Robustness of Average-Case Meta-Complexity via Pseudorandomness

Rahul Ilango
ilangorahul@gmail.com
MIT
Massachusetts, USA

Hanlin Ren*
h4n1in.r3n@gmail.com
University of Oxford
Oxford, UK

Rahul Santhanam
rahul.santhanam@cs.ox.ac.uk
University of Oxford
Oxford, UK

ABSTRACT

We show broad equivalences in the average-case complexity of many different meta-complexity problems, including Kolmogorov complexity, time-bounded Kolmogorov complexity, and the Minimum Circuit Size Problem. These results hold for a wide range of parameters (various thresholds, approximation gaps, weak or strong average-case hardness, etc.) and complexity notions, showing the theory of meta-complexity is very *robust* in the average-case setting.

Our results are shown by establishing new and generic connections between meta-complexity and the theory of pseudorandomness and one-way functions. Using these connections, we give the first unconditional characterization of one-way functions based on the average-case hardness of the Minimum Circuit Size Problem. We also give a surprising and clean characterization of one-way functions based on the average-case hardness of (the worst-case uncomputable) Kolmogorov complexity. Moreover, the latter is the first characterization of one-way functions based on the average-case hardness of a fixed problem on *any* samplable distribution.

We give various applications of these results to the foundations of cryptography and the theory of meta-complexity. For example, we show that the average-case hardness of deciding k-SAT or Clique on any samplable distribution of high enough entropy implies the existence of one-way functions. We also use our results to unconditionally solve various meta-complexity problems in CZK (computational zero-knowledge) on average, and give implications of our results for the classic question of proving NP-hardness for the Minimum Circuit Size Problem.

CCS CONCEPTS

• **Theory of computation → Pseudorandomness and derandomization**; *Problems, reductions and completeness*; Circuit complexity; *Cryptographic primitives.*

KEYWORDS

meta-complexity, average-case complexity, coding theorem, minimum circuit size problem, Kolmogorov complexity, one-way functions

*Part of this work was done when Hanlin Ren was affiliated with Tsinghua University.

STOC '22, June 20–24, 2022, Rome, Italy

ACM ISBN 978-1-4503-9264-8/22/06.
https://doi.org/10.1145/3519935.3520051

ACM Reference Format:
Rahul Ilango, Hanlin Ren, and Rahul Santhanam. 2022. Robustness of Average-Case Meta-Complexity via Pseudorandomness. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing (STOC '22), June 20–24, 2022, Rome, Italy*. ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3519935.3520051

1 INTRODUCTION

In general, *meta-complexity* studies the complexity of computing various complexity measures, such as the Kolmogorov complexity of a string and the circuit complexity of a function given by its truth table. The study of these problems dates back decades [55], but in recent years there has been a surge of interest in the area, with several breakthrough results being shown. These include NP-hardness and ETH-hardness results for various complexity measures [24–26, 40], "hardness magnification" results showing that even weak lower bounds for some of these measures can lead to breakthrough complexity separations [13, 14, 43, 46, 48], and connections with learning [12, 47] and proof complexity [50]. Allender's 2020 survey [2] summarizes much of this work.

In this paper, we are interested in *average-case meta-complexity*, an area in which some of the most exciting recent progress has happened. For example, in a sequence of works [20–22], Hirahara uses meta-complexity to present new non-black-box worst-case to average-case reductions for problems in NP and PH, and several recent works [4, 38, 41, 52] characterize the existence of one-way functions based on the average-case hardness of meta-complexity problems on the uniform distribution.

We address two fundamental questions:

Question 1: Is there a characterization of one-way functions by some natural average-case complexity assumption about MCSP? (MCSP is the *minimum circuit size problem* [32]: given a Boolean function F represented as its truth table, decide whether F can be computed by a circuit of size at most s, where s is a complexity threshold parameter.)

Question 2: Is average-case meta-complexity *robust* in the sense that it does not depend critically on the parameters of the meta-complexity problem (such as the complexity threshold or approximation gap), the precise notion of average-case hardness, or even the precise complexity measure being computed?

Question 1 has been the focus of much work since the seminal paper of Kabanets and Cai [32], who build on the Natural Proofs framework [51] to show that MCSP is hard if one-way functions exist. Question 1 essentially asks if there is a converse to this fact. Santhanam [53] showed an equivalence between the zero-error average-case hardness of MCSP on the uniform distribution and the existence of one-way functions under a certain assumption on optimal pseudo-random generators, but there has been no progress on establishing that assumption, and we seek an unconditional result. The recent work of Liu and Pass [38] gives a characterization of one-way functions based on the average-case hardness of polynomial-time-bounded Kolmogorov complexity on the uniform distribution, but their technique does not generalize easily to other complexity measures such as circuit size. More recent follow-up works [4, 52] do give implications from average-case hardness assumptions on MCSP and its variants to the existence of one-way functions, but these implications are not known to be equivalences.

In this work, we give the first such (unconditional) equivalence.

Theorem 1.1 (Informal version of Theorem 2.1). *One-way functions exist if and only if there is a "locally samplable" distribution on truth tables on which "approximating" circuit complexity is intractable.*

We discuss the features of this theorem in more detail in the results section (in particular, the notion of locally samplable and the precise degree of approximation). However, we would like to specifically highlight that the techniques we use to prove Theorem 1.1 are very different from those used in previous connections between meta-complexity and one-way functions. Moreover, while the proof of Theorem 1.1 is technically involved, it has a quite intuitive high-level approach.

In fact, only after proving Theorem 1.1 did we realize that this simple high-level approach is actually *much more* broadly applicable. Indeed, the approach applies to almost any reasonable complexity measure on distributions that satisfies a "coding theorem" with respect to that complexity measure (in hindsight, the bulk of the proof of Theorem 1.1 is actually spent proving this coding condition).

Theorem 1.2 (Informal version of Theorem 2.2). *Let C be a "nice" complexity measure, and let D be an efficiently samplable distribution that has a "coding theorem" with respect to C. Then one-way functions exist if approximating C is hard on average on D.*

One way to interpret this result is as an intriguing converse to theorems showing that if one-way functions exist, then many meta-complexity problems are hard on average. Those results follow roughly because if one-way functions exist, one can generate low-complexity "pseudorandom" strings that are computationally indistinguishable from random strings, which have high complexity [3, 17, 19, 32, 51]. A major conceptual takeaway from this paper is that in many settings there is, in fact, a general method for going in the reverse direction, i.e., from the average-case hardness of approximating a complexity measure to the existence of one-way functions. Moreover, as we will discuss later, the high-level approach for doing this is surprisingly simple.

