Advanced Cryptography CS 655

Week 3:

- Memory-Tight Reductions
- RSA-FDH
- Memory-Tightness

Memory-Tight Reductions

Benedikt Auerbach¹ David Cash² Manuel Fersch¹ Eike Kiltz¹

¹Horst Görtz Institute for IT Security, Ruhr University Bochum, Germany

² Rutgers University, New Jersey, USA

Crypto 2017, Santa Barbara August 21, 2017



Tightness and Memory-Tightness

Cryptographic Reductions and Tightness



Reduction is *tight* if **Time**(A_P) \approx **Time**(A_S) and **Succ**(A_P) \approx **Succ**(A_S)

Benedikt Auerbach: Memory-Tight Reductions

4/23



Time-success trade-off plot for algorithms solving problem **P**



Time-success trade-off plot for algorithms solving problem **P**





5/23

What about memory?

Resources of adversary

- Running time
- Success probability
- Memory consumption

Benedikt Auerbach: Memory-Tight Reductions

Contributions

- Raise awareness for memory usage in reductions
- Propose tools for achieving memory-tightness
- Concrete application: memory-tight reduction for RSA-FDH
- Impossibility of memory-tight reductions for certain problems

7/23

Time-Memory Trade-offs

Time-Memory Trade-offs

- Some problems are harder with less memory
 - in particular many lattice / coding problems
- Needs to be taken into account in reductions
- Concrete example: Learning Parity with Noise (LPN)

Benedikt Auerbach: Memory-Tight Reductions

Adding Memory-Consumption



Adding Memory-Consumption



Adding Memory-Consumption



For simplicity: consider algorithms with constant success probability





Algorithm 2: BKW algorithm for dimension λ Mem $\approx 2^{\lambda/\log \lambda}$, Time $\approx 2^{\lambda/\log \lambda}$





Time-memory trade-off plot for LPN [EKM, 11:00 Lotte Lehman Hall]

Memory-Tight Reductions



Memory-Tight Reductions



(Time-)Tight but not memory-tight reduction

Memory-Tight Reductions



Memory-Sensitive Problems



Memory-sensitive:

- LPN
- Shortest Vector Problem
- 3 collision resistance
- DLP in finite fields
- Factoring

Might change in future



Not memory-sensitive:

Mem

- Collision resistance
- Preimage resistance
- DLP over elliptic curves

Benedikt Auerbach: Memory-Tight Reductions

Typical Non-Memory-Tight Reductions

Example 1: Random Oracle Simulation



Worst case: $Mem(A_P) \approx Time(A_S)$, while $Mem(A_S)$ small

Example 2: Unforgeability of Signatures



Worst case: $Mem(A_P) \approx Time(A_S)$, while $Mem(A_S)$ small

Recap

- Currently memory often ignored in reductions
- Many existing reductions not memory-tight
 - Worst case $Mem(A_P) \approx Time(A_S)$ while $Mem(A_S)$ small
- Particularly problematic for memory-sensitive problems

Achieving Memory-Tightness

Example 1: Random Oracle Simulation



RO simulation via lazy sampling



Memory efficient RO simulation [Bernstein 2011]

Example 2: Unforgeability of Signatures



Usual simulation of unforgeability game



Memory-efficient simulation of unforgeability game



Memory-efficient simulation of unforgeability game



Important: Coins of A_P and A_S have to be stored memory-efficiently



Unforgeability and Multi-Unforgeability


Unforgeability and Multi-Unforgeability





Tight but not memory-tight reduction (store queries)

g h

R



Memory-tight reduction with lower success probability (guess forgery)

g h

R

Unforgeability and Multi-Unforgeability



Memory-tight reduction with higher running time (rewind adversary)

Unforgeability and Multi-Unforgeability

Reductions from mForge to Forge (for *q* adversarial queries)

	Time	Succ	Mem
store queries	1	1	q
guess forgery	1	1/q	1
rewind adversary	q	1	1

Lower Bounds

Theorem

A certain class of black box reductions from mForge to Forge can not be simultaneously tight and memory-tight.

