Cryptography from Sublinear-Time Average-Case Hardness of Time-Bounded Kolmogorov Complexity

Yanyi Liu
Cornell Tech

Rafael Pass
Cornell Tech &
Tel Aviv University

Properties of Meta-Complexity Problems

- Let **C** be a complexity measure
- **MCP[s]** = {**x** | **C**(**x**) <= s(|**x**|)}

Question: Does hardness of **MCP[n/4]** => hardness of **MCP[n/2]**?
Or, **MCP[polylog]** is hard for **TIME[n]** <=> **MCP[n/2]** is hard for **TIME[2^{n^ϵ}]**?

[RS21]: ∃ oracle world in which **MCSP[n/4]** is hard but **MCSP[n/2]** is easy.

Main Result: Yes (for both)!
(when considering appropriate notions for K^t and avg-case hardness)
Our proof goes through the notion of **OWFs**.

One-way Functions (OWF) [Diffie-Hellman'76]

A function **f** that is

- **Easy to compute**: can be computed in poly time
- **Hard to invert**: no PPT can invert it

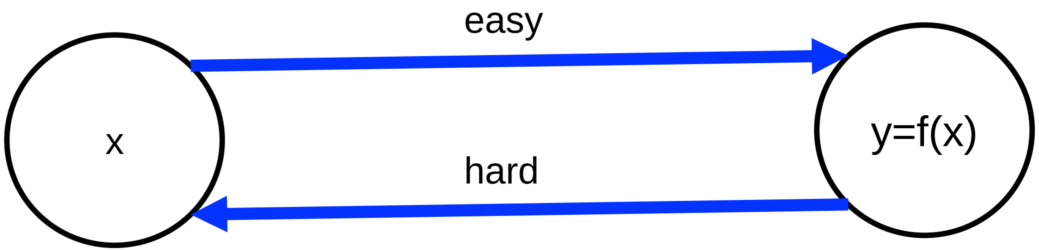

Ex [Factoring]: use x to pick 2 random "large" primes p,q, and output y = p* q

One-way Functions (OWF) [Diffie-Hellman'76]

A function **f** that is

- **Easy to compute**: can be computed in poly time
- **Hard to invert**: no PPT can invert it

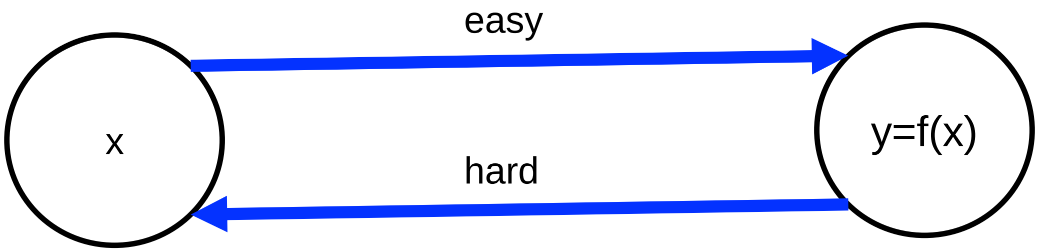

Definition 2.1. *Let $f : \{0,1\}^* \to \{0,1\}^*$ be a polynomial-time computable function. f is said to be a* one-way function (OWF) *if for every* PPT *algorithm $\mathcal{A}$, there exists a negligible function μ such that for all $n \in \mathbb{N}$,*

$$\Pr[x \leftarrow \{0,1\}^n; y = f(x) : A(1^n, y) \in f^{-1}(f(x))] \le \mu(n)$$

One-way Functions (OWF) [Diffie-Hellman'76]

A function **f** that is

- **Easy to compute**: can be computed in poly time
- **Hard to invert**: no PPT can invert it

OWF both necessary [IL'89] and sufficient for:

- Private-key encryption [GM84,HILL99]
- Pseudorandom generators [HILL99]
- Digital signatures [Rompel90]
- Authentication schemes [FS90]
- Pseudorandom functions [GGM84]
- Commitment schemes [Naor90]
- Coin-tossing [Blum'84]
- …

Not included:

public-key encryption, OT, obfuscation

Whether OWF exists is the most important problem in Cryptography

Characterization of OWFs [LP'20]

For every polynomial $t(n)>1.1n$:
OWFs exist iff **$MK^tP[n - O(\log n)]$** is mildly hard-on-average

$MK^tP[s]$: the set of strings x with t-bounded Kolmogorov complexity at most $s(|x|)$

Today: what happens when the threshold changes?