Theorem 1.2 allows us to make progress on the aforementioned Question 2. An unfortunate theme throughout theorems in meta-complexity is that results are often fragile with respect to the precise notion of meta-complexity. Proofs and results can depend on the precise notion of complexity (e.g. what gates are allowed in circuits for MCSP), or settings of a size threshold. To illustrate the importance of this "lack of robustness" consider that in recent years researchers have shown all of the following (informally stated):

- we know certain variants of MCSP and time-bounded Kolmogorov complexity are NP-complete [24–26, 40],
- we know certain variants of MCSP and time-bounded Kolmogorov complexity have worst-case to average-case reductions [20, 21, 53], and
- we know time-bounded Kolmogorov complexity is average-case hard if and only if one-way functions exist [38].[1]

Now, if all these results held for the same meta-complexity problem, then we could compose them to get an amazing implication: that the worst-case hardness of NP is equivalent to the existence of one-way functions[2]! This would solve a major open question in the foundations of cryptography, and rule out two of Impagliazzo's five possible worlds of average-case complexity: Heuristica and Pessiland [28]. Thus there is a strong motivation to prove "more robust" versions of the results mentioned above.

Making progress on the aforementioned robustness issue, we use Theorem 1.2 to prove a broad equivalence on the average-case complexity of meta-complexity problems through connections to one-way functions. These problems include approximation versions of MCSP, (unrestricted) Kolmogorov complexity, and time-bounded Kolmogorov complexity. (For technical reasons, in order to include time-bounded Kolmogorov complexity in our equivalence, we need to switch to the infinitely-often setting and assume complexity-theoretic derandomization.)

Theorem 1.3 (Informal version). *All of the following are equivalent.*

1. *Infinitely-often one-way functions exist.*
2. *There exists a locally samplable distribution on which "approximating" circuit complexity is intractable infinitely often.*
3. *There exists a samplable distribution on which "approximating" K-complexity is intractable infinitely often.*
4. *There exists a samplable distribution on which "approximating" $\mathsf{K}^{\mathrm{poly}(n)}$-complexity is intractable infinitely often (assuming a standard derandomization assumption, namely that linear exponential time requires linear-exponential-sized circuits).*

Perhaps the main feature of this theorem is the *robustness* of these equivalences. Although we have stated these results informally, they hold for a wide range of complexity thresholds, approximation gaps, and average-case error tolerances (i.e., for weak average-case hardness and strong average-case hardness). We discuss this in greater detail in Section 2.

Finally, we specifically highlight the equivalence between Item 1 and Item 3, as we find it to be both surprising and elegant: one can characterize the existence of cryptography by the average-case complexity of approximating Kolmogorov complexity (a worst-case uncomputable problem!).

Theorem 1.4 (Equivalence between Item 1 and Item 3 restated). *One-way functions exist if and only if there exists a samplable distribution on which "approximating" K-complexity is intractable.*[3]

[1]Unlike the previous two items, this was only known for variants of time-bounded Kolmogorov complexity and not for variants of MCSP.

[2]In order to actually get this implication, we would need the same notion of average-case hardness in the second and third item. Current results give zero-error average-case hardness for the second item, but require bounded-error average-case hardness for the third.

[3]This theorem holds in both infinitely-often and almost-everywhere settings and does not need the complexity-theoretic derandomization assumption.

Perhaps surprisingly, the proof of this result is quite simple and relies only on basic facts about Kolmogorov complexity and one-way functions. To our knowledge, however, this result was never shown previously (we later discuss some possible reasons why in Section 3.1). Indeed, we view the simplicity of the proof as a significant feature of this result.

We also note that this gives the first characterization of the existence of one-way functions by the average-case hardness of a problem on *any* samplable distribution. Previous characterizations such as Levin's universal one-way function [36, 37] and the results of Liu and Pass [38] require average-case hardness on a *specific* distribution in order to infer the existence of one-way functions.

In summary, we make three main contributions:

(1) Giving the first unconditional characterization of the existence of one-way functions based on MCSP.
(2) Showing the robustness of average-case meta-complexity with respect to many different parameters (e.g. size thresholds, approximation gaps, error tolerances) and notions (e.g. Kolmogorov complexity, time-bounded Kolmogorov complexity, and circuit size).
(3) Giving a simple and elegant characterization of one-way functions based on (unrestricted) Kolmogorov complexity. Moreover, this is the first characterization of the existence of one-way functions by the average-case hardness of a problem on any samplable distribution.

In the next section, we describe our results more formally and mention some additional applications to cryptography and meta-complexity as well.

2 OUR RESULTS

2.1 Equivalences

We give the first equivalence between the average-case hardness of MCSP and the existence of one-way functions. In order to do this, we consider a samplability notion that we believe is interesting in its own right: *local samplability*[4]. A t-local sampler is a sampler that, in order to compute a given bit of its output, runs in time t with random access to its input[5]. Here t is typically some sub-linear function. Several natural distributions, such as the uniform distribution and distributions induced by pseudorandom function generators [17], are t-locally samplable for small t. Further motivating the notion is the fact that for most pairs L, L' of natural NP-complete problems, there is a *local* reduction from L to L', and hence the hardness of L with respect to some locally samplable distribution translates to the hardness of L' also with respect to some locally samplable distribution.

In what follows, let GapMCSP$[s, c]$ denote the promise problem of distinguishing between truth tables of Boolean functions with circuit complexity at most s and Boolean functions with circuit complexity at least c. We say that a problem is *weakly average-case hard* on a distribution $\mathcal{D}$ if every probabilistic polynomial-time algorithm fails to solve it with probability $1/n^{O(1)}$ (over $\mathcal{D}$ and the randomness of the algorithm) on almost all input lengths, and we say that a problem is *strongly average-case hard* on a distribution $\mathcal{D}$ if every probabilistic polynomial-time algorithm fails to solve it with probability $1/2 - 1/n^{\omega(1)}$ on almost all input lengths.

Theorem 2.1. *The following are equivalent:*

(1) One-way functions exist.
(2) For some constant $\delta > 0$ and $s = \Omega(n^\delta)$, there is an (n^δ)-locally samplable distribution $\mathcal{D}$ such that GapMCSP$[s, sn^{5\delta}]$ *is weakly average-case hard on $\mathcal{D}$.*
(3) For every constant $\delta > 0$, there is an (n^δ)-locally samplable distribution $\mathcal{D}$ such that GapMCSP$[n^\delta, o(\frac{n}{\log n})]$ *is strongly average-case hard on $\mathcal{D}$.*

Theorem 2.1 shows that in the setting of average-case complexity with respect to locally samplable distributions, GapMCSP is very robust with respect to the complexity parameters s, c as well as the error tolerance (i.e., weak or strong average-case hardness).