- Proof uses techniques from streaming algorithms
- Similar results for multi-collision to collision resistance

Lower Bounds

Theorem

A certain class of black box reductions from mForge to Forge can not be simultaneously tight and memory-tight.



Conclusions and Future Work

- Memory usage is ignored but affects security.
- Many reductions are easily fixed...
- ... but some seem inherently loose, including some widely used implicitly.

Future work:

- Give memory-tight reductions for some constructions (e.g. Hashed ElGamal).
- Prove lower bounds in less restrictive models.

ia.cr/2017/675

Signature Experiment (Sig – $forge_{A,\Pi}(n)$)



Signature Experiment (Sig – forge, – (n))

Formally, let $\Pi = (Gen, Sign, Vrfy)$ denote the signature scheme, call the experiment Sig – forge_{A, Π}(n)

We say that Π is existentially unforgeable under an adaptive chosen message attack (or just secure) if for all PPT adversaries A, there is a negligible function μ such that $\Pr[\text{Sig} - \text{forge}_{A,\Pi}(n) = 1] \le \mu(n)$

Existential Unforgeability

- Limitation: Does not prevent replay attacks
 - $\sigma \leftarrow \operatorname{Sign}_{sk}("Pay Bob \$50", R)$
 - If this is a problem then you can include timestamp in signature
- Unforgeability: does rule out the possibility attacker modifies a signature
- Plain RSA signatures are malleable (does not satisfy our security notion)
- **Remark:** By design signatures cannot hide all information about message *m*
 - Public Verification \rightarrow Attacker can easily distinguish between a signature for m₁ and m₂

Plain RSA Signatures

- Plain RSA
- Public Key (pk): N = pq, e such that $GCD(e, \phi(N)) = 1$
 - $\phi(N) = (p-1)(q-1)$ for distinct primes p and q
- Secret Key (sk): N, d such that ed=1 mod $\phi(N)$

$$\operatorname{Sign}_{sk}(m) = m^d \mod N$$
$$\operatorname{Vrfy}_{pk}(m, \sigma) = \begin{cases} 1 & if \ m = [\sigma^e \mod N] \\ 0 & otherwise \end{cases}$$

• Verification Works because $\left[\operatorname{Sign}_{sk}(m)^{e} \mod N\right] = \left[m^{ed} \mod N\right] = \left[m^{\left[ed \mod \phi(N)\right]} \mod N\right] = m$

No Message Attack

- **Goal:** Generate a forgery using only the public key
 - No intercepted signatures required
- Public Key (pk): N = pq, e such that $GCD(e, \phi(N)) = 1$ • $\phi(N) = (p-1)(q-1)$ for distinct primes p and q
- Pick random $\sigma \in \mathbb{Z}_{_{N}}^{*}$
- Set $m = [\sigma^e \mod N]$.
- Output (m, σ)

$$\operatorname{Vrfy}_{pk}(m,\sigma) = \begin{cases} 1 & if \ m = [\sigma^e \ mod \ N] \\ 0 & otherwise \end{cases}$$

- Full Domain Hash: $H: \{0,1\}^* \to \mathbb{Z}_N$
- Given a message $m \in \{0,1\}^*$

$$\sigma = \operatorname{Sign}_{sk}(m) = H(m)^d \mod N$$

Theorem 12.7: RSA-FDH is a secure signature scheme assuming that the RSA-Inversion problem is hard and H is modeled as a random oracle.

Remark: The range of H (e.g., SHA3) may be shorter than \mathbb{Z}_N .

Solution: Repeated application of H e.g., $H'(m) = H(1|m) \dots H(k|m) \mod N$

- Full Domain Hash: $H: \{0,1\}^* \to \mathbb{Z}_N^*$
- Given a message $m \in \{0,1\}^*$

$$\sigma = \operatorname{Sign}_{sk}(m) = H(m)^d \bmod N$$

Theorem 12.7: RSA-FDH is a secure signature scheme assuming that the RSA-Inversion problem is hard and H is modeled as a random oracle.