Time-Bounded Kolmogorov Complexity

Give a truthtable $\mathbf{x} \in \{0,1\}^n$ of a Boolean function, what is the size of the smallest "program" that computes $\mathbf{x}$?

When "program" = time-bounded TMs

- t-time-bounded Kolmogorov Complexity [Kol'68, Ko'86, Sip'83, Har'83, AKB+06]
- There are many ways to define time-bounded Kolmogorov complexity.
  We here consider the "**local compression**" version.

When "program" = circuits

- Minimum Circuit Size problem (MCSP) [KC'00, Tra'84]

Time-Bounded Kolmogorov Complexity

Give a truthtable $\mathbf{x} \in \{0,1\}^n$ of a Boolean function, what is the size of the smallest "program" that computes $\mathbf{x}$?

When "program" = time-bounded TMs

- t-time-bounded Kolmogorov Complexity [Kol'68, Ko'86, Sip'83, Har'83, AKB+06]

$\mathbf{K^t(x)}$ = length of the shortest program Π such that Π computes truthtable $\mathbf{x}$ within time $\mathbf{t(|\Pi|)}$

Fix a universal TM U, and a running time bound t. We are looking for the length of the shortest program Π s.t. $U(\Pi(i),1^{t(|\Pi|)}) = x_i$, $\forall$ i <= |x|.

$\mathbf{MK^tP[s]}$: $\{x \mid K^t(x) <= s(|x|)\}$

Characterization of OWFs [LP'20]

OWFs exist iff **MK^{t}P[n - O(log n)]** is mildly hard-on-average

What happens if the threshold **s << n**?

Does (avg-case) hardness of **MK^{t}P[poly log n]** imply OWFs?
Is **MK^{t}P[n - O(log n)]** harder than **MK^{t}P[poly log n]** (on avg)?

Main Theorem

Our main theorem demonstrates that for an *appropriate* notion of mild avg-case hardness:

1. **Subexp-secure OWF $\Leftrightarrow$ mild avg-case hardness of MK^{t}P[polylog n] w.r.t. sublinear algorithms (running in time n^ε, $\varepsilon<1$)**

2. **Qpoly-secure OWF $\Leftrightarrow$ mild avg-case hardness of MK^{t}P[$2^{O(\sqrt{\log n})}$] w.r.t. sublinear algorithms**

Proving the existence of **Subexp OWFs** is equivalent to proving a **sublinear** avg-case lowerbound

Main Theorem

Our main theorem demonstrates that for an *appropriate* notion of mild avg-case hardness:

1. **Subexp-secure OWF $\Longleftrightarrow$ mild avg-case hardness of MK^{t}P[polylog n] w.r.t. sublinear algorithms (running in time n^{ε}, ε<1)**

2. **Qpoly-secure OWF $\Longleftrightarrow$ mild avg-case hardness of MK^{t}P[$2^{O(\sqrt{\log n})}$] w.r.t. sublinear algorithms**

Unconditional Lower Bounds

- **MK^{t}P[polylog n] is worst-case hard for sublinear time algorithms.**
- **MK^{t}P[n-log n] is mildly avg-case hard for Dtime(t/n^3)**

Main Theorem

Our main theorem demonstrates that for an *appropriate* notion of mild avg-case hardness:

1. **Subexp-secure OWF $\Leftrightarrow$ mild avg-case hardness of MK^{t}P[polylog n] w.r.t. sublinear algorithms (running in time n^ε, ε<1)**

2. **Qpoly-secure OWF $\Leftrightarrow$ mild avg-case hardness of MK^{t}P[$2^{O(\sqrt{\log n})}$] w.r.t. sublinear algorithms**

Threshold **s** for the **MK^{t}P[s]** captures the **quantitative** hardness of OWFs (or **MK^{t}P[n - O(log n)]**)
Smaller s $\Leftrightarrow$ "Easier" **MK^{t}P[s]**

The brute-force attacker running time "ratio" remains the same.

Avg-case Hardness for Sparse Languages

Observation: $|\mathbf{MK^tP[s]} \cap \{0,1\}^n| \approx 2^{s(n)}$.

