On One-Way Functions from NP-Complete Problems

Yanyi Liu
Cornell Tech, New York, NY, USA

Rafael Pass
Cornell Tech, New York, NY, USA
Tel-Aviv University, Israel

Abstract

We present the first natural NP-complete problem whose average-case hardness w.r.t. the uniform distribution over instances is *equivalent* to the existence of one-way functions (OWFs). The problem, which originated in the 1960s, is the *Conditional Time-Bounded Kolmogorov Complexity Problem*: let $K^t(x \mid z)$ be the length of the shortest "program" that, given the "auxiliary input" z, outputs the string x within time $t(|x|)$, and let $\mathsf{McK}^t\mathsf{P}[\zeta]$ be the set of strings (x, z, k) where $|z| = \zeta(|x|)$, $|k| = \log |x|$ and $K^t(x \mid z) < k$, where, for our purposes, a "program" is defined as a RAM machine.

Our main result shows that for every polynomial $t(n) \geq n^2$, there exists some polynomial ζ such that $\mathsf{McK}^t\mathsf{P}[\zeta]$ is NP-complete. We additionally extend the result of Liu-Pass (FOCS'20) to show that for every polynomial $t(n) \geq 1.1n$, and every polynomial $\zeta(\cdot)$, mild average-case hardness of $\mathsf{McK}^t\mathsf{P}[\zeta]$ is equivalent to the existence of OWFs. Taken together, these results provide the following crisp characterization of what is required to base OWFs on $\mathsf{NP} \not\subseteq \mathsf{BPP}$:

> *There exists concrete polynomials t, ζ such that "Basing OWFs on $\mathsf{NP} \not\subseteq \mathsf{BPP}$" is equivalent to providing a "worst-case to (mild) average-case reduction for $\mathsf{McK}^t\mathsf{P}[\zeta]$".*

In other words, the "holy-grail" of Cryptography (i.e., basing OWFs on $\mathsf{NP} \not\subseteq \mathsf{BPP}$) is equivalent to a basic question in algorithmic information theory.

As an independent contribution, we show that our NP-completeness result can be used to shed new light on the feasibility of the *polynomial-time bounded symmetry of information* assertion (Kolmogorov'68).

2012 ACM Subject Classification Theory of computation → Computational complexity and cryptography

Keywords and phrases One-way Functions, NP-Completeness, Kolmogorov Complexity

Digital Object Identifier 10.4230/LIPIcs.CCC.2022.36

Related Version *Full Version*: `https://eccc.weizmann.ac.il/report/2021/059/`

Funding *Rafael Pass*: Work partially done while being on a sabbatical at Tel-Aviv University. Supported in part by NSF Award SATC-1704788, NSF Award RI-1703846, AFOSR Award FA9550-18-1-0267, and a JP Morgan Faculty Award. This material is based upon work supported by DARPA under Agreement No. HR00110C0086. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Government or DARPA.

Acknowledgements We are very grateful to Vinod Vaikuntanathan and Rahul Ilango for helpful comments on an earlier version of this paper. We thank the anonymous reviewers for their helpful comments.

37th Computational Complexity Conference (CCC 2022).
Editor: Shachar Lovett; Article No. 36; pp. 36:1–36:24
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

A one-way function (OWF) [15] is a function f that can be efficiently computed (in polynomial time), yet no probabilistic polynomial-time (PPT) algorithm can invert f with inverse polynomial probability for infinitely many input lengths n. Whether OWFs exist is unequivocally the most important open problem in Cryptography: OWFs are both necessary [34] and sufficient for many of the most central cryptographic primitives and protocols (e.g., pseudorandom generators [10, 24], pseudorandom functions [19], private-key encryption [20], digital signatures [56], commitment schemes [51], identification protocols [16], coin-flipping protocols [9], and more). These primitives and protocols are often referred to as *private-key primitives*, or "Minicrypt" primitives [33] as they exclude the notable task of public-key encryption [15, 55].

While many candidate constructions of OWFs are known – most notably based on factoring [55], the discrete logarithm problem [15], or the hardness of lattice problems [1] – the question of whether OWFs can be based on some "standard" complexity-theoretic assumption is mostly wide open. Indeed, a central open problem – often referred to as the "holygrail of cryptography" – originating in the seminal work of Diffie and Hellman [15] is whether the existence of OWFs can be based on the assumption that $\mathsf{NP} \not\subseteq \mathsf{BPP}$.[1] So far, however, most results in the literature have been negative. Notably, starting with the work by Brassard [13] in 1983, a long sequence of works have shown various types of black-box separations between *restricted* types of OWFs (e.g., one-way permutations) and NP-hardness (see e.g., [13, 12, 3, 53, 21, 47, 22, 11]). We emphasize, however, that these results only show limited separations: they either consider restricted types of one-way functions, or restricted classes of black-box reductions. Thus, even w.r.t. black-box reductions, the question of whether OWFs can be based on the assumption that $\mathsf{NP} \not\subseteq \mathsf{BPP}$, is wide open. In this work, our focus is on providing a complexity-theoretic characterization of *exactly* what is required for basing OWFs on $\mathsf{NP} \not\subseteq \mathsf{BPP}$:

> *Is there a simple complexity-theoretic characterization of what is required for basing OWFs on the assumption that* $\mathsf{NP} \not\subseteq \mathsf{BPP}$*?*

We believe that having a crisp complexity-theoretic characterization will be useful both for obtaining more meaningful separation results, and towards the goal of eventually getting a construction of OWFs based on $\mathsf{NP} \not\subseteq \mathsf{BPP}$.

Towards Characterizing the Possibility of Basing OWFs on NP $\not\subseteq$ BPP. A first step towards answering the above question is implied by a recent work by Liu and Pass [43]; they demonstrated the first natural NP-language whose average-case hardness characterizes the existence of OWFs. In more detail, they demonstrated that for any polynomial $t(n) \geq 1.1n$, OWFs exist if and only the *t-time-bounded Kolmogorov complexity problem*, $\mathsf{MK}^t\mathsf{P}$, is mildly hard-on-average, where a language L is said to be *mildly hard-on-average* if there exists some polynomial $p(\cdot)$ such that no PPT heuristic $\mathcal{H}$ can decide L with probability $1 - 1/p(n)$ over random n-bit instances for infinitely many input lengths n. (We provide more details on the definition of $\mathsf{MK}^t\mathsf{P}$ below.) $\mathsf{MK}^{\mathsf{poly}}\mathsf{P}$ is contained in NP, but it is unknown whether this problem (which has been studied since the 1960s) is NP-complete. Indeed, this is one of the

[1] Or more precisely, whether OWFs can be based on the assumption that $\mathsf{NP} \not\subseteq \mathsf{ioBPP}$ since the definition of OWFs requires "almost everywhere" hardness. For convenience, in the introduction we are ignoring this issue.

long-standing open problems in algorithmic information-theory [38]. A simple corollary of the result from [43] (as far as we know, this has not been previously observed) is that basing OWFs on $\mathsf{NP} \not\subseteq \mathsf{BPP}$ is *equivalent* to (1) proving that $\mathsf{MK}^{\mathsf{poly}}\mathsf{P}$ is NP-complete (perhaps with a non-constructive reduction), and (2) providing a worst-case to average-case reduction for $\mathsf{MK}^{\mathsf{poly}}\mathsf{P}$.[2] To see why this is the case, note that if (1) and (2) are satisfied, then by the result of [43], we have directly based OWFs on $\mathsf{NP} \not\subseteq \mathsf{BPP}$. For the converse direction, note that if OWFs can be based on $\mathsf{NP} \not\subseteq \mathsf{BPP}$, then $\mathsf{NP} \not\subseteq \mathsf{BPP}$ implies the existence of OWF, which by [43] implies that $\mathsf{MK}^{\mathsf{poly}}\mathsf{P}$ is average-case hard (and thus also worst-case hard); thus we have that (1) must hold. To see that (2) also holds, note that since $\mathsf{MK}^{\mathsf{poly}}\mathsf{P} \in \mathsf{NP}$, it follows that $\mathsf{MK}^{\mathsf{poly}}\mathsf{P} \not\subseteq \mathsf{BPP}$ implies that $\mathsf{NP} \not\subseteq \mathsf{BPP}$, which in turn implies OWF (by assumption) which in turn by [43] implies that $\mathsf{MK}^{\mathsf{poly}}\mathsf{P}$ is average-case hard, and thus (2) follows.

The above discussion, however, leaves open the question of whether a crisper characterization can be obtained. In particular, if one can come up with a *natural* NP-complete language L whose average-case hardness is equivalent to the existence of OWFs, then the question of whether OWFs can be based on $\mathsf{NP} \not\subseteq \mathsf{BPP}$ would be equivalent to the question of whether there exists a worst-case to average-case reduction for this particular problem. This thus begs the question whether there exists some natural NP-*complete* language that characterizes the existence of OWFs:

> *Does there exist some "natural"* NP*-complete language L such that OWFs exist iff L is hard-on-average?*

This question was recently raised (but not solved) in a paper by Allender et al [5] (and without the above motivation). We note that "naturality" of the language L is *key* for this question to make sense: It is easy to modify $\mathsf{MK}^{\mathsf{poly}}\mathsf{P}$ into a new "artificial" language L' which is both NP-complete, yet (mild) average-case hardness of L' is equivalent to mild average-case hardness of $\mathsf{MK}^{\mathsf{poly}}\mathsf{P}$ (and thus equivalent to the existence of OWFs).[3] But such an artificial problem would have no relevance to the central question that concerns us (i.e., providing a crisp characterization of when OWF can be based on $\mathsf{NP} \not\subseteq \mathsf{BPP}$).

There is a long history of work on trying to base OWFs on average-case hardness of NP-complete problems, starting with the work of Merkle and Hellman [50]. While the original attempts failed to produce secure schemes (see [52] for a survey), more recent approaches pioneered by Impaglizzo and Naor [35], Ajtai [1] and Ajtai and Dwork [2] produced not just OWFs but also more advanced cryptographic primitives (such as collision-resistant hash functions and public-key encryption) based on well-founded average-case hardness assumptions on the subset sum problem (which is NP-complete). However, it is not known whether the existence of OWFs implies average-case hardness of the subset sum problem (i.e., they only have a one-sided implication).

In this work, we identify the first natural NP-complete language L – *time-bounded Conditional Kolmogorov-complexity* [62, 42, 59, 48] – such that mild average-case hardness of L (with respect to the uniform distribution on instances) is equivalent to the existence

[2] We emphasize that we need a worst-case to *2-sided* error average-case reduction. Hirahara's elegant work [25] makes partial progress on this question by presenting a worst-case (approximate) to *errorless* average-case reduction; errorless average-case hardness does not suffice for [43].

[3] Simply consider the language L' of $2n$-bit instances $x||y$ where $x, y \in \{0,1\}^n$, and either (a) $x = 0^n$ and $y \in \mathsf{SAT}$, or (b) $x \neq 0^n$ and $y \in \mathsf{MK}^{\mathsf{poly}}\mathsf{P}$. In other words, L' is a combination of SAT and $\mathsf{MK}^{\mathsf{poly}}\mathsf{P}$, so clearly this language is NP-complete, but when considering uniform statements, we only hit SAT instances with negligible probability, and thus this language behaves essentially just like $\mathsf{MK}^{\mathsf{poly}}\mathsf{P}$ on average.

of OWFs. As a consequence, we get that basing OWFs on $\mathsf{NP} \not\subseteq \mathsf{BPP}$ is equivalent to providing a worst-case to average case reduction for this particular problem, yielding a simple complexity-theoretic characterization of exactly what it takes to base OWFs on $\mathsf{NP} \not\subseteq \mathsf{BPP}$.

