COL876: SPECIAL TOPICS IN FORMAL METHODS

Formal verification of security protocols

Lecture 11, 19 October 2023

RECAP: COMPUTATIONAL SOUNDNESS

- Want to map symbolic terms to distributions over strings
- Map symbolic attacks to non-negligible adversary advantage
- Need to keep track of adversary "view"
- "What can an adversary learn from an encrypted term?" "Patterns"
- Equivalence of patterns == Indistinguishability of ciphertexts

PATTERNS FROM TERMS

- $P, Q := i \mid k \mid (P, Q) \mid \{P\}_k \mid \square$ where $i \in \{0, 1\}$ and $k \in$ Keys
- Given a set of keys T and a term M, pat(M, T) gives the pattern that an attacker can see in M if he has access to T
  - Inductive definition; two cases for encryption
- $\mathrm{pat}(M) = \mathrm{pat}(M, \{ k \in \mathrm{Keys} \mid M \vdash k \})$; $M \equiv N$ iff $\mathrm{pat}(M) = \mathrm{pat}(N)$
- $M \cong N$ iff $M \equiv N\sigma$ for some bijection σ on Keys

PATTERNS FROM TERMS: EXAMPLES

- $0 \cong 0$ and $0 \ncong 1$ and $\{0\}_k \cong \{1\}_k$ and $\{0\}_k \cong \{1\}_{k'}$
- $(k, \{0\}_k) \ncong (k, \{1\}_k)$, but $(k, \{ (\{0\}_{k'}, 0) \}_k) \cong (k, \{ (\{1\}_{k'}, 0) \}_k)$.
- $(\{0\}_k, \{0\}_k) \cong (\{0\}_k, \{1\}_k)$ Cannot identify identical plaintexts
- $(\{0\}_k, \{1\}_k) \cong (\{0\}_k, \{1\}_{k'})$ Cannot identify whether same key is used
- $\{((1, 0), (0, 1))\}_k \cong \{0\}_k$ Length of plaintext is not revealed

INITIAL ASSOCIATIONS

- Given an encryption scheme $\Pi = (K, E, D)$, associate to a term M a distribution on strings $M(\Pi, \eta)$; lift to collection $M(\Pi)$
- Define an algorithm Conv which works over terms as follows:
  - Map each key k occurring in M to a string of bits $\tau(k)$ using $K(\eta)$
  - Map constants 0 and 1 in the term algebra to their bitstrings
  - Lift easily to pairs; for $M = senc(M', k)$, map it to $E(M'(\Pi, \eta), \tau(k))$
  - Tag every bitstring with its type: "key", "bool", "pair", "ciphertext"

RECAP: ENCRYPTION SCHEMES

- An encryption scheme Π, is a triple of PTIME algorithms (K, E, D) parametrized by η
  - K is the key generation algorithm
    - input: parameter, coins output: key
  - E is the encryption algorithm
    - input: key, string, coins output: ciphertext
  - D is the decryption algorithm
    - input: key, string output: plaintext
  - D(k, E(k, m, r)) = m if m is a valid plaintext, ∅ otherwise

RECAP: NEGLIGIBLE ADVANTAGE

- Probabilistic PTIME adversary A
- A function f: N $\rightarrow$ R is negligible if, for all c > 0, there exists an N_c such that $f(\eta) \leq \eta^{-c}$ for all $\eta \geq N_c$.
- $adv(\eta) := Pr[\, x \leftarrow D \mid A(\eta, x) = 1\,] - Pr[\, x \leftarrow D' \mid A(\eta, x) = 1\,]$
- We say D and D' are indistinguishable ($D \approx D'$) if for every probabilistic PTIME adversary A, $adv(\eta)$ is negligible

EQUIVALENCE IMPLIES INDISTINGUISHABILITY

- $M \cong N$ implies $M(\Pi) \approx N(\Pi)$
- $0 \cong 0$, so $0(\Pi) \approx 0(\Pi)$. Both ensembles put all the probability mass on <0, "bool">
- $\{0\}_k \cong \{1\}_k$, so $\{0\}_k(\Pi) \approx \{1\}_k(\Pi)$
  - Non-trivial; depends heavily on our assumptions about type-0 security of the encryption scheme