Proof Sketch: Given an RSA-Inversion challenge $c = r^e \mod N$ (r is unknown) we will simulate the signature attacker. WLOG assume attacker always queries $H(m_i)$ before $Sign_{sk}(m_i)$

- 1. Whenever the attacker queries $H(m_i)$ we can pick a random $r_i \in \mathbb{Z}_N^*$ and program $H(m_i) = r_i^e \mod N$ so that $H(m_i)^d = r_i \mod N$.
- 2. Whenever the attacker queries $\text{Sign}_{sk}(\mathbf{m}_i)$ we can simply return $\mathbf{r}_i \in \mathbb{Z}_N^*$
- 3. Exception: Pick a random query index $j \le q_{hash}$ and program $H(m_i) = c \times r_i^e \mod N$ instead of $H(m_i) = r_i^e \mod N$
- 4. If the attacker forges a signature σ for m_i we can win the RSA-Inversion game by computing $r = \sigma \times r_i^{-1}$ since $\sigma = \text{Sign}_{sk}(m) = (c \times r_i^e)^d = (r \times r_i)^{ed} = rr_i \mod N$.

Analysis: If signature forgery attacker wins with probability f(n) we win RSA-inversion game with probability $f(n)/q_{hash}$ where q_{hash} is the number of queries to the random oracle.

- Full Domain Hash: $H: \{0,1\}^* \to \mathbb{Z}_N^*$
- Given a message $m \in \{0,1\}^*$

$$\sigma = \operatorname{Sign}_{sk}(m) = H(m)^d \mod N$$

Theorem 12.7 (Concrete): Suppose that any attacker running in time at most t'(n) wins RSA-Inversion game with probability at most $\varepsilon'(n)$ then the RSA-FDH is $(t(n), q_H, q_{Sig}, \varepsilon(n))$ -secure i.e., any attacker running in time $t(n) = t'(n) - (q_H + q_{Sig} + 1)O(poly(n))$ and making at most q_H (resp. q_{Sig}) queries to the random oracle (resp. signature oracle) wins with probability at most $\varepsilon(n) \le 4q_{Sig}\varepsilon'(n)$

Theorem 12.7 (Concrete): Suppose that any attacker running in time at most t'(n) wins RSA-Inversion game with probability at most $\varepsilon'(n)$ then the RSA-FDH is $(t(n), q_H, q_{Sig}, \varepsilon(n))$ -secure i.e., any attacker running in time $t(n) = t'(n) - (q_H + q_{Sig} + 1)O(poly(n))$ and making at most q_H (resp. q_{Sig}) queries to the random oracle (resp. signature oracle) wins with probability at most $\varepsilon(n) \leq 4q_{Sig}\varepsilon'(n)$

Proof Idea: Whenever the attacker queries $H(m_i)$ we can pick a random $r_i \in \mathbb{Z}_N^*$ and flip a biased coin $\Pr[\text{heads}] = \left(1 - \frac{1}{1 + q_{Sig}}\right)^{-1}$

- 1. Heads: program $H(m_i) = r_i^e \mod N$ so that $H(m_i)^d = r_i \mod N$.
- 2. Tails: and program $H(m_i) = c \times r_i^e \mod N$
 - 1. If attacker queries signing oracle on this message we will need to abort

$$\Pr[no \text{ abort}] = \Pr[\text{heads}]^{q_{Sig}} = \left(1 - \frac{1}{1 + q_{Sig}}\right)^{q_{Sig}} \approx \frac{1}{e}$$

Theorem 12.7 (Concrete): Suppose that any attacker running in time at most t'(n) wins RSA-Inversion game with probability at most $\varepsilon'(n)$ then the RSA-FDH is $(t(n), q_H, q_{Sig}, \varepsilon(n))$ -secure i.e., any attacker running in time $t(n) = t'(n) - (q_H + q_{Sig} + 1)O(poly(n))$ and making at most q_H (resp. q_{Sig}) queries to the random oracle (resp. signature oracle) wins with probability at most $\varepsilon(n) \leq 4q_{Sig}\varepsilon'(n)$

Proof Idea: Whenever the attacker queries $H(m_i)$ we pick a random $r_i \in \mathbb{Z}_N^*$ and flip a biased coin $Pr[heads] = \left(1 - \frac{1}{1+q_{Sig}}\right)$

