Message Authentication

15-859I
Spring 2003

Motivation

- Suppose Alice is an ATM and Bob is a Bank, and Alice sends Bob messages about transactions over a public channel.
- Bob would like to know that when he receives a message saying “credit $128 to Carol’s Account – Alice”, it originates from the ATM. Bob is concerned with the *authenticity* of the message.
- He also wants to know that Carol has not modified the message from “credit $16 to Carol’s Account – Alice.” This concerns the *integrity* of the message.

Authentication and Encryption

- Should we expect to get good message authentication via encryption? i.e., is it enough to guarantee authenticity of M by transmitting $E_K(M)$?
- No! e.g., if E_K is CTR from lecture 6, then it is easy for Carol to change $E_K(16)$ to $E_K(128)$ via $E_K(128) = E_K(16) \oplus 144$.
- In general, good encryption does not necessarily imply integrity.

Message Authentication Codes (MACs)

- Formally: A MAC is a trio of algorithms (G,T,V) such that:
- $G(1^k)$ generates a k-bit key K.
- $T_K(M)$ generates a L(k)-bit tag σ
- $V_K(M, \sigma)$ verifies the tag σ for the message M.
- Require that for all K, M, random choices or states of T, $V_K(M,T_K(M)) = 1$.

If T_K is deterministic and stateless, V_K is trivial.

Security of MACs

- The adversary’s goal might be to sign some specific message m.
- We want it to be hard to produce *any* (M,σ) pair such that $V_K(M,\sigma)=1$.
- This should be true even if the adversary has seen several $M', T_K(M')$ pairs
- Should be conservative: Allow the adversary to choose the M’ messsages.

Existential Unforgeability

- Notation: let $Q(A^O)$ denote the list of oracle queries that A makes with O as its oracle.
- Define the chosen-message attack (cma) advantage of A against MAC = (G,T,V) by:

$$\mathrm{Adv}^{\mathrm{cma}}_{A,MAC}(k) = \Pr[V_K(M,\sigma)=1 \wedge M \notin L :$$
$$K \leftarrow G(1^k), (M,\sigma) \leftarrow A^{T_K}(1^k), L = Q(A^{T_K}(1^k))]$$

- Say that MAC is existentially unforgeable under chosen message attack if every ppt A has negligible advantage.

MAC Insecurity

- For a fixed security parameter k, define the Insecurity of MAC=(G,T,V) against time-t adversaries which make q queries with total message length l by:

$$\mathrm{InSec}_{MAC}^{\text{uf-cma}}(t,q,l) = \max_{A:A(t,q,l)} \left\{ \mathrm{Adv}_{A,MAC}^{\text{cma}}(k) \right\}$$

PRFs are good MACs

- Let $F : K \times \{0,1\}^d \rightarrow \{0,1\}^s$ be a function family. Then if F_K is pseudorandom, F_K is a good MAC for the message space $\{0,1\}^d$:

$$\mathrm{InSec}_F^{\text{uf-cma}}(t,q,dq) \le \mathrm{InSec}_F^{\text{prf}}(t',q) + 2^{-s}$$

- Proof: Let A be a chosen-message forger for F as a MAC. We show how to construct a PRF distinguisher D for F that has almost the same advantage as A, and runs in the same time.

PRFs are good MACs

$D^f(1^k)$:

Run $A(1^k)$: respond to q with f(q), add q to Q

Set (M,σ) = output of A.

If $f(M) = \sigma$ and $M \notin Q$, return 1, else return 0.

Notice: $\Pr[D^F(1^k) = 1] = \mathrm{Adv}_{A,F}^{cma}(k)$. And since M was never queried, $\Pr[D^{\mathcal{F}}(1^k) = 1] \le 1/2^s$.

So $\mathrm{Adv}_{A,F}^{prf}(k) \ge \mathrm{Adv}_{A,F}^{cma}(k) - 2^{-s}$

Almost-XOR-Universal$_2$ (AXU$_2$) Hash Functions

- Let $H: K \times D \rightarrow \{0,1\}^L$ be a family of functions. Define the XOR-2-Universality of H by

$$\mathrm{Adv}^{\text{xuh}}(H) = \max_{a_1,a_2 \in D, b \in \{0,1\}^L} \{\Pr_K[H_K(a_1) \oplus H_K(a_2) = b]\}$$

- We say H is ε-almost-XOR-Universal$_2$ (ε-AXU$_2$) if $\mathrm{Adv}^{xuh}(H) \le \varepsilon$.
- H is XOR-Universal$_2$ if it is $1/2^L$-AXU$_2$.
- Notice that Pairwise-independent hash functions are XOR-Universal$_2$.