Only after proving Theorem 2.1 did we realize that the high-level approach works for not only MCSP but also many different meta-complexity problems. The essential condition we need is a *coding theorem* [34]: if there is an efficiently samplable distribution that generates a string x w.p. at least p, then the complexity of x is at most $\log(1/p) + O(\log |x|)$. The notion of "efficiently samplable" and "complexity" may vary here; for every definition of them that satisfies the coding theorem (with some mild restrictions on the definition of "complexity"), hardness of approximating the "complexity" of a string under an "efficiently samplable" distribution imply one-way functions.

Theorem 2.2. *Let C be a "nice" complexity measure*[6] *and S be a class of polynomial-time samplable distributions. Suppose the following coding theorem holds:*

- *For every string $x \in \{0, 1\}^n$ that is sampled with probability p from a distribution $D \in S$, we have $C(x) \le \log(1/p) + O(\log n)$.*

Let $\Delta = \omega(\log n)$ be the approximation gap. If there is a distribution in S on which it is weakly average-case hard to approximate C within an additive factor of Δ, then one-way functions exist.

For example, we will prove a coding theorem where "complexity" is interpreted as the circuit complexity of a truth table and "efficiently samplable" means locally samplable.[7] Given the coding theorem, the implication from the second item to the first item in Theorem 2.1 essentially follows from the machinery of Theorem 2.2.

Building on Theorem 2.2 and the well-known coding theorem for Kolmogorov complexity, we give a surprising characterization of the existence of one-way functions by the average-case hardness of *unbounded* Kolmogorov complexity. Below, let GapK$[s, c]$ denote the promise problem of distinguishing between strings of Kolmogorov complexity at most s and of Kolmogorov complexity at least c.

Theorem 2.3. *The following are equivalent:*

(1) One-way functions exist.

[4]Indeed, the notion of local samplability has been studied previously. See for example [15].
[5]Consequently, a t-local sampler can access at most t bits of its input while computing a specific output bit.
[6]See the full version for the exact definition of "nice" complexity measures.
[7]Actually, since the coding theorem for MCSP is not tight and circuit complexity is not "nice enough", we only show the existence of one-way functions assuming hardness of GapMCSP under a *multiplicative* gap, instead of an *additive* gap as in Theorem 2.2.

(2) For some $s = n^{\Omega(1)}$ and $\Delta = \omega(\log(n))$, there is a samplable distribution $\mathcal{D}$ such that $\mathsf{GapK}[s, s+\Delta]$ *is weakly average-case hard on $\mathcal{D}$.*
(3) For every $\epsilon > 0$, there is a samplable distribution $\mathcal{D}$ such that $\mathsf{GapK}[n^{\epsilon}, n - \omega(\log(n))]$ *is strongly average-case hard on $\mathcal{D}$.*

The fact that the average-case hardness of approximating Kolmogorov complexity implies the existence of one-way functions might seem especially surprising, given that candidate one-way functions are usually defined based on problems in NP, while Kolmogorov complexity is uncomputable. While most constructions of one-way functions based on the average-case hardness of some computational problem use the structure of the problem to define the one-way function and argue security based on the distributional average-case hardness, our construction does the reverse[8]: the one-way function is defined based on the distribution, while the proof of security exploits the structure of the problem assumed to be hard, i.e., Kolmogorov complexity.

A natural question is whether there is a version of Theorem 2.3 where the equivalence involves the hardness of a meta-complexity problem known to be in NP, such as the K^t problem considered in [38].[9] We are able to achieve this with a more complicated proof, but only under a complexity-theoretic derandomization assumption and for infinitely-often one-way functions.

Theorem 2.4. *Assume that* $\mathsf{E} \notin \mathsf{ioSIZE}[2^{.01n}]$. *The following are equivalent:*

(1) Infinitely-often one-way functions exist.
(2) For some $s = n^{\Omega(1)}$ and $\Delta = \omega(\log(n))$, there is a samplable distribution $\mathcal{D}$ such that for every large enough polynomial τ, $\mathsf{GapK}^{\tau}[s, s+\Delta]$ *is weakly average-case hard on $\mathcal{D}$ on infinitely many input lengths.*
(3) For every $\epsilon > 0$, there is a samplable distribution $\mathcal{D}$ such that for every large enough polynomial τ, $\mathsf{GapK}^{\tau}[n^{\epsilon}, n - \omega(\log(n))]$ *is strongly average-case hard on $\mathcal{D}$ on infinitely many input lengths.*

As can be seen from the statement of Theorems 2.1, 2.3 and 2.4, our results are extremely robust with respect to the parameters of the underlying meta-complexity problem. For example, it follows from Theorem 2.3 that the average-case complexity of $\mathsf{GapK}[s, c]$ under samplable distributions remain the same regardless of the complexity threshold (s could be any polynomial), the gap parameter ($c - s$ could be as small as $\log^2 n$ or as large as $n - n^{0.1}$), and the error tolerance (weak or strong average-case hardness).

Finally, by strengthening the samplability assumption in Theorem 2.3 and using the influential localization technique of [9], we can also characterize one-way functions computable in NC^0.[10]

Theorem 2.5. *The following are equivalent:*

(1) There are one-way functions computable in NC^0.
(2) For some $s = n^{\Omega(1)}$ and $\Delta = \omega(\log(n))$, there is a logspace-samplable distribution $\mathcal{D}$ such that $\mathsf{GapK}[s, s+\Delta]$ *is weakly average-case hard on $\mathcal{D}$.*
(3) There is an NC^0*-samplable distribution $\mathcal{D}$ such that* $\mathsf{GapK}[n - n^{0.99}, n - \omega(\log(n))]$ *is strongly average-case hard on $\mathcal{D}$.*

2.2 Applications

Our new connections between the hardness of meta-complexity problems on samplable distributions and the existence of one-way functions have several applications in cryptography and the theory of meta-complexity.

One-way functions from hardness of SAT *and* Clique. First, they allow us to show that one-way functions follow from the average-case hardness of NP-complete problems such as SAT and Clique under more general assumptions than were known before. Specifically, average-case hardness of SAT or Clique on *any* samplable distribution of *high enough entropy* implies the existence of one-way functions. Candidate one-way functions based on SAT and Clique are often based on very specific distributions, which lead to hardness assumptions that are not very robust. By showing that one-wayness can be derived from more general classes of distributions, we make progress towards basing one-way functions simply on the average-case hardness of NP.

Below, the entropy *deficiency* of a distribution on m bits is the difference between m and the entropy.

