Fixing Cracks in the Concrete: Random Oracles with Auxiliary Input, Revisited

Yevgeniy Dodis New York University

Joint work with Siyao Guo (Simons Institute, UC Berkeley) Jonathan Katz (University of Maryland)

Hash Functions Are Ubiquitous



- OWFs
- PRG/PRFs
- MACs
- CRHFs
- KDFs

How to assess the <u>best possible</u> concrete security for each application?

Random Oracle Model Methodology [BR93]

Random Function O:[N]->[M]

A: T queries

Clean proofs, Precise bounds: e.g., OWFs/MACs T/min(N,M), PRFs/PRGs T/N, CRHFs T²/M, etc.

Simple Proof Techniques: programmability, lazy sampling, distinguishing-to-extraction, etc.

<u>Practical heuristics</u>: for "natural" applications, Security in ROM = Security in "standard model" (for the best instantiation of O)

Random Oracle Model Methodology [BR93]



<u>Practical heuristics</u>: for "natural" applications, Security in ROM = Security in "standard model" (for the best instantiation of O)

Random Oracle Model Methodology [BR93]



<u>Practical heuristics</u>: for "natural" applications, Security in ROM = Security in "standard model" (for the best instantiation of O)



Random Oracle Standard Model

CRHF: T^2/M

OWF: T/N

[N]->[N]

1 (rainbow tables: in time N^{2/3} [Hel80])

PRG: T/N

1/N^{1/2}

(constant time [AGHP92])

Non-uniform Cracks in the Concrete

Random Oracle Standard Model

CRHF: T^2/M 1

 OWF:
 T/N
 1

 [N]->[N]
 (rainbow tables;

 in time N^{2/3} [Hel80])

PRG: T/N

1/N^{1/2} (constant time [AGHP92])

Security

Non-Uniform Adversaries

- Modeled as families of circuits
 - Can "hardwire" arbitrary (bounded) "advice" before attacking the system
 - <u>Preprocessing</u>: special case of "computable" advice (corresponds to potentially implementable attack)
- Why/how did this become "standard" model?
 - Uniform model too weak (e.g., attacker can focus on a given security parameter in advance)
 - Sometimes preprocessing realistic (rainbow tables!)
 - Seems critical for protocol composition (i.e., ZK)
 - Wlog, deterministic attacker (P/poly=BPP/poly)
 - (Non-uniform) hardness vs randomness : non-uniform lower bounds => derandomization

Non-uniform Cracks in the Concrete

Securit

Can we "salvage" ROM methodology and be consistent with <u>non-uniform</u> attackers? [Unr07] YES: ROM with Auxiliary-Input (ROM-AI)

The ROM methodology is blatantly false for most <u>natural</u> and <u>widely deployed</u> applications when:

- Preprocessing is allowed
- The standard model adversary is non-uniform

<u>Practical heuristics</u>: for "natural" applications, Security in ROM = Security in "standard model" (for the best instantiation of O)

Fixing Cracks in the Concrete



$$A = (A_0, A_1)$$

$$A_0: S \text{ bits } A_0: C \text{ bits } A_1: T \text{ queries}$$

- A₀: computationally <u>unbounded</u>, gets entire RO, and passes S bits of O-dependent advice to A₁
 - Becomes non-uniform advice when O instantiated
 - Separating S and T for more accurate time-space tradeoffs (e.g., for RAM attackers vs. circuits)
- ROM vs. standard model "separations" disappear! Concrete bounds in <u>ROM-AI</u> are meaningful against standard model <u>non-uniform</u> attackers!

Fixing Cracks in the Concrete

ROM-AI Random function O:[N]->[M]

 $A = (A_0, A_1)$ $A_0: S bits \overset{\mathcal{O}_{V_i}}{\longrightarrow} A_1: T queries$

<u>ROM-AI methodology</u>: for "natural" applications, Security in ROM-<u>AI</u> = Security in "standard model" against <u>non-uniform</u> attackers (for the best instantiation of O)

Security against any preprocessing attacks

Handling Auxiliary Input?