ε-AXU$_2$ Hash Families for large domains

- Let $h: K \times \{0,1\}^{2L} \rightarrow \{0,1\}^L$ be ε AXU$_2$. Define the family $H : K^n \times \{0,1\}^{2^n L} \rightarrow \{0,1\}^L$ as follows:
- $H_{K1\ldots Kn}(a_1,\ldots,a_{2n}) = H_{K2\ldots Kn}(h_{K1}(a_1,a_2),h_{K1}(a_3,a_4),\ldots,h_{K1}(a_{2n-1},a_{2n}))$; $H_K(a_1,a_2) = h_K(a_1,a_2)$
- Claim: H is (nε)-AXU$_2$.
- Proof: If the inputs to h_{Kn} are not the same, then the xor-probability is at most ε. The probability that the inputs to h_{Kn} are the same is at most ε given that not all inputs to $h_{K(n-1)}$ are the same, and so on.

AXU$_2$ - MAC

- Let $F: \mathcal{K} \times \{0,1\}^l \rightarrow \{0,1\}^L$ be a PRF. Let $H : K \times D \rightarrow \{0,1\}^L$ be ε-AXU$_2$. Define the MAC $\mathcal{C\text{-}UHM}$ as follows:
- G: Select $K \leftarrow \mathcal{K}$, $\kappa \leftarrow K$. Return (K, κ)
- $T_{(K,\kappa)}(M)$ =
  - Let $x = H_\kappa(M)$; $\tau = F_K(ctr) \oplus x$; $\sigma=(ctr,\tau)$
  - Set ctr = ctr+1
  - Return σ
- $V_{(K,\kappa)}(M,(s,\tau)) = 1$ iff $F_K(s) \oplus \tau = H_\kappa(M)$.

C-UHM theorem

- Theorem: for any $q \le 2^l$,
  $InSec_{C\text{-}UHM}^{uf\text{-}cma}(t,q,l) \le \varepsilon + InSec_F^{prf}(t',q+1)$
- Proof: Let A be any MAC adversary. Suppose we choose $f \leftarrow \mathcal{F}$ and run A against UHM instantiated with f in place of F_K. Denote the queries that A makes by M_i, $1 \le i \le q$; denote the responses by $\sigma_i = (i,\tau_i)$
- Finally, A returns some message $M \notin \{M_1,\ldots,M_q\}$, and a tag (s,τ)

C-UHM Theorem, continued.

Let NEW be the event that s > q-1, that is, the s returned by A was not a value input to f in C-UHM. Let OLD be the event s < q.

Claim 1: $Pr[V_\kappa^f(M,(s,\tau)) = 1 \mid OLD] \le \varepsilon$

Proof:

$Pr[V_\kappa^f(M,(s,\tau)) = 1] =$

$Pr[H_\kappa(M) \oplus H_\kappa(M_s) = \tau \oplus \tau_s] \le \varepsilon.$

Claim 2: $Pr[V_\kappa^f(M,(s,\tau)) = 1 \mid NEW] \le 2^{-L}.$

Proof: $= Pr[H_\kappa(M) \oplus \tau = f(s)] = 2^{-L}.$

C-UHM Theorem, continued.

Thus:

$Pr[V_\kappa^f(M,(s,\tau)) = 1] =$

$Pr[V_\kappa^f(M,(s,\tau)) = 1 \mid OLD]\, Pr[OLD] +$

$Pr[V_\kappa^f(M,(s,\tau)) = 1 \mid NEW](1-Pr[OLD])$

$\le \varepsilon q + 2^{-L}(1-q)$

$\le \varepsilon q + \varepsilon(1-q) = \varepsilon.$

The theorem follows, since we can distinguish F_K from f by trying to use A to forge a MAC and then checking if A was successful.

