

A Note on Zero-Knowledge for \mathbf{NP} and One-Way Functions

Yanyi Liu *

Noam Mazor †

Rafael Pass ‡

June 12, 2025

Abstract

We present a simple alternative exposition of the recent result of Hirahara and Nanashima (STOC'24) showing that one-way functions exist if (1) every language in \mathbf{NP} has a zero-knowledge proof/argument (i.e., $\mathbf{NP} \subseteq \mathbf{ZKA}$) and (2) \mathbf{ZKA} contains non-trivial languages (i.e., $\mathbf{ZKA} \not\subseteq \mathbf{iP/poly}$). Our presentation does not rely on meta-complexity and we hope it may be useful for didactic purposes, and may make extensions/generalization easier.

Indeed, using our modular approach, we can directly extend the result to hold also in the *uniform setting* (i.e., assuming $\mathbf{NP} \not\subseteq \mathbf{iOBPP}$). Additionally, we show that the same result holds for (imperfect) \mathbf{iO} for 3CNF, or Witness Encryption for \mathbf{NP} , solving an open problem from Komargodski et al (FOCS '14).

*Cornell Tech. E-mail: y12866@cornell.edu. Research partly supported by NSF CNS-2149305.

†Tel Aviv University. E-mail: noammaz@gmail.com. Research partly supported by NSF CNS-2149305 and DARPA under Agreement No. HR00110C0086.

‡Cornell Tech, Technion and Tel-Aviv University. E-mail: rafael@cs.cornell.edu. Supported in part by NSF Award CNS 2149305, AFOSR Award FA9550-23-1-0387, AFOSR Award FA9550-23-1-0312, AFOSR Award FA9550-24-1-0267, ISF Award 2338/23 and ERC Advanced Grant KolmoCrypt - 101142322. 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, the AFOSR, the European Union or the European Research Council Executive Agency.

1 Introduction

Zero-knowledge (ZK) proofs, introduced by Goldwasser, Micali and Rackoff [GMR89] are paradoxical constructs that enables a Prover to convince a Verifier that some instance x belongs to a language L without revealing any additional information. In the case where the soundness condition only holds against polynomial-size attacker, the proof systems is instead referred to as a *zero-knowledge argument* [BCC88]; let **ZKA** denote the class of languages having ZK arguments. While classic results in the early 1990's established that, assuming the existence of *one-way functions* (OWFs), every language in **NP** has a ZK proof [GMW91; HILL99; Nao91], it remains an open problem whether the existence of *non-trivial* ZK proofs/arguments also imply the existence of OWFs.

Seminal results by Ostrovsky [Ost91], and Ostrovsky and Wigderson (OW) [OW93] from the 1990's, however, show that non-trivial ZK “almost” implies OWFs. In particular, they show that if **ZKA** $\not\subseteq$ **ioP/poly**, then a relaxation of OWFs, referred to as a *auxiliary-input* OWF (ai-OWF) exist.¹ Additionally OW shows that if **ZKA** contains a language that is *average-case hard*, then (standard) OWFs exist.²

Thus, the only “gap” is between ai-OWFs and (standard) OWFs, or between worst-case and average-case hardness for a language in **ZKA**. A recent elegant paper by Hirahara and Nanashima (HN) [HN24] closes this gap under the assumption that all of **NP** (or even just a specific meta-complexity language) has zero-knowledge arguments (i.e., **NP** \subseteq **ZKA**). In other words, they show:

Theorem 1.1 (Main, reproving [HN24]). *Assume that $\mathbf{NP} \subseteq \mathbf{ZKA}$, and that $\mathbf{ZKA} \not\subseteq \mathbf{ioP/poly}$. Then one-way functions exist.*

In this note, we provide a somewhat alternative presentation of the proof of the HN result. The proof elements are very similar, but our exposition is more direct and we dispense of the use of meta-complexity. As such, we hope that their results becomes easier to appreciate (without any background on meta-complexity).