Theorem 2.6. *Given an integer k, let $\Delta \geq 2^{k+3}$ be a large enough integer, and let $t : \mathbb{N} \to \mathbb{N}$ be any function such that $t(n) = \omega(\log n)$. If*

- k-SAT *on Δn clauses is strongly average-case hard w.r.t. some samplable (resp. logspace-samplable) distribution $\mathcal{D}$ with entropy deficiency at most $\Delta n / 2^{k+1}$, or*
- t-Clique *is strongly average-case hard w.r.t. some samplable (resp. logspace-samplable) distribution $\mathcal{D}$ with entropy deficiency at most $0.99\binom{t}{2}$,*

then one-way functions (resp. one-way functions computable in NC^0*) exist.*

It is natural to wonder if the hardness assumptions in Theorem 2.6 are reasonable. We show that in fact, the hardness assumption for Clique follows from the well-studied Planted Clique Hypothesis [5, 31, 33], while the hardness assumption for SAT follows from pseudo-randomness of random local functions (often referred to as"Goldreich's PRG") [7, 8, 16]. Thus our assumptions generalize hypotheses that have been intensively studied.

Unconditional CZK *protocols for meta-complexity.* Turning to the theory of meta-complexity, our characterizations imply *unconditional* average-case simulations of the corresponding meta-complexity problems in CZK (Computational Zero Knowledge) infinitely often. As far as we are aware, these are the first natural examples of approximation problems shown to be *unconditionally* in CZK (on average) without also being shown to be in SZK (Statistical Zero Knowledge).

Below, we say that a problem is infinitely often in CZK on a distribution $\mathcal{D}$ if for each $k > 0$, there is a CZK protocol that is

[8] There are precedents for this, such as [29]. In [29], they prove that if NP is hard over some samplable distribution, then NP is also hard over the uniform distribution. One of the cases they consider is that the sampler itself already implements a one-way function; in this case, a hard NP language over the uniform distribution follows easily.
[9] Intuitively K^t is an easier problem than K, therefore basing one-way functions on hardness of K^t should be *easier* than basing one-way functions on hardness of K. This is not the case, though, for the proof techniques we use.
[10] NC^0 is the class of (multi-output) functions computable by uniform circuits such that each output bit is connected to a constant number of input bits. It was proved in [9] that logspace-computable one-way functions exist if and only if NC^0-computable one-way functions exist.

correct with probability at least $1 - 1/n^k$ on inputs sampled from $\mathcal{D}$, for infinitely many n.

Theorem 2.7. *For every* $s : \mathbb{N} \to \mathbb{N}$ *such that* $s(n) = n^{\Omega(1)}$ *and for every samplable distribution* $\mathcal{D}$*,* $\mathsf{GapK}[s, s + \omega(\log(n))]$ *is infinitely often in* CZK *on* $\mathcal{D}$*.*

For every $\delta > 0$*,* $s = \Omega(n^\delta)$*, and* (n^δ)*-locally samplable distribution* $\mathcal{D}$*,* $\mathsf{GapMCSP}[s, sn^{5\delta}]$ *is infinitely often in* CZK *on* $\mathcal{D}$*.*

We note that the proof of Theorem 2.7 builds on two common themes in previous work on CZK: connections with one-way functions and win-win arguments (see e.g. [49, 56]).

*Non-*NP*-hardness of* GapMCSP *under randomized local reductions.* Finally, we use our results to shed some light on the long-standing open question of whether MCSP is NP-complete. Based on the assumption that one-way functions in NC^0 exist, we rule out NP-hardness of GapMCSP under randomized local reductions. To the best of our knowledge, this is the first piece of evidence against *randomized* reductions from SAT to GapMCSP. We note that Murray and Williams [44] unconditionally ruled out NP-hardness of MCSP under deterministic local reductions.

Theorem 2.8. *Suppose there are one-way functions computable in* NC^0*. Then for each* $\delta > 0$ *and* $s = \Omega(n^\delta)$*, there are no randomized* (n^δ)*-local reductions from* SAT *to* $\mathsf{GapMCSP}[s, sn^{4\delta}]$*.*

3 TECHNIQUES

Here we discuss the main ideas used in our proofs. We restrict ourselves to high-level arguments in this section, and do not delve too deeply into the choice of parameters.

3.1 Meta-Complexity and One-Way Functions

In this subsection, we outline how to prove the equivalences between the average-case hardness of approximating various complexity measures and the existence of one-way functions.

Reverse Directions: Average-Case Hardness from One-Way Functions. The reverse directions of these equivalences are relatively straightforward given previous work. Suppose one-way functions exist, and we wish to show that, for example, GapK and GapMCSP are strongly hard on average. By [17, 19], for each $\epsilon > 0$ there are pseudo-random generators with seed length n^ϵ computable in polynomial time such that each output of the generator, when interpreted as the truth table of a function, has circuit size $n^{O(\epsilon)}$. This also implies that every output of the generator has Kolmogorov complexity at most $n^{O(\epsilon)}$. On the other hand, a random string x has $\mathsf{K}(x)$ close to n and circuit size close to $n/\log(n)$ with high probability. Thus, we can consider the samplable distribution $\mathcal{D}$ that generates a uniformly random string with probability $1/2$ and a uniformly random output of the pseudo-random generator with probability $1/2$. Any algorithm for GapK or GapMCSP that has a noticeable advantage over random could be used to distinguish the uniform distribution from the pseudo-random distribution, contradicting the pseudo-randomness assumption. In order to get average-case hardness on NC^0-samplable distributions from NC^0 one-way functions, we use [18] rather than [19] to build an NC^0-computable pseudo-random generator from the NC^0-computable one-way function.

Forward Directions: One-Way Functions from Average-Case Hardness. Our main technical contribution towards proving these equivalences is giving a generic approach for showing how the average-case hardness of approximating a complexity measure on specific distributions can imply the existence of one-way functions. Let $C : \{0,1\}^\star \to \mathbb{N}$ be a complexity measure. Let D_n be some efficiently samplable distribution on n-bit strings. Let $t \in \mathbb{N}$ be some threshold. We work in the contrapositive, i.e., we show that if no one-way functions exist, then one can efficiently distinguish whether a string has C-complexity at most t or C-complexity much larger than t on average (with two-sided error) over D_n. To do this we introduce some notation: for a string $y \in \{0,1\}^n$, let p_y denote the probability that y is sampled from D_n.

Our framework works by showing the following.

(1) **If p_y is low, then y has high complexity (on average over D_n).** By a union bound, if we sample y from D_n, the probability that $p_y \le q$ and y has complexity at most k is at most
$$|\{y \in \{0,1\}^n : C(y) \le k\}| \cdot q.$$
Assuming our complexity measure C has the property that the number of low-complexity strings is relatively small, the above quantity will be small, if q is also small. Consequently, we get that when p_y is low, y has high complexity on average.

(2) **Coding Theorem: If p_y is high, then y has low complexity.** Intuitively, if p_y is large, then to describe y (in the information-theoretic setting), one should need roughly $\log(1/p_y)$ bits. We say a coding theorem holds for a complexity measure C and a distribution D_n if for all y, we have that $C(y)$ is upper bounded by roughly $\log(1/p_y)$. Assuming we have a coding theorem (which is a non-trivial task in of itself, especially if the complexity measure is resource-constrained!), we get that if p_y is small, then y has low complexity.