- 1. Heads: program $H(m_i) = r_i^e \mod N$ so that $H(m_i)^d = r_i \mod N$.
- 2. Tails: and program $H(m_i) = c \times r_i^e \mod N$
 - 1. If attacker queries signing oracle on this message we will need to abort

$$\Pr[no \text{ abort}] = \Pr[\text{heads}]^{q_{Sig}} = \left(1 - \frac{1}{1 + q_{Sig}}\right)^{q_{Sig}} \approx \frac{1}{e}$$

Reducing Memory Usage in Reduction

Prior reduction is not memory-tight since we need to remember that $H(m_i) = r_i^e \mod N$ for each query.

Solution: Set $r_i = H$ (*K*, m_i) where K is secret key used in the reduction. Program H(m_i) as before. flip a biased coin Pr[heads] = $\left(1 - \frac{1}{1+q_{Sig}}\right)$

- 1. Heads: program $H(m_i) = r_i^e \mod N$ so that $H(m_i)^d = r_i \mod N$.
- 2. Tails: and program $H(m_i) = c \times r_i^e \mod N$
 - 1. If attacker queries signing oracle on this message we will need to abort

$$\Pr[no \text{ abort}] = \Pr[\text{heads}]^{q_{Sig}} = \left(1 - \frac{1}{1 + q_{Sig}}\right)^{q_{Sig}} \approx \frac{1}{e}$$

Idea 1: If the attacker does not query H (K, ...) then the reduction is unchanged.

Idea 2: If the attacker forges a signature σ for m_i we will hope that we programmed $H(m_i) = c \times r_i^e \mod N$ (tails) and compute $r_i = H^-(K, m_i)$ $r = \sigma \times r_i^{-1}$ since $\sigma = \text{Sign}_{sk}(m) = (c \times r_i^e)^d = (r \times r_i)^{ed} = rr_i \mod N$.

CPA-Security Game (Single Message Version)



Random bit b $K \leftarrow Gen(1^n)$

 $(t(n), q(n), \varepsilon(n))$ -secure if any attacker A running in time t and making at most q queries wins with probability at most $\frac{1}{2} + \varepsilon(n)$

Recall: Week 1 Reduction

 $\operatorname{Enc}_{k}(m) = \langle r, F_{k}(r) \oplus m \rangle$

 $\operatorname{Dec}_{k}(\langle r, s \rangle) = F_{k}(r) \oplus s$

For any attacker A running in time t(n) and making at most q(n) encryption queries we have

$$\Pr[\operatorname{Priv} K_{A,\Pi}^{cpa1} = 1] \leq \frac{1}{2} + \frac{q(n)}{2^n} + \mu(n, t'(n), q(n))$$

$$\underset{\text{Collision with initial challenge}}{\operatorname{PRF Security}}$$

CPA-Security Game (Left-Right) m_0^1, m_1^1





Random bit b $K \leftarrow Gen(1^n)$

 $(t(n), q(n), \varepsilon(n))$ -secure if any attacker A running in time t and making at most q encryption queries wins with probability at most $\frac{1}{2} + \varepsilon(n)$

Example: Left-Right Security

 $\operatorname{Enc}_{k}(m) = \langle r, F_{k}(r) \oplus m \rangle$

 $\operatorname{Dec}_{k}(\langle r, s \rangle) = F_{k}(r) \oplus s$

For any attacker A running in time t(n) and making at most q(n) encryption queries we have

$$\Pr\left[\operatorname{Priv}K_{A,\Pi}^{\operatorname{cpal}R}=1\right] \leq \frac{1}{2} + \frac{\binom{q(n)}{2}}{2^n} + \mu(n,t'(n),q(n))$$

Probability there <u>exists</u> a nonce collision

Example: Left-Right Security

F

6

Question: Suppose the attacker is space bounded and that $S \ll q(n)$ so that the attacker cannot store all of the random nonces. Can we prove a tighter security bound?