<u>Problem</u>: conditioned on S-bit "leakage", values of random oracle are not random and independent

	Traditional ROM	ROM-AI
Lazy Sampling		×
Programmability		×
Distinguishing- to-Extraction		×





Handling Auxiliary Input: Pre-sampling [Unruh07]

 <u>Intuition</u>: conditioned on S-bit leakage, only "few" values of O(x) are "heavily biased"

 $-A_0$ can pass these "few" values as advice to A_1

- The rest can be re-sampled fresh and independent of the leakage!
 - Lazy sampling, programmability, etc. all come back as long as avoid the "few" pre-sampled points





- Pre^o can depend on S-bit "leakage" z about O
 P is a free parameter optimized later (see below)
- But R is random and independent on z
- How big is ε?

[Unr07]: ε < (ST/P)^{1/2}

S bits about O, then T queries



Our Motivating Question



Exact security for basic primitives? (critical for <u>ROM-AI methodology</u>!)

Our Results



- Unruh's "pre-sampling conjecture" <u>false</u>
 - For many apps (OWFs, MACs, etc.), pre-sampling (as defined above) cannot give tight bounds
- New technique: Compression
 - Apply to get nearly tight ROM-AI bounds for OWFs, MACs, PRGs, PRFs, (salted) CRHFs
 - Bounds much weaker than traditional ROM (because there <u>are</u> non-uniform attacks!)
- Salting provably defeats preprocessing!
 - Long-enough salt \Rightarrow ROM-AI-sec. \approx ROM-sec.
 - Possible way to reconcile theory and practice!



Our Results



- Salting provably defeats preprocessing!
 - Long-enough salt \Rightarrow ROM-AI-sec. \approx ROM-sec.
 - Possible way to reconcile theory and practice!

S bits, T queries |Pre^o|=P

Random Function O:[N]->[N]

3≈

Random Function R:[N]->[N] | Pre⁰

• [Unr07]: ε < (ST/P)^{1/2}

– Can't get negl(n) security with P = poly(n) \otimes

- <u>Conj</u>: can get ε = negl(n) for P = poly(n) \odot

- Our result: $\varepsilon > \Omega(ST/P)$ \leftarrow tight! [CDGS17] - Unruh's conjecture false (in this generality)
- Is it enough to prove tight security?





Our New Technique





Compression Paradigm [GT00,DTT10]! High advantage \Rightarrow Compressing RO RO is impossible to compress \Rightarrow Exact security bounds

<u>Challenge</u>: need to compress by more than **S** bits!

(salted) CRHFs

$$Pr[A_1^{O}(A_0(O), a) = (x, x') \text{ s.t. } x \neq x', O(a, x) = O(a, x')]$$
Number of salts a
$$= O(S/K + T^2/M)$$

<u>**Optimal</u>**: can hardwire S collisions inside advice $A_0(O)$!</u>

<u>Idea</u>: compress O(a,x') using indices $i,j \in [T]$ and O(a,x)

of saved bits: $|# of a succeeds| \times (logM - 2logT)$ = $\epsilon K \times \frac{1}{2} M/T^2$

 $S > \epsilon K \log (\epsilon M/eT^2)$

The Order Issue

Consider 2 salts: $O(a_1, x_1) = O(a_1, x_1')$; $O(a_2, x_2) = O(a_2, x_2')$

Ideally, compress both $O(a_1, x_1')$ and $O(a_2, x_2')$

<u>Problem</u>: what if A(z, a₁) would query O(a₂,x₂')?? (not so crazy because of advice z...)

