, A commits to it using some standard commitment schemes. Now the decommitment phase will also require the decommitment to the commitment of root in addition to all the necessary nodes to verify the root.

One concern with Merkle commitments is that the committer may not commit to all tree leaves, but rather pick some intermediate node in the tree randomly as a starting point. How do we check that this is not the case. One option is to ask the committer to turn y into an error correcting code $C(y)$ which can recover y from, say $\frac{1}{4}C(y)$ errors. The verifier then checks $O(\log |X'|)$ leaves which guarantee with high probability all but $1 - \frac{1}{\text{poly}(|X'|)}$ leaves exist. Thus, in the proof, given a root of the Merkle tree, one can rewind the prover and get all but say 1% of the leaves and then use error correcting program to recover X' completely.

2.3 Constructing efficient ZK-proofs

We will use the PCP theorem as a black-box to construct efficient ZK-proofs for all languages in **NP**. To do so, we will construct an efficient ZK-proof for 3-SAT. As 3-SAT is **NP**-complete, the result will follow.

Following is an efficient strategy to prove $X \in 3\text{-SAT}$. The prover is given the witness to prove the above statement, whereas both the prover and the verifier get the statement as

³Giving y_2 is necessary. But the real block y_2 of y can be prevented from revealing by inserting a random block of size k next to each real block of y . Thus, we will be exposing only one block of y during decommit phase. This, however, is not needed for our goal of constructing efficient ZK-proofs.

their input. the protocol proceeds as follows:

- Prover P computes the string y using PCP-reductions. Let, $y = \text{PCP}(X)$.
- P executes with V the commit protocol of Merkle commitment scheme, to commit to y .
- V chooses $O(\log(|X'|))$ positions of clauses of X' uniformly at random and sends them to P .
- P now decommits to all those blocks y that contain an assignment to a clause asked by V , using the decommit protocol of Merkle commitments.
- V verifies that the root is correctly computed for each of the clauses and that all the clauses are satisfied by the obtained assignment.

Notice that given all the tests succeed, above proof convinces the verifier that the statement is true. Prover cannot change his mind on the variable assignments once it has sent the root to the verifier (and by previous discussion the prover has to commit to a large fraction of the PCP proof). To cheat, the prover must create collisions on the hash function. Because of the PCP theorem, only logarithmic number of clauses suffice to reduce the error probability down to a small value.

Making the protocol ZK. The protocol presented above is *not* zero knowledge. This is because the verifier learns the assignment to the variables. Furthermore, while learning the path for verifying the root of Merkle(PCP(X)), it might also learn some other part of the witness w' . Thus the protocol is not zero-knowledge. However, notice that using the commitment schemes and normal zero-knowledge proof systems, we can make above protocol zero-knowledge as follows.

Instead of actually decommitting, P provides a zero-knowledge proof of the following statement using standard zero-knowledge proof techniques: “Had the prover decommitted, the verifier would have accepted”. This makes the protocol computationally zero-knowledge.

Communication overhead. We now measure the communication complexity of the protocol. Notice that sending h , and the root of the Merkle tree takes only constant number of bits. Sending the random clauses to the prover takes $O(\log(|X|))$ bits. Committing to the assignments and nodes of Merkle tree again takes $O(\log(|X|))$ bits. Proving to the verifier that the committed values are correct requires a zero-knowledge proof on a statement of size $O(\log(|X|))$ bits. Thus, it will take only $O(\text{poly}(\log(|X|)))$ bits. Thus, overall communication complexity is $O(\text{poly}(k) + \text{poly}(\log(|X|)))$.

3 Barak's Non-Black-Box Technique

In traditional zero-knowledge proofs, the simulator has only black-box access to the verifier. That is, it is not given access to the code of the verifier. Barak showed that if the simulator is given access to the code of the verifier (also called non-black-box access to the verifier), then some of the goals that are impossible to achieve in traditional (black-box) zero-knowledge model, become possible to achieve.

To give a concrete example, let us define Arthur-Merlin type proofs. We call a zero-knowledge proof to be of Arthur-Merlin type if the coins of the verifier are all *public* and the verifier is restricted to only send his coin flips. Goldreich and Krawczyk showed that there do not exist constant round black-box ZK-proofs of Arthur-Merlin type for non-trivial (i.e. outside **BPP**) **NP** languages. Barak showed that constant-round Arthur-Merlin type ZK-proofs for non-trivial languages in **NP** exist in the non-black-box model.