1.1 Our Results

Before describing our results in detail, let us first briefly recall the notion of Time-bounded Kolmogorov Complexity and the result of [43] that we will be relying on.

Time-bounded Kolmogorov Complexity and OWFs. What makes the string 1212121212 1212121 less random than 60484850668340357492? The notion of *Kolmogorov complexity* (K-complexity) [58, 40, 14] from the field of algorithmic information theory provides an elegant method for measuring the amount of "randomness" in individual strings: The K-complexity of a string is the length of the shortest program (to be run on some fixed universal Turing machine U) that outputs the string x. The notion of *$t(\cdot)$-time-bounded Kolmogorov Complexity (K^t-complexity)* is a computationally-restricted version of K-complexity: $K^t(x)$ is defined as the length of the shortest program that outputs the string x within time $t(|x|)$. As surveyed by Trakhtenbrot [59], the problem of efficiently determining the K^t-complexity for $t(n) = \mathsf{poly}(n)$ predates the theory of NP-completeness and was studied in the Soviet Union since the 60s as a candidate for a problem that requires "brute-force search". The modern complexity-theoretic study of this problem goes back to Sipser [57], Ko [37] and Hartmanis [23]. Let $\mathsf{MK}^{t(\cdot)}\mathsf{P}$ denote the decisional $t(n)$-time bounded Kolmogorov complexity problem; namely, the language of pairs (x, k) where $|k| = \lceil \log |x| \rceil$ and $K^t(x) \le k$.

As mentioned above, Liu and Pass [43] demonstrated that for every polynomial $t(n) \ge 1.1n$, mild average-case hardness of $\mathsf{MK}^t\mathsf{P}$ is equivalent to the existence of OWFs. But as mentioned, it is not known whether $\mathsf{MK}^t\mathsf{P}$ is NP-complete (for any polynomial t). Towards getting a characterization of OWFs based on average-case hardness of an NP-complete problem, we will consider a generalization of $\mathsf{MK}^t\mathsf{P}$ based on *conditional* Kolmogorov complexity.

Conditional Time-bounded Kolmogorov Complexity. The *$t(\cdot)$-time-bounded Conditional Kolmogorov Complexity* [62, 42, 59, 48] of a string x conditioned on the string z – denoted $K^t(x \mid z)$ – is the length of the shortest program that, given the "auxiliary input" z, outputs the string x within time $t(|x|)$. More formally,

$$K^t(x \mid z) = \min_{\Pi \in \{0,1\}^*} \{|\Pi| : U(\Pi(z), 1^{t(|x|)}) = x\},$$

where U is a universal Turing machine, and we let $U(\Pi(z), 1^t)$ denote the output of the program Π on input z after t steps. Whereas the notion of a "program" typically is taken to be a Turing machine, in this work we focus on the setting where a program is taken to be a RAM-machine – namely Π is now allowed to be a RAM-machine that can make Random Access queries into the auxiliary string z. Let $\mathsf{McK}^{t(\cdot)}\mathsf{P}[\zeta(\cdot)]$ denote the decisional $t(\cdot)$-time-bounded $\zeta(\cdot)$-conditional Kolmogorov complexity problem; namely, the language of triples (x, z, k) where $|z| = \zeta(|x|)$, $|k| = \lceil \log |x| \rceil$ and $K^t(x \mid z) \le k$. Whereas conditional (time-bounded) Kolmogorov complexity has been studied for decades (see e.g., [48]), it has also remained an open question to determine whether this problem is NP-complete.[4]

[4] We remark, however, that as far as we know, we are the first to consider this problem w.r.t. RAM programs as opposed to Turing machines. In our view, this RAM version of the problem is as natural (if not more) than the "standard" TM version.

We observe that the result of [43] extends, with only relatively minor modifications in the proof, also to conditional Kolmogorov complexity: We show that for every polynomial $t(\cdot) \geq 1.1n$, and every polynomial $\zeta(\cdot)$, mild average-case hardness of $\mathsf{McK}^t\mathsf{P}[\zeta]$ is equivalent to the existence of OWFs.

▶ **Theorem 1.1** (closely following [43]). *For every polynomial $t(n) \geq 1.1n$, every polynomial $\zeta(\cdot)$, mild average-case hardness of $\mathsf{McK}^t\mathsf{P}[\zeta]$ is equivalent to the existence of OWFs.*

So, if we could show that $\mathsf{McK}^t\mathsf{P}[\zeta]$ is NP-complete for some polynomials t, ζ, we would be done. Our main theorem does exactly this.

▶ **Theorem 1.2** (Main Theorem). *For every polynomial $t(n) \geq n^2$, there exists some polynomial $\zeta(\cdot)$, such that $\mathsf{McK}^t\mathsf{P}[\zeta]$ is NP-complete (under randomized polynomial-time reductions).*

Let us emphasize that the combination of Theorem 1.1 and Theorem 1.2, for instance, yields the following crisp characterization of the "holygrail" of Cryptography:

> *There exists (concrete) polynomials t, ζ such that "Basing OWFs on $\mathsf{NP} \not\subseteq \mathsf{BPP}$" is equivalent to the existene of a worst-case to (mild) average-case reduction for $\mathsf{McK}^t\mathsf{P}[\zeta]$.*

In other words, the "holy grail" of Cryptography is equivalent to a basic question in algorithmic information theory. Furthermore, let us point out that for the unconditional time-bounded Kolmogorov complexity problem $\mathsf{MK}^{\mathsf{poly}}\mathsf{P}$, some partial[5] worst-case to average-case reductions are known [25], so this gives us hope that a full worst-case to average-case reduction may be possible also for $\mathsf{McK}^t\mathsf{P}$.

As we shall discuss shortly, Theorem 1.2 is also interesting in its own right and has other direct applications: we show how to shed new light on a long-standing open problem regarding *symmetry of information* [62] for the setting of time-bounded Kolmogorov complexity.

Let us emphasize that for the NP-completeness result to hold, it is imperative that our notion of conditional Kolmogorov complexity views programs as *RAM-machines* (as opposed to *Turing machines*). We leave it as an intriguing open problem to determine whether the "standard" conditional time-bounded Kolmogorov complexity (where interpreting a program as a Turing machine) is also NP-complete.

We proceed to providing a proof overview of the main theorem (i.e. Theorem 1.2).

1.2 Proof Overview

We first note that it directly follows that for all polynomials t, ζ, $\mathsf{McK}^t\mathsf{P}[\zeta] \in \mathsf{NP}$ – the witness for an instance (x, z, k) is simply a RAM program Π such that $|\Pi| \leq k$ and $\Pi(z)$ generates x within $t(|x|)$ steps. We turn to discussing how to prove that there exist polynomials t, ζ, such that $\mathsf{McK}^t\mathsf{P}[\zeta]$ is NP-hard. On a high-level, our approach will start off by using the recent breakthrough approach by Ilango [29, 30] showing NP-hardness of an *oracle*-variant of the *circuit minimization problem (MCSP)* [36] – that is, the problem of, given the truth table of a boolean function, determining the size of the smallest circuit that computes the function – and next extend it to deal with the conditional Kolmogorov complexity problem by appropriately embedding the "oracles" used in the construction of [29] in the auxiliary input.

[5] The reduction only shows so-called *errorless*, as opposed to 2-sided error, average-case hardness.

In more detail, following [8, 28, 29, 30, 32], we will embed an (approximate) Bounded Set Cover instance into an $\mathsf{McK}^t\mathsf{P}[\zeta]$ instance; the approximate Bounded Set Cover problem is known to be NP-complete [60]. Recall that in the Bounded Set Cover problem, we are given a collection of sets $S_1, S_2, \ldots S_r$, each of which is a *constant*-size subset of the universe $U = [n]$ and the goal is to find a minimal set of indexes, s, such that $\cup_{i \in s} S_i = [n]$ (i.e., finding the minimal collection $\mathcal{S}$ of sets S_i that cover $[n]$). We start off by generalizing an idea from [29, 30, 32] and replace the universe $U = [n]$ with n random strings $A_i \in \{0,1\}^m$, where $m(n)$ is some sufficiently large polynomial (in the formal proof $m(n) = n^3$). Roughly speaking, the rationale for doing this is that a set cover, intuitively, should give a succinct (proportional to the size of the set cover) way to generate the random string $A = A_1||A_2|| \ldots ||A_n$ if we have *oracle access* to the sets S_i – we simply need to specify the sets in the set cover and can then reconstruct the union of these sets. This construction was used in [29] to prove NP-hardness of the oracle-version of the MCSP problem – the sets S_i were simply placed into the oracle. ([30] provides a more elaborate construction that also shows NP-hardness of a conditional variant of the MCSP problem; we will, however, not rely on that extension.)

To convert the above set-cover instance into a conditional Kolmogorov complexity problem, our new idea will be to place the description of the sets S_i (each of which consists of some set of strings A_i) at *random locations* in the auxiliary string z and to make sure z is very long (yet still only of polynomial length), and consider the conditional Kolmogorov complexity problem of computing $K^t(A \mid z)$ where $A = A_1||A_2|| \ldots ||A_n$. Conceptually, one can view this approach as a way to *obfuscate* the oracle used in [29] and placing the obfuscation in z. Intuitively, since we are placing the descriptions of the sets S_i at random locations in z, a time-bounded algorithm can only access S_i if it "knows" the random location where it has been put, and thus, intuitively, we can view z as an information-theoretic obfuscation of gates that compute these descriptions.[6]

If there exists a set cover s of size ℓ, $K^t(A \mid z)$ should be no more than $\ell O(\log n) + O(1)$, by considering the program that simply hardcodes the location in z of the descriptions of the sets S_i for $i \in s$. The harder part is showing that if $K^t(A \mid z) \leq \ell O(\log n)$ then there exists a set-cover of size $O(\ell)$. Relying on the intuition that z acts as an obfuscation of the description of the sets $\{S_i\}$, the intuition for why this holds is that if z is sufficiently longer than $t(|x|)$, and the descriptions of the sets are put into random positions of z, any program with running-time $t(|x|)$ that reconstructs the string $A = A_1||A_2|| \ldots ||A_n$ (which with overwhelming probability has high Kolmogorov complexity) must "know" the location in the auxiliary string z of sets $\{S_i\}_{i \in s}$ in some set cover s, in the sense that by running this program, those positions can be "reconstructed". In a bit more detail, by running this program and looking at the memory access queries made by the program into z, we must be hitting the locations where the sets have been put. But since these locations are random (by construction of z), the program needs to basically "hard-code" them, or else we would be able to compress the indexes of these locations, but these indexes have high Kolmogorov complexity as they were picked at random, which is a contradiction.

The reader may note that, perhaps curiously, we are using an argument based on Kolmogorov complexity to formalize the statement that $K^t(A \mid z) \approx \ell O(\log n)$. In more detail, we are relying on a Kolmogorov-complexity style compression argument to formalize

[6] This intuition is somewhat misleading: since z is only of polynomial length, the "obfuscation" only works with respect to a-priori time-bounded attackers (that can only explore a small fraction of z) and can only have inverse polynomial security. But in our context, such a relaxed notion of security suffices.

that z acts as a good "obfuscation" of the description of the sets $\{S_i\}$. This proof technique bears similarities to the proof technique pioneered by Gennero and Trevisan [17] in the context of proving that a random permutation is one-way w.r.t. polynomial-size circuits.