EQUIVALENCE IMPLIES INDISTINGUISHABILITY

- Let M and N be terms* and Π an encryption scheme*. If $M \cong$ N, then M(Π) ≈ N(Π).
- Overall steps:
  - Assume M and N are pattern equivalent.
  - Rename keys
  - "Hybrid patterns" M_i and N_i to form a chain between the renamed versions of M and N to maintain pattern equivalence
  - Define ensembles for each M_i and N_i, final ensembles M'(Π) and N'(Π)
  - Want to show that any adversary advantage between M'(Π) and N'(Π) is negligible
  - Assume not; Contradict the type-0 security of Π

KEY RENAMING

- Want to modify M and N so that keys encrypt other keys in a systematic manner
- Rename so that:
  - M and N have l recoverable keys $j_1, j_2, \ldots, j_l$
  - M and N have some hidden keys
    - M has m hidden keys $k_1, k_2, \ldots, k_m$
    - N has n hidden keys $k_1, k_2, \ldots, k_n$
    - k_p encrypts k_q only when $p \geq q$
- Can do this because terms do not have key cycles; a "deeper" key gets a smaller index
- Get terms M' and N' after this renaming

HYBRID PATTERNS

- $M_0, M_1, \ldots, M_m$ and $N_0, N_1, \ldots, N_n$ to form chain from M' to N'
  - $M_i = \text{pat}(M', \text{recoverable}(M') \cup \{k_1, k_2, \ldots, k_i\})$
  - $N_j = \text{pat}(N', \text{recoverable}(N') \cup \{k_1, k_2, \ldots, k_j\})$
  - $M_0 = \text{pat}(M')$ and $M_m = M'$ and $N_0 = \text{pat}(N')$ and $N_n = N'$
- M_i and N_i are the patterns the adversary could see in M' and N' if they had access to (hitherto hidden) keys k_1 through k_i
- Acyclicity: these keys do not give access to other keys k_j where $j > i$

DEFINING ENSEMBLES

- We map each M_0, M_1, … , M_m and N_0, N_1, …, N_n to an ensemble
- Lift the Conv algorithm to work over patterns, not just terms
  - Generate a new fixed key $\tau(k_0)$ using $K(\eta)$
  - Map $\Box$ to $E(\mathbf{0}, \tau(k_0))$, tag with "ciphertext"
  - $\tau(k_0)$ is only for use with $\Box$

ADVERSARY ADVANTAGE

- We know that M(Π) ≈ M'(Π) and N(Π) ≈ N'(Π) (only keys have been renamed). Want to show that M'(Π) ≈ N'(Π)
- Assume there is an adversary A who can distinguish between M'(Π) and N'(Π) with non-negligible advantage
- $\lambda(\eta) = \Pr[\, y \leftarrow M'(\Pi) \mid A(\eta, y) = 1 \,] - \Pr[\, y \leftarrow N'(\Pi) \mid A(\eta, y) = 1 \,]$
- For some constant c and infinite set S, $\lambda(\eta) > \eta^{-c}$ for all $\eta \in S$.

ADVERSARY ADVANTAGE

- We define the following for $0 \leq i \leq m$ and $1 \leq j \leq n$:
  - $p_i(\eta) = \Pr[\, y \leftarrow M_i(\Pi, \eta) \mid A(\eta, y) = 1 \,]$
  - $q_j(\eta) = \Pr[\, y \leftarrow N_j(\Pi, \eta) \mid A(\eta, y) = 1 \,]$
- Since $M' = M_m$ and $N' = N_n$, $\lambda = p_m - q_n$. Also $p_0 = q_0$. So,

$$\lambda = (p_m - p_{m-1}) + (p_{m-1} - p_{m-2}) + \ldots + (p_1 - p_0) + (q_0 - q_1) + (q_1 - q_2) + \ldots + (q_{n-1} - q_n)$$

- Have m+n quantities that add up to λ. Triangle inequality: There is either
  - $1 \leq i \leq m$ s.t. $p_i - p_{i-1} \geq \lambda/(m + n)$, or
  - $1 \leq j \leq n$ s.t. $q_j - q_{j-1} \geq \lambda/(m + n)$
- A suitable i or j exists for each such η, and since we have finite, fixed summands, there is some i or j that works for infinitely many η.