Additionally, our more modular viewpoint enables also directly obtaining a uniform version of their result:³

Theorem 1.2 (Main, uniform setting). *Assume that $\mathbf{NP} \subseteq \mathbf{ZKA}$, and that $\mathbf{ZKA} \not\subseteq \mathbf{ioBPP}$. Then uniformly-secure one-way functions exist.*

Proof Overview The proof proceeds in two steps:

1. **Errorless average-case hardness of ZKA implies OWFs.** As mentioned, OW already showed that average-case hardness of **ZKA** implies OWFs. We observe (following the approach in [HN24]) that essentially a direct combination of a characterization of Vadhan [Vad06] together with the results of Ostrovsky [Ost91] yields the stronger statement that *errorless*⁴ average-case

¹Roughly speaking, an ai-OWF is family of functions f_i such that no polynomial-size attacker A can invert *every* function in the family. That is, there exists some i on which A fails to invert f_i .

²In essence, when just assuming worst-case hardness of some language in **ZKA**, the index i is selected as an instance in the language on which the attacker will fail to decide the language, whereas in the case of average-case hardness, the index can be efficiently sampled so we get a standard OWF.

³We remark that for this result, ZK still needs to hold for *all* inputs (i.e., with a non-uniform notion of security).

⁴That is, we consider hardness against algorithms that either give the right answer or \perp , and that only output \perp with small probability.

hardness suffices (i.e., that $\mathbf{ZKA} \not\subseteq \mathbf{ioAvgBPP/poly}$).⁵ In essence, the reason why errorless average-case hardness suffices is that given a statement x , the reduction provided in these earlier works either correctly decides x , or fails to invert some ai-OWF candidate f_x ; but since the latter event is checkable, the reduction can easily be made errorless.

2. **Auxiliary-input OWFs imply average-case hardness of \mathbf{NP} .** We observe that the PRG construction of HILL [HILL99] yields an average-case hardness language in \mathbf{NP} even when instantiated with an ai-OWF: The language consists of all strings in the range of the PRG, and average-case hardness holds over the uniform distribution. The argument (which can be traced back to [Hir18], and made explicit for general PRGs in [LP21b]), is simple: an errorless algorithm for the uniform distribution needs to output 1 with high probability when given a uniform sample (since with high probability those strings are not in the range of the PRG), and output either 0 or \perp when given a sample in the range of the PRG, so simply interpret \perp as a 0 and we have a PRG distinguisher that works with high probability, which using the HILL reduction can be turned into an inverter for the (ai)-OWFs.⁶

So, by OW, non-trivial \mathcal{ZK} implies ai-OWF, which by step (2), and the assumption that $\mathbf{NP} \subseteq \mathbf{ZKA}$ implies that $\mathbf{ZKA} \not\subseteq \mathbf{ioAvgBPP/poly}$ which by step (1) implies OWFs. (Note that this results also hold just for honest-verifier zero-knowledge; in any case, by [OV07], a language has zero-knowledge argument if and only if it has honest-verifier zero-knowledge argument.)

1.1 Corollaries to iO and WE

We note an interesting new corollary of Theorem 1.1 (i.e., the main result of [HN24]) with respect to *indistinguishability obfuscation* (iO) [Bar+12] and *witness encryption for \mathbf{NP}* (WE) [GGSW13]. [Kom+14] showed that (imperfect) iO for polynomial-size circuits and the assumption that \mathbf{NP} is worst-case hard yield the existence of one-way functions. They further showed that iO for just the class of 3CNF formulas, or even WE for \mathbf{NP} (which is implied by iO for 3CNFs) together with *average-case hardness* of \mathbf{NP} yield one-way functions, but left open the problem of whether worst-case hardness of \mathbf{NP} suffices. We observe that Theorem 1.1, and Theorem 1.2 can be used to solve this—in fact, to get a strict improvement, our uniform strengthening of the result of [HN24] (i.e., Theorem 1.2) is crucial.

In particular, [Kom+14] show that these primitives can be used to construct statistical zero-knowledge arguments for \mathbf{NP} (see Theorem 5.2 and the following remark in [Kom+14]); next, by relying on Theorem 1.1, we can conclude the existence of one-way functions assuming just worst-case hardness of \mathbf{NP} .⁷

Corollary 1.3. *Assume that $\mathbf{NP} \not\subseteq \mathbf{ioP/poly}$. If there exists an (imperfect) iO for 3CNF formulas, or WE for \mathbf{NP} , then one-way functions exist.*

Using Theorem 1.2 we derive similar results in the uniform setting.

⁵Coincidentally, OW actually stated their result assuming that $\mathbf{ZKA} \not\subseteq \mathbf{ioAvgBPP/poly}$, but while the notation *Avg* typically denotes errorless average-case hardness in the literature, they defined it as two-sided error average-case hardness, so this is a strict strengthening of their result.