Let us end by noting that the above proof outline oversimplifies and misses several crucial details that make the actual proof quite a bit more complicated.

1.3 Applications to Polynomial-time Bounded Symmetry of Information

The celebrated *symmetry of information* theorem by Kolmogorov and Levin from 1967 [62] states that for all strings $x, y \in \{0,1\}^*$:

$$K(xy) = K(y) + K(x \mid y) \pm O(\log |xy|)$$

where xy denotes the concatenation of the strings x, y. The proof of this theorem, however, involves a computationally expensive exhaustive search through all strings of lengths $|x|, |y|$. The question of whether a *polynomial-time bounded* version of information symmetry holds, where K-complexity is replaced by K^t-complexity for polynomials t, has remained an open problem. We refer to the assertion that there exists some constant $0 < \epsilon \leq 1$ such that for all sufficiently large polynomials t, all $x, y \in \{0,1\}^*$ (of polynomially-related length), it holds that

$$K^{t^\epsilon}(xy) \geq K^t(y) + K^t(x \mid y) - O(\log |xy|)$$

$$K^{t^{1/\epsilon}}(xy) \leq K^t(y) + K^t(x \mid y) + O(\log |xy|)$$

as the *polynomial-time symmetry of information assertion (*polySOI*)*. The question of whether a polynomial-time symmetry of information assertion holds goes back to a work by Kolmogorov from 1968 [39]; as retold by Levin [41]:

> *Kolmogorov suggested at the time [39] that this information symmetry theorem may be a good test case to prove that for some tasks exhaustive search cannot be avoided (in today's terms,* $\mathsf{P} \neq \mathsf{NP}$*)*

The first formal complexity-theoretic investigation of this question goes back to works by Longpré and Mocas [48] and Longpré and Watanabe [49]. They consider a *length-restricted* version of the polySOI where we additionally add the requirement that $|x| = |y|$ and show that (1) the length-restricted polySOI assertion holds if $\mathsf{NP} = \mathsf{P}$, and (2) the length-restricted polySOI assertion is false if one-way functions exist. We here focus our attention on the above "length-unrestricted" version of the polySOI assertion, where x and y can be of arbitrary polynomially-related lengths. We demonstrate, as a corollary of the techniques behind the proof of Theorem 1.2, that the (length unrestricted) polySOI assertion is *unconditionally* false.

▶ **Theorem 1.3.** *The (length-unrestricted) polynomial-time symmetry of information assertion is false.*

Let us provide a brief overview of how this theorem is proven, and the instrumental role that the proof of Theorem 1.2 plays in this proof. The proof, roughly speaking, proceeds in the following steps:

- **Step 1: Polynomial increase in running-time can only decrease $K^t(x)$ by $O(\log |x|)$.** We first show that polySOI implies that if we polynomially increase the running-time bound, then this cannot have a significant effect on K^t; more precisely, for large enough t, it holds that for all constants c, $K^t(x)$ cannot be more than $O(\log |x|)$ more than $K^{t^c}(x)$. Intuitively, this follows from polySOI by letting y be a sufficiently long all "dummy" string, 0^m. We note that this step inherently relies on us considering a length *unrestricted* polySOI.

- **Step 2: Showing that a "strong" polySOI assertion holds.** We next observe that combining Step 1 with polySOI yields a strong form of the polySOI assertion where $\epsilon = 1$; i.e., it holds that

$$K^t(xy) = K^t(y) + K^t(x \mid y) \pm O(\log |xy|)$$

- **Step 3: Polynomial increase in running-time can only decrease *conditional* $K^t(x|z)$ by $O(\log |xz|)$.** We then combine the "strong" polySOI assertion from Step 2 with Step 1 to show that an analog of Step 1 holds also with respect to conditional time-bounded Kolmogorov complexity. More precisely, for large enough t, it holds that for all constants c, $K^t(x \mid z)$ cannot be more than $O(\log |xz|)$ larger than $K^{t^c}(x \mid z)$.
- **Step 4: Contradicting the construction in Theorem 1.2.** We finally observe that the statement of Step 3 contradicts the construction in the proof of Theorem 1.2. In particular, in the proof of Theorem 1.2, we showed strings x, z where the condition time-bounded Kolmogorov complexity $K^t(x \mid z)$ is large (roughly the size of the set-cover). However, this relied on t being sufficiently small so that the program cannot read all of z. If t is large, so that the program can read all of z, then $K^t(x \mid z)$ is tiny (x can be trivially reconstructed from z). This contradicts the statement of Step 3.

Let us emphasize that Theorem 1.3 is incomparable to the result of [49]: [49] consider a weaker "length-restricted" polynomial-time symmetry of information assertion (where $|x| = |y|$), whereas we are considering a "length-unrestricted" version. While [49] show that the length-restricted polySOI indeed holds if NP = P, we show that the length-unrestricted version is unconditionally false. We do, however, note that the "standard" (non time-bounded) symmetry of information theorem of [62] consider the length unrestricted case, in analogy with what we do.

Furthermore, for our result to hold (and in contrast to [49]), it is crucial that we consider RAM-programs as the model of computation, as opposed to the (more standard in the context of time-bounded Kolmogorov complexity) notion of TM-programs; nevertheless, in our eyes, considering RAM programs is equally motivated (if not more) than TM programs.

1.4 Related Works

As mentioned above, there has been a recent sequence of surprising works proving NP-hardness results for variants of the MCSP problem [8, 28, 29, 30, 32]; in particular, as mentioned, Ilango [30] proves that a conditional version of the MCSP problem is NP-hard. As observed already in [59], and further explored in [4], the MCSP problem is closely related to the time-bounded Kolmogorov complexity problem – intuitively, the two problems capture the same concept, but using a different model of computation – but a formal reduction between these problems is not known so these results do not directly extend to the setting we consider. (However, as mentioned above, the starting point for our approach is the result of [29] showing NP-hardness for an oracle version of the MCSP problem.)

A recent result by Hirahara [26] directly addresses conditional time-bounded Kolmogorov complexity and shows NP-hardness for a variant of this problem, $\mathsf{McK}^{\mathsf{poly}}\mathsf{P}^{\mathsf{SAT}}$, where the program has access to a SAT-oracle. (The $\mathsf{McK}^{\mathsf{poly}}\mathsf{P}^{\mathsf{SAT}}$ problem, however, is not known to be in NP, but is in $\mathsf{NP}^{\mathsf{NP}}$, so NP-completeness is not shown).

An intriguing recent paper by Allender et al [5] presented a natural NP-complete problem L – a sparse variant of the MCSP problem – such that average-case hardness of L was claimed to imply the existence of OWFs; the authors also claimed a "weak" converse of this implication – that the existence of OWFs implies a very weak, so-called "non-trivial", notion of average-case hardness of the language[7]; unfortunately, an error was found in the paper. Concurrently and independently from the current work, the authors of [5] show how to repair the issues in their proof and present a different NP-complete language whose average-case hardness implies the existence of OWFs, and for which the same weak converse holds.[8] While their original posting [5] – which inspired the current work – attempted to base OWFs on the average-case hardness of a sparse version of the MCSP problem, their new paper [6] instead bases OWFs on average-case hardness of a conditional Kolmogorov complexity style problem, just as in the current work. Their conditional Kolmogorov complexity problem differs from ours in several aspects: (1) whereas we consider conditional Kolmogorov complexity w.r.t. RAM programs, [6] considers it w.r.t. Turing machines with "oracle-access" to the auxiliary input z; and (2) instead of considering a time-bounded version of conditional Kolmogorov complexity (as we do), [6] instead *charge* for running-time in their notion of Kolmogorov complexity, following the KT notion of [4]. Due to these differences, NP-completeness of their problem follows essentially directly from the NP-completeness results of [29] (whereas we have to work a lot harder, as explained above). However, due to these differences, they only manage to show a one-directional implication between average-case hardness of their problem and OWFs (and only a weak converse in the other direction), whereas we establish an *equivalence* between average-case hardness of $\mathsf{McK}^t\mathsf{P}[\zeta]$ (for any polynomials $t(n) > 1.1n, \zeta(\cdot)$) and OWFs.

Subsequent to the initial posting [44] of this paper,[9] [7] have shown, based on the results in [54] that (mild) average-case hardness of the NP-complete problem considered in [6] is equivalent to the existence of OWFs computable in log space; their work thus provides an elegant characterization of what it means to base OWFs computable in logspace on $\mathsf{NP} \not\subseteq \mathsf{BPP}$.

After the initial posting of this paper, we were informed by Rahul Ilango [31] that he had independently also shown NP-completeness of some conditional time-bounded Kolmogorov complexity problem, but without writing down the results. Indeed, as far as we can tell, our paper is the first to present any type of NP-completeness results for Kolmogorov complexity problems.

Resource bounded notions of conditional Kolmogorov complexity are useful also in other (related) contexts. In a companion paper to the current work [46], we rely on a notion of *space-bounded* conditional Kolmogorov complexity (defined similarly to the time-bounded notion of conditional Kolmogorov complexity used in the current paper) to characterize OWFs in NC^0; alternative characterizations without relying on condition Kolmogorov complexity were provided in [54].

In [46], we also identify a problem whose (infinitely-often) average-case hardness w.r.t. error-less heuristics is equivalent to $\mathsf{EXP} \neq \mathsf{BPP}$ (i.e., the problem is EXP-average-case complete w.r.t. errorless heuristics), yet (two-sided error) average-case hardness of this problem is equivalent to the existence of OWFs; related results were also obtained in [54].

[7] Roughly speaking, that average-case hardness holds for an inverse *exponential*, as opposed to inverse polynomial, fraction of inputs.

[8] The papers were submitted to on ECCC/Eprint within one day of each other.

[9] The initial posting [44] did not contain the results on polynomial-time symmetry of information

Taken together, the current work and [46, 54], demonstrate that the existence of OWFs can be characterized through the average-case hardness of both NP-complete (this work) and EXP-complete ([46, 54]) languages.

It has been recently shown in [27, 18] that some other variant of symmetry of information holds under the assumption that NP is easy on average. We mention that our result (proved in Theorem 1.3) combined with the aforementioned result in [27, 18] does *not* prove that NP is hard on average unconditionally (due to the difference in formulating symmetry of information). In the form used in [27, 18], they require that symmetry of information holds with respect to individual running time bounds $t \in \mathbb{N}$ that are not polynomially-related to $|x|$ (or $|y|$), whereas we consider polynomially-related running time bounds and allow a (polynomially) larger running time bound when the string is longer. Our proof does not work in the setting of [27, 18] for technical reasons.

1.5 Outline

We will provide the formalizations and proofs of Theorem 1.2 in Section 3. We refer the reader to the full version [45] for formal treatments of Theorem 1.1 and 1.3.

2 Preliminaries

We let $[n]$ denote the set $\{1, 2, \ldots, n\}$ for any integer $n \in \mathbb{N}$. For any two strings x, y, let $x||y$ denote the concatenation of x and y; whenever it is clear from context, we sometimes also use xy to denote the concatenation of x and y. In this work, we sometimes consider strings that contain a special symbol $\perp$ (besides 0 and 1). We will use the following standard encoding scheme – which we refer to as simple the *standard encoding scheme* $\mathsf{enc}_\perp$) to transform a string that may contain $\perp$ into a binary string: $\mathsf{enc}_\perp(x)$, of a string $x \in \{0, 1, \perp\}^*$ is a $2|x|$-bit binary string where we replace each bit in x by 00 for 0, 01 for 1, and 11 for $\perp$.