R-UHM

- Let $F: \mathcal{K} \times \{0,1\}^l \to \{0,1\}^L$ be a PRF. Let $H : K \times D \to \{0,1\}^L$ be ε-AXU_2. Define the MAC $\mathcal{R\text{-}UHM}$ as follows:
- G: Select $K \leftarrow \mathcal{K}$, $\kappa \leftarrow K$. Return (K, κ)
- $T_{(K,\kappa)}(M) =$
  - Choose $s \leftarrow \{0,1\}^l$.
  - Let $x = H_\kappa(M)$; $\tau = F_K(s) \oplus x$
  - Return $\sigma=(s,\tau)$
- $V_{(K,\kappa)}(M,(s,\tau)) = 1$ iff $F_K(s) \oplus \tau = H_\kappa(M)$.

R-UHM theorem

- Theorem: for any $q \le 2^{l/2}$,
  $InSec_{R\text{-}UHM}^{uf\text{-}cma}(t,q,l) \le \varepsilon + InSec_F^{prf}(t',q+1) + q(q-1)/2^{l+1}$
- Proof: Consider the same experiment as before. Clearly when there are no collisions in the values $s_1,\ldots,s_q$, the same argument upper bounds the success probability of A. And when there is a collision, the success probability of A is at most 1. The probability of a collision is $q(q-1)/2^{l+1}$.

CBC-MAC

- Let $F: \{0,1\}^k \times \{0,1\}^l \to \{0,1\}^l$ be a PRF. Define the MAC $F^{(m)}$ on ml-bit messages as follows:
- $T_K(x_1,\ldots,x_m) =$
  - Let $y_0 = 0^l$
  - For i = 1,…,m
    $y_i = F_K(M_i \oplus y_{i-1})$
  - return y_m
- Theorem:
  $Insec_{F(m)}^{prf}(t,q) \le Insec_F^{prf}(t+O(qml),qm) + 3q^2m^2/2^{l+1}.$
- So $F^{(m)}$ is a secure MAC if F is a secure PRF.

CBC-MAC Proof

- Lemma 1: If $f \leftarrow \mathcal{F}_{l,l}$ then
  $\mathrm{Insec}_{f^{(m)}}^{prf}(q) \le 3m^2q^2/2^{l+1}$
- Consider the 2^l-ary tree of depth m. A sequence of strings $X = (x_1,\ldots,x_n) \in \{0,1\}^{nl}$ uniquely specifies a node in this tree.
- Let f: $\{0,1\}^l \rightarrow \{0,1\}^l$, denote the labeling of a sequence $x_1\ldots x_n$ by $Z_f() = 0^l$, $Y_f(x_1,\ldots,x_n) = x_n \oplus Z_f(x_1,\ldots,x_{n-1})$, $Z_f(X)=f(Y_f(X))$
- Call $(X_1,\ldots,X_n)$ a query sequence if every X_i has parent either the root or X_j for some j < i.

CBC-MAC proof

- Consider an (unbounded) adversary A trying to distinguish $f^{(m)}$ from a sample from $\mathcal{F}_{ml,l}$ with q queries. We let A make qm queries $X_1\ldots X_{qm}$ in the form of a query tree, and whenever X_i is at depth m, A learns $Z_f(X_i)$.
- We let $\mathcal{Z}_n$ be the depth-m labels A has learned after the n^{th} query, and $V_n = (X_1,\ldots,X_n; \mathcal{Z}_n)$ denotes the View of A after the n^{th} query.
- If the labeling Z_f is collision-free, then A's view is identical to its view on a random function from ml bits to l bits. So we only need to bound the probability of a collision in Z_f.

CBC-MAC Proof

- Lemma 2: Let Z_n^1 and Z_n^2 be collision-free output labelings consistent with a depth-m labeling $\mathcal{Z}_n$. Then:

$$\Pr[Z_n = Z_n^1 \mid V_n = (X_1,\ldots,X_n; \mathcal{Z}_n)] = \Pr[Z_n = Z_n^2 \mid V_n = (X_1,\ldots,X_n; \mathcal{Z}_n)].$$

- Proof: By induction. Obviously true for n=1, since the first node in a query tree is not at depth m.