When $s(n) = n/2$, the trivial outputting NO heuristic succeeds w.p. $1\text{-}2^{-n/2}$.
so $MK^tP[s]$ is trivially **easy-on-average.**

Our notion: require hardness **conditioned** on both YES and NO**.**

μ-heuristic* H for L: H succeeds on at least a $1\text{-}\mu(n^*)$ fraction of **YES** instances, and at least a $1\text{-}\mu(n^*)$ fraction of **NO** instances, where **$n^* = \log|L\cap \{0,1\}^n|$**

Avg-case* hardness: We say that $\mathbf{MK^tP[s]}$ is **mildly hard on average* (HoA*)** if there exists a poly p such that $MK^tP[s]$ does not have (1/p)-heuristic* w.r.t. infinitely many input length.

Main Theorem

"Nice" classes: We say that **F** is a "nice" class of time-bounds if (a) all T $\in$ **F**, T is strictly increasing and (b) all T $\in$ **F**, T is closed under poly-composition (T $\in$ **F** => $T(n^{\varepsilon})^{\varepsilon} \in$ **F**).

Main THM: Let F be a "nice" class of super-polynomial (but subexp) functions.
Let t(n) >= 1.1n be a polynomial. The following are equivalent:

1. $\exists T \in F$ s.t. **T-Hard secure OWF** exists
2. $\exists T \in F$ **$MK^tP[T^{-1}]$** is mildly HoA* w.r.t. **sublinear** (i.e time n^{ε}) algorithms.
3. $\exists T \in F$ **$MK^tP[n/2]$** is mildly HoA* w.r.t. **T**-time algorithms.

Corr:

1. **Quasipoly-secure OWF** $\Leftrightarrow$ **$MK^tP[2^{O(\sqrt{\log n})}]$ is sublinear mild-HOA***
2. **Subexp-secure OWF** $\Leftrightarrow$ **MK^tP[polylog n] is sublinear mild-HOA***

Related Work

- **Hardness Magnification**[OS18, MMW19, CT19, OPS19, CMMW19, Oli19, CJW19, CHO+20]
  - weak lower bounds => breakthrough separations
  - compress an instance in a sparse language into another much shorter instance
  - Our result: **hardness magnification for OWFs**

- **Fine-grained Complexity**[BRSV17, BRSV18, GR18, LLW19, BABB19, DLW20]
  - fine-grained lower bounds => fine-grained OWFs (secure w.r.t. a-priori bounded poly-time attacker)
  - Our result: **sublinear lower bounds <=> super-poly hard OWFs**

Main Theorem

Main THM: Let F be a “nice” class of super-polynomial (but subexp) functions.
Let t(n) >= 1.1n be a polynomial. The following are equivalent:

1. $\exists T \in F$ s.t. **T-Hard secure OWF** exists
2. $\exists T \in F$ **$MK^tP[T^{-1}]$** is mildly HoA* w.r.t. **sublinear** (i.e time n^{ε}) algorithms.
3. $\exists T \in F$ **$MK^tP[n/2]$** is mildly HoA* w.r.t. T-time algorithms.

Today: (1) ⇔ (2)
(2) ⇔ (3) follows from the same type of argument (in the paper)

Let F be a "nice" class of super-polynomial (but subexp) functions.
Let $t(n) >= 1.1n$ be a polynomial.

Theorem 1:

Assume that $\exists s \in F^{-1}$ **$MK^tP[s]$** is mildly HoA* w.r.t. **sublinear** algorithms. Then $\exists T \in F$ s.t. **T-Hard secure OWF exists**

Theorem 2:

Assume that $\exists T \in F$ s.t. T-Hard secure OWF exists. Then $\exists T \in F$ $MK^tP[T^{-1}]$ is mildly HoA* w.r.t. sublinear (i.e time n^{ε}) algorithms.

Theorem 1

Assume that ∃s ∈ F^{-1} **$MK^tP[s]$** is mildly HoA* w.r.t. **sublinear** algorithms.
Then ∃T ∈ F s.t. **T-Hard secure OWF exists**

For simplicity, we focus our attention on $F = F_{subexp}$ and s = polylog n.
That is: **$MK^tP[s]$** is mildly HoA* w.r.t. **sublinear** attackers => **Subexp OWFs**

Theorem 1

Assume that $\exists s \in F^{-1}$ **$MK^tP[s]$** is mildly HoA* w.r.t. **sublinear** algorithms.
Then $\exists T \in F$ s.t. **T-Hard secure OWF exists**

By Yao's hardness amplification Lemma [Yao'82], it suffices to construct a weak **T**-hard OWF.