$$\Pr\left[\operatorname{Priv}_{A,\Pi}^{\operatorname{cpal}_{R}}=1\right] \leq \frac{1}{2} + \frac{\binom{q(n)}{2}}{2^{n}} + \mu(n,t'(n),q(n))$$

Probability there exists a nonce collision

On the Streaming Indistinguishability of a Random Permutation and a Random Function

> Itai Dinur Ben-Gurion University



Eurocrypt 2020

"Switching Lemma" for Random Permutation\Function

- Classical problem: adversary A tries to distinguish a random permutation P:[N]->[N] from random function F:[N]->[N] with Q queries
- "Switching Lemma": A has advantage bounded by O(Q²/N)
 - $|\Pr[A^{P(.)} = 1] \Pr[A^{F(.)} = 1]| \in O(Q^2/N)$
- Widely used to establish concrete security of cryptosystems up to **birthday bound** of $Q = \sqrt{N}$
 - E.g., modes of operation (counter-mode)

oracle

$$q_i \downarrow x_i = P(q_i)$$

 $A \quad or F(q_i)$

"Switching Lemma" for Random Permutation\Function

- "Switching Lemma": A has advantage bounded by O(Q²/N)
 - $|\Pr[A^{P(.)} = 1] \Pr[A^{F(.)} = 1]| \in O(Q^2/N)$
- Matching algorithm: store the Q query outputs and look for collision (F(q_i)= F(q_j) for q_i ≠q_j)



Memory-Restricted Adversaries

- Algorithm requires **memory** ≈**Q** bits
- What about **memory-restricted** adversaries?
- Use cycle detection algorithm to obtain optimal O(Q²/N) advantage with ≈log(N) memory
- Requires **adaptive queries** to primitive
- What if adversary with S memory bits only given stream of Q elements produced by random function/permutation?
- Considered by Jaeger and Tessaro at EUROCRYPT 2019 [JT'19]

oracle
$$X_1$$
 X_2 X_3 ... X_{Q-1} X_Q $x_i = P(i)$ S or F(i) $-$ A

Streaming Switching Lemma [JT'19]

- "Streaming switching lemma" [JT'19]: adversary with S bits of memory with (1-pass) access to stream of Q elements from **random permutation\function** has distinguishing advantage of **at most** $\sqrt{Q \cdot S/N}$
- Application: better **security bounds** against **memoryrestricted** adversaries for some modes of operation

Streaming Switching Lemma [JT'19]

- Application: better **security bounds** against **memory- restricted** adversaries for some modes of operation
- AES-based counter-mode:
- m_i encrypted to $(r_i, c_i = AES_K(r_i) \bigoplus m_i)$ for uniform r_i
- Eavesdropping adversary sees stream (r₁, c₁), (r₂, c₂),...
- **Replace** AES by random P + **apply** streaming switching lemma (several times):
- show (r₁, c₁), (r₂, c₂),... Indistinguishable from
- $(r_i, \alpha_i), (r_i, \alpha_i), \dots$ for uniform α_i

Streaming Switching Lemma

 "Streaming switching lemma" [JT'19]: adversary with S bits of memory with access to stream of Q elements from random permutation\function has distinguishing

advantage of **at most** $\sqrt{Q} \cdot S/N$

- Application: if *S* is limited, counter-mode secure beyond birthday bound
- Limitations of [JS'19]:
- 1) Proof based on unproven combinatorial **conjecture**
- 2) Bound $\sqrt{Q \cdot S/N}$ not tight when $Q \cdot S \ll N$
 - E.g., when S = Q, bound is $\sqrt{Q^2/N}$, but (original) switching lemma gives Q^2/N

New Streaming Switching Lemma

- In this work: overcome limitations
- New streaming switching lemma bound $O(\log Q \cdot Q \cdot S/N)$
- **Tight** (up to poly-log factors):
 - Algorithm: store **first S** elements and look for collision with *Q* elements
 - Advantage: $\approx Q \cdot S/N$
- Note: when S = Q, we get (original) switching lemma