Lemma 2 proof, con't

- Two cases for n>1:
  - X_n at depth < m. Then the view is independent of $Z(X_n)$. Thus $\Pr[Z_n = Z_n^i \mid V_n]$

$= \Pr[Z_n=Z_n^i|V_{n-1}]$

$= \Pr[Z_{n-1}=Z_{n-1}^i|V_{n-1}]\Pr[Z_n(X_n) = Z_n^i(X_n)|Z_{n-1}=Z_{n-1}^i, V_{n-1}]$

$= \Pr[Z_{n-1}=Z_{n-1}^i|V_{n-1}]2^{-l}$

These are equal for i=1,2 by IH.

  - X at depth m. Then $\Pr[Z_n = Z_n^i|V_n]$

$= \Pr[Z_{n-1} = Z_{n-1}^i \mid V_{n-1}, Z_n(X_n) = z]$

$= \Pr[Z_n(X_n) = z|Z_{n-1}=Z_{n-1}^i,V_{n-1}]\Pr[Z_{n-1}=Z_{n-1}^i|V_{n-1}]/\Pr[Z_n(X_n) = z]$

$= 2^{-l}\Pr[Z_{n-1}=Z_{n-1}^i|V_{n-1}]/\Pr[Z_n(X_n) = z]$

Lemma 3

- Lemma 3: Let CF(Z) denote the event that Z is collision-free. Let $\Pr_n[E]$ denote the quantity $\Pr[E \mid V_n=(X_1,\ldots,X_n;\mathcal{Z}_n), CF(Z_n)]$. Let $n^2/4 + n - 1 \le 2^l/2$ Let $(x_1\ldots x_i) \in \{X_1,\ldots,X_m\}$ and i<m; let z_S be a collision-free label of the nodes in $S = \{X_1,\ldots,X_n\} \setminus \{(x_1\ldots x_i)\}$ consistent with $\mathcal{Z}_n$. Then
  (1) For any $(x_1\ldots x_i x_{i+1}) \in S$, any $y^* \in \{0,1\}^l$:
  $\Pr_n[Y_n(x_1\ldots x_i x_{i+1}) = y^* \mid Z_n^S = z_S] \le 2\,2^{-l}$
  (2) For any $z^* \in \{0,1\}^l$, $\Pr_n[Z_n(x_1\ldots x_i)=z^*|Z_n^S=z_S] \le 2\,2^{-l}$.

Proof of Lemma 3(1)

Let $y \in \{0,1\}^l$ be some fixed string. Define the labeling $Z_{z,y}(X_j) = z_S(X_j)$ if $X_j \ne x_1\ldots x_i$, and $y \oplus x_{i+1}$ otherwise. Let $Y_{z,y}$ be the labeling induced by $Z_{z,y}$:

$Y_{z,y}(X_j) = y_S(X_j)$ if $X_j \notin \mathrm{children}(x_1\ldots x_i)$
$y \oplus x_{i+1} \oplus x'_{i+1}$ if $X_j = x_1\ldots x_i x'_{i+1}$.

Let $\mathcal{Y}(z_S)$ be the set of all strings y such that $Z_{z,y}$ is collision-free. $y \notin \mathcal{Y}(z_S)$ iff either:

- $y \oplus x_{i+1} \in \{z_S(X_j) : 0 < j < n+1$ and $X_j \ne (x_1\ldots x_i)\}$; or
- for some x'_{i+1}, $y \oplus x_{i+1} \oplus x'_{i+1} \in \{y_S(X_j), X_j \notin \mathrm{children}(x_1\ldots x_i)$, and $0 < j < n+1\}$

Thus $|\{0,1\}^l \setminus \mathcal{Y}(z_S)| \le (n-1) + (n-s)(s) \le n-1 + n^2/4 \le 2^l/2$.

This proves (1).

Proof of Lemma 3(2)

Let $z \in \{0,1\}^l$ be some fixed string. Define the labeling $Z_{z,y}(X_j) = z_S(X_j)$ if $X_j \neq x_1 \ldots x_i$, and z otherwise. Let Y_z be the labeling induced by $Z_{,z}$:

$Y_z(X_j) =$ $y_S(X_j)$ if $X_j \notin \text{children}(x_1 \ldots x_i)$
$z \oplus x'_{i+1}$ if $X_j = x_1 \ldots x_i x'_{i+1}$.