<u>Solution</u>: run A on all salts a where he succeeds, and keep track of "salt-specific" indices i_a , j_a for the first collision (which exists!) on all such a's



ROM-AI Bounds for Basic Primitives

Primitive	Our ROM-AI Bound
OWF*	ST/N
PRG	(ST/N) ^{1/2}
PRF	(ST/N) ^{1/2}
MAC*	ST/N

Always Better than Pre-sampling $\ensuremath{\textcircled{\odot}}$

Primitive	Our ROM-AI Bound	Pre-Sampling
OWF*	ST/N	(ST/N) ^{1/3}
PRG	(ST/N) ^{1/2}	(ST/N) ^{1/3}
PRF	(ST/N) ^{1/2}	(ST/N) ^{1/3}
MAC*	ST/N	(ST/N) ^{1/3}

Nearly Tight ©

Primitive	Our ROM-AI Bound	Best Attack
OWF*	ST/N	$Min(ST/N, (S^2T/N^2)^{1/3})$
PRG	(ST/N) ^{1/2}	(ST/N) ^{1/2}
PRF	(ST/N) ^{1/2}	(ST/N) ^{1/2}
MAC*	ST/N	$Min(ST/N, (S^2T/N^2)^{1/3})$

But Much Weaker than ROM 😕

Primitive	Our ROM-AI Bound	Traditional ROM
OWF*	ST/N	T/N
PRG	(ST/N) ^{1/2}	T/N
PRF	(ST/N) ^{1/2}	T/N
MAC*	ST/N	T/N

But Much Weaker than ROM 😣



Maybe we can all live in peace?



How to Defeat Preprocessing?

Extensively used in theory and practice: Saw the magic for CRHFs already!



Chose random public salt after preprocessing; Prepend as input to O

Security Bounds for Salting $O: [K] \times [N] \rightarrow [M]$

Primitive	Salted ROM-AI Bound	Traditional ROM
OWF*	T/N + ST/ <mark>K</mark> N	T/N
PRG	T/N + (ST/KN) ^{1/2}	T/N
PRF	T/N + (ST/KN) ^{1/2}	T/N
MAC*	T/N + ST/ <mark>K</mark> N	T/N
CRHF	T²/M + S/K	T ² /M

Security Bounds for Salting $O: [K] \times [N] \rightarrow [M]$

Primitive	Salted ROM-AI Bound	Traditional ROM
OWF*	T/N +	T/N
PRG	T/N + (ST, 1) ^{1/2}	T/N
PRF	T/N + (S1) ^{1/2}	T/N
MAC*	T/N + S (N	T/N
CRHF	T ² /M + S K	T ² /M

Salting Provably helps!



At most n bits of salt yield ≈ <u>same</u> security in ROM with auxiliary input as without auxiliary input



n-bit salt provably defeats preprocessing



Summary





Summary

- ROM-AI is the new "cool kid" in town !
 - -very clean: just S and T!
 - addresses both theory (non-uniformity) and practice (preprocessing)
 - non-trivial, yet elegant proofs
 - 1000's of ROM papers need re-evaluation !
- Obfuscation without the mess !



Your proposal is written with clarity and conviction. Send it up to legal for obfuscation.

Follow-Up Work [CDGS17]

- Optimal Pre-Sampling Error ST/P
 - Improves $(ST/P)^{1/2}$ [Unruh07]
 - Gives tight bounds for indistinguishability apps
- New pre-sampling for unpredictability apps
 - Matches compression for all current apps
- Salting generically defeats preprocessing
- Random Permutation and Ideal Cipher with Auxiliary Input

Limitation of Pre-sampling

Random Function O:[N]-> {0,1}

≈ε

Random Function R:[N]-> {0,1} | Pre⁰

 $A=(A_0,A_1)$ " |Pre|=P

 $\Pr[A_1^{O}(A_0(O))=1] - \Pr[A_1^{R|Pre}(A_0(O))=1] > 1/24P$

> 1/2 + 1/3L^{1/2} <=1/2 + P/2L L= 4P²+1

 $A_0(O) = Marj(O_1, ..., O_L)$ $A_1 = 1$ if $A_0(O) = O_i$ where i~[L]

Extending to large S,T

Extending to large T: xor first



Extending to large S: repeat on disjoint inputs