⁶In fact, essentially the same argument shows errorless average-case hardness of the *Minimum time-bounded Kolmogorov complexity problem* [Ko86], see Appendix A. We note that HN used a similar but more complicated argument to show average-case hardness of the Minimum Circuit Size problem [KC00] relying on the construction of a PRF [GGM84] from OWFs.

⁷In contrast, whereas [Kom+14] required average-case hardness, their result also applied in the uniform setting.

Corollary 1.4. *Assume that $\mathbf{NP} \not\subseteq \mathbf{ioBPP}$. If there exists an (imperfect) iO for 3CNF formulas, or WE for \mathbf{NP} , then uniformly-secure one-way functions exist.*

In fact, it suffices to assume WE for MCSP (using [HN24]) or $\mathbf{MK}^t\mathbf{P}$ (by Corollary A.3.)

2 Preliminaries and Definitions

Let \mathbf{ZKA} denote set of languages having zero-knowledge arguments [GMR89; BCC88], and let \mathbf{SZKP} denote the set of languages having statistical zero-knowledge proofs (see e.g., [SV97]). We proceed to defining the notion of errorless average-case hardness.

Definition 2.1 ($\mathbf{ioAvgBPP/poly}$). *A pair $(\mathcal{L}, \mathcal{D})$ of a language \mathcal{L} and a samplable distribution $\mathcal{D} = \{\mathcal{D}_n\}_{n \in \mathbb{N}}$ is in $\mathbf{ioAvgBPP/poly}$ if there exists a non-uniform PPT A such that the following holds for infinitely many $n \in \mathbb{N}$:*

- For every $x \in \text{Supp}(\mathcal{D}_n)$, $\Pr[A(x) \in \{\perp, \mathcal{L}(x)\}] \geq 0.9$
- $\Pr_{x \leftarrow \mathcal{D}_n}[A(x) = \perp] \leq 1/4$.

We recall that notion of an ai-OWF:

Definition 2.2 (ai-OWF). *A function family $\{f_a: \{0, 1\}^{m(n)} \rightarrow \{0, 1\}^{m(n)}\}_{a \in \{0, 1\}^n}$ is an auxiliary-input one-way function (ai-OWF) if for every non-uniform PPT \mathcal{A} there exists a negligible function μ such that for every $n \in \mathbb{N}$, there exists some $a \in \{0, 1\}^n$ such that*

$$\Pr_{x \leftarrow \{0, 1\}^{m(n)}}[\mathcal{A}(a, f_a(x)) \in f_a^{-1}(f_a(x))] \leq \mu(n) \quad (1)$$

If a is simply 0^n , we refer to the family as simply a one-way function (OWF). We say that $\{f_a\}$ is almost-everywhere hard on a set $\mathcal{I} \subset \{0, 1\}^$ if Equation (1) holds for every $a \in \mathcal{I} \cap \{0, 1\}^n$. We say that $\{f_a\}$ is uniformly-secure if the above only holds for every uniform PPT \mathcal{A} .*

Finally, let us recall the classic result by Ostrovksy and Wigderson [OW93]:

Theorem 2.3 (\mathcal{ZK} to ai-OWF, [OW93]).

- Assume that $\mathbf{ZKA} \not\subseteq \mathbf{ioP/poly}$. Then ai-OWF exists.
- Assume that $\mathbf{ZKA} \not\subseteq \mathbf{ioBPP}$. Then uniformly-secure ai-OWF exists.

3 Proof of Main Theorem

3.1 Step 1: $\mathbf{ZKA} \not\subseteq \mathbf{ioAvgBPP/poly} \Rightarrow \mathbf{OWFs}$

Recall that OV showed that *two-sided* error average-case hardness of \mathbf{ZKA} implies OWF; we here show the same by starting with just *errorless* average-case hardness of \mathbf{ZKA} .⁸

Lemma 3.1 (Errorless Avg-Hardness of \mathcal{ZK} to OWFs).

⁸This result also easily follows from two theorem statements in [HN24], but was not explicit stated as far as we can tell.

- Assume that there exists a language $\mathcal{L} \in \mathbf{ZKA}$ and samplable distribution \mathcal{D} such that $(\mathcal{L}, \mathcal{D}) \notin \mathbf{ioAvgBPP/poly}$. Then OWFs exist.
- Assume that there exists a language $\mathcal{L} \in \mathbf{ZKA}$ and samplable distribution \mathcal{D} such that $(\mathcal{L}, \mathcal{D}) \notin \mathbf{ioAvgBPP}$. Then uniformly-secure OWFs exist.