(3) **If one can approximate p_y, then one can approximate C-complexity on D_n on average.** Combining (1) and (2), we get that if we are able to approximate p_y given y on average over D_n, then we can approximate C-complexity on D_n on average: simply output "high complexity" if p_y is low and "low-complexity" if p_y is high.

(4) **If one-way functions do not exist, one can approximate p_y.** Because the distribution D_n is efficiently samplable by some algorithm A, one can use the non-existence of one-way functions and hashing ideas to approximately count the number of pre-images of y under A [29, 30]. This gives an efficient approximation of p_y on average over D_n.

(5) **If one-way functions do not exist, one can approximate C-complexity on D_n on average.** This is by combining (3) and (4).

We now emphasize which parts of the above argument depend on the choice of complexity measure and distribution. Parts (3) and (5) hold as long as the remaining parts hold. For part (1), we need that the complexity measure C has relatively few low-complexity strings (a bound of the form 2^k or even $2^{k \log k}$ on the total number of strings of complexity k is fine for us). This condition seems to hold for all natural complexity measures. Next, part (4) of the argument holds as long as D_n is polynomial-time samplable.

Finally, part (2) is the most delicate part to prove, and the one we need to work the hardest to show throughout this paper. In the case of Kolmogorov complexity, it is relatively easy to show that if D_n is an efficiently samplable distribution, then $C(y) \leq \log(1/p_y) + O(\log n)$. However, in general, showing a coding theorem for a class of distributions and a specific complexity measure can be difficult. Showing Coding Theorems. As mentioned previously, in the case of (unrestricted) Kolmogorov complexity, showing the corresponding coding theorem is straightforward. However, in both the case of MCSP and time-bounded Kolmogorov complexity we need to work harder.

For time-bounded Kolmogorov complexity, we were not able to get an *unconditional* coding theorem for samplable distributions. Instead, we prove an *average-case* coding theorem under the assumptions that *one-way functions do not exist* and *complexity-theoretic derandomization holds*. Luckily, this suffices for our purposes. The details are somewhat intricate, but the main idea (at a high level) is this: suppose y is a high complexity string and, for contradiction, p_y is large. To achieve a contradiction, we want to come up with a small, efficient description of y. We will do this by specifying a small hash v of y such that this hash is unique among the (not too many) strings z such that p_z is large. We then use the non-existence of one-way functions, in two different ways[11], and a derandomization assumption to show that given v, one can deterministically recover y (on average).

For MCSP, we show an (unconditional, worst-case) coding theorem on *locally-samplable* distributions.[12] In particular, we show the contrapositive, that is, we show that if y is the truth table of a function with high circuit complexity, then p_y is low, where p_y is the probability that y is sampled from $\mathcal{D}$. The key idea is to "reveal" bits of the input to the sampler used to compute y in stages. Each stage reveals a small number of bits of the input to the sampler, by its locality. If after a small number of stages, all bits of y can be correctly computed by an approximate majority over random choices of the unrevealed bits, then we can argue that we get a small circuit for y. Suppose this is not the case. Then there is some bit of y for which random choices to the unrevealed bits give the wrong answer with probability $\geq 1/3$. In this case, we can argue that p_y must decrease by a factor of $2/3$. If y has large circuit complexity, the number of stages in this process must be high, and hence p_y must be low.

Why were these results not shown earlier? Looking at the simplicity of our high-level approach for going from approximating average-case hardness of complexity measures to the existence of one-way functions, it is natural to ask why these results were not shown earlier. Indeed, all the pseudorandomness machinery we use was developed in the early 90s, and developing stronger connections between one-way functions and meta-complexity has been a longstanding question. It is especially surprising that there is such a simple proof showing that the existence of one-way functions is equivalent to the average-case hardness of approximating (unrestricted) Kolmogorov complexity.

We suggest some heuristic reasons for why these results took so long to discover. For one, it seems somewhat hard to believe that there could be a relationship between an uncomputable problem and the existence of (efficiently computable) one-way functions. Indeed, for Kolmogorov complexity to even be computable one needs to consider the simultaneous restriction to approximation and two-sided error on samplable distributions.

Another contributing factor is that, counter-intuitively, it seems *easier* to prove an equivalence via our framework when the complexity measure is *more powerful* (i.e. like Kolmogorov complexity), since it is easier to prove coding theorems in this setting. As a result, while one might think one is starting with "simpler cases" like MCSP, in fact, those cases are more difficult to work with in our framework.

Finally, exciting recent positive results in this area, especially that of Liu and Pass [38], has led the community to look more closely at connections between Kolmogorov complexity (and its variants) and cryptography.

3.2 Applications

Theorem 2.6: One-way functions from hardness of SAT *and* Clique. Our proof of Theorem 2.6 is inspired by a zero-error average-case reduction from SAT to computing KT complexity[13] in [23]. The idea is that random k-CNF formulas are incompressible, while k-CNFs with satisfying assignments can be compressed if they are long enough. A similar idea gives a zero-error reduction from Clique to computing KT.

Here we adapt these ideas to the bounded-error average-case setting. When considering bounded-error average-case complexity, it is no longer the case that the uniform distribution is a reasonable one to consider for k-SAT, since answering "Unsatisfiable" works with overwhelmingly high probability. However, it is still reasonable to expect average-case hardness on distributions with high entropy. We show that if the distribution has high enough entropy, then there are bounded-error reductions from k-SAT and Clique to approximating Kolmogorov complexity. The reductions themselves are the simplest possible, namely the identity reduction! However, the proof that they work requires the compressibility argument from [23] as well as the fact that high entropy distributions must place noticeable probability on strings of high Kolmogorov complexity. We thereby get a reduction from computing SAT or Clique on average with noticeable advantage over random on a samplable distribution $\mathcal{D}$ with high enough entropy to computing GapK with all but inverse polynomial probability on $\mathcal{D}$. The robustness of GapK on average is crucial to our argument, as the reduction needs the algorithm for GapK to be correct w.p. $1 - 1/\text{poly}(n)$.

To show that our average-case assumptions are reasonable, we show that they are implied by well-studied hardness assumptions in average-case complexity and cryptography, namely the Planted Clique Hypothesis for Clique and the pseudorandomness of random local functions for k-SAT.

[11] The reason why the K^t result only holds infinitely often is that we need to invert one-way functions twice, and we need to make sure the input lengths where our two inverters succeed line up.

[12] We know that MCSP is hard to approximate even on locally-samplable distributions if one-way functions exist, so for the purpose of proving an equivalence, it suffices to consider locally-samplable distributions.

[13] KT complexity is a meta-complexity notion defined in [1] that is closely related to circuit complexity.

