Advanced Cryptography — Midterm Exam

Serge Vaudenay

13.5.2014

- duration: 3h00
- documents are allowed
- a pocket calculator is allowed
- communication devices are not allowed
- the exam invigilators will *not* answer any technical question during the exam
- readability and style of writing will be part of the grade

1 Cryptosystem based on Matrices

We define a new cryptosystem. Let p be a large prime number. Let $a \in \mathbf{Z}_p$ and $b \in \mathbf{Z}_p^*$ be arbitrary such that $a^2 + b^2 \in \mathbf{Z}_p^*$. Let G be the matrix $G = \begin{pmatrix} a & b \\ -b & a \end{pmatrix}$. The public parameters are given by (p, a, b). We let $x \in \mathbf{Z}$ be a secret key and let $Y = G^x$ be a public key. To encrypt $m \in \mathbf{Z}_p$, we pick a random integer r, compute $U = G^r$, $V = Y^r$, and $w = V_{1,1} + m \bmod p$ (where $V_{1,1}$ is the upper left coefficient of V). The ciphertext is the pair (U, w).

Q.1 Explain how to decrypt.

Q.2 Let $J = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$. Show that the following properties are equivalent.

- $p \bmod 4 = 1$;
- there is an invertible 2×2 matrix P with coefficients in $\mathbf{Z}_p$ such that $P^{-1}JP$ is a diagonal matrix;
- there is an invertible 2×2 matrix P with coefficients in $\mathbf{Z}_p$ such that $P^{-1}GP$ is a diagonal matrix.

HINT: we recall that a 2×2 matrix is diagonalizable if and only if it has two different eigenvalues or it is already diagonalized.

Q.3 For $p \bmod 4 = 1$, show that the key recovery problem reduces to the discrete logarithm problem in $\mathbf{Z}_p^*$.

Q.4 For $p \bmod 4 = 3$, we define $\mathbf{K} = \mathbf{Z}_p[x]/(x^2 + 1)$, the field of $\mathbf{Z}_p$ extended with a root θ of $x^2 + 1$. By working with matrices with coefficients in $\mathbf{K}$, show that the key recovery problem reduces to the discrete logarithm problem in $\mathbf{K}^*$.

HINT: show that J is diagonalizable as a matrix with coefficients in $\mathbf{K}$.

Q.5 In general, give a positive integer q such that G^q is the identity matrix.

2 Predicate Encryption

We define a new cryptographic primitive called *predicate encryption*. We consider a predicate P. A predicate encryption for P is defined by four algorithms:

$\mathsf{Setup}(1^\lambda) \to (\mathsf{pp}, \mathsf{msk})$**:** (probabilistic) given a security parameter λ, it generates a key pair where msk is the master key (secret) of the authority and pp is the public parameter, which is distributed to all participants.

$\mathsf{Keygen}(\mathsf{msk}, k) \to \mathsf{sk}$**:** (probabilistic) given a key k for the predicate $P(k, .)$, the authority generates a secret key sk for a participant Bob.

$\mathsf{Enc}(\mathsf{pp}, \mathsf{ind}, m) \to c$**:** (probabilistic) given a value ind called *index* and a message m, Alice generates a ciphertext c. Note that this is independent of k and Bob.

$\mathsf{Dec}(\mathsf{sk}, c)$**:** (deterministic) given the ciphertext and the secret sk, this decryption algorithm yields m if $P(k, \mathsf{ind})$ is true and $\perp$ otherwise. (I.e., this is the correctness property of the primitive.)

The security of this primitive specifies that from c, Bob (holding sk) does not learn m if $P(k, \mathsf{ind})$ is false.

Q.1 In *identity-based encryption* (IBE), there is an authority holding a master key, distributing some public parameters to everyone, and giving a secret key to each user. We want that if, e.g., Alice is offline and cannot connect to retrieve the public key of Bob to any public directory, she can still encrypt a message which can only be decrypted by Bob (and the authority). More precisely, we have

$\mathsf{IBE.Setup}(1^\lambda) \to (\mathsf{pp}, \mathsf{msk})$**:** generate the public parameters and the master key.

$\mathsf{IBE.Keygen}(\mathsf{msk}, \mathsf{id}) \to \mathsf{sk}$**:** given the identity of Bob, generate his secret key.

$\mathsf{IBE.Enc}(\mathsf{pp}, \mathsf{id}, m) \to c$**:** encrypt a message m for a given identity.

$\mathsf{IBE.Dec}(\mathsf{sk}, c) \to m$**:** decrypt the message given the correct secret key.

We want that a user who does not hold the correct sk learns nothing about m.

By well choosing a predicate, construct one IBE scheme with the above syntax based on predicate encryption.

Give an argument for the security.