Towards proving this, we will rely on the following characterization of \mathbf{ZKA} of Vadhan [Vad06].

Definition 3.2 (SZK/OWF [Vad06]). A promise problem $\Pi = (\Pi_{\mathcal{Y}}, \Pi_{\mathcal{N}})$ satisfies the SZKP/OWF Condition if there is $\mathcal{I} \subseteq \Pi_{\mathcal{Y}} \cup \Pi_{\mathcal{N}}$ such that:

- The promise problem $(\Pi_{\mathcal{Y}} \setminus \mathcal{I}, \Pi_{\mathcal{N}} \setminus \mathcal{I})$ is in **SZKP**.
- There exists an auxiliary-input one-way function which is almost everywhere hard on \mathcal{I} .

Theorem 3.3 ([OV07]). If $(\mathcal{Y}, \mathcal{N}) \in \mathbf{ZKA}$ then $(\mathcal{Y}, \mathcal{N})$ satisfies the SZKP/OWF Condition.

We will also rely on the following version of the results by Ostrovsky (as explicitly stated in [Vad06].)

Theorem 3.4 ([Ost91] (c.f. [Vad06, Theorem 7.5])). Let $\Pi = (\Pi_{\mathcal{Y}}, \Pi_{\mathcal{N}}) \in \mathbf{SZKP}$. Then there exists a function family $\{f_x\}_{x \in \{0,1\}^*}$ and an oracle-aided PPT R , such that for every $x \in \Pi_{\mathcal{Y}} \cup \Pi_{\mathcal{N}}$ and any algorithm A that inverts f_x with probability at least 0.01, $\Pr[R^A(x) = \Pi(x)] \geq 0.99$.

We now turn to the proof of Lemma 3.1.

Proof of Lemma 3.1. We prove the first item in Lemma 3.1. Let $(\Pi = (\mathcal{L}, \bar{\mathcal{L}}), \mathcal{D}) \in (\mathbf{ZKA}) \setminus (\mathbf{ioAvgBPP/poly})$. Let \mathcal{I} and $\{h_x\}_{x \in \{0,1\}^*}$ be the set and the auxiliary-input one-way function promised by Theorem 3.3 and the SZK/OWF condition for Π . Let $\{g_x\}_{x \in \{0,1\}^*}$ and R be the auxiliary-input one-way function and reduction promised by Theorem 3.4 for $\Pi' = (\Pi_{\mathcal{Y}} \setminus \mathcal{I}, \Pi_{\mathcal{N}} \setminus \mathcal{I})$.

Let $f(r, y_1, y_2) = \mathcal{D}(r) || h_{\mathcal{D}(r)}(y_1) || g_{\mathcal{D}(r)}(y_2)$. We claim that f is a weak one-way function. Indeed, assume towards a contradiction that an efficient algorithm A inverts f with probability 0.99 for infinitely many input lengths n , and fix such large enough n . Let A' be the algorithm that uses A to invert g_x : given an input $x, g_x(y)$, A' samples y_1 and executes $A(x || h_x(y_1) || g_x(y))$ to get a pre-image of $g_x(y)$. Let B be the algorithm that given input x , samples random y_1, y_2 , and runs A on $z = x || h_x(y_1) || g_x(y_2)$. If A failed in inverting f on z , B outputs \perp . Otherwise, B outputs $R^{A'}(x)$. We next show that B contradicts the assumption that $\Pi \notin \mathbf{ioAvgBPP/poly}$.

Observe that for every r such that $\mathcal{D}(r) \in \mathcal{I}$, A' can only invert $g_{\mathcal{D}(r)}$ with negligible probability and consequently, A inverts $f(r, \cdot)$ only with negligible probability. Thus B outputs \perp on every (sufficiently large) such input with probability 0.99. On the other hand, for every r such that $\mathcal{D}(r) \notin \mathcal{I}$, B outputs either \perp or the right answer with probability at least 0.99. Thus B outputs the wrong answer with probability at most 0.01 for any x . This implies the first item in Definition 2.1.