We have seen the definition of zero-knowledge interactive proof systems in previous lectures. We remind the definition of zero-knowledge below. Other properties (completeness and soundness) were also defined in previous lectures (skipped here).

Zero Knowledge. $\forall x \in L$, where $L \in \mathbf{NP}$, $\forall V^*$, $\exists \mathcal{S}_{V^*}(x)$ such that the simulator \mathcal{S} can simulate the view of V^* interacting with the real prover such that no PPT distinguisher can distinguish the simulated view from the real view with non-negligible probability.

Now we present Barak's protocol, which is a constant round Arthur-Merlin type non-black-box zero-knowledge proof system for proving membership of x in L . Denote by $\text{EZK}(x)$, the communication efficient ZK proof of membership of x in L which we described in the previous section. Let $(\text{Com}, \text{DCom})$ denote a secure commitment scheme with Com being the commitment procedure and DCom being the decommitment procedure. Let h be a collision resistant hash function. If T and a are two strings, then the notation $T(a)$ denotes the output of a Turing machine described by the string T and run on input a . Finally, let 0^n denote a string of n contiguous zeros. Following is the Barak's protocol:

- V sends h to P .
- Let $z = 0^n$. P computes $\alpha = \text{Com}(z)$ and sends α to V .
- V chooses a random string r uniformly, and sends it to P .
- Let $y \equiv$ "Either $z(\alpha) = r$ or $x \in L$ ". Here $z = \text{DCom}(\alpha)$. Prover P gives an $\text{EZK}(y)$ to V .

We require that $r \gg \alpha$ so that there is enough entropy and hence the prover cannot predict r in the real execution. The intuition behind Barak's protocol is that because r is hard

to guess, P cannot come up with a proof for $z(\alpha) = r$ with non-negligible probability and hence it cannot cheat the verifier unless it knows a way to prove that $x \in L$. Whereas on the other hand, the simulator can simulate a transcript by setting $z = V^*$ (as it has access to the **code** and the coins of V^* and all the side information that V^* has)⁴. Because now $z = V^*$, we will always have $z(\alpha) = r$ and the simulator can generate a proof of this, and hence it can simulate indistinguishable transcripts. Notice that the description of z might be arbitrarily long (even of exponential size). Because of this, we require efficient zero-knowledge proofs which use Merkle commitments. Also notice that this is the first zero-knowledge proof in which the simulator does not require rewinding.

⁴That is, the simulator commits $\alpha = \text{Com}(V^*, V^*$'s random coins, and all the side information V^* has)


Lecture 16

1 Software Protection

In this lecture¹ we consider the problem of software protection. Informally, we show a method of simulating an arbitrary program by a program which runs on a CPU having access to a random oracle and a few protected registers such that no information other than the input-output behavior of the program can be deduced.

The cryptographic tool that we are using is pseudo-random functions of [Goldreich, Goldwasser, Micali] which were discussed in the previous lecture. Pseudo-random functions are used in this work for both (private-key) encryption and to simulate a random oracle (see previous lecture for definitions).

1.1 The probabilistic RAM Model



Probabilistic RAM Model

The probabilistic RAM consists of a CPU which has some protected registers and has access to a random oracle. The CPU is connected to a random access memory (RAM) whose contents can be read or altered by the adversary (or the software pirate). The adversary can also snoop on the address bus of the CPU and figure out the order in which the RAM is accessed. However there is no way by which one can find the contents of the CPU's register at a given point in time.

1.2 Protection of Program's Data

1.2.1 Protection against non-tempering adversaries

A non-tempering adversary can read the RAM and watch the communication between the RAM and the CPU, but it can't modify or alter the RAM contents. The first step is to encrypt the program and its data so that no one else copy or use it. This can be done by encrypting the contents of the RAM. One can replace the bit b_i stored at the i th location by $b'_i = F(i) \oplus b_i$,

¹Scribe notes taken by: Rahul Garg, September 27, 29, 1994

where $F(\cdot)$ is the random function computed using the oracle. The program reads b'_i from the RAM and computes $b_i = b'_i \oplus F(i)$ to compute the actual value of the data. Since $F(\cdot)$ is a random function, the contents of the RAM after this transformation are random to anyone who doesn't know $F(\cdot)$. The next step is to simulate a read (or a write) by a (dummy) read followed by a write at the same memory location. This would hide whether the access was a read or a write. (Thus the most frequently used RISC instructions "load" and "store" would become indistinguishable from each other).

