Scrypt is Maximally Memory Hard







Krzysztof Pietrzak IST Austria



Leonid Reyzin Boston University





Password hashing: Store a hash of a password + salt





Traditional hardness metric: <u>Time complexity</u> (e.g., PKCS #5)



Honest users

(General-purpose CPU)



cost per **F** eval = **C**

better parallelization, pipelining for speedup; lower energy costs ...

Adversaries (ASICs)



 $cost per F eval = C' \ll C$

Can we enforce $C' \approx C$?

Idea: Memory costs (e.g., on-chip area, access time, \$-cost) are platform independent

Memory-hard functions (MHFs)[Percival, 2009]



large memory



Small memory \Rightarrow <u>slow</u>

evaluation of MHF F

Memory-hardness, more precisely

Goal: Maximize <u>cumulative memory complexity (CMC)</u> for any (possibly **parallelized**) strategy to evaluate **F.**



 $CMC = \sum_{t=0}^{T} Memory(t)$

[Alwen and Serbinenko, STOC '15]

Memory-hardness

PHC call for submissions



The Password Hashing Competition (PHC) organizers solicit proposals from any interested party for candidate password hashing schemes, to be considered for inclusion in a portfolio of schemes suitable for widespread adoption, and covering a broad range of applications.

Memory-hardness was defacto requirement for PHC

- Security
 - Cryptographic security: the function should behave as a random function (random-looking output, oneway, collision resistant, immune to length extension, etc.).
 - Speed-up or other efficiency improvement (e.g., in terms of memory usage per password tested) of cracking-optimized implementations (checking multiple sets of inputs in parallel, and doing so in a CPU's native code) compared to implementations intended for password validation should be minimal.
 - Speed-up or other efficiency improvement (e.g., in terms of area-time product per password tested) of cracking-optimized ASIC, FPGA, and GPU implementations (checking multiple sets of inputs in parallel) compared to CPU implementations intended for password validation should be minimal.
 - Resilience to side-channel attacks (timing attacks, leakages, etc.). In particular, information should not leak on a password's length.

Many memory-hard candidates: Argon2d, Argon2i, Scrypt,

Lyra2, Balloon hashing, Catena, Yescrypt,

Can we build provably memory-hard functions?

Towards optimal memory hardness

Previous <u>provably</u> MHFs [AS15,BCS16,ABP17] are iMHFs: <u>data-independent</u> memory access patterns!





Two issues raised by Alwen and Blocki:

(1) No iMHF achieves <u>optimal</u> memory hardness.

(2) <u>Practical</u> iMHFs are even less memory hard for parallel evaluation strategies.

Can datadependence help?

This paper: Scrypt is optimally memory hard

- Very first conjectured MHF: Proposed by Colin Percival in 2009
- Used within PoWs in Litecoin
- Inspired the design of Argon2d one of the winners of Password Hashing Competition
- Covered by RFC 7914

Take home message: Very first example of function with provably optimal memory hardness.

+ it is practical, already in use, and relatively simple

Finding such proof has been a surprisingly hard problem:

- [Percival, 2009] is incorrect
- [ACKKPT16] only gave restricted result

No iMHF achieves optimal memory hardness





Roadmap

The Scrypt function Definition, memory-hardness intuition, and challenges

2. Optimal memory hardness of Scrypt

3. Conclusions



Core of Scrypt: ROMix

Modeled as a random oracle! \longrightarrow H: A Salsa20 based "hash function" with output length w.



n: a tunable parameter. e.g., $n = 2^{14}$, w = 1 KB

ROMix: Answering challenges



this!

2. Need to answer <u>all challenges</u> to complete the evaluation



Abstract 2nd phase: challenges are **H**-dependent random!

For all round $j = 0, \dots, n-1$:



<u>Adversary's goal:</u> Reduce its own CMC for answering all challenges!