2.1 Set Cover

Let n be an integer and $S_1, S_2, \ldots, S_\ell, T$ be sets $\subseteq [n]$. We say that the sets $S_1, S_2, \ldots, S_\ell$ cover T if $T \subseteq S_1 \cup S_2 \cup \ldots \cup S_\ell$. Let $\mathcal{S}$ be a collection of sets. We define $\mathsf{cover}(T, \mathcal{S})$ to be the minimum number of sets in $\mathcal{S}$ necessary to cover T.

We recall the γ-Bounded Set Cover Problem:

- Input: $(1^n, 1^\ell, \mathcal{S})$ where n, ℓ are integers $\in \mathbb{N}$ and $\mathcal{S} = \{S_1, S_2, \ldots, S_r\}$ is a collection of subsets $\subseteq [n]$. It is guaranteed that all the sets in $\mathcal{S}$ covers $[n]$ together and for all $i \in [r]$, $|S_i| \leq \gamma$.
- Decide: Is $\mathsf{cover}([n], \mathcal{S}) \leq \ell$.

We also consider the approximate version of the γ-Bounded Set Cover problem. The α-approximate γ-Bounded Set Cover Problem is a promise problem $(\Pi_{\mathsf{yes}}, \Pi_{\mathsf{no}})$ where Π_{yes} contains $(1^n, 1^\ell, \mathcal{S})$ such that $\mathsf{cover}([n], \mathcal{S}) \leq \ell$ and Π_{no} consists of $(1^n, 1^\ell, \mathcal{S})$ such that $\mathsf{cover}([n], \mathcal{S}) > \alpha \cdot \ell$.

Trevisan [60] showed that approximating the γ-Bounded Set Cover Problem within a constant factor is NP-hard:

▶ **Theorem 2.1** ([60]). *For every constant $\alpha \geq 1$, there exists a constant $\gamma \in \mathbb{N}$ such that the α-approximate γ-Bounded Set Cover Problem is* NP*-hard. More concretely, for any language $L \in$* NP*, there exists a polynomial-time algorithm R such that on input $x \in L$, $R(x)$ outputs an instance in Π_{yes}; on input $x \notin L$, $R(x)$ outputs an instance in Π_{no}, where $(\Pi_{\mathsf{yes}}, \Pi_{\mathsf{no}})$ denotes the α-approximate γ-Bounded Set Cover Problem.*

2.2 The RAM Model

A RAM program $\Pi = (M, y)$ consists of a CPU "next-step" Turing machine M, and some initial input $y \in \{0,1\}^*$. Let $\mathsf{state} = 0$ be an initial state. The execution of this RAM program Π on input $z \in \{0,1\}^*$ (which may be empty) proceeds as follows.

- At initialization, the memory is set to $y||\bot||z$, and the "read bit" b^{read} is set to $\bot$. (For simplicity, we assume that each memory position contains a symbol $\in \{0,1,\bot\}$.[10] We assume that the memory is of infinite length and the rest of the positions in the memory are filled with $\bot$.)
- At each CPU step, M receives as input $\mathsf{state} \in \{0,1\}^*$, the most recently read bit b^{read}, and outputs a new state $\mathsf{state}' \in \{0,1\}^*$, a read position i^{read}, a write position i^{write} and some bit b^{write} (to be written to position i^{write}).[11]
- The execution of this step replaces state with state', sets b^{read} to the content of memory position i^{read}, and replaces the content of memory position i^{write} by b^{write}.
- When $\mathsf{state} = \varepsilon$ (i.e., the empty string), the computation ends and the output of the of the computation is defined as the content of the memory tape up to the symbol $\bot$.[12]
- The running time of Π is defined to be the sum of the running time of M in all CPU steps.

Note that any polynomial-time Turing machine can be simulated by a polynomial-time RAM program by simply copying the content of the memory into state, next letting M run the original Turing machine using state as its tape, and finally copying the content of state back into the memory.

2.3 Time-bounded Conditional Kolmogorov Complexity

We introduce the notion of time-bounded conditional Kolmogorov complexity with respect to RAM programs. Roughly speaking, the *t-time-bounded Kolmogorov complexity*, $K^t(x \mid z)$, of a string $x \in \{0,1\}^*$ conditioned on a string $z \in \{0,1\}^*$ is the length of the shortest RAM program $\Pi = (M, y)$ such that $\Pi(z)$ outputs x in $t(|x|)$ steps.

Let U be some fixed Universal Turing machine that can emulate any RAM program Π with polynomial overhead. Let $U(\Pi(z), 1^t)$ denote the output of $\Pi(z)$ when emulated on U for t steps. We now define the notion of t-time-bounded conditional Kolmogorov complexity.

▶ **Definition 2.2.** *Let t be a polynomial. For all $x \in \{0,1\}^*$ and $z \in \{0,1\}^*$, define*

$$K^t(x \mid z) = \min_{\Pi \in \{0,1\}^*} \{|\Pi| : U(\Pi(z), 1^{t(|x|)}) = x\}$$

where $|\Pi|$ is referred to as the description length *of Π. When there is no time bound, we define*

$$K(x \mid z) = \min_{\Pi \in \{0,1\}^*} \{|\Pi| : U(\Pi(z), 1^{t'}) = x \text{ for some finite } t'\}$$

[10] When we implement this, we always use the standard encoding scheme, $\mathsf{enc}_\bot$. We also note that the string y and z can never contain the symbol $\bot$ (since they exclusively consist of 0s and 1s). When we load y and z into the memory, instead of storing y and z directly, we store the standard encoding of y and z (where 0 becomes 00 and 1 becomes 01).

[11] Formally, the inputs and outputs of M are separated by the $\bot$ symbol so that state can be of variable length.

[12] In a real execution, the content of the memory is encoded by the standard encoding scheme. The output of the computation is then defined by the *decoded* content of the memory.

We also consider the decisional variant of the minimum t-time-bounded conditional Kolmogorov complexity problem. Let t, ζ be two polynomials, and let $\mathsf{McK}^t\mathsf{P}[\zeta]$ denote the language of triples (x, z, k), having the property that $K^t(x \mid z) \leq k$, where $z \in \{0,1\}^{\zeta(|x|)}$ and $k \in \{0,1\}^{\lceil \log n \rceil}$.

We note that for any string $z \in \{0,1\}^*, x \in \{0,1\}^*$, for any polynomial $t(\cdot)$, $K^t(x \mid z)$, is always upper bounded by $|x| + O(1)$.

▶ **Fact 2.3.** *There exists a constant $c \in \mathbb{N}$ such that for all polynomial $t(\cdot)$, for all string $z \in \{0,1\}^*, x \in \{0,1\}^*$, $K^t(x \mid z) \leq |x| + c$.*

Proof. Consider the RAM program $\Pi = (M, x)$ where M is a Turing machine that directly sets $\mathsf{state} = \varepsilon$. Note that in the execution of Π, x will be put into the memory and Π will halt immediately. Thus Π will output the string x. Note that M is a constant-size machine, so the description length of Π is at most $|x| + c$ for some constant c. ◀

We finally remark that for any polynomials $t(\cdot), \zeta(\cdot)$, $\mathsf{McK}^t\mathsf{P}[\zeta] \in \mathsf{NP}$.

▷ Claim 2.4. For all polynomials $t(\cdot), \zeta(\cdot)$, $\mathsf{McK}^t\mathsf{P}[\zeta] \in \mathsf{NP}$.

Proof. On input an instance $(x, z, k) \in \mathsf{McK}^t\mathsf{P}[\zeta]$, and a witness Π, checking if $|\Pi| \leq k$, $|z| = \zeta(|x|)$ and $U(\Pi(z), 1^{t(|x|)}) = x$ can be done in polynomial time. ◁

2.4 One-way Functions

We recall the definition of one-way functions [15]. Roughly speaking, a function f is one-way if it is polynomial-time computable, but hard to invert for PPT attackers.

▶ **Definition 2.5.** *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 $n \in \mathbb{N}$,*

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

We may also consider a weaker notion of a *weak one-way function* [61], where we only require all PPT attackers to fail with probability noticeably bounded away from 1:

▶ **Definition 2.6.** *Let $f : \{0,1\}^* \to \{0,1\}^*$ be a polynomial-time computable function. f is said to be a* α-weak one-way function (α-weak OWF) *if for every* PPT *algorithm $\mathcal{A}$, for all sufficiently large $n \in N$,*

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

2.5 Average-case Hard Languages

We turn to defining what it means for a language to be average-case hard (for PPT algorithms). We will be considering languages that are only defined on some input lengths (such as $\mathsf{McK}^t\mathsf{P}[\zeta]$). We say that a language L is *defined over inputs lengths* $s(\cdot)$ if $L \subseteq \cup_{n \in \mathbb{N}} \{0,1\}^{s(n)}$. For concreteness, note that $\mathsf{McK}^t\mathsf{P}[\zeta]$ is defined on input lengths $s(n) = n + \zeta(n) + \lceil \log n \rceil$.

We now turn to defining average-case hardness.

▶ **Definition 2.7.** *We say that a language L defined over inputs lengths $s(\cdot)$ is $\alpha(\cdot)$* hard-on-average (α-HoA) *if for all* PPT *heuristic $\mathcal{H}$, for all sufficiently large $n \in \mathbb{N}$,*

$$\Pr[x \leftarrow \{0,1\}^{s(n)} : \mathcal{H}(x) = L(x)] < 1 - \alpha(n)$$

In other words, there does not exist a PPT "heuristic" $\mathcal{H}$ that decides L with probability $1 - \alpha(n)$ on infinitely many input lengths $n \in N$ over which L is defined.

We refer to a language L as being *mildly HoA* if there exists a polynomial $p(\cdot) > 0$ such that L is $\frac{1}{p(\cdot)}$-HoA.

3 NP-Hardness of $\mathsf{McK}^t\mathsf{P}[\zeta]$

In this section, we prove our main theorem: We show that there exists a reduction from the approximate γ-Bounded Set Cover Problem to $\mathsf{McK}^t\mathsf{P}[\zeta]$ when t, ζ are sufficiently large.

▶ **Theorem 3.1.** *For all polynomial $t(n) \geq n^2$, there exists a polynomial $\zeta(n)$ such that $\mathsf{McK}^t\mathsf{P}[\zeta]$ is NP-hard under many-one randomized polynomial-time reductions.*

Proof. The theorem follows from Proposition 3.3 and Proposition 3.4 (stated and proved in Section 3.2), and Theorem 2.1. ◀

In fact, we note that the reduction only has one-sided errors:

▶ **Theorem 3.2.** *For all polynomial $t(n) \geq n^2$, if there exists a polynomial $\zeta(n)$ such that $\mathsf{McK}^t\mathsf{P}[\zeta] \in \mathsf{coRP}$, then $\mathsf{NP} \subseteq \mathsf{coRP}$.*

Proof. By Proposition 3.3, our reduction succeeds with probability 1 on YES instances. By Proposition 3.4, our reduction succeeds with high probability ($\geq \frac{1}{2}$) on NO instances. Finally, the corollary follows from Theorem 2.1. ◀

3.1 A Reduction from the γ-Bounded Set Cover Problem to $\mathsf{McK}^t\mathsf{P}$

Let γ be a constant, let $t(n) \geq n^2$ be a polynomial, and consider $\zeta(n) = (t(n))^4 n^{2\gamma}$. We will show that there exists a randomized reduction from the γ-Bounded Set Cover Problem to $\mathsf{McK}^t\mathsf{P}[\zeta]$.