Next, to see that the second item holds, observe that for any r such that

$$\Pr_y[A(f(r, y) \in f^{-1}(f(r, y)))] \geq 0.9,$$

B outputs a (non- \perp) right answer with probability at least 0.9 (probability of running B) - 0.01 (probability that B outputs an incorrect answer) = 0.89. Moreover, by an averaging argument,

$$\Pr_r[\Pr_y[A(f(r, y) \in f^{-1}(f(r, y)))] \geq 0.9] \geq 0.9.$$

(since otherwise, it cannot be that A inverts f with probability 0.99). Thus, B outputs the right answer with probability at least $0.89 \cdot 0.9 \geq 3/4$.

The second item of Lemma 3.1 follows in identically the same way but just considering a uniform algorithm A . \square

3.2 Step 2: ai-OWF \Rightarrow NP $\not\subseteq$ ioAvgBPP/poly

We turn to observing that ai-OWFs imply (errorless) average-case hardness of **NP**; this observation may be folklore in the community but, as far as we can tell, was first explicitly stated in [HN24] using a somewhat more complicated proof for a stronger statement (in particular, they proved not only that **NP** is average-case hard but also that the particular MCSP problem [KC00] is so).⁹

Lemma 3.5 (ai-OWF to AvgBPP hardness).

- Assume that ai-OWF exists. Then there exists a language $\mathcal{L} \in \mathbf{NP}$ such that $(\mathcal{L}, \{U_{m(n)}\}_{n \in \mathbb{N}}) \notin \mathbf{ioAvgBPP/poly}$ for some $m \in \text{poly}$.
- Assume that uniformly-secure ai-OWF exists. Then there exists a language $\mathcal{L} \in \mathbf{NP}$ such that $(\mathcal{L}, \{U_{m(n)}\}_{n \in \mathbb{N}}) \notin \mathbf{ioAvgBPP}$ for some $m \in \text{poly}$.

Proof. We prove the first item of Lemma 3.5. The second item follows in identically the same way.

Let $\{f_x\}_{x \in \{0,1\}^*}$ be an ai-OWF. For every $x \in \{0,1\}^*$, let $G_x: \{0,1\}^{m(|x|)} \rightarrow \{0,1\}^{2m(|x|)}$ be the PRG construction of HILL from f_x with $m \in \text{poly}$ such that $m(|x|) > 2|x|$. [HILL99]. By [HILL99] it holds that (1) $G_x(s)$ can be efficiently computed given x, s , and (2) there is a reduction from distinguishing the output of G_x from uniform and inverting f_x . Let

$$\mathcal{L}_{\text{HILL}} = \left\{ y: \exists x \in \{0,1\}^*, s \in \{0,1\}^{m(|x|)} \text{ s.t. } G_x(s) = y \right\},$$

and let $\mathcal{D} = \{\mathcal{D}_n = U_{2m(n)}\}_{n \in \mathbb{N}}$. We claim that $(\mathcal{L}_{\text{HILL}}, \mathcal{D}) \notin \mathbf{ioAvgBPP/poly}$. To see this, assume toward contradiction that $(\mathcal{L}_{\text{HILL}}, \mathcal{D}) \in \mathbf{ioAvgBPP/poly}$, and let A be the algorithm that decides $\mathcal{L}_{\text{HILL}}$ with good probability over the \mathcal{D}_n for infinite many n 's. We show that A can be used to invert $\{f_x\}$ for every $x \in \{0,1\}^n$ and for infinitely many n 's.

Indeed, fix n such that A succeed on \mathcal{D}_n , and fix $x \in \{0,1\}^n$. Observe that $G_x(s) \in \mathcal{L}_{\text{HILL}}$ for any $s \in \{0,1\}^{m(n)}$. Thus, A outputs 1 or \perp on x with probability at least 0.9 on every output of G_x . On the other hand,

$$|\mathcal{L}_{\text{HILL}} \cap \{0,1\}^{2m(n)}| \leq 2^n \cdot 2^{m(n)} \leq 2^{2m(n)-n}.$$

Therefore, it holds that $\Pr_{y \leftarrow U_{2m(n)}}[y \in \mathcal{L}_{\text{HILL}}] \leq 2^{-n}$, and thus

$$\Pr_{y \leftarrow U_{2m(n)}}[A(y) \neq 0] \leq 2^{-n} + 0.1 + 1/4 \leq 1/2.$$

We get that A distinguishes the output of G_x from random with advantage at least $0.9 - 0.5 = 0.4$. \square

⁹As mentioned, our proof directly extends also to showing that the MK^tP problem [Ko86] also is average-case hard—see Appendix A for more details—but this is not of relevance for proving the main result.