CMC = Cumulative Memory Complexity = $\sum_{t=0}^{T}$ Memory(t)

Round game – Naïve strategy





Round game – Memory-less strategy



Previous two strategies are special cases: consistent memory size



More general strategy: memory consumption can vary a lot

e.g., forget values, re-compute afterward



Memory hardness: intuition



Single-shot memory-time trade-off

Simplifying assumption: upon learning challenge C_j , adversary only stores p of the values X_0, \ldots, X_{n-1}



Fact: Avg-distance from X_{C_i} to closest stored X_i preceding X_{C_i} is n/2p

Regardless of parallelism, as computation of *X*-values is inherently sequential! Expected time to answer the challenge is n/2p

≈|memory|

How to translate this intuition into a memoryhardness proof for ROMix?

Three technical barriers:

 Adversary stores <u>arbitrary information</u> e.g., XOR of X_i values, halves of X_i, reconstruct information adaptively on challenges, etc.
 [ACKKPT16] considered restricted strategies and exhibited round games where general storing strategies can help!

Focus on

1 and **2**

Memory

2. Memory variation during computation single-shot memory-time trade-off not enough! [ACKKPT16] only shows CMC = $\Omega(\frac{n^2w}{\log^2(n)})$

3. H-dependent challenges, as opposed to truly random see the paper!

Roadmap

- 1. The Scrypt function
- 2. Optimal memory hardness of Scrypt Model, theorem, and proof approach

3. Conclusions



The parallel random oracle model

[Alwen and Serbinenko, STOC '15]



At each step: Adv asks one batch of parallel H queries + performs unbounded computation

Goal of adv: minimize $CMC = \sum_{i=1}^{T} |S_i|$

Main Theorem.

For any adversary A evaluating ROMix,

$$\mathsf{CMC}(\mathsf{A}) \geq \frac{1}{25} \cdot n^2 \cdot (w - 4 \cdot \log(n))$$

w/ overwhelming probability over the choice of **H**.

The $4\log(n)$ loss is inherent in the proof.

 $\Omega(n^2w)$ clearly <u>best possible</u> for any construction making *n* queries to **H**.

Naïve strategy: Make n calls, remember all outputs

Proof strategy: step 1



Memory-time trade-off \Rightarrow lower bound on memory The memory-time trade-off holds true for adv storing <u>X-values</u> even if the adv stores <u>arbitrary</u> information!

Single-shot memory-time trade-off



Z: <u>arbitrary</u> computation on **H**-outputs Goal: Lower bound |Z|

 E.g., pre-computation of H's entries, XOR of {X_i} values, halves of X_i

[ACKKPT16]: computation on H-outputs can help in some round games

[This result]: computation on H-outputs cannot help for Scrypt!

Single-shot memory-time trade-off



Lemma. For all A, for most **H**, if $|Z| \approx pw$ bits $\Pr_{C} \left[T > \frac{n}{2p} \right] > \frac{1}{2}$ **Lemma.** For all A, for most **H**, if $|Z| \approx pw$ bits $\Pr_{C} \left[T > \frac{n}{2p} \right] > \frac{1}{2}$

Proof idea:

If adversary $A^{H}(Z, C)$ answers too <u>fast</u> for most challenges C

 $A^{\mathbf{H}}(\mathbf{Z}, \mathbf{C})$ can output or query many X_i values w/o querying **H** first

Can compress the oracle **H** using state **Z**

Cannot be true for too many H: random oracle is incompressible

[Dwork, Naor and Wee, Crypto'05], [Alwen and Serbinenko, STOC '15]

Proof strategy: step 2

Technical barriers:

- 1. Adversary stores <u>arbitrary information</u> Single-shot memory-time trade-off for <u>arbitrary adv</u>
- 2. Memory variation during computation



Single-shot memory-time trade-off Generalize



Optimal CMC lower bound for the round game



memory-time trade-off \Rightarrow memory lower bound for the step <u>right before</u> the challenge is revealed

CMC = ???



Similar trade-off holds for every step before challenge is revealed



By adding lower bounds over rounds from 0 to n - 1, we have $CMC = \Omega(n^2w)$

Roadmap

- 1. The Scrypt function
- 2. Optimal memory hardness of Scrypt

3. Conclusions



Summary



- Scrypt is maximally memory hard
 - First <u>optimal</u> memory-hardness proof.
 - Validates a practical MHF design.

- Open problem
 - Optimal memory hardness proof for Argon2d?

Thank you! – Merci! https://eprint.iacr.org/2016/989