$CC \rightarrow Streaming$

- Main idea: **reduce** from **communication complexity** (**CC**) problem (with **strong lower bounds**) to streaming
- General reduction framework from **one-way** CC problem:
 - Alice, Bob solve CC problem given access to streaming algorithm:
 - View concatenated inputs as stream
 - Alice **simulates** streaming algorithm on her input, **passes state** to Bob which continues simulation, outputs result



$\mathsf{CC} \rightarrow \mathsf{Streaming}$

- Streaming algorithm with memory S gives one-way communication protocol with communication cost S (and same advantage)
- Lower bound on cost of communication protocol → lower bound on memory of streaming algorithm



Reduction Attempt for Random Permutation\Function

- Attempt: CC problem each player gets Q/2 elements, chosen using ran permutation\function
- Useless: CC problem is easy
 - E.g., if $Q > \sqrt{N}$, players can **trivially distinguish** between permutation\function with **no communication**
 - Each player has **unlimited resources** and can detect a collision locally



Reduction Attempt for Random Permutation\Function

- General restriction: in **hard CC problem** joint distributions for Alice and Bob's inputs should have **identical marginals**
 - Alice and Bob should have same local view
- **Impossible** when considering rand permutation\function distributions
- Solution: use **hybrid argument**
 - Consider intermediate hybrid distributions between random permutation and random function
 - Prove indistinguishability of **neighboring hybrid distributions** by reduction from CC
Hybrid Argument

Attempt: define **Q** hybrids games





w\o replacement w replacement

- (Standard) hybrid argument far from tight
 - (Distinguishing advantage) x (num of hybrids) too large

Improved Hybrid Argument

- Main idea: break dependency between halves
- Denote 1st sequence by $x_1, x_2, ..., x_{Q/2}, y_1, y_2, ..., y_{Q/2}$
- 1st distribution: elements chosen using (same) permutation
- 1st intermediate hybrid: $x_1, x_2, ..., x_{Q/2}$ and $y_1, y_2, ..., y_{Q/2}$ chosen using independent permutations
- Reduction from (one-way) CC:
- Alice gets 1st half of sequence, Bob gets 2nd half (decide if they obtain same or independent permutations)
 - Marginals are identical

Permutation Dependence

- (one way) CC problem **permutation dependence** (**PDEP**):
- Alice and Bob decide if their inputs were drawn using **same** or **independent** permutations
- **PDEP** to **streaming** reduction:



UDISJ-> PDEP

- Communication **cost** \ **advantage** tradeoff for **PDEP**?
- Reduction from (unique) **disjointness (UDISJ)**
 - Each player receives a set of size n (domain size O(n)), need to decide if sets intersect or disjoint
- Theorem (informal)[BM'13, GW'14]: if Alice and Bob communicate c bits for DISJ (UDISJ) in the worst case, their max advantage is O(c/n)
 - Even when given access to **public randomness**





UDISJ-> PDEP



- Theorem (informal): there is a public coin local reduction that converts a UDISJ instance of size n=N/Q to a PDEP instance of size Q
 - Shorter inputs harder from PDEP, but easier for UDISJ
- Overall: UDISJ -> PDEP-> streaming bounds max advantage for hybrid game by O(c/n) = O(S/(N/Q)) = O(Q · S/N)

The Full Hybrid Argument

- Once dependency between 2 halves broken:
 - Continue recursively (tree structure)
- 2'nd level: 2 games of distinguishing stream distributions on Q/2 elements
- Final distribution: Q elements divided into Q independent permutations == random function
- Max advantage for each level: $O(Q \cdot S/N)$
- **Total** max advantage: $O(\log Q \cdot Q \cdot S/N)$



Conclusions

- New streaming switching lemma bound $O(\log Q \cdot Q \cdot S/N)$
- **Tight** up to poly-log factors
- Reduction from CC to streaming uses unconventional hybrid argument
- Standard streaming problems defined in **worst case setting**
 - Gives freedom to choose hard distributions for CC problem
- In our (cryptographic) setting streams distributions fixed
 - Hybrid argument reduction applicable to more problems?
- Extension: **multi-pass** streaming switching lemma
 - Streaming alg allowed multiple passes over data

Thanks for your attention!