Given an instance $(1^n, 1^\ell, \mathcal{S})$ where $\mathcal{S} = \{S_1, S_2, \ldots, S_r\}$ of the γ-Bounded Set Cover Problem, we proceed as follows:

- Let $m = n^3$; for each $i \in [n]$, sample a random string $A_i \in \{0,1\}^m$, and consider the length-$(n \times m)$ concatenation $A = A_1||A_2||\ldots||A_n$ of the sampled strings. Think of A_i as a randomized encoding of the element i in the Set Cover problem. See Figure 1 for an illustration of these strings.

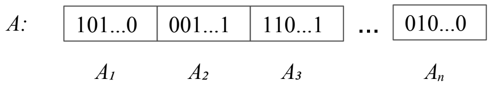

Figure 1 An illustrative example for the string A.

- For each $i \in [r]$, we construct a "gadget" string $W_i \in (\{0,1\}^m)^n$ (for set S_i). We partition W_i into n blocks $W_{i,1}, W_{i,2}, \ldots, W_{i,n}$ where each block is of size m. We let $W_{i,j} = A_j$ if $j \in S_i$, and otherwise $W_{i,j} = 0^m$. In other words, W_i reveals the strings A_j for all $j \in S_i$; think of W_i as a randomized encoding of the set S_i. See Figure 2 for an illustration of there strings.
- Let $\lambda = 4 \log r + 4 \log t(nm)$. For each $i \in [r]$, we sample a "key" $k_i \in \{0,1\}^\lambda$ for W_i. For simplicity, we assume that the sampled keys are distinct with each other. (If this is not the case, the reduction just aborts; since this happens only with negligible probability we may ignore this event in the analysis.)
- We are finally ready to describe the "auxiliary input" z. The idea is to hide the gadgets $\{W_i\}$ in z at random locations specified by the keys so as to ensure that the only way for a t-time bounded program to recover W_i is to essentially hard-code the key k_i as part of its description. In more detail, we consider a string z of length $2^\lambda \times n \times m$; partition z

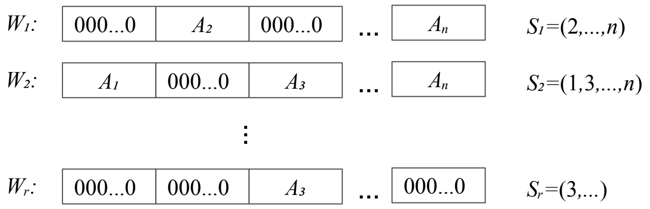

Figure 2 An illustrative example for the gadget strings W_i. Note that if we have a Set Cover $(i_1, i_2, \dots, i_\ell)$, then the bitwise OR of the strings W_{i_1}, W_{i_1}, $\dots W_{i_\ell}$ equals A.

into 2^λ blocks $z_{0^\lambda}, z_{0^{\lambda-1}1}, \dots, z_{1^{\lambda-1}0}, z_{1^\lambda}$ where for all $p \in \{0,1\}^\lambda$, $|z_p| = n \times m$. For all $p \in \{0,1\}^{2^\lambda}$, let $z_p = W_i$ if $p = k_i$ for some $i \in [r]$, and otherwise, let $z_p = 0^{n \times m}$. See Figure 3 for an illustration of these strings.

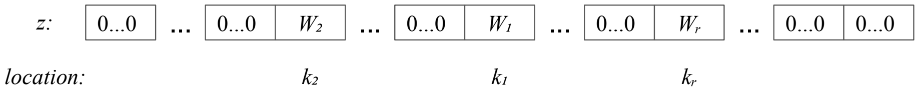

Figure 3 An illustrative example of the "auxiliary" input z.

- Finally, the reduction will output YES if $K^t(A \mid z) \le 2\lambda\ell$. Note that the length of z is upper bounded by $\zeta(|A|)$, and thus this is a syntactically valid reduction to an $\mathsf{McK}^t\mathsf{P}[\zeta]$ instance.

We turn to analyzing the success probability of the reduction.

3.2 Analyzing the Reduction

We will prove that the above reduction gives us a 4-approximation of the γ-Bounded Set Cover Problem. We first show that if $[n]$ can be covered by a small number ($\le \ell$) of sets, the time-bounded Kolmogorov complexity of A conditioned on the string z will be small ($\le 2\lambda\ell$): the program computing A simply needs to hard-code the keys k_i corresponding to the ℓ sets in the set cover; it can look into z at the positions specified by the keys and output the bitwise OR of the content of those positions.

▶ **Proposition 3.3.** *If* $\mathsf{cover}([n], \mathcal{S}) \le \ell$ *then* $K^t(A \mid z) \le K^{t'}(A \mid z) \le 2\lambda\ell$ *where* $t'(n) = n^2$.

Proof. Let $S_{i_1}, S_{i_2}, \dots, S_{i_\ell}$ be the ℓ sets in $\mathcal{S}$ that cover $[n]$. (Since the sets are γ-Bounded, it follows that $\ell \ge n/\gamma$.) Let Π be a RAM program with n, m, λ and the keys $k_{i_1}, \dots, k_{i_\ell}$ hardwired in it. For each $j \in [\ell]$, Π first reads $W'_{i_j} = z_{k_{i_j}}$ from the k_{i_j}-th block of the string z (where $|z_{k_{i_j}}| = n \times m$). (Recall that z is partitioned into 2^λ blocks and each block is of size $n \times m$.) Π then obtains $W'_{i_1}, \dots, W'_{i_\ell}$ and Π simply outputs

$$W'_{i_1} \vee W'_{i_2} \vee \dots \vee W'_{i_\ell}$$

where $\vee$ denotes the bitwise OR for binary strings.

We first show that Π indeed outputs the string A. Note that by the construction of string z, it holds that

$$(W'_{i_1}, \dots, W'_{i_\ell}) = (z_{k_{i_1}}, \dots, z_{k_{i_\ell}}) = (W_{i_1}, \dots, W_{i_\ell}).$$

Recall that in the construction of the gadget string W_{i_j} (for each $j \in [\ell]$), W_{i_j} is partitioned into n blocks $W_{i_j,1}, \ldots, W_{i_j,n}$. And for each block $b \in [n]$, $W_{i_j,b} = A_b$ if $b \in S_{i_j}$, and otherwise $W_{i_j,b} = 0^m$. Since the sets $S_{i_1}, \ldots, S_{i_\ell}$ cover $[n]$, for all $b \in [n]$, there exists an index j such that the b-th block of the gadget string W_{i_j} matches A_b. Thus, $W_{i_1} \vee W_{i_2} \vee \ldots \vee W_{i_\ell} = A$.

We then show that Π can be described within $2\lambda\ell$ bits. Recall that Π contains the values n, m, λ (which takes $O(\log n)$ bits to describe), the keys $k_{i_1}, \ldots, k_{i_\ell}$ (which takes $\lambda\ell$ bits), and the code of Π (which takes $O(1)$ bits). We will provide a more fine-grained analysis in the full version [45] to show that the code of Π is of constant-bit length in the RAM model. Thus, Π can be represented using $\lambda\ell + O(\log n) \leq 2\lambda\ell$ bits.

Finally, note that Π runs in time $O(\ell nm \mathsf{poly} \log n) \leq (nm)^2 = t'(|A|) \leq t(|A|)$ (since in each CPU step, the CPU next-step machine takes $O(\mathsf{poly} \log n)$ time). (We refer the reader to the full version [45] for a more detailed running time analysis.) Thus, we conclude that $K^{t'}(A \mid z) \leq K^t(A \mid z) \leq 2\lambda\ell$. ◀

The key part of the analysis is showing that if $K^t(A \mid z) \leq 2\lambda\ell$ then $\mathsf{cover}([n], \mathcal{S}) \leq 4\ell$:

▶ **Proposition 3.4.** *With probability at least $1 - 2/n$ over the random choice of $k_1, k_2, \ldots, k_r$ (which determines z) and A, it holds that if $K^t(A \mid z) \leq 2\lambda\ell$ then $\mathsf{cover}([n], \mathcal{S}) \leq 4\ell$.*

The proof of Proposition 3.4 is provided in Section 3.3. Proposition 3.3 together with Proposition 3.4 concludes that our reduction achieves a 4-approximation.

3.3 Proof of Proposition 3.4

Let Π be a RAM program such that $|\Pi| \leq 2\lambda\ell$ and $\Pi(z)$ prints A in $\leq t = t(|A|)$ CPU steps (where t is the running time bound associated with the problem $\mathsf{McK}^t\mathsf{P}[\zeta]$). The existence of such Π is implied by the assumption that $K^t(A \mid z) \leq 2\lambda\ell$. We will now show how to use Π, z to extract out a Set Cover of size 4ℓ. Towards this, recall that when executing $\Pi(z)$, in each CPU step, $\Pi(z)$ will read one bit from the memory. Let

$$q_1, q_2, \ldots, q_t$$

be the memory positions that $\Pi(z)$ reads in the execution of $\Pi(z)$ (such that in CPU step i, $\Pi(z)$ reads the content of memory position q_i). Note that the string z will be stored in the memory of $\Pi(z)$, and we are interested in the memory positions where the string z is stored. So, we let d be the memory position such that z is stored from position d to position $d + |z| - 1$. In addition, most of the bits in z are just zeros and $z_{k_1}, z_{k_2}, \ldots, z_{k_r}$ are the only informative blocks. (Recall that z is partitioned into 2^λ blocks of size $n \times m$.) Thus, let

$$p_i = \lfloor (q_i - d)/(n \times m) \rfloor$$

be the index of the block in z from which $\Pi(z)$ reads one bit in CPU step i. When p_i matches some key k_j, $z_{p_i} = z_{k_j} = W_j$. When p_i does not match any of the keys, $z_{p_i} = 0^{n \times m}$.[13]

We say that $\Pi(z)$ makes a $\mathsf{useful\ access}$ to the string z in CPU step i if there exists $j \in [r]$ such that $p_i = k_j$ and for all $i' < i$, $p_i \neq p_{i'}$. In other words, $\Pi(z)$ makes a $\mathsf{useful\ access}$ when it first reads some bit in the block z_{k_j} for some $j \in [r]$. We say that $\Pi(z)$ hits some block z_p if in some CPU step i, $\Pi(z)$ reads one bit from z_p.

[13] Here we discuss the string z constructed by the reduction, instead of the one stored in the memory. (So Π can not manipulate values in z.) Thus, when p_i is out of the range (e.g., $p_i < 0$), it still holds that $z_{p_i} = 0^{n \times m}$.

Bounding the number of useful accesses. We first present an upper-bound on the number of useful accesses. The following central claim shows that if the number of useful accesses is large, then the Kolmogorov complexity of Keys must be small.