Weak T-hard OWF: "mild-HoA version" of a OWF:
efficient function f s.t. no T-time algorithms can invert f w.p. **1-1/p(n)** for inf many n, for some poly p(n)>0.

Let **t** be a (polynomial) time-bound (the time-bound from the K-complexity problem)

OWF **f** construction

- Use the input to sample a random length ℓ <= n and a length-ℓ program $\mathbf{\Pi}$
- For i ϵ **[2n]**, let y_i = output of $\mathbf{\Pi}$(**i**) after t(ℓ) steps. (y_i is a single bit.)
- Output $y_1\ y_2 \ldots\ y_{2n-1}\ y_{2n}$

Assume for contradiction that f is not a weak T-hard OWF.
That is, there exists a **subexp-time attacker A** that inverts f w.h.p**.**

We construct a **sublinear-time heuristic H** (using A) that **decides MK^{t}P[s] w.h.p.,** which concludes that **MK^{t}P[s]** is not mildly HoA, a contradiction.

Given an instance **x** $\in\{0, 1\}^n$, need to **decide** whether $\mathbf{K^t(x) <= s(n).}$

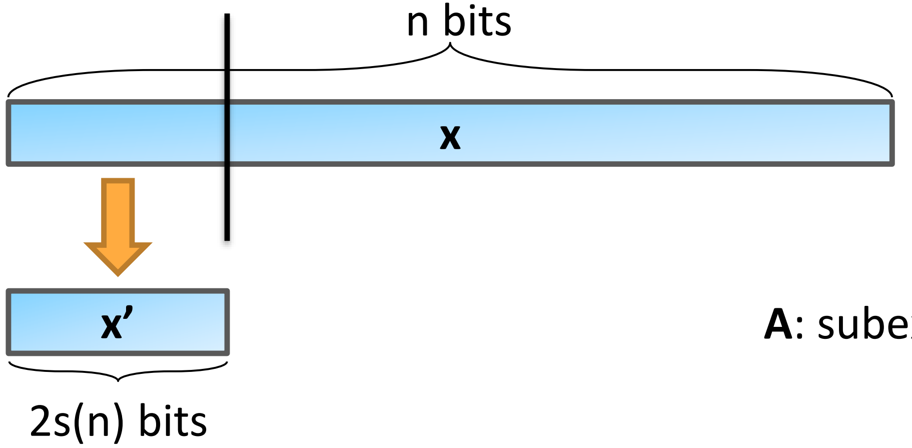

A: subexp-time heuristic for **f**.

H(x) first truncates **x** to 2s(n) bits (and gets **x'**) and outputs 1 if **A(x')** outputs a $\mathbf{K^t}$**-witness** for **x'** of length <= s(n).

Although **A** runs in subexp-time, **x'** is so short that **H** just runs in sublinear time in **|x|** (since s(n) = poly log n).

If x is a YES Instance for MK^{t}P[s]

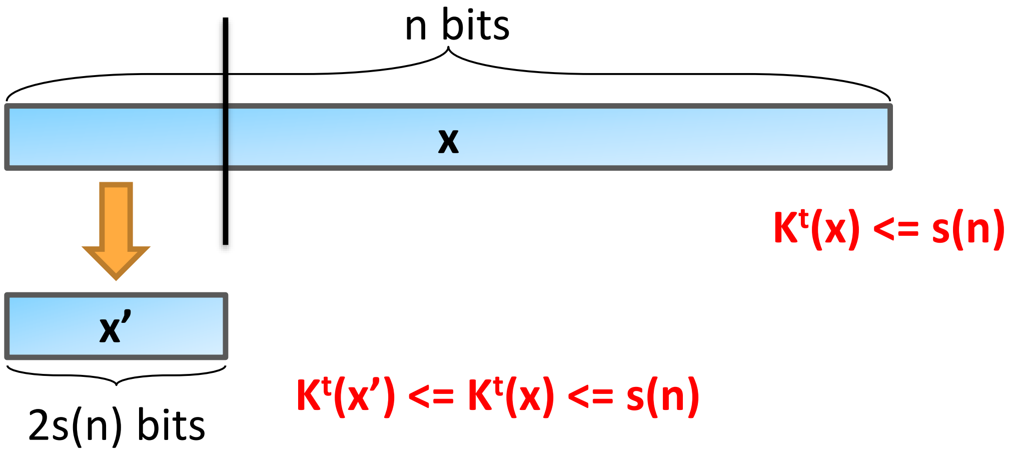

A(x') should be able to find a witness of **x'**

We have to argue that **H** succeeds with high probability.