3.3 Concluding the Proof of Theorem 1.1 and Theorem 1.2

Proof of Theorem 1.1. Assume that $\mathbf{ZKA} \not\subseteq \mathbf{ioP/poly}$. By Theorem 2.3 ai-OWF exist. By Lemma 3.5, there exists $\mathcal{L} \in \mathbf{NP}$ and samplable distribution \mathcal{D} such that $(\mathcal{L}, \mathcal{D}) \notin \mathbf{ioAvgBPP/poly}$. Finally, since by assumption $\mathbf{NP} \subseteq \mathbf{ZKA}$, we get by Lemma 3.1 that OWFs exist. \square

Proof of Theorem 1.2. The proof is identical to that of Theorem Theorem 1.1 except that we use the second part of Theorem Theorem 2.3 to conclude the existence of uniformly-secure ai-OWF. For concreteness:

Assume that $\mathbf{ZKA} \not\subseteq \mathbf{ioBPP}$. By Theorem 2.3 uniformly-secure ai-OWF exist. By Lemma 3.5, there exists $\mathcal{L} \in \mathbf{NP}$ and samplable distribution \mathcal{D} such that $(\mathcal{L}, \mathcal{D}) \notin \mathbf{ioAvgBPP}$. Finally, since by assumption $\mathbf{NP} \subseteq \mathbf{ZKA}$, we get by Lemma 3.1 that OWFs exist. \square

References

- [Bar+12] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. “On the (im)possibility of obfuscating programs”. In: *J. ACM* 59.2 (2012), 6:1–6:48. DOI: 10.1145/2160158.2160159. URL: <https://doi.org/10.1145/2160158.2160159> (cit. on p. 3).
- [BCC88] Gilles Brassard, David Chaum, and Claude Crépeau. “Minimum Disclosure Proofs of Knowledge”. In: *Journal of Computer and System Sciences* (1988), pp. 156–189 (cit. on pp. 2, 4).
- [GGM84] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. “On the Cryptographic Applications of Random Functions”. In: *Advances in Cryptology: Proceedings of CRYPTO 84*. 1984, pp. 276–288 (cit. on p. 3).
- [GGSW13] Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. “Witness encryption and its applications”. In: *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*. 2013, pp. 467–476 (cit. on p. 3).
- [GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. “The Knowledge Complexity of Interactive Proof Systems”. In: *SIAM Journal on Computing* (1989). Preliminary version in *STOC’85*, pp. 186–208 (cit. on pp. 2, 4).
- [GMW91] Oded Goldreich, Silvio Micali, and Avi Wigderson. “Proofs that Yield Nothing But Their Validity or All Languages in NP Have Zero-Knowledge Proof Systems”. In: *Journal of the ACM* (1991). Preliminary version in *FOCS’86*, pp. 691–729 (cit. on p. 2).
- [HILL99] Johan Hastad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. “A pseudorandom generator from any one-way function”. In: *SIAM Journal on Computing* (1999), pp. 1364–1396 (cit. on pp. 2, 3, 6).
- [Hir18] Shuichi Hirahara. “Non-black-box worst-case to average-case reductions within NP”. In: *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2018, pp. 247–258 (cit. on pp. 3, 9).
- [HN24] Shuichi Hirahara and Mikito Nanashima. “One-Way Functions and Zero Knowledge”. In: *STOC*. 2024 (cit. on pp. 2–4, 6).