1.2.2 Protection against tempering adversaries

A tempering adversary is allowed to change the contents of RAM. Such an adversary might try to change the RAM contents to get additional information about the working of the program. To secure against such adversaries store of a value v at a location i is simulated as store of a pair $(v, F(v))$. While reading back the value, a pair (v, y) is read and $F(v)$ is computed to check if $y = F(v)$. An adversary which wants to replace v by v' should also replace $F(v)$ by $F(v')$, so that the simulator doesn't come to know that the RAM contents have been altered. Since $F(\cdot)$ is a random function known only to the CPU the adversary can't compute $F(v')$ and hence can't change the RAM contents to store a fake value. However it can still copy a $(v, F(v))$ pair from one RAM location to another and clobber the program's memory. The program can be protected from this by storing a value v at location i as $(v, F(i.v))$ where $i.v$ represents concatenation of i and v . The adversary could still alter a location by placing a value which was stored there in the past. This can be avoided by storing a version number with the value. The version number is incremented each time the location is updated. Thus a value could be stored at location i as $(v, \text{version}, F(v.i.\text{version}))$ (we discuss later how appropriate version number is checked).

With the above encryption a program and its data can be hidden securely. Even after this an adversary might be able to get some information about the working of a program by carefully analyzing the order in which the program accesses its memory. The final milestone in achieving the goal of software protection is to hide the access pattern of a program. In other words the programs have to be transformed in such a way that the access patterns of any two program taking the same running time are indistinguishable from each other.

This problem is formalized as efficient simulation of any RAM program on an oblivious RAM. A machine is oblivious if the sequence in which it accesses memory locations is equivalent for any two inputs with same running time. For example an oblivious Turing Machine is one for which the movement of heads on the tapes is identical for each computation. In 1979 Pippenger and Fischer showed how a two tape oblivious Turing Machine can simulate, on-line, a one tape Turing Machine, with a logarithmic slowdown in the running time. An analogue result for the RAM model of computation is presented.

1.3 The restricted problem

Consider a simpler problem for restricted set of programs which access each location of the RAM exactly once. A solution to the problem can be attained by shuffling the memory locations. After the memory is shuffled (randomly permuted), the accesses to the RAM would appear to be like a random permutation and hence would be oblivious. The same idea can be extended to the general case by organizing the memory into hierarchy.

Lets assume that the program accesses each memory location exactly once. If the memory is permuted randomly before the accesses are made then the accesses will appear to be a random permutation. Thus the access pattern of any two programs will be indistinguishable. The problem reduces to the problem of permuting the memory obliviously.

Assume that the original program accesses n words of the RAM. Allocate $4n \cdot O(\log n)$ words of memory and organize it into $4n$ buckets of size $O(\log n)$ each. Now the original memory is hashed into these buckets using the random oracle to compute the hash function. The hash function $h_s(i)$ can be defined as $F(i.s) \bmod 4n$, where s is chosen randomly.


Notice that n items are stored into a hash table with $4n$ entries according to a random oracle. Hence the probability (taken over a choice of randomly chosen function) that the number of entries in a bucket will exceed $O(\log n)$ (bucket overflow) is negligible.

If one can somehow transfer the original RAM contents into the data structure described above, then the access patterns can be hidden. An access to a memory location i can be translated into searching for the appropriate tuple in bucket $h_s(i)$. Any two distinct memory locations are assigned independent buckets and each memory location is accessed exactly once, thus the access pattern (i.e. the bucket access pattern) is independent of the memory access pattern. Whenever a bucket is accessed, each of its element is examined, as a result the memory accesses looks as if one is randomly choosing a bucket and scanning it sequentially.

The final problem is of hashing the elements into buckets obliviously (oblivious hash). This problem can be solved if one can perform sorting of elements obliviously. We first describe the a scheme to perform oblivious sorting and then we show a scheme to do oblivious hash.

