CS60088 Foundations of Cryptography, Spring 2013–2014

Mid-Semester Test

19–February–2014 CSE-107, 2:00–4:00pm Maximum marks: 35

Roll no: ____________ **Name:** ______________________________

[*Write your answers in the question paper itself. Be brief and precise. Answer all questions.*]

1. So far, we have discussed signature schemes *with appendix*, in which the signature is appended to the message. In a signature scheme *with message recovery*, only the signature is presented to the verifier. The verification algorithm recovers the message. If the message is supposed to contain some redundancy (like it is an English text), then the verification of a forged signature recovers a message which is expected with high probability not to contain the redundancy.

Nyberg and Rueppel (Eurocrypt 1994) propose an ElGamal-like signature with message recovery. We work in a finite field $\mathbb{F}_p$ with an element g of large prime order q dividing $p-1$. The signature on a message m is pair (r, s) generated as follows:

$$\begin{aligned} r &\equiv mg^l \pmod p \text{ for a randomly chosen } l \in \mathbb{Z}_q, \\ r' &= r \text{ rem } q, \\ s &\equiv -l - r'x \pmod q. \end{aligned}$$

(a) How can the message m be recovered from the signature (r, s)? **(5)**

Solution $m \equiv rg^{-l} \equiv rg^{s+r'x} \equiv rg^s y^{r'} \pmod p$, where $r' = r$ rem q.

(b) Show how existential forgery is possible for the Nyberg–Rueppel scheme. **(5)**

Solution If (r, s) is a valid Nyberg–Rueppel signature on a message m, then $(r, s+t)$ is again a Nyberg–Rueppel signature on the message $mg^t \pmod p$ for any t.

Moreover, if we start with any arbitrary signature (r, s), the message-recovery algorithm outputs an m on which (r, s) is a valid signature.

2. Sakurai and Takagi (2001) propose a semantically secure RSA-like encryption scheme. Their scheme is based upon earlier works by Paillier (Eurocrypt 1999) and by Catalano, Gennaro, Howgrave-Graham and Nguyen (CCS 2001). A modulus $n = pq$ is chosen with suitably large primes p, q. The public key of Alice is (n and) an e (usually small for efficiency) coprime to $\phi(n)$, and her private key is $d \equiv e^{-1} \pmod{\phi(n)}$. The encryption involves arithmetic modulo n^2. For a message $m \in \mathbb{Z}_n$, a uniformly random $r \in \mathbb{Z}_n^*$ is chosen, and the ciphertext is computed as $c \equiv r^e(1+mn) \pmod{n^2}$.

(a) Prove that for every $c \in \mathbb{Z}_{n^2}^*$, there exist unique $r \in \mathbb{Z}_n^*$ and $m \in \mathbb{Z}_n$ such that $c \equiv r^e(1+mn) \pmod{n^2}$. **(5)**

Solution Consider the function $f : \mathbb{Z}_n^* \times \mathbb{Z}_n \to \mathbb{Z}_{n^2}^*$ that maps (r,m) to $c \equiv r^e(1+mn) \pmod{n^2}$. We show that f is a bijection. Let $c = f(r,m) = f(\rho,\mu)$. Modulo n, we have $r^e \equiv \rho^e \equiv c \pmod{n}$. Exponentiation to the d-th power gives $r \equiv \rho \equiv c^d \pmod{n}$. But both r and ρ are from $\mathbb{Z}_n^*$, so $r = \rho$. This in turn implies that $m = \mu$, that is, the function f is injective. Since both the domain and the range of f have the same size, namely $\phi(n^2) = p(p-1)q(q-1) = n\phi(n)$, we conclude that f is a bijection.

(b) Explain how decryption is carried out in the Sakurai–Takagi scheme. **(5)**

Solution Alice first recovers r by RSA decryption. The ciphertext c reduced modulo n is $r^e \pmod n$. Its d-th power exponentiation is $c^d \equiv r \pmod n$. Then, r^e is computed modulo n^2. But $\gcd(r,n) = 1$, so r^e is invertible modulo n^2 too, and m is retrieved as $\left(c(r^e)^{-1} - 1 \pmod{n^2}\right)/n$.

(c) Demonstrate that the Sakurai–Takagi scheme is not secure against chosen-ciphertext attacks. **(5)**

Solution Sakurai–Takagi encryption is additively homomorphic. More precisely, $f(r,m)f(\rho,m) = f(r\rho, m+\mu)$. In order to decrypt a target ciphertext c corresponding to the plaintext message m, Malice multiplies c by the ciphertext $c' = f(\rho,\mu)$ for randomly chosen ρ and μ. But then, cc' is a random element of $\mathbb{Z}^*_{n^2}$, so Alice decrypts cc' to reveal the message $m + m'$ to Malice. Since m' was chosen by Malice, he obtains the target plaintext m.

(d) Prove that Sakurai–Takagi encryption is all-or-nothing secure against passive attacks if and only if the RSA assumption holds. **(5)**

Solution If the RSA assumption does not hold, Malice (a passive eavesdropper in the part) can compute $r \pmod{n}$ by RSA decryption of r^d. He then recovers m as in Part (b).