Problem 1: x does not have the right distribution (similar to [LP20])

Problem 2: The "truncating mapping" is **not** one-to-one: many YES-instances x could lead to the same x'.

If x is a YES Instance for $MK^tP[s]$

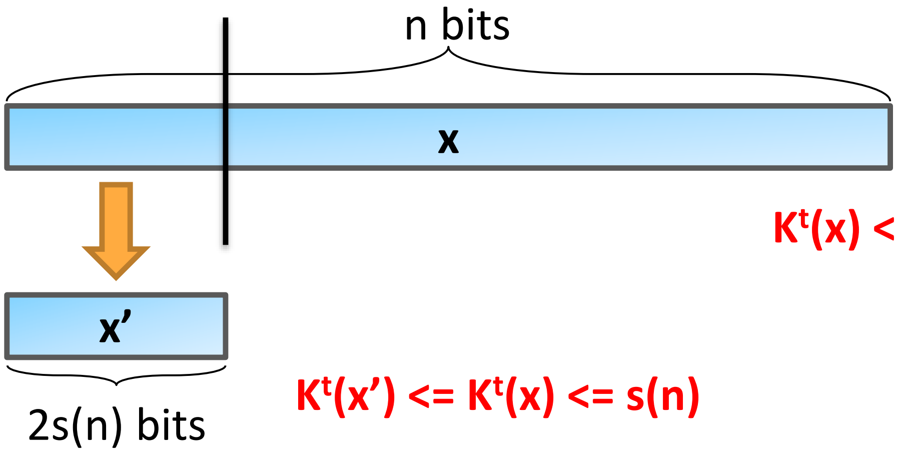

$K^t(x) <= s(n)$

$K^t(x') <= K^t(x) <= s(n)$

A(x') should be able to find a witness of **x'**

Key observation: The **more** YES-instances that are mapped to **x'**, the **larger** the probability mass **x'** has in the OWF experiment!

Can next use similar analysis to [LP'20].

If x is a NO Instance for $MK^tP[s]$

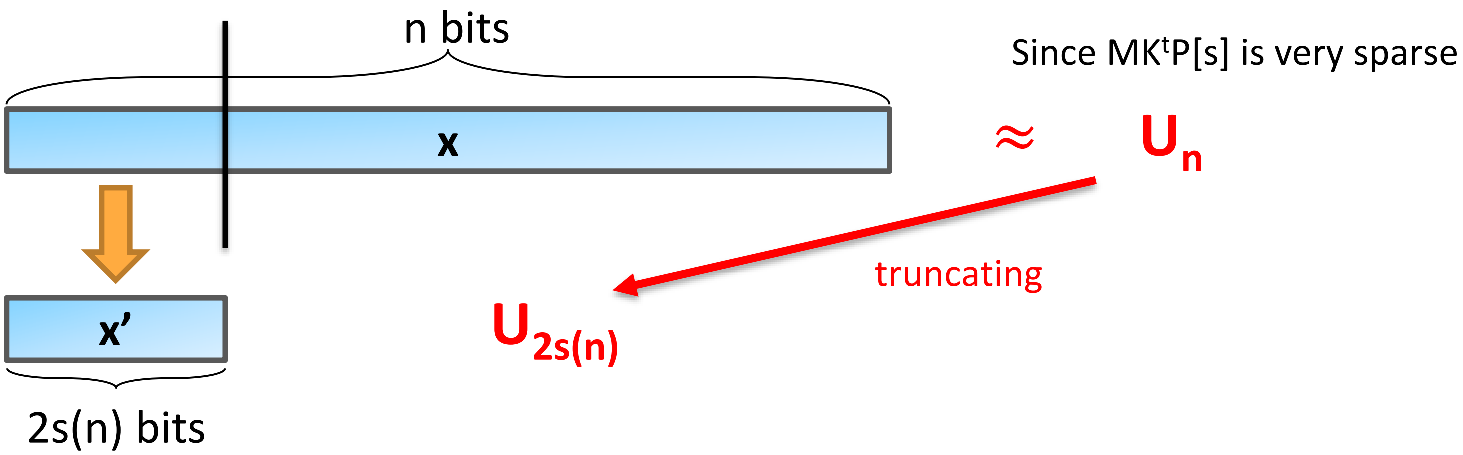

If **x** is a **NO** instance, after being truncated, **x'** could become a **YES** instance.
But we only need to show **H** works on a random NO instance.