Let $\mathcal{Z}(z_S)$ be the set of all strings z such that Z_z is collision-free. $z \notin \mathcal{Z}(z_S)$ iff either:

- $z \in \{z_S(X_j) : 0 < j < n+1 \text{ and } X_j \neq (x_1 \ldots x_i)\}$; or
- for some x'_{i+1}, $z \oplus x'_{i+1} \in \{y_S(X_j), X_j \notin \text{children}(x_1 \ldots x_i), \text{ and } 0 < j < n+1\}$

Thus $|\{0,1\}^l \setminus \mathcal{Y}(z_S)| \leq (n-1) + (n-s)(s) \leq n-1 + n^2/4 \leq 2^l/2$.

This proves (2).

Lemma 4: Pr[not CF(Z)]

- Let $n^2/4 + n-1 < 2^l/2$. Let $X_1 \ldots X_n$ be a query sequence and $\mathcal{Z}$ be the labeling of depth-m nodes. Then
$$\Pr[\text{not } CF(Z_{n+1}) \mid V_n, CF(Z_n)] \leq 3n\, 2^{-l}.$$
- Proof: Denote $\Pr[E | V_n, CF(Z_n)]$ by $\Pr_n[E]$.
- Case 1: X_{n+1} is at depth 1. Then let $X_{n+1} = x_1^*$. $Y(X_{n+1}) = x_1^*$ by definition. Now for each $1 \leq t \leq n$,
$$\Pr_n[Y_n(X_t) = x_1^*] \leq 2\, 2^{-1}.$$
- This is because if X_t is at level 1, $\Pr[Y(X_t) = x_1^*] = 0$. Otherwise X_t is at depth at least 2, and is the child of some $(x_1 \ldots x_i) \in \{X_1 \ldots X_n\}$ and so the equation follows because of lemma 3.
- Then $\Pr_n[\text{not } CF(Z_n)] \leq \Pr_n[x_1^* \in \{Y_n(X_1), \ldots, Y_n(X_n)\}] + \Pr_n[Z_{n+1}(X_{n+1}) \in \{Z_n(X_1) \ldots Z_n(X_n)\} \mid x_1^* \notin \{Y_n(X_1), \ldots, Y_n(X_n)\}] \leq 2n/2^l + n/2^l = 3n/2^l$.

Lemma 4: Case 2

- Case 2: $X_{n+1} = x_1 \ldots x_i x_{i+1}$, $i > 0$, is the child of some $x_1 \ldots x_i \in \{X_1, \ldots, X_n\}$. Let $S = \{X_1, \ldots, X_n\} \setminus \{x_1 \ldots x_i\}$. Notice that for any $X_t \in \{X_1, \ldots, X_n\}$,
$$\Pr_n[Y_{n+1}(X_{n+1}) = Y_n(X_t)] \leq 2/2^l.$$
Since if X_{n+1} and X_t are siblings, the probability is 0, and otherwise any collision free labeling z_S determines $Y_n(X_t)$; thus
$\Pr_n[Y_{n+1}(X_{n+1}) = Y_n(X_t)]$
$= \Sigma_z \Pr_n[Y_{n+1}(X_{n+1}) = Y_n(X_t) \mid Z_n^S = z_S] \Pr[Z_n^S = z_S]$
$= \Sigma_z \Pr_n[Z_n(x_1 \ldots x_i) = Y_n(X_t) \oplus x_{i+1} | Z_n^S = z_S] \Pr[Z_n^S = z_S]$
$\leq 2/2^l \, \Sigma_z \Pr[Z_n^S = z_S] \leq 2/2^l$.
This gives us that
$\Pr_n[\text{not } CF(Z_{n+1})] \leq \Pr_n[Y_{n+1}(X_{n+1}) \in \{Y_n(X_1), \ldots, Y_n(X_n)\}] +$
$\Pr_n[Z_{n+1}(X_{n+1}) \in \{Z_n(X_1) \ldots Z_n(X_n)\} | Y_{n+1}(X_{n+1}) \notin \{Y_n(X_1), \ldots, Y_n(X_n)\}]$
$\leq 2n/2^l + n/2^l = 3n/2^l$.

Pr[CF(Z)]

- So
$\Pr[\text{not } CF(Z)] \leq \Sigma_n \Pr[\text{not } CF(Z_n) \mid CF(Z_{n-1})]$
$\leq 3/2^l \, (qm)(qm-1)/2$
$= 3/2 \, q^2 m^2 / 2^l$.