▷ Claim 3.5. Let $\mathsf{Keys} = k_1||k_2||\ldots||k_r$ be the concatenation of $k_1, k_2, \ldots, k_r$. If $\Pi(z)$ makes α (or more) useful accesses to the string z, then

$$K(\mathsf{Keys} \mid A, \mathcal{S}) \leq |\Pi| + (r-\alpha)\lambda + \alpha(\log t + \log r) + O(\log n)$$

We defer the proof of Claim 3.5 to Section 3.4

We observe that since Keys are picked at random, their (conditional) Kolmogorov complexity is high.

▷ Claim 3.6. For all $A \in \{0,1\}^{n\times m}$, with probability $1 - 1/n$ (over the random choice of Keys), it holds that

$$K(\mathsf{Keys} \mid A, \mathcal{S}) \geq |\mathsf{Keys}| - \log n = r\lambda - \log n.$$

Proof. Note that the total number of RAM programs with description length $< r\lambda - \log n$ is at most $2^{r\lambda - \log n} \leq \frac{2^{r\lambda}}{n}$, while the total number of the choices of Keys is $2^{r\lambda}$); thus the claim follows. ◁

By combining Claim 3.5 and Claim 3.6 we get the following bound on the number of useful accesses.

▶ **Corollary 3.7.** *With probability* $1 - 1/n$ *over the random choice of* Keys*, if* $|\Pi| \leq 2\lambda\ell$*, it holds that* $\Pi(z)$ *makes at most* 4ℓ useful accesses

Proof. Assume not. Then by Claim 3.5,

$$\begin{aligned}
K(\mathsf{Keys} \mid A, \mathcal{S}) &\leq |\Pi| + (r-4\ell)\lambda + 4\ell(\log t + \log r) + O(\log n)\\
&\leq 2\lambda\ell + r\lambda - 4\lambda\ell + 4\ell(\log t + \log r) + O(\log n)\\
&\leq r\lambda - (2\lambda\ell - 4\ell(\log t + \log r) - O(\log n))\\
&\leq r\lambda - (2\cdot 4(\log t + \log r)\cdot\ell - 4\ell(\log t + \log r) - O(\log n))\\
&\leq r\lambda - (\frac{n}{\gamma} - O(\log n))\\
&< r\lambda - \log n
\end{aligned}$$

which contradicts Claim 3.6. ◀

Extracting a small Set Cover. We now turn to showing that we can extract a Set Cover from Π, z which is bounded in size by the number of useful accesses. We first show that if $\Pi(z)$ manages to output the string A, yet does not make useful accesses such that the union of all the blocks that are hit by $\Pi(z)$ equal A, then the Kolmogorov complexity of A must be small.

▷ Claim 3.8. Assume that

- $\Pi(z)$ makes α useful accesses;
- $\Pi(z)$ outputs the string A.
- $z_{p_1} \vee z_{p_2} \vee \ldots \vee z_{p_t} \neq A$; [14]

[14] When $p_i < 0$ or $p_i \geq 2^\lambda$, we assume that z_{p_i} is an all-zero string and $z_{p_i} = 0^{n\times m}$.

Then,

$$K(A \mid \mathcal{S}) \leq |\Pi| + (n-1)m + \alpha(\log t + \log r) + O(\log n)$$

We defer the proof of Claim 3.8 to Section 3.4.

We observe that since A is a random string, its (conditional) Kolmogorov complexity must be high.

▷ Claim 3.9. With probability $1 - 1/n$ (over the random choice of A), it holds that

$$K(A \mid \mathcal{S}) \geq |A| - \log n \geq nm - \log n.$$

Proof. Note that the total number of RAM programs with description length $< nm - \log n$ is at most $2^{nm-\log n} \leq \frac{2^{nm}}{n}$ (while the total number of the choices of A is 2^{nm}); thus the claim follows. ◁

Combining Claim 3.8 and Claim 3.9, we conclude that the union of all the blocks hit by $\Pi(z)$ must equal A (provided that $\Pi(z)$ prints the string A and makes at most 4ℓ useful accesses).

▶ **Corollary 3.10.** *With probability $1 - 1/n$ over the random choice of A, if $|\Pi| \leq 2\lambda\ell$, $\Pi(z) = A$, and $\Pi(z)$ makes at most 4ℓ useful accesses, it holds that $z_{p_1} \vee z_{p_2} \vee \ldots \vee z_{p_t} = A$.*

Proof. Assume not. Then by Claim 3.8,

$$\begin{aligned} K(A \mid \mathcal{S}) &\leq |\Pi| + (n-1)m + 4\ell(\log t + \log r) + O(\log n) \\ &\leq 2\lambda\ell + (n-1)m + 4\ell(\log t + \log r) + O(\log n) \\ &= nm - (m - (2\lambda\ell + 4\ell(\log t + \log r) + O(\log n))) \\ &< nm - \log n \quad (\text{since } m = n^3, \lambda \leq n, \ell \leq n) \end{aligned}$$

which contradicts to Claim 3.9. ◀

We finally show that if the union of all the blocks hit by $\Pi(z)$ matches A, then we can extract out a Set Cover whose size is bounded by the number of useful accesses $\Pi(z)$ made.

▷ Claim 3.11. If $\Pi(z)$ makes at most 4ℓ useful accesses and $z_{p_1} \vee z_{p_2} \vee \ldots \vee z_{p_t} = A$, then $\mathsf{cover}([n], \mathcal{S}) \leq 4\ell$.

Proof. Let α be the number of useful accesses made by $\Pi(z)$. Let

$$i_1, i_2, \ldots, i_\alpha$$

be the CPU steps when $\Pi(z)$ makes a useful access; that is, $i_1, i_2, \ldots, i_\alpha$ is a sequence of CPU step indices such that for each $l \in [\alpha]$, $\Pi(z)$ will make a useful access in CPU step i_l. (Recall that except for $z_{k_1}, \ldots, z_{k_r}$, the blocks in the string z are all-zero strings.) Recall that $\Pi(z)$ makes a useful access when it first reads some bit in the block z_{k_j} for some $j \in [r]$.) Thus, by the definition of useful access, it follows that

$$z_{p_{i_1}} \vee z_{p_{i_2}} \vee \ldots \vee z_{p_{i_\alpha}} = z_{p_1} \vee z_{p_2} \vee \ldots \vee z_{p_t} = A$$

Since $\Pi(z)$ makes a useful access in CPU step i_l, p_{i_l} must equal some key. We let $j_1, j_2, \ldots, j_\alpha \in [r]$ be a sequence of indices of the keys such that

$$(p_{i_1}, p_{i_2}, \ldots, p_{i_\alpha}) = (k_{j_1}, k_{j_2}, \ldots, k_{j_\alpha})$$

Note that (by the construction of the string z)

$$(W_{j_1}, W_{j_2}, \ldots, W_{j_\alpha}) = (z_{p_{i_1}}, z_{p_{i_2}}, \ldots, z_{p_{i_\alpha}})$$

Thus, it follows that

$$W_{j_1} \vee W_{j_2} \vee \ldots \vee W_{j_\alpha} = z_{p_{i_1}} \vee z_{p_{i_2}} \vee \ldots \vee z_{p_{i_\alpha}} = A$$

Finally, we argue that $S_{j_1}, S_{j_2}, \ldots, S_{j_\alpha}$ cover $[n]$, which concludes the proof (since $\alpha \leq 4\ell$). We recall that for each $l \in [\alpha]$, W_{j_l} is the gadget string for S_{j_l}. Furthermore, W_{j_l} is partitioned into n blocks, $W_{j_l,1}, W_{j_l,2}, \ldots, W_{j_l,n}$. For each block $b \in [n]$, $W_{j_l,b} = A_b$ if $b \in S_{j_l}$, and otherwise $W_{j_l,b} = 0^m$. Since $W_{j_1} \vee W_{j_2} \vee \ldots \vee W_{j_\alpha} = A$, it follows that for all blocks $b \in [n]$,

$$W_{j_1,b} \vee W_{j_2,b} \vee \ldots \vee W_{j_\alpha,b} = A_b.$$

Thus, for all $b \in [n]$, there must exist $l \in [\alpha]$ such that $b \in S_{j_l}$. We conclude that the sets $S_{j_1}, S_{j_2}, \ldots, S_{j_\alpha}$ indeed cover $[n]$. ◁

We can now conclude the proof of Proposition 3.4:

Proof of Proposition 3.4. By Corollary 3.7, with probability $1-1/n$, $\Pi(z)$ makes at most 4ℓ useful accesses. By Corollary 3.10, with probability $1-1/n$, it holds that $z_{p_1} \vee z_{p_2} \vee \ldots \vee z_{p_t} = A$. Finally by Claim 3.11, it holds that $\mathsf{cover}([n], \mathcal{S}) \leq 4\ell$, which happens with probability at least $1 - 2/n$ (by a union bound). ◀

3.4 Proof of Claim 3.5 and Claim 3.8

In both Claim 3.5 and Claim 3.8, the goal is to compress some strings (either Keys or A) provided that $\Pi(z)$ prints A. Towards doing this, we need be able to find a short representation of the information needed to perform the execution of $\Pi(z)$. Towards this, it will be helpful to track when $\Pi(z)$ makes a useful access. Furthermore, note that every useful access corresponds to some key k_j such that z_{k_j} stores the gadgets W_j of the set S_j. For each such useful access, we will also track this "key index" j. As we shall see, given Π, A and S, as well as the sequence of CPU steps and key indexes (of useful accesses), the whole execution of $\Pi(z)$ can be emulated without having access to z. In fact, as we shall formalize now, we actually do not even need the full content of A and S, but rather just the gadgets W_j corresponding to the sets hit by the useful accesses.

To formalize this, let $t = t(|A|)$ be the maximum number of CPU steps that $\Pi(z)$ can run, and let α be some integer bounded by the number of useful accesses made by $\Pi(z)$. We refer a pair of sequences of CPU steps and key indexes $\omega = ((i_1, i_2, \ldots, i_\alpha), (j_1, j_2, \ldots, j_\alpha)) \in [t]^\alpha \times [r]^\alpha$ as a configuration. We say that $\Pi(z)$ *matches* ω if the first time $\Pi(z)$ makes a useful access is in CPU step i_1 and $\Pi(z)$ reads one bit from the block $z_{k_{j_1}}$ (and recall that $z_{k_{j_1}} = W_{j_1}$), and the second time $\Pi(z)$ makes a useful access is in CPU step i_2 and $\Pi(z)$ reads one bit from the block $z_{k_{j_2}}$, and so on.

▶ **Lemma 3.12.** *Let $\alpha \in \mathbb{N}$, and $\omega = ((i_1, \ldots, i_\alpha), (j_1, \ldots, j_\alpha))$ be a* configuration *in $[t]^\alpha \times [r]^\alpha$. If $\Pi(z)$ matches ω then one can emulate $\Pi(z)$ for i_α CPU steps using the code of Π, the configuration ω, and $W_{j_1}, W_{j_2}, \ldots, W_{j_\alpha}$ (without having access to z).*

Proof. We now describe how to emulate the execution of $\Pi(z)$ for i_α steps using the code of Π, the configuration ω, and $W_{j_1}, W_{j_2}, \ldots, W_{j_\alpha}$. Recall that d is the memory position where z starts at; that is, z is stored in memory positions d to $d + |z| - 1$.

