Public Key Cryptography and Chosen Ciphertext Attacks

15-859I
Spring 2003

Chosen Ciphertext Attacks

- As with symmetric crypto, we are concerned stronger attacks than chosen plaintext.
- One example: the *lunchtime* attack:
- Charlie sneaks into Bob's office and gets his PC to decrypt a bunch of "test" ciphertexts
- When Alice sends Bob a message, Charlie attempts to use the decryptions to get some information about Alice's message.

Formal "Lunchtime CC Attack"

- Let S=(G,E,D) be a PKC.
- Let (F,T) be PPT. For bit b, define
  $p^b_{F,T}(k) = Pr[T(y,s)=1 : (e,d) \leftarrow G(1^k); (m_0,m_1,s) \leftarrow F^{D(d,.)}(e); y \leftarrow E(e,m_b)]$
- Define the advantage of A=(F,T) by
  $Adv^{lcc}_{A,S}(k) = |p^0_A(k) - p^1_A(k)|$
- S is (lunchtime) chosen-ciphertext secure if every PPT A has negligible lcc advantage.

Building LCC-secure PKC

- Let PKC S = (G,E,D) be semantically secure.
- We will need a *Noninteractive Zero Knowledge Certifying Proof System* for the language
  $L = \{ (e_1,e_2,c_1,c_2) : \exists\ b,r_1,r_2\ .\ c_i = E(e_i,b;r_i) \}$
  We call L the language of "consistent encryptions"

Noninteractive Zero Knowledge

- A NIZKPS for L is a triple (P,V,U) such that:
  - Completeness: $\forall x \in L$, w witness,
    $Pr[V(p,x,R)=1 : R \leftarrow U, p \leftarrow P(x,w,R)] = 1$
  - Soundness: $\forall p, x \notin L$,poly,
    $Pr[V(p,x,R)=1: R \leftarrow U] < 1/poly(|x|)$
  - Simulation: $\exists$ PPT S: $\forall$ poly,PPT A, $x \in L$:
    $| Pr[A(S(x))=1] - Pr[A(p,x,R)=1] | < 1/poly(|x|)$

R is called the "Common Random String"

A NIZK Proof "Gives no information" about $x \in L$.

Additional Required Properties

We require a NIZKPS (P,V,U) satisfying:

- Strong Soundness: $\forall y \notin L, p \in \{0,1\}^*$,poly
  $Pr[V(p,y,R)=1 : R \leftarrow U(|y|)] < 1/poly(|x|)$
- Validity: $\forall$ PPT C, poly:
  $Pr[V(q,y,R)=1 \ \&\ y \notin L \ \&\ x \in L: x \leftarrow C(1^k); (p,x,R) \leftarrow S(x); (q,y) \leftarrow C(x,R)] < 1/poly(k)$
- Recognizability: $\exists$PPT M recognizing (p,y,R): V(p,y,R)=1 for $y \notin L$ when $R \leftarrow S(x)$, $x \in L$.
  1-negligible probability of detecting such *invalid words*, negligible probability of false positive.

Call a NIZKPS satisfying these a *certifying system*.

Certifying Systems and NIZKPS

- Certifying Systems for L exist iff NIZKPS exist for L [NY]
- NIZKPS exist for every L ∈ NP if trapdoor OWPs exist. [BMF]
- [KP] give a NIZKPS that has proof and random string length $O(Lk^2)$ for a satisfiable of size L and tOWP on k-bit domain.
- [NY] transformation to Certifying System multiplies proof length, random string length by $O(k^2)$.

LCC-Secure PKC

- Let S=(G,E,D) be semantically secure, and let (P,V,U) be a certifying proof system for the language of consistent encryptions.
- $G'(1^k)$: draw $(e_1,d_1) \leftarrow G(1^k)$, $(e_2,d_2) \leftarrow G(1^k)$, $R \leftarrow U(k)$. Output $((e_1,e_2,R), (d_1,d_2))$
- $E'((e_1,e_2,R),b)$: choose r_1,r_2. $c_1 \leftarrow E(e_1,b;r_1)$, $c_2 \leftarrow E(e_2,b;r_2)$, $p \leftarrow P((e_1,e_2,c_1,c_2),(b,r_1,r_2),R)$ return (c_1,c_2,p)
- $D'((d_1,d_2,R),c_1,c_2,p)$ = if $V(p,(e_1,e_2,c_1,c_2),R)=0$ then $\perp$ else $D(d_1,c_1)$

Security Proof: Sketch

- Suppose A=(F,T) distinguishes E'(0) from E'(1) in LCC attack. We will write a chosen plaintext attacker with similar advantage.
- Idea: we get e_1, E(b). Choose an e_2 and b' and use simulator to generate an R so we can prove that $w = (e_1,e_2, E(e_1,b), E(e_2,b'))$ is consistent.
- Run F on (e_1,e_2,R), use d_2 to decrypt.
- Run T on w, state from F.

Formal CPA attack

- Write chosen-plaintext attacker B as follows:
- $B(e_1,c_1= E(e_1,b), e_2, c_2=E(e_2,b_2),d_i)$ =
  - $(p,R) \leftarrow S(e_1,e_2,c_1,c_2)$
  - Run F; on query check p, decrypt with d_i
  - $(s,0,1) = F(e_1,e_2,R)$
  - return $T(s,c_1,c_2,p)$.

Potential Problems

1. F, on R chosen by simulator could query the decryption of an inconsistent ciphertext.
   - Validity Lemma: If F submits an invalid word with non-negligible probability δ, it can CPA distinguish E(0) from E(1) with probability δ/2.
2. We sometimes run T on an inconsistent encryption…
   - Distinguisher Lemma: If F submits only valid words and (F,T) has advantage ε, then there is a CPA attacker with advantage ε/4.

Validity Lemma:

- Suppose F submits an invalid word with non-negligible probability δ. Consider the following distinguisher for S=(G,E,D):
- $T'(e_1,c_1=E(e_1,b))$ =
  - Choose $(e_2,d_2) \leftarrow G(1^k)$;
  - Run $B(e_1,c_1,e_2,E(e_2,0),d_2)$
  - If F queries invalid word, guess b=1
  - else, flip a coin.
- If b=0, S and U are indistinguishable, so Pr[output 1] < Pr[say valid word invalid] < δ/4.
- If b=1, F submits invalid query with probability δ, we detect with probability 1-negligible, so Pr[output 1] > 3δ/4.

Distinguisher Lemma

- Let $P_{bb'} = Pr[B(e_1,E(e_1,b),b')=1]$
- Lemma 1: $|P_{00} – P_{11}| > \varepsilon/2$.
- Corollary: either $|P_{00}-P_{01}| > \varepsilon/4$ or $|P_{01}-P_{11}| > \varepsilon/4$.
- So, to distinguish E(0) from E(1), choose $(e',d') \leftarrow G(1^k)$; in the first case we run B(e',E(e',0),e,E(e,b),d'), and in the second case B(e,E(e,b),e',E(e',1),d')

Lemma 1 Proof

- Let P'_{00} and P'_{11} denote the probability of outputting 1 when R is chosen according to U rather than S, and p is chosen from P(R,.).
- Since we assume (F,T) has LCC advantage ε, $|P'_{00} - P'_{11}| > \varepsilon$.
- By assumption on S, $|P_{00} – P'_{00}|$ and $|P_{11}-P'_{11}|$ must be negligible. (Otherwise we have a distinguisher for S)
- So we must have $|P_{00}-P_{11}| > \varepsilon/2$.

Conclusion

- If trapdoor OWPs exist, then LCC-secure PKC exists.