- [KC00] 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*. 2000, pp. 73–79 (cit. on pp. 3, 6).
- [Ko86] Ker-I Ko. “On the Notion of Infinite Pseudorandom Sequences”. In: *Theor. Comput. Sci.* 48.3 (1986), pp. 9–33. DOI: 10.1016/0304-3975(86)90081-2. URL: [https://doi.org/10.1016/0304-3975\(86\)90081-2](https://doi.org/10.1016/0304-3975(86)90081-2) (cit. on pp. 3, 6, 9).
- [Kol68] A. N. Kolmogorov. “Three approaches to the quantitative definition of information”. In: *International Journal of Computer Mathematics* 2.1-4 (1968), pp. 157–168 (cit. on p. 9).
- [Kom+14] Ilan Komargodski, Tal Moran, Moni Naor, Rafael Pass, Alon Rosen, and Eylon Yogev. “One-way functions and (im) perfect obfuscation”. In: *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*. IEEE. 2014, pp. 374–383 (cit. on p. 3).
- [LP21a] Yanyi Liu and Rafael Pass. “Cryptography from sublinear-time average-case hardness of time-bounded Kolmogorov complexity”. In: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 2021, pp. 722–735 (cit. on p. 9).
- [LP21b] Yanyi Liu and Rafael Pass. “On the Possibility of Basing Cryptography on $\text{EXP} \neq \text{BPP}$ ”. In: *CRYPTO*. 2021 (cit. on p. 3).
- [Nao91] Moni Naor. “Bit Commitment Using Pseudorandomness”. In: *Journal of Cryptology* (1991), pp. 151–158 (cit. on p. 2).
- [Ost91] Rafail Ostrovsky. “One-Way Functions, Hard on Average Problems, and Statistical Zero-Knowledge Proofs”. In: *Proceedings of the 6th Annual Structure in Complexity Theory Conference*. IEEE Computer Society, 1991, pp. 133–138 (cit. on pp. 2, 5).
- [OV07] Shien Jin Ong and Salil Vadhan. “Zero Knowledge and Soundness are Symmetric”. In: 2007, pp. 187–209 (cit. on pp. 3, 5).
- [OW93] Rafail Ostrovsky and Avi Wigderson. “One-Way Functions are Essential for Non-Trivial Zero-Knowledge”. In: *Proceedings of the 2nd Israel Symposium on Theory of Computing Systems*. IEEE Computer Society, 1993, pp. 3–17 (cit. on pp. 2, 4).
- [SV97] Amit Sahai and Salil P. Vadhan. “A Complete Promise Problem for Statistical Zero-Knowledge”. In: *focs8*. 1997, pp. 448–457 (cit. on p. 4).
- [Vad06] Salil P. Vadhan. “An Unconditional Study of Computational Zero Knowledge”. In: *SIAM Journal on Computing* (2006), pp. 1160–1214 (cit. on pp. 2, 5).

A On Average-case Hardness of MK^tP

We here briefly observe that ai-OWF imply average-case hardness of the Minimum Kolmogorov complexity problem, MK^tP [Kol68; Ko86]; this result may be folklore in the community but as far as we know has not been explicitly stated (although a close variant of it can be found in [LP21a], following the approach in [Hir18]).

Recall that for a threshold $s = s(n)$, $\text{MK}^t\text{P}[s]$ is the language of all strings $x \in \{0,1\}^*$ with $K^t(x) \leq s(|x|)$. Similarly, $\text{MK}^t\text{P}[s_0, s_1]$ is the promise problem in which the Yes instances are the strings with $K^t(x) \leq s_0(|x|)$, and the No instances are the strings with $K^t(x) > s_1(|x|)$. We prove the following theorem.

Theorem A.1. *Let $\epsilon > 0$ be a constant. Assuming that ai-OWF exist, $(\text{MK}^t\text{P}[\epsilon n, (1-\epsilon)n], U_{m(n)}) \notin \text{ioAvgBPP/poly}$ for any $t(n) \geq n^{1+\epsilon}$ and some $m \in \text{poly}$.*

Proof. Assume toward contradiction that $(\text{MK}^t\text{P}[\epsilon n, (1-\epsilon)n], U_{m(n)}) \in \text{ioAvgBPP/poly}$ for any $m \in \text{poly}$. We claim that there is no ai-OWF. Indeed, we can use the MK^tP solver to invert any function family $\{f_x\}_{x \in \{0,1\}^*}$ on every x of length n , for infinite many n 's.

To see that, we use each function f_x to construct a PRG $G_x: \{0,1\}^{\epsilon \cdot m(n)/2} \rightarrow \{0,1\}^{m(n)}$ for some polynomial $m(n)$, such that (1) $\epsilon \cdot m(n)/2 \geq 2n$, (2) $G_x(z)$ can be computed in time at most $(m(n))^{1+\epsilon}$ given x and an input z , and (3) there is a reduction from distinguishing the output of G_x from uniform and inverting f_x .

Let A be the zero-error algorithm that for $(\text{MK}^t\text{P}[\epsilon n, (1-\epsilon)n], U_{m(n)})$. Using the family $\{G_x\}$ and the algorithm A we can invert $\{f_x\}$ almost everywhere by the observation that A distinguish between