Theorem 2.7: Unconditional CZK *protocols for meta-complexity.* The proof of Theorem 2.7 uses a win-win argument:

- It is well known that if one-way functions exist, then CZK = IP = PSPACE [10, 54]. In this case, since p_y can be computed in polynomial space for any string y sampled from the distribution $\mathcal{D}$, we have that GapK is in CZK on average.
- Suppose, on the other hand, that one-way functions don't exist. Then by Theorem 2.3, GapK is infinitely often in probabilistic polynomial time on $\mathcal{D}$. Since CZK trivially contains probabilistic polynomial time, GapK is infinitely often in CZK on $\mathcal{D}$ in this case as well.

A similar argument works for GapMCSP on locally samplable distributions, using Theorem 2.1 instead of Theorem 2.3.

*Theorem 2.8: Non-*NP*-hardness of* GapMCSP *under randomized local reductions.* Finally, to prove Theorem 2.8, we first show that if a language L has randomized local reductions to GapMCSP, then L is easy on average over locally samplable distributions. The main ingredient of this proof is showing that GapMCSP is easy on average over a locally samplable distribution when given the randomness of the sampler, rather than just its output. This argument is similar to the argument that p_y is low for strings y of high circuit complexity sampled by a local sampler $\mathcal{D}$. We observe that under the assumption that there are one-way functions in NC^0, k-SAT is average-case hard on some locally samplable distribution, and combining this with the lemma about randomized local reductions concludes the proof.

4 RELATED WORK

There have been several works relating one-way functions to non-cryptographic notions. Impagliazzo and Levin [29] show that one-way functions exist if and only if "universal extrapolation" does not, where universal extrapolation is a generic procedure to sample from continuations of the output of some samplable process. Some of the ideas we use are similar to theirs, though there does not seem to be a formal connection between the results. Blum, Furst, Kearns, and Lipton [11] relate the existence of one-way functions to an average-case notion of learning. Oliveira and Santhanam [47] show that exponentially hard (non-uniform) one-way functions exist if and only if non-trivial (non-uniform) learning is hard.

More recently, there has been a number of papers considering the average-case hardness of meta-complexity problems on the uniform distribution and relating it to one-way functions. Santhanam [53] showed that under an assumption on universal succinct pseudorandom distributions, MCSP is zero-error hard on average on the uniform distribution if and only if one-way functions exist. By considering K^t rather than MCSP and bounded-error hardness rather than zero-error hardness, Liu and Pass [38] gave an amazing unconditional equivalence. Characterizations of NC^0 cryptography by meta-complexity over the uniform distribution are given in [41, 52]. An implication for one-way functions from the average-case hardness of the conditional KT-complexity problem is given in [4]. [40] give a natural NP-complete problem whose average-case hardness on the uniform distribution is equivalent to the existence of one-way functions.

This paper is an updated version of [27], which appeared with a different title and slightly different exposition. Soon after [27] appeared online, Liu and Pass [39] published a note where they observed that the argument in the proof of Theorem 2.6 generalizes to give one-way functions from the average-case hardness of any sparse enough language with respect to a distribution of high enough entropy. This is a very interesting observation. However, we would like to stress their result only generalizes Theorem 2.6 and does not seem applicable to most of our other results. In particular, the key phenomenon behind our results is not just sparsity, but rather the interplay between complexity measures and coding theorems for distributions as captured by our Theorem 2.2.

Indeed, a major part of our work is spent proving coding theorems for various complexity measures and distributions. We remark that Antunes and Fortnow [6, Lemma 3.2] also proved a coding theorem for K^t under complexity-theoretic derandomization assumptions. Compared with their results, our coding theorem uses a weaker derandomization assumption[14], but our coding theorem also requires the non-existence of infinitely-often one-way functions and only works on the average case. Recently, Lu and Oliveira [42] showed a coding theorem for rKt (a randomized version of Levin's Kt complexity [35, 45]).

ACKNOWLEDGMENTS

We thank Igor Carboni Oliveira for collaborating on this project at an early stage, and Shuichi Hirahara and Ryan Williams for useful discussions. Hanlin Ren is supported in part by the Zhongguancun Haihua Institute for Frontier Information Technology.

REFERENCES

[1] Eric Allender. 2001. When Worlds Collide: Derandomization, Lower Bounds, and Kolmogorov Complexity. In *FST TCS 2001: Foundations of Software Technology and Theoretical Computer Science, 21st Conference, Bangalore, India, December 13-15, 2001, Proceedings (Lecture Notes in Computer Science, Vol. 2245)*. Springer, 1–15. https://doi.org/10.1007/3-540-45294-X_1

[2] Eric Allender. 2020. The New Complexity Landscape Around Circuit Minimization. In *Language and Automata Theory and Applications - 14th International Conference, LATA 2020, Milan, Italy, March 4-6, 2020, Proceedings (Lecture Notes in Computer Science, Vol. 12038)*. Springer, 3–16.

[3] Eric Allender, Harry Buhrman, Michal Koucký, Dieter van Melkebeek, and Detlef Ronneburger. 2006. Power from Random Strings. *SIAM J. Comput.* 35, 6 (2006), 1467–1493. https://doi.org/10.1137/050628994

[4] Eric Allender, Mahdi Cheraghchi, Dimitrios Myrisiotis, Harsha Tirumala, and Ilya Volkovich. 2021. One-Way Functions and a Conditional Variant of MKTP. In *41st IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2021, December 15-17, 2021, Virtual Conference (LIPIcs, Vol. 213)*, Mikolaj Bojanczyk and Chandra Chekuri (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 7:1–7:19. https://doi.org/10.4230/LIPIcs.FSTTCS.2021.7

[5] Noga Alon, Michael Krivelevich, and Benny Sudakov. 1998. Finding a Large Hidden Clique in a Random Graph. In *Proceedings of the Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, 25-27 January 1998, San Francisco, California, USA*. ACM/SIAM, 594–598. http://dl.acm.org/citation.cfm?id=314613.315014

[6] Luis Filipe Coelho Antunes and Lance Fortnow. 2009. Worst-Case Running Times for Average-Case Algorithms. In *Proceedings of the 24th Annual IEEE Conference on Computational Complexity, CCC 2009, Paris, France, 15-18 July 2009*. IEEE Computer Society, 298–303. https://doi.org/10.1109/CCC.2009.12

[7] Benny Applebaum. 2013. Pseudorandom Generators with Long Stretch and Low Locality from Random Local One-Way Functions. *SIAM J. Comput.* 42, 5 (2013), 2008–2037. https://doi.org/10.1137/120884857

[8] Benny Applebaum. 2016. Cryptographic Hardness of Random Local Functions - Survey. *Comput. Complex.* 25, 3 (2016), 667–722. https://doi.org/10.1007/s00037-015-0121-8

[9] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. 2006. Cryptography in NC^0. *SIAM J. Comput.* 36, 4 (2006), 845–888. https://doi.org/10.1137/S0097539705446950

[14] We only assume that linear exponential time requires exponential-size circuits, while Antunes and Fortnow assumed that linear exponential time requires exponential-size circuits *even with* Σ_2^p *oracle gates*.