Given the code of Π, we start to emulate $\Pi(z)$ with the content of memory positions $d, d+1, \ldots, d+|z|-1$ (which are supposed to store z) set to 0. In the simulation, we keep track of all memory positions that $\Pi(z)$ has written to. In each CPU step i, if i matches some value in $\{i_1, i_2, \ldots, i_\alpha\}$ (and suppose $i = i_l$), we proceed as follows:

- Let q_i be the memory position which Π will read from in CPU step i and proceed as follows.
- Let $p_i = \lfloor (q_i - d)/(n \times m) \rfloor$. Put the string $W_{j_l} \in \{0,1\}^{n \times m}$ into the memory from position $d + p_i \times nm$ to position $d + p_i \times nm + nm - 1$, with the following exception: If Π has ever previously written into a memory between position $d + p_i \times nm$ and $d + p_i \times nm + nm - 1$, we keep those bits unchanged.
- Finally, let Π will read the bit from the memory (just as if the string z had been there), and we continue to emulate the execution of $\Pi(z)$ in the rest of CPU step i.

If i does not appear in $\{i_1, i_2, \ldots, i_\alpha\}$, we simply emulate the execution honestly. When $i = i_\alpha$, we stop to emulate $\Pi(z)$.

We argue, by induction, that the above procedure perfectly emulates the execution of $\Pi(z)$ in the first i_α CPU steps. For the base case, we consider CPU step $i = 0$, in which $\Pi(z)$ has not started yet, so the statement is trivially true. For any $i \leq i_\alpha$, we now assume that in all the steps $\leq i-1$, our simulation perfectly emulates $\Pi(z)$, and we will prove that also in CPU step i, the simulation does so as well. First note that if, in CPU step i, Π attempts to read from a memory position q_i that has (1) previously been written or read from, (2) the memory position is not within the range $[d, d+|z|-1]$, or (3) the memory access to q_i is not a useful access, then the induction step directly follows from the induction hypothesis and the fact that the step is performed in exactly the same way in the simulation as in the real execution. We thus only need to consider the case when the memory access to q_i is a useful access. But whenever this happens, by the induction hypothesis, the simulation will produce exactly the same content in the block of z where q_i is contained, as in the real execution of $\Pi(z)$. It thus follows that also this step is perfectly emulated.

Thus, we conclude that $\Pi(z)$ can be emulated for i_α steps using the code of Π, the configuration ω, and $W_{j_1}, W_{j_2}, \ldots, W_{j_\alpha}$. ◀

We are now ready to prove Claim 3.5, which we restate for the convenience of the reader.

▷ Claim 3.13 (Claim 3.5, restated). Let $\mathsf{Keys} = k_1||k_2||\ldots||k_r$ be the concatenation of $k_1, k_2, \ldots, k_r$. If $\Pi(z)$ makes α (or more) useful accesses to the string z, then

$$K(\mathsf{Keys} \mid A, \mathcal{S}) \leq |\Pi| + (r-\alpha)\lambda + \alpha(\log t + \log r) + O(\log n)$$

Proof. If $\Pi(z)$ makes at least α useful accesses, $\Pi(z)$ must match some configuration

$$\omega = ((i_1, i_2, \ldots, i_\alpha), (j_1, j_2, \ldots, j_\alpha))$$

where $\omega \in [t]^\alpha \times [r]^\alpha$. We let $\{j'_1, j'_2, \ldots, j'_{r-\alpha}\} = [r] - \{j_1, j_2, \ldots, j_\alpha\}$ be the set of key indices that do *not* appear in ω.

We consider the following program Π' that prints the string $\mathsf{Keys} = k_1||k_2||\ldots||k_r$ with the string A and the collection of sets $\mathcal{S}$ as auxiliary information. Π' has the values n, m, λ, α, t, r hardwired in it, and the code of Π' also includes the configuration ω, the code of Π, and the $r-\alpha$ keys $k_{j'_1}, k_{j'_2}, \ldots, k_{j'_{r-\alpha}}$. Π' first computes $W_{j_1}, W_{j_2}, \ldots, W_{j_\alpha}$ from A and $\mathcal{S}$. Π' then emulates the execution of $\Pi(z)$ (using the code of Π, the configuration ω, and $W_{j_1}, W_{j_2}, \ldots, W_{j_\alpha}$, using the method described in Lemma 3.12) for i_α CPU steps (recall that i_α is the CPU step when $\Pi(z)$ makes its α'th useful access). Let d be the index such that z is initially stored in the memory from position d to the position $d+|z|-1$. Let

$$q_1, q_2, \ldots, q_{i_\alpha}$$

be the memory positions that $\Pi(z)$ reads (such that in CPU step i, $\Pi(z)$ reads one bit from memory position q_i) in the first i_α CPU steps. We will decode $k_{j_1}, k_{j_2}, \ldots, k_{j_\alpha}$ from $q_1, q_2, \ldots, q_{i_\alpha}$ as follows: For each $i \leq i_\alpha$, let

$$p_i = \lfloor (q_i - d)/(n \times m) \rfloor$$

Since $\Pi(z)$ matches ω, it follows that

$$p_{i_1} = k_{j_1}, p_{i_2} = k_{j_2}, \ldots, p_{i_\alpha} = k_{j_\alpha}$$

Thus, Π' has access to $k_{j'_1}, k_{j'_2}, \ldots, k_{j'_{r-\alpha}}$ (hardwired) and can compute $k_{j_1}, k_{j_2}, \ldots, k_{j_\alpha}$ as specified above. Thus, Π' can recover and output the string $\mathsf{Key} = k_1 || k_2 || \ldots || k_r$.

Finally, we show that the description length of Π' is at most $|\Pi| + (r-\alpha)\lambda + \alpha(\log t + \log r) + O(\log n)$. To describe Π', we require:

- $|\Pi|$ bits to store the code of Π;
- $(r-\alpha)\lambda$ bits to store the $r-\alpha$ keys $k_{j'_1}, k_{j'_2}, \ldots, k_{j'_{r-\alpha}}$;
- $\alpha(\log t + \log r)$ bits to store the configuration ω.
- $O(\log n)$ bits to store the values n, m, λ, α, t, r.
- $O(1)$ bits to describe the CPU next-step machine.

Thus, the description length of Π' is at most $|\Pi| + (r-\alpha)\lambda + \alpha(\log t + \log r) + O(\log n)$, and from this we conclude that

$$K(\mathsf{Keys} \mid A, \mathcal{S}) \leq |\Pi| + (r-\alpha)\lambda + \alpha(\log t + \log r) + O(\log n)$$

which completes the proof. ◁

We next proceed to prove Claim 3.8, which we first restate:

▷ Claim 3.14 (Claim 3.8, restated). Assume that

- $\Pi(z)$ makes α useful accesses;
- $\Pi(z)$ outputs the string A.
- $z_{p_1} \vee z_{p_2} \vee \ldots \vee z_{p_t} \neq A$; [15]

Then,

$$K(A \mid \mathcal{S}) \leq |\Pi| + (n-1)m + \alpha(\log t + \log r) + O(\log n)$$

Proof. Consider some Π, z satisfying the pre-conditions of the claim. Since $\Pi(z)$ has the property that

$$z_{p_1} \vee z_{p_2} \vee \ldots \vee z_{p_t} \neq A,$$

and recalling that each z_{p_i} is divided n m-size blocks, $z_{p_i,1}, \ldots, z_{p_i,n}$, it follows that there exists a block index $b \in [n]$ such that for each block $z_{p_i} \in \{0,1\}^{n \times m}$ that $\Pi(z)$ reads, $z_{p_i,b} = 0^m$. In addition, note that $\Pi(z)$ makes α useful accesses, so $\Pi(z)$ must match some configuration

$$\omega = ((i_1, i_2, \ldots, i_\alpha), (j_1, j_2, \ldots, j_\alpha))$$

where $\omega \in [t]^\alpha \times [r]^\alpha$. Since $\Pi(z)$ matches ω, we know that

$$(W_{j_1}, W_{j_2}, \ldots, W_{j_\alpha}) = (z_{p_{i_1}}, z_{p_{i_2}}, \ldots, z_{p_{i_\alpha}})$$

[15] When $p_i < 0$ or $p_i \geq 2^\lambda$, we assume that z_{p_i} is an all-zero string and $z_{p_i} = 0^{n \times m}$.

Thus,

$$W_{j_1} \vee W_{j_2} \vee \ldots \vee W_{j_\alpha} \neq A$$

It follows that for all $l \in [\alpha]$, $W_{j_l,b} = 0^m$. From this, we can conclude that the gadget strings $W_{j_1}, W_{j_2}, \ldots, W_{j_\alpha}$ can be constructed from $\mathcal{S}$ and all randomized encodings $A_1, \ldots, A_{b-1}, A_{b+1}, \ldots, A_n$ *excluding* A_b.

Based on this observation, let us show how to construct a program Π' that outputs the string A given $\mathcal{S}$ as auxiliary information. The program Π' embeds the values n, m, λ, α, r, t, the value of b, the code of Π, the configuration ω, and strings $A_1, \ldots, A_{b-1}, A_{b+1}, \ldots, A_n$ into its code. Π' first computes $W_{j_1}, W_{j_2}, \ldots, W_{j_\alpha}$ from $A_1, \ldots, A_{b-1}, A_{b+1}, \ldots, A_n$ and $\mathcal{S}$. Π' then simulates the execution of $\Pi(z)$ using the code of Π, the configuration ω, and the gadget strings $W_{j_1}, W_{j_2}, \ldots, W_{j_\alpha}$ (making use of Lemma 3.12), and finally outputs whatever $\Pi(z)$ outputs. Note that since $\Pi(z)$ makes exactly α useful accesses, Π' can emulate $\Pi(z)$ all the way until it terminates. Furthemore, recall that by assumption $\Pi(z)$ outputs A, so Π' will do so as well.

We finally show that the description length of Π' is at most $|\Pi| + (n-1)m + \alpha(\log t + \log r) + O(\log n)$. To see this, note that to specify Π', we require:

- $|\Pi|$ bits to include the code of Π;
- $(n-1)m$ bits to store strings $A_1, \ldots, A_{b-1}, A_{b+1}, \ldots, A_n$;
- $\alpha(\log t + \log r)$ bits to save the configuration ω.
- $O(\log n)$ bits to strore the values n, m, λ, α, r, t, b
- $O(1)$ bits to implement the CPU next-step machine:

Thus, we have that the description length of Π' is at most $|\Pi| + (n-1)m + \alpha(\log t + \log r) + O(\log n)$. From this we conclude that

$$K(A \mid \mathcal{S}) \leq |\Pi| + (n-1)m + \alpha(\log t + \log r) + O(\log n).$$

which proves the claim. ◁

References

1 Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In Gary L. Miller, editor, *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996*, pages 99–108. ACM, 1996. doi:10.1145/237814.237838.

2 Miklós Ajtai and Cynthia Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In Frank Thomson Leighton and Peter W. Shor, editors, *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing, El Paso, Texas, USA, May 4-6, 1997*, pages 284–293. ACM, 1997. doi:10.1145/258533.258604.

3 Adi Akavia, Oded Goldreich, Shafi Goldwasser, and Dana Moshkovitz. On basing one-way functions on NP-hardness. In *STOC '06*, pages 701–710, 2006. doi:10.1145/1132516.1132614.

4 Eric Allender, Harry Buhrman, Michal Kouckỳ, Dieter Van Melkebeek, and Detlef Ronneburger. Power from random strings. *SIAM Journal on Computing*, 35(6):1467–1493, 2006.