It follows that **A** can rarely find K^t-witness of length <= s(n)

Let F be a “nice” class of super-polynomial (but subexp) functions.
Let $t(n) >= 1.1n$ be a polynomial.

Theorem 1:

Assume that $\exists s \in F^{-1}$ **$MK^tP[s]$** is mildly HoA* w.r.t. **sublinear** algorithms. Then $\exists T \in F$ s.t. **T-Hard secure OWF exists**

Theorem 2:

Assume that $\exists T \in F$ s.t. T-Hard secure OWF exists. Then $\exists T \in F$ $MK^tP[T^{-1}]$ is mildly HoA* w.r.t. sublinear (i.e time n^{ε}) algorithms.

Let F be a “nice” class of super-polynomial (but subexp) functions.
Let $t(n) >= 1.1n$ be a polynomial.

Theorem 1:

Assume that $\exists s \in F^{-1}$ **$MK^tP[s]$** is mildly HoA* w.r.t. **sublinear** algorithms. Then $\exists T \in F$ s.t. **T-Hard secure OWF exists**

Theorem 2:

Assume that $\exists T \in F$ s.t. **T-Hard secure OWF exists**. Then $\exists T \in F$ **$MK^tP[T^{-1}]$** is mildly HoA* w.r.t. sublinear (i.e time n^{ε}) algorithms.

Theorem 2

Assume that ∃T ∈ F s.t. **T-Hard secure OWF exists**.
Then ∃s ∈ F^{-1} **$MK^tP[s]$ is mildly HoA*** w.r.t. sublinear (i.e time n^{ε}) algorithms.

cond-EP PRG w/ **sublinear** stretch $G:\{0,1\}^n \to \{0,1\}^{n+n^\epsilon}$

- **Pseudorandomness:** $G(U_n \mid E)$ indistinguishable from U_{n+n^ϵ}
- **Entropy-preserving:** $[G(U_n \mid E)]_n$ has **Shannon** entropy n-O(log n)

Lemma 1: **cond-EP PRG** w/ sublinear stretch => **$MK^tP[s]$** is mildly HoA*
(passes through PRFs)

Lemma 2: **cond-EP PRG** w/ sublinear stretch from OWF

cond-EP PRG from OWFs

Lemma: OWFs => cond-EP PRG w/ sublinear stretch

[LP'20]: OWFs => cond-EP PRG w/ **logarithmic** stretch
Why not do repeated applications **G(…G(G(-))…)**

This does **not** give us a cond-EP PRG w/ sublinear stretch directly.

We here give a new construction of a cond-EP PRG with sublinear stretch.

cond-EP PRG from OWFs

Lemma: OWFs => cond-EP PRG w/ sublinear stretch

Proof:

- OWFs => PRGs [HILL'99]
- Let **G** $\{0,1\}^n \to \{0,1\}^{2n}$ be a PRG. Sample **i** as a guess of "degeneracy" of **G(x)**.
- If **i** is correct, given **G(x)**, **x** **has min-entropy i**.
- Applying hash functions as extractors

$$G^*(i, x, h1) = h1, G(x), [h1(x)]_{i-O(\log n)}$$

- Not pseudorandom: length **i** is leaked.

$$G'(i, x, h1, h2) = h1, h2, [h2(\underbrace{G(x), [h1(x)]_{i-O(\log n)}}_{\text{Entropy (roughly) } n - O(\log n)})]_{3n/2}$$

cond-EP PRG from OWFs

$\mathbf{G'(i, x, h1, h2) = h1, h2, [h2(G(x), [h1(x)]_{i-O(\log n)})]_{3n/2}}$

Pseudorandomness:

- We need to show that if ∃ **D'** that breaks **G'** conditioned on **i being correct**, then ∃ **D** that breaks **G**
- **Bad news**: D does **not** know **i**
- Guessing i does not work, as the distinguisher can be very bad when the guess is incorrect.

- A central contribution is dealing with this issue.

Conclusion

1. **Subexp-OWFs** ⇔ proving an avg-case lowerbound w.r.t. **sublinear** attackers.

2. The threshold **s(n)**, for the **MK^{t}P[s]** problem captures the **quantitative** hardness of OWFs.

3. Technically: the notion of a **“EP-PRG with large stretch”** plays a central role.

Open question: can we characterize exponential OWF!

Thank You