1.3.1 Oblivious Sorting

The idea is to implement Batcher's Sorting Network which consists of a number of 2-element sorters as shown in the figure. A 2 element sorter takes two inputs and it outputs them in sorted order. To sort n elements, a network of $(\log n)^2$ such 2-element sorters can be built. Since the network consists of 2-element sorters only, an implementation of this in software will only consist of operations which read two a-priory predetermined memory location and write them back after swapping them (if needed). Thus the order in which the memory locations are accessed by the program will be fixed in advance and be independent of the data being sorted.



4x4 Sorting Network using 2x2 Sorters

1.4 Oblivious Hashing

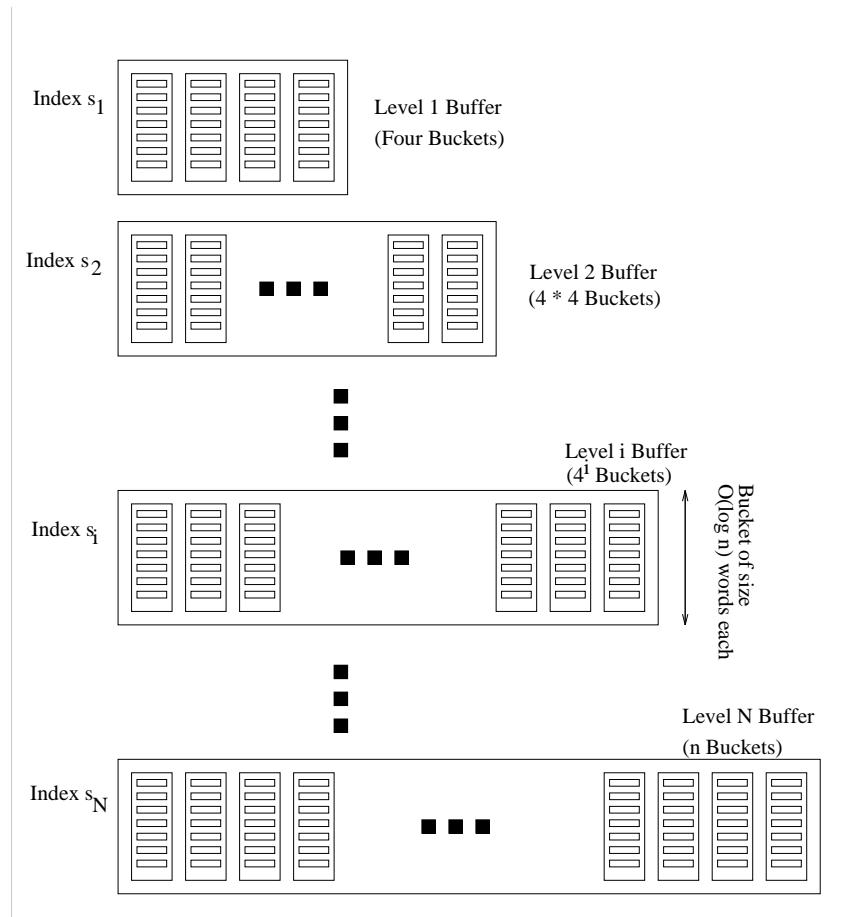
Now that we know how to perform oblivious sorting, we can make its use to perform oblivious hash. The oblivious sorting is used to to the following.

- Sorting elements within a bucket to group together the empty and full entries. Empty entries are tagged 0 and full entries are tagged 1, the sorting is done on tags.
- Sorting elements within two buckets to move all the non empty entries into a single bucket.
- Sorting buckets as individual elements.

First, all the elements of the RAM are tagged by their hash function and then, $4n$ dummy elements corresponding to the $4n$ possible tag values are introduced to make sure that each bucket has at least one element after hashing. The elements are sorted according to their tags and all the elements with the same tag value are put into the same bucket (using oblivious sort in way as described above). The buckets are then sorted according to tags of their elements to make sure that the i th bucket contains all the elements having a tag value of i . This is always ensured because dummy elements make sure that bucket corresponding to each tag value will contain at least one element. Thus an element of the RAM stored at location i is now moved to the bucket numbered $h_s(i)$.