ADVERSARY ADVANTAGE

- Let i be such an index. There exists an infinite set $S' \subseteq S$ s.t. $p_i(\eta) - p_{i-1}(\eta) \geq \lambda(\eta)/(m + n)$ for each $\eta \in S'$.

- A suitable i or j exists for each such η, and since we have finite, fixed summands, there is some i or j that works for infinitely many η.

- Using this adversary A, we want to construct a computational adversary A_0 who violates the type-0 security of Π.

ADVERSARY A_o

- A_o generates $\tau(k)$ using $K(\eta)$ for every k in M'
- It then runs an algorithm called Conv2 on M' (coming up) and obtains a result y
- It then calls A using the parameter η and y, and returns the result.
- A_o (and Conv2) has access to two oracles f and g, instantiated either as
  - $f = \mathcal{E}_{K_i}(.)$ for $K_i \leftarrow K(\eta)$; $g = \mathcal{E}_{K_o}(.)$ for $K_o \leftarrow K(\eta)$, or
  - $f = \mathcal{E}_{K_o}(.)$ for $K_o \leftarrow K(\eta)$; $g = \mathcal{E}_{K_o}(.)$ for $K_o \leftarrow K(\eta)$

ALGORITHM CONV2

- Conv2 same as Conv except for encryptions; everything tagged as earlier
- For encryptions of the form $\{M^*\}_k$
  - If $k \in \{j_1, \ldots, j_l, k_1, \ldots, k_{i-1}\}$, map to E(Conv2($M^*$), k)
  - If $k = k_i$, map to f(Conv2(M^*))
  - If k in $\{k_{i+1}, \ldots, k_m\}$, map to g($\emptyset$)
- Encryption under a recoverable key k corresponds to encryption under the associated key $\tau(k)$.
- Encryption under a hidden key from $\{k_1, \ldots, k_{i-1}\}$ also corresponds to encryption under the associated key $\tau(k)$.
- Encryption under a hidden key in $\{k_{i+1}, \ldots, k_m\}$ results in $\emptyset$ encrypted under K_0.

CONTRADICTING TYPE-0 SECURITY

- We have
  - $p_i(\eta) = \Pr[\, K_i, K_0 \leftarrow K(\eta) \mid A_0^{\mathcal{E}_{K_i}(.),\, \mathcal{E}_{K_0}(.)}(\eta) = 1 \,]$
  - $p_{i-1}(\eta) = \Pr[\, K_0 \leftarrow K(\eta) \mid A_0^{\mathcal{E}_{K_0}(.),\, \mathcal{E}_{K_0}(.)}(\eta) = 1 \,]$
- Conv2(M’) returns a sample from
  - $M_i(\Pi)$ when $f = \mathcal{E}_{K_i}(.)$ and $g = \mathcal{E}_{K_0}(.)$, and
  - $M_{i-1}(\Pi)$ when $f = \mathcal{E}_{K_0}(\mathtt{0})$ and $g = \mathcal{E}_{K_0}(\mathtt{0})$
- For p_i, encryption under the hidden key k_i corresponds to encryption under K_i
- For p_{i-1}, encryption under k_i results in $\mathtt{0}$ encrypted under K_0.

CONTRADICTING TYPE-0 SECURITY

- Therefore, for infinitely many values of η, we get

$$\mathrm{adv}(\eta) \text{ for } A_0 \text{ is } p_i(\eta) - p_{i-1}(\eta)$$

$$\geq \lambda(\eta)/(m + n)$$

$$> \eta^{-c}/(m+n)$$

$$> \eta^{-(c+1)}$$