Q.2 In *ciphertext-policy attribute-based encryption* (CP-ABE), there is an authority holding a master key, distributing some public parameters to everyone, and giving a secret key to each user. Each user has list $z = (z_1, \ldots, z_n)$ of attributes associated to a semantic. (E.g., if member of EPFL or not, if MSc student or faculty member or admin staff, if registered to the Advanced Cryptography class, etc.) Each message is encrypted with a formula φ. It can only be decrypted for holders of attributes satisfying the formula φ. (E.g., expressing that the message can only be decrypted by MSc students or admin staff of EPFL who registered to Advanced Cryptography.) More precisely, we have

$\mathsf{CPABE.Setup}(1^\lambda) \to (\mathsf{pp}, \mathsf{msk})$**:** generate the public parameters and the master key.

$\mathsf{CPABE.Keygen}(\mathsf{msk}, z) \to \mathsf{sk}$**:** given the attributes of Bob, generate his secret key.

$\mathsf{CPABE.Enc}(\mathsf{pp}, \varphi, m) \to c$**:** encrypt a message m for a formula φ.

$\mathsf{CPABE.Dec}(\mathsf{sk}, c)$**:** obtain m is $\varphi(z)$ holds and $\perp$ otherwise.

We want that a user who holds attributes not satisfying φ learns nothing about m.

By well choosing a predicate, construct one CP-ABE scheme with the above syntax based on predicate encryption.

Give an argument for the security.

Q.3 Given a modulus N and a length n, we assume that, when $\mathsf{ind} = (a_1, \ldots, a_n)$ and $k = (b_1, \ldots, b_n)$ are in $\mathbf{Z}_N^n$, we have a predicate encryption scheme for inner product (IP), i.e., for the predicate

$$P(k, \mathsf{ind}) \iff a_1 b_1 + \cdots + a_n b_n \equiv 0 \pmod{N}$$

We denote this scheme by $(\mathsf{IP.Setup}_N^n, \mathsf{IP.Keygen}_N^n, \mathsf{IP.Enc}_N^n, \mathsf{IP.Dec}_N^n)$.

Based on an IP scheme with the above syntax, construct one predicate encryption scheme where users are associated to a variable x and messages are encrypted with a polynomial f of bounded degree d, and the decryption works if x is a root of f. More precisely, construct a scheme $(\mathsf{POL.Setup}_N^d, \mathsf{POL.Keygen}_N^d, \mathsf{POL.Enc}_N^d, \mathsf{POL.Dec}_N^d)$ in which $\mathsf{POL.Keygen}_N^d(\mathsf{msk}, x)$ gives sk, $\mathsf{POL.Enc}_N^d(\mathsf{pp}, f, m)$ gives c, and $\mathsf{POL.Dec}_N^d(\mathsf{sk}, c) = m$ if and only if $f(x) = 0$.

Q.4 We now give several variables to the participants. Extend the previous construction to multivariate polynomials.

Q.5 We consider a predicate $P(a, b)$ which is a CNF (conjunctive normal form) of terms of form $a_i = b_j$. I.e., $P(a, b) = \bigwedge_u \bigvee_v (a_{i_{u,v}} = b_{j_{u,v}})$. ($\vee$ is a notation for the OR and $\wedge$ is a notation for the AND.) Given some (secret) random r_u, we consider the polynomial $f(a, b) = \sum_u r_u \prod_v (a_{i_{u,v}} - b_{j_{u,v}})$.

Q.5a If $P(a, b)$ is true, show that $f(a, b) = 0$.

Q.5b Given a and b fixed such that $P(a, b)$ is false, show that $\Pr[f(a, b) = 0]$ is small, over the distribution of the r_u's.
HINT: assume that N is prime.

Q.5c From the IP scheme, show that we can construct a predicate encryption scheme for the predicate P. How large is n in the above construction?

3 Distribution Fitting

Let p be a prime number and $\ell \leq \log_2 p$. Let $r = p \bmod 2^\ell$. Let $V \in_U \mathbf{Z}_p$ be uniformly distributed. Let $X = V \bmod 2^\ell$. We want to distinguish the distribution of X from the uniform distribution over $\{0, \ldots, 2^\ell - 1\}$.

Q.1 Compute the distribution of X: depending on $x \in \{0, \ldots, 2^\ell - 1\}$, provide a formula to compute $\Pr[X = x]$ in terms of x, r, p, ℓ.

Q.2 Given a single sample, compute the best advantage of a distinguisher to distinguish the distribution of X from a uniform one.

Q.3 We assume that p is selected arbitrarily among prime numbers of k bits. With ℓ fixed and in the worst case for p, how large should k be so that the best advantage given a single sample is lower than $2^{-\ell}$?

Q.4 Assume the identified condition is satisfied, approximate the Chernoff information by the squared Euclidean imbalance and estimate the number of samples needed to distinguish the distribution of X from a uniform one.