Conversely, suppose that there is an oracle ST that, upon the input of $c \in \mathbb{Z}_{n^2}^*$, returns the plaintext m (with non-negligible advantage). This in turn implies that $r^e \equiv c(1+mn)^{-1} \pmod{n^2}$ can be computed easily. Let us see how this oracle can be used to solve the RSA problem.

Let $c \equiv r^e \pmod{n}$ be the challenge RSA ciphertext that Malice wants to decrypt. Malice queries the oracle ST with c (treated as an element of $\mathbb{Z}_{n^2}^*$) as input. Let $c = f(\rho, \mu)$ (see Part (a)). Malice then computes $\lambda \equiv \rho^e \pmod{n^2}$ as mentioned in the last paragraph. Since $c \equiv r^e \equiv \rho^e \pmod{n}$, and $r, \rho \in \mathbb{Z}_n^*$, we have $r = \rho$, that is, Malice has computed $\lambda \equiv r^e \pmod{n^2}$. Indeed, Malice could have sent any ciphertext of the form $c + kn$ to the oracle to compute this λ.

Malice then makes a second query to the oracle ST, this time with $c/2^e \pmod{n}$ (again treated as an element of $\mathbb{Z}_{n^2}^*$) as input. This enables Malice to compute $\lambda' \equiv (r')^e \pmod{n^2}$, where $r' \equiv r/2 \pmod{n}$.

If r is even, we have $r' = r/2$, so $2^e \lambda' \equiv \lambda \pmod{n^2}$. If n is odd, then $r' = (r+n)/2$ so that $2^e \lambda' \equiv (r+n)^e \equiv r^e + er^{e-1}n \pmod{n^2}$ (by the binomial theorem). By construction, $er^{e-1} \in \mathbb{Z}_n^*$, that is, $2^e \lambda' \not\equiv \lambda \pmod{n^2}$. This means that by checking whether the congruence $2^e \lambda' \equiv \lambda \pmod{n^2}$ holds, Malice can decide the LSB of r, that is, the RSA parity oracle exists. This in turn implies that the RSA assumption does not hold.

(e) The (decisional) e-th power residuosity problem is the determination of whether an $x \in \mathbb{Z}_{n^2}^*$ can be expressed as $x \equiv r^e \pmod{n^2}$ for some $r \in \mathbb{Z}_n^*$. Prove that Sakurai–Takagi encryption is semantically secure (that is, IND-CPA secure) if and only if the e-th power residuosity problem is intractable. **(5)**

Solution First, suppose that a decider (an oracle) for the e-th power residuosity problem exists. Using this, Malice can win the IND-CPA game as follows. Malice chooses $m_0 = 0$ and $m_1 = 1$. We have $c_0 \equiv r_0^e \pmod{n^2}$ and $c_1 \equiv r_1^e(1+n) \pmod{n^2}$ for some $r_0, r_1 \in \mathbb{Z}_n^*$. Clearly, c_0 is an e-th power residue. If $c_1 \equiv r_2^e \pmod{n^2}$ for some $r_2 \in \mathbb{Z}_n^*$, then $c_1 = f(r_1, 1) = f(r_2, 0)$, a contradiction to the bijectivity of f (see Part (a)). Therefore, c_1 not an e-th power residue. Thus, the decider for the e-th power residuosity problem straightaway reveals to Malice the random choice between c_0 and c_1 made by Alice's encryption oracle, that is, Sakurai–Takagi encryption is not semantically secure.

Conversely, suppose that Sakurai–Takagi encryption is not semantically secure, that is, an oracle ST-CPA exists that, given m_0, m_1, c^*, reveals to Malice (with non-negligible advantage) whether c^* comes from m_0 or from m_1. Using this oracle, the e-th power residuosity problem can be solved as follows.

Suppose that we want to decide whether $z \in_U \mathbb{Z}_{n^2}^*$ is an e-th power residue or not. For randomly chosen m_0, m_1 (of the same length), we query ST-CPA with $m_0, m_1, c^* \equiv z(1+m_1 n) \pmod{n^2}$ as input. If $z \equiv r^e \pmod{n^2}$ is an e-th power residue, then $c^* = f(r, m_1)$, and ST-CPA responds by outputting the bit 1 (with non-negligible advantage). On the other hand, if z is not an e-th power residue, then $z \equiv r^e(1+mn)$ for some $r \in \mathbb{Z}_n^*$ and $m \in \mathbb{Z}_n$ with $m \neq 0$ (recall that the function f is a bijection). But then, $c^* = z(1+m_1 n)$ is an encryption of $m_1 + m \pmod{n}$. Now, $m \neq 0$, so $m \neq m_1$. Moreover, for a random z, we have $m \equiv m_0 - m_1 \pmod{n}$ with only negligible probability. Therefore, c^* is the encryption of neither m_0 nor m_1. In this case, ST-CPA may report failure or outputs 0 or 1 randomly.

In both the cases, if ST-CPA outputs 0 or 1, we output the same bit. If ST-CPA outputs failure, then we output 0. This gives us a Monte-Carlo-type randomized algorithm to solve the e-th power residuosity problem.

For leftover answers and rough work

For leftover answers and rough work