[10] Michael Ben-Or, Oded Goldreich, Shafi Goldwasser, Johan Håstad, Joe Kilian, Silvio Micali, and Phillip Rogaway. 1988. Everything Provable is Provable in Zero-Knowledge. In *CRYPTO '88 (Lecture Notes in Computer Science, Vol. 403)*. Springer, 37–56. https://doi.org/10.1007/0-387-34799-2_4

[11] Avrim Blum, Merrick L. Furst, Michael J. Kearns, and Richard J. Lipton. 1993. Cryptographic Primitives Based on Hard Learning Problems. In *Advances in Cryptology - CRYPTO '93, 13th Annual International Cryptology Conference, Santa Barbara, California, USA, August 22-26, 1993, Proceedings (Lecture Notes in Computer Science, Vol. 773)*. Springer, 278–291. https://doi.org/10.1007/3-540-48329-2_24

[12] Marco L. Carmosino, Russell Impagliazzo, Valentine Kabanets, and Antonina Kolokolova. 2016. Learning Algorithms from Natural Proofs. In *31st Conference on Computational Complexity, CCC 2016, May 29 to June 1, 2016, Tokyo, Japan (LIPIcs, Vol. 50)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 10:1–10:24. https://doi.org/10.4230/LIPIcs.CCC.2016.10

[13] Lijie Chen, Ce Jin, and R. Ryan Williams. 2019. Hardness Magnification for all Sparse NP Languages. In *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, David Zuckerman (Ed.). IEEE Computer Society, 1240–1255. https://doi.org/10.1109/FOCS.2019.00077

[14] Lijie Chen, Ce Jin, and R. Ryan Williams. 2020. Sharp threshold results for computational complexity. In *Proccedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy (Eds.). ACM, 1335–1348. https://doi.org/10.1145/3357713.3384283

[15] Anindya De and Thomas Watson. 2012. Extractors and Lower Bounds for Locally Samplable Sources. *ACM Trans. Comput. Theory* 4, 1, Article 3 (March 2012), 21 pages. https://doi.org/10.1145/2141938.2141941

[16] Oded Goldreich. 2000. Candidate One-Way Functions Based on Expander Graphs. *Electron. Colloquium Comput. Complex.* 7, 90 (2000). http://eccc.hpi-web.de/eccc-reports/2000/TR00-090/index.html

[17] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. 1986. How to construct random functions. *J. ACM* 33, 4 (1986), 792–807. https://doi.org/10.1145/6490.6503

[18] Iftach Haitner, Omer Reingold, and Salil P. Vadhan. 2013. Efficiency Improvements in Constructing Pseudorandom Generators from One-Way Functions. *SIAM J. Comput.* 42, 3 (2013), 1405–1430. https://doi.org/10.1137/100814421

[19] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. 1999. A Pseudorandom Generator from any One-way Function. *SIAM J. Comput.* 28, 4 (1999), 1364–1396. https://doi.org/10.1137/S0097539793244708

[20] Shuichi Hirahara. 2018. Non-Black-Box Worst-Case to Average-Case Reductions within NP. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*. IEEE Computer Society, 247–258. https://doi.org/10.1109/FOCS.2018.00032

[21] Shuichi Hirahara. 2020. Characterizing Average-Case Complexity of PH by Worst-Case Meta-Complexity. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*. IEEE, 50–60. https://doi.org/10.1109/FOCS46700.2020.00014

[22] Shuichi Hirahara. 2021. Average-case hardness of NP from exponential worst-case hardness assumptions. In *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, Samir Khuller and Virginia Vassilevska Williams (Eds.). ACM, 292–302. https://doi.org/10.1145/3406325.3451065

[23] Shuichi Hirahara and Rahul Santhanam. 2017. On the Average-Case Complexity of MCSP and Its Variants. In *32nd Computational Complexity Conference, CCC 2017, July 6-9, 2017, Riga, Latvia (LIPIcs, Vol. 79)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 7:1–7:20. https://doi.org/10.4230/LIPIcs.CCC.2017.7

[24] Rahul Ilango. 2020. Approaching MCSP from Above and Below: Hardness for a Conditional Variant and $AC^0[p]$. In *11th Innovations in Theoretical Computer Science Conference, ITCS 2020, January 12-14, 2020, Seattle, Washington, USA (LIPIcs, Vol. 151)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 34:1–34:26. https://doi.org/10.4230/LIPIcs.ITCS.2020.34

[25] Rahul Ilango. 2020. Constant Depth Formula and Partial Function Versions of MCSP are Hard. In *61st IEEE Annual Symposium on Foundations of Computer Science*. IEEE, 424–433.

[26] Rahul Ilango, Bruno Loff, and Igor Carboni Oliveira. 2020. NP-Hardness of Circuit Minimization for Multi-Output Functions. In *35th Computational Complexity Conference, CCC 2020, July 28-31, 2020, Saarbrücken, Germany (Virtual Conference) (LIPIcs, Vol. 169)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 22:1–22:36. https://doi.org/10.4230/LIPIcs.CCC.2020.22

[27] Rahul Ilango, Hanlin Ren, and Rahul Santhanam. 2021. Hardness on any Samplable Distribution Suffices: New Characterizations of One-Way Functions by Meta-Complexity. *Electron. Colloquium Comput. Complex.* (2021), 82. https://eccc.weizmann.ac.il/report/2021/082

[28] Russell Impagliazzo. 1995. A Personal View of Average-Case Complexity. In *Proceedings of the Tenth Annual Structure in Complexity Theory Conference, Minneapolis, Minnesota, USA, June 19-22, 1995*. IEEE Computer Society, 134–147.

[29] Russell Impagliazzo and Leonid A. Levin. 1990. No Better Ways to Generate Hard NP Instances than Picking Uniformly at Random. In *31st Annual Symposium on Foundations of Computer Science, St. Louis, Missouri, USA, October 22-24, 1990, Volume II*. IEEE Computer Society, 812–821. https://doi.org/10.1109/FSCS.1990.89604

[30] Russell Impagliazzo and Michael Luby. 1989. One-way Functions are Essential for Complexity Based Cryptography (Extended Abstract). In *30th Annual Symposium on Foundations of Computer Science, Research Triangle Park, North Carolina, USA, 30 October - 1 November 1989*. IEEE Computer Society, 230–235. https://doi.org/10.1109/SFCS.1989.63483

[31] Mark Jerrum. 1992. Large Cliques Elude the Metropolis Process. *Random Struct. Algorithms* 3, 4 (1992), 347–360. https://doi.org/10.1002/rsa.3240030402

[32] Valentine Kabanets and Jin-Yi Cai. 2000. Circuit minimization problem. In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, May 21-23, 2000, Portland, OR, USA*. ACM, 73–79. https://doi.org/10.1145/335305.335314

[33] Ludek Kucera. 1995. Expected Complexity of Graph Partitioning Problems. *Discret. Appl. Math.* 57, 2-3 (1995), 193–212. https://doi.org/10.1016/0166-218X(94)00103-K

[34] Leonid A. Levin. 1974. Laws of information conservation (nongrowth) and aspects of the foundation of probability theory. *Problemy Peredachi Informatsii* 10, 3 (1974), 30–35.