1.5 Simulation on an Oblivious RAM

Instead of a single hash table introduced to solve the simple problem, several hash tables of different sizes are used in the general case. There are N hash tables and each table associated with an index s_i (see figure). These hash tables are called “buffers” which are obviously hashed at different intervals (smaller buffers are hashed more often than larger ones). The idea is to ensure that for each buffer, no element in that buffer is accessed more than once in between two consecutive oblivious hashes, thus reducing the general problem to the restricted problem already solved.



The Hierarchical organization of RAM

Initially memory is allocated for $N = (1 + \lceil \log_4 n \rceil)$ (n is the length of the input) distinct buffers, where i th level buffer is a hash table of size 4^i buckets and each bucket is of size $\log n$. For each buffer of level i , a random seed s_i is picked and a hash function $h_{s_i}(\cdot)$, where $h_{s_i}(v) = F(s_i.v) \bmod 4^i$, is associated with it. All the buffers are empty in the beginning. The program begins with copying (i.e. hashing) the program and input into the biggest, N th level buffer. During the simulation if the program refers to a location v , the first level 1 buffer is completely scanned to find it. If it is not found in the first level buffer, then $h_{s_2}(v)$ bucket of the next level buffer is scanned. If it is still not found there then $h_{s_3}(v)$ bucket of the next level buffer is scanned and so on. Note that address of each element is stored along-with the element, so one can compare the address to find whether the an element in the bucket is actually the one which is needed. After the required element is found, the buffers of remaining level are randomly scanned till the N th buffer. The value is put into level 1 buffer after the access. After some time the buffer at level 1 may start overflowing. To avoid this potential overflow, we move the contents of buffer level i to level $i + 1$ after every 4^{i-1} accesses. The contents of buffer at level i and $i + 1$ are obviously hashed into level $i + 1$ according to a newly chosen random seed s'_i . Thus, each time when a buffer at level i is at most a quarter full, it is emptied into a higher level. If the higher level does not exist, it is allocated.

Notice that at round 4^i , the total number of entries in the first i buffers is at most 4^i and

hence merging their contents into i th buffer (of size 4^{i+1} buckets) can have only a quarter of the entries. When the contents of a lower buffer are moved into the larger buffer, a new hash function is picked and the elements are rehashed according to this new hash function. Thus the collisions which occurred in the smaller buffer do not influence the collisions which may occur in the bigger buffer. That is, the collisions do not accumulate.

1.5.1 Obliviousness of Access Pattern

From the restricted problem, we know that as long as any memory reference is made at most once, the choice of the hash bucket accessed appears completely random to the adversary.

The crucial observation is that for all buffers other than 1, the program never looks for the same memory location more than once, in between hashing stage of that buffer. After a tuple is found at some level, it is moved to the first level. Thus subsequent searches for the same location will find the required tuple in smaller levels unless the tuple migrates to the same level again. Each time something is migrated to a same level, the level is rehashed according to a new randomly selected hash function. The bucket in which a tuple would go will be independent of its previous bucket. Thus, for all buffers $i \geq 2$ and location v , whenever a new s_i is picked, $h_{s_i}(v)$ is computed at most once and hence the access patterns at each level appear to be random to the adversary.

1.5.2 Overhead of Oblivious Simulation

Let t be the running time of the program. At the end of simulation, there would be at most $N = O(\log t)$ buffers. For each memory access in the original program, the transformed program scans all the buckets of level 1 buffer and then scans a single bucket of size $O(\log t)$ on every other level, resulting into $O((\log t)^2)$ accesses for each original access. In addition, during simulation level i and level $i - 1$ buffers are hashed into level i buffer a total of 4^{n-i+1} times, $i = 2, \dots, n$. Since the joint size of level $i - 1$ and level i buffers is $O(4^i)$ buckets of size $O(\log t)$, it takes $O(4^i \cdot \log t \cdot \log(4^i \cdot \log t))$ steps to obviously hash them. Thus the total number of steps needed to perform all the hashing of all the buffers is equal to :

$$\sum_{i=2}^N 4^{N-i+1} \cdot 4^i \cdot \log t \cdot \log(4^i \cdot \log t) = O(t \cdot (\log t)^3)$$

Therefore the overhead in the oblivious simulation of an arbitrary RAM program is a factor of $O((\log t)^3)$ (with a negligible probability of failure of the simulation).