5 Eric Allender, Mahdi Cheraghchi, Dimitrios Myrisiotis, Harsha Tirumala, and Ilya Volkovich. One-way functions and a conditional variant of MKTP. *Electron. Colloquium Comput. Complex.*, 28:9, 2021. Revision 2; October 19, 2021.

6 Eric Allender, Mahdi Cheraghchi, Dimitrios Myrisiotis, Harsha Tirumala, and Ilya Volkovich. One-way functions and a conditional variant of MKTP. *Electron. Colloquium Comput. Complex.*, 28:9, 2021. Revision 1; April 18, 2021.

7 Eric Allender, Mahdi Cheraghchi, Dimitrios Myrisiotis, Harsha Tirumala, and Ilya Volkovich. One-way functions and a conditional variant of MKTP. *Electron. Colloquium Comput. Complex.*, 28:9, 2021.

8 Eric Allender, Lisa Hellerstein, Paul McCabe, Toniann Pitassi, and Michael E. Saks. Minimizing disjunctive normal form formulas and ac^0 circuits given a truth table. *SIAM J. Comput.*, 38(1):63–84, 2008. doi:10.1137/060664537.

9 Manuel Blum. Coin flipping by telephone - A protocol for solving impossible problems. In *COMPCON'82, Digest of Papers, Twenty-Fourth IEEE Computer Society International Conference, San Francisco, California, USA, February 22-25, 1982*, pages 133–137. IEEE Computer Society, 1982.

10 Manuel Blum and Silvio Micali. How to generate cryptographically strong sequences of pseudo-random bits. *SIAM Journal on Computing*, 13(4):850–864, 1984.

11 Andrej Bogdanov and Christina Brzuska. On basing size-verifiable one-way functions on NP-hardness. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part I*, volume 9014 of *Lecture Notes in Computer Science*, pages 1–6. Springer, 2015.

12 Andrej Bogdanov and Luca Trevisan. On worst-case to average-case reductions for np problems. In *FOCS '03*, pages 308–317, 2003.

13 Gilles Brassard. Relativized cryptography. *IEEE Transactions on Information Theory*, 29(6):877–893, 1983.

14 Gregory J. Chaitin. On the simplicity and speed of programs for computing infinite sets of natural numbers. *J. ACM*, 16(3):407–422, 1969.

15 Whitfield Diffie and Martin Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654, 1976.

16 Uriel Feige and Adi Shamir. Witness indistinguishable and witness hiding protocols. In *STOC '90*, pages 416–426, 1990. doi:10.1145/100216.100272.

17 Rosario Gennaro and Luca Trevisan. Lower bounds on the efficiency of generic cryptographic constructions. In *41st Annual Symposium on Foundations of Computer Science, FOCS 2000, 12-14 November 2000, Redondo Beach, California, USA*, pages 305–313. IEEE Computer Society, 2000.

18 Halley Goldberg and Valentine Kabanets. A simpler proof of the worst-case to average-case reduction for polynomial hierarchy via symmetry of information. *Electronic Colloquium on Computational Complexity (ECCC)*, 2022. URL: https://eccc.weizmann.ac.il/report/2022/007.

19 Oded Goldreich, Shafi Goldwasser, and Silvio Micali. On the cryptographic applications of random functions. In *CRYPTO*, pages 276–288, 1984.

20 Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *J. Comput. Syst. Sci.*, 28(2):270–299, 1984.

21 S. Dov Gordon, Hoeteck Wee, David Xiao, and Arkady Yerukhimovich. On the round complexity of zero-knowledge proofs based on one-way permutations. In *LATINCRYPT*, pages 189–204, 2010.

22 Iftach Haitner, Mohammad Mahmoody, and David Xiao. A new sampling protocol and applications to basing cryptographic primitives on the hardness of NP. In *IEEE Conference on Computational Complexity*, pages 76–87, 2010.

23 J. Hartmanis. Generalized kolmogorov complexity and the structure of feasible computations. In *24th Annual Symposium on Foundations of Computer Science (sfcs 1983)*, pages 439–445, November 1983. doi:10.1109/SFCS.1983.21.

24 Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.

25 Shuichi Hirahara. Non-black-box worst-case to average-case reductions within NP. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018*, pages 247–258, 2018.

26 Shuichi Hirahara. Unexpected hardness results for kolmogorov complexity under uniform reductions. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy, editors, *Proccedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, pages 1038–1051. ACM, 2020.

27 Shuichi Hirahara. Symmetry of information in heuristica. Manuscript, 2021.

28 Shuichi Hirahara, Igor Carboni Oliveira, and Rahul Santhanam. Np-hardness of minimum circuit size problem for OR-AND-MOD circuits. In Rocco A. Servedio, editor, *33rd Computational Complexity Conference, CCC 2018, June 22-24, 2018, San Diego, CA, USA*, volume 102 of *LIPIcs*, pages 5:1–5:31. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.CCC.2018.5.

29 Rahul Ilango. AC^0[p] lower bounds and NP-hardness for variants of MCSP. *Electron. Colloquium Comput. Complex.*, 26:21, 2019. URL: https://eccc.weizmann.ac.il/report/2019/021.

30 Rahul Ilango. 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*, pages 34:1–34:26, 2020.

31 Rahul Ilango. Personal communication, 2021.

32 Rahul Ilango, Bruno Loff, and Igor Carboni Oliveira. NP-hardness of circuit minimization for multi-output functions. In *35th Computational Complexity Conference, CCC 2020*, pages 22:1–22:36, 2020.

33 Russell Impagliazzo. A personal view of average-case complexity. In *Structure in Complexity Theory '95*, pages 134–147, 1995.

34 Russell Impagliazzo and Michael Luby. 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*, pages 230–235, 1989.

35 Russell Impagliazzo and Moni Naor. Efficient cryptographic schemes provably as secure as subset sum. In *30th Annual Symposium on Foundations of Computer Science, Research Triangle Park, North Carolina, USA, 30 October - 1 November 1989*, pages 236–241. IEEE Computer Society, 1989.

36 Valentine Kabanets and Jin-yi Cai. Circuit minimization problem. In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, May 21-23, 2000, Portland, OR, USA*, pages 73–79, 2000.

37 Ker-I Ko. On the notion of infinite pseudorandom sequences. *Theor. Comput. Sci.*, 48(3):9–33, 1986. doi:10.1016/0304-3975(86)90081-2.

38 Ker-I Ko. On the complexity of learning minimum time-bounded turing machines. *SIAM J. Comput.*, 20(5):962–986, 1991.

39 A. N. Kolmogorov. Several theorems about algorithmic entropy and algorithmic amount of information (a talk at a moscow math. soc. meeting 10/31/67). *An abstract in Usp. Mat. Nauk*, 23(2):201, 1968.

40 A. N. Kolmogorov. Three approaches to the quantitative definition of information. *International Journal of Computer Mathematics*, 2(1-4):157–168, 1968.

41 L. A. Levin. The tale of one-way functions. *Problems of Information Transmission*, 39(1):92–103, 2003. doi:10.1023/A:1023634616182.

42 Leonid A. Levin. Universal search problems (russian), translated into English by BA Trakhtenbrot in [59]. *Problems of Information Transmission*, 9(3):265–266, 1973.

43 Yanyi Liu and Rafael Pass. 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*, pages 1243–1254. IEEE, 2020.

44 Yanyi Liu and Rafael Pass. On one-way functions from np-complete problems. Cryptology ePrint Archive, Report 2021/513, 2021. ; received on April 19, 2021. URL: https://ia.cr/2021/513.

45 Yanyi Liu and Rafael Pass. On one-way functions from np-complete problems. *Electron. Colloquium Comput. Complex.*, 28:59, 2021. URL: https://eccc.weizmann.ac.il/report/2021/059.

46 Yanyi Liu and Rafael Pass. On the possibility of basing cryptography on EXP $\neq$ BPP. In *CRYPTO*, 2021.

47 Noam Livne. On the construction of one-way functions from average case hardness. In *ICS*, pages 301–309. Citeseer, 2010.

48 Luc Longpré and Sarah Mocas. Symmetry of information and one-way functions. In Wen-Lian Hsu and Richard C. T. Lee, editors, *ISA '91 Algorithms, 2nd International Symposium on Algorithms, Taipei, Republic of China, December 16-18, 1991, Proceedings*, volume 557 of *Lecture Notes in Computer Science*, pages 308–315. Springer, 1991. doi:10.1007/3-540-54945-5_75.

49 Luc Longpré and Osamu Watanabe. On symmetry of information and polynomial time invertibility. In *Algorithms and Computation*, pages 410–419, Berlin, Heidelberg, 1992. Springer Berlin Heidelberg.

50 R. Merkle and M. Hellman. Hiding information and signatures in trapdoor knapsacks. *IEEE Transactions on Information Theory*, 24(5):525–530, 1978.

51 Moni Naor. Bit commitment using pseudorandomness. *J. Cryptology*, 4(2):151–158, 1991.

52 A. M. Odlyzko. The rise and fall of knapsack cryptosystems. In *In Cryptology and Computational Number Theory*, pages 75–88. A.M.S, 1990.

53 Rafael Pass. Parallel repetition of zero-knowledge proofs and the possibility of basing cryptography on NP-hardness. In *21st Annual IEEE Conference on Computational Complexity (CCC 2006), 16-20 July 2006, Prague, Czech Republic*, pages 96–110. IEEE Computer Society, 2006.

54 Hanlin Ren and Rahul Santhanam. Hardness of KT characterizes parallel cryptography. *Electron. Colloquium Comput. Complex.*, 28:57, 2021. URL: https://eccc.weizmann.ac.il/report/2021/057.

55 Ronald L. Rivest, Adi Shamir, and Leonard M. Adleman. A method for obtaining digital signatures and public-key cryptosystems (reprint). *Commun. ACM*, 26(1):96–99, 1983. doi:10.1145/357980.358017.

56 John Rompel. One-way functions are necessary and sufficient for secure signatures. In *STOC*, pages 387–394, 1990.

57 Michael Sipser. A complexity theoretic approach to randomness. In *Proceedings of the 15th Annual ACM Symposium on Theory of Computing, 25-27 April, 1983, Boston, Massachusetts, USA*, pages 330–335. ACM, 1983.

58 R.J. Solomonoff. A formal theory of inductive inference. part i. *Information and Control*, 7(1):1–22, 1964. doi:10.1016/S0019-9958(64)90223-2.

59 Boris A Trakhtenbrot. A survey of Russian approaches to perebor (brute-force searches) algorithms. *Annals of the History of Computing*, 6(4):384–400, 1984.

60 Luca Trevisan. Non-approximability results for optimization problems on bounded degree instances. In Jeffrey Scott Vitter, Paul G. Spirakis, and Mihalis Yannakakis, editors, *Proceedings on 33rd Annual ACM Symposium on Theory of Computing, July 6-8, 2001, Heraklion, Crete, Greece*, pages 453–461. ACM, 2001. doi:10.1145/380752.380839.

61 Andrew Chi-Chih Yao. Theory and applications of trapdoor functions (extended abstract). In *23rd Annual Symposium on Foundations of Computer Science, Chicago, Illinois, USA, 3-5 November 1982*, pages 80–91, 1982.

62 A. K. Zvonkin and L. A. Levin. the Complexity of Finite Objects and the Development of the Concepts of Information and Randomness by Means of the Theory of Algorithms. *Russian Mathematical Surveys*, 25(6):83–124, December 1970. doi:10.1070/RM1970v025n06ABEH001269.