[35] Leonid A. Levin. 1984. Randomness Conservation Inequalities; Information and Independence in Mathematical Theories. *Inf. Control.* 61, 1 (1984), 15–37. https://doi.org/10.1016/S0019-9958(84)80060-1

[36] Leonid A. Levin. 1987. One-way functions and pseudorandom generators. *Comb.* 7, 4 (1987), 357–363. https://doi.org/10.1007/BF02579323

[37] Leonid A. Levin. 2003. The Tale of One-Way Functions. *Probl. Inf. Transm.* 39, 1 (2003), 92–103. https://doi.org/10.1023/A%3A1023634616182

[38] Yanyi Liu and Rafael Pass. 2020. On One-way Functions and Kolmogorov Complexity. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*. IEEE, 1243–1254. https://doi.org/10.1109/FOCS46700.2020.00118

[39] Yanyi Liu and Rafael Pass. 2021. A Note on One-way Functions and Sparse Languages. *Electron. Colloquium Comput. Complex.* (2021), 92.

[40] Yanyi Liu and Rafael Pass. 2021. On One-way Functions from NP-Complete Problems. *Electron. Colloquium Comput. Complex.* 28 (2021), 59. https://eccc.weizmann.ac.il/report/2021/059

[41] Yanyi Liu and Rafael Pass. 2021. On the Possibility of Basing Cryptography on EXP ≠ BPP. In *Advances in Cryptology - CRYPTO 2021 - 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16-20, 2021, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 12825)*, Tal Malkin and Chris Peikert (Eds.). Springer, 11–40. https://doi.org/10.1007/978-3-030-84242-0_2

[42] Zhenjian Lu and Igor Carboni Oliveira. 2021. An Efficient Coding Theorem via Probabilistic Representations and Its Applications. In *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, July 12-16, 2021, Glasgow, Scotland (Virtual Conference) (LIPIcs, Vol. 198)*, Nikhil Bansal, Emanuela Merelli, and James Worrell (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 94:1–94:20. https://doi.org/10.4230/LIPIcs.ICALP.2021.94

[43] Dylan M. McKay, Cody D. Murray, and R. Ryan Williams. 2019. Weak lower bounds on resource-bounded compression imply strong separations of complexity classes. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*. ACM, 1215–1225. https://doi.org/10.1145/3313276.3316396

[44] Cody D. Murray and R. Ryan Williams. 2017. On the (Non) NP-Hardness of Computing Circuit Complexity. *Theory Comput.* 13, 1 (2017), 1–22. https://doi.org/10.4086/toc.2017.v013a004

[45] Igor Carboni Oliveira. 2019. Randomness and Intractability in Kolmogorov Complexity. In *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019, Patras, Greece (LIPIcs, Vol. 132)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 32:1–32:14. https://doi.org/10.4230/LIPIcs.ICALP.2019.32

[46] Igor Carboni Oliveira, Ján Pich, and Rahul Santhanam. 2019. Hardness Magnification near State-Of-The-Art Lower Bounds. In *34th Computational Complexity Conference, CCC 2019, July 18-20, 2019, New Brunswick, NJ, USA (LIPIcs, Vol. 137)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 27:1–27:29. https://doi.org/10.4230/LIPIcs.CCC.2019.27

[47] Igor Carboni Oliveira and Rahul Santhanam. 2017. Conspiracies Between Learning Algorithms, Circuit Lower Bounds, and Pseudorandomness. In *32nd Computational Complexity Conference, CCC 2017, July 6-9, 2017, Riga, Latvia (LIPIcs, Vol. 79)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 18:1–18:49. https://doi.org/10.4230/LIPIcs.CCC.2017.18

[48] Igor Carboni Oliveira and Rahul Santhanam. 2018. Hardness Magnification for Natural Problems. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*. IEEE Computer Society, 65–76. https://doi.org/10.1109/FOCS.2018.00016

[49] Rafail Ostrovsky and Avi Wigderson. 1993. One-Way Fuctions are Essential for Non-Trivial Zero-Knowledge. In *Second Israel Symposium on Theory of Computing Systems, ISTCS 1993, Natanya, Israel, June 7-9, 1993, Proceedings*. IEEE Computer Society, 3–17. https://doi.org/10.1109/ISTCS.1993.253489

[50] Ján Pich and Rahul Santhanam. 2019. Why are Proof Complexity Lower Bounds Hard?. In *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS*

2019, Baltimore, Maryland, USA, November 9-12, 2019. IEEE Computer Society, 1305–1324. https://doi.org/10.1109/FOCS.2019.00080

[51] Alexander A. Razborov and Steven Rudich. 1997. Natural Proofs. *J. Comput. Syst. Sci.* 55, 1 (1997), 24–35. https://doi.org/10.1006/jcss.1997.1494

[52] Hanlin Ren and Rahul Santhanam. 2021. Hardness of KT Characterizes Parallel Cryptography. In *36th Computational Complexity Conference, CCC 2021, July 20-23, 2021, Toronto, Ontario, Canada (Virtual Conference) (LIPIcs, Vol. 200)*, Valentine Kabanets (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 35:1–35:58. https://doi.org/10.4230/LIPIcs.CCC.2021.35

[53] Rahul Santhanam. 2020. Pseudorandomness and the Minimum Circuit Size Problem. In *11th Innovations in Theoretical Computer Science Conference, ITCS 2020, January 12-14, 2020, Seattle, Washington, USA (LIPIcs, Vol. 151)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 68:1–68:26. https://doi.org/10.4230/LIPIcs.ITCS.2020.68

[54] Adi Shamir. 1992. IP = PSPACE. *J. ACM* 39, 4 (1992), 869–877. https://doi.org/10.1145/146585.146609

[55] Boris A. Trakhtenbrot. 1984. A Survey of Russian Approaches to Perebor (Brute-Force Searches) Algorithms. *IEEE Ann. Hist. Comput.* 6, 4 (1984), 384–400. https://doi.org/10.1109/MAHC.1984.10036

[56] Salil P. Vadhan. 2006. An Unconditional Study of Computational Zero Knowledge. *SIAM J. Comput.* 36, 4 (2006), 1160–1214. https://doi.org/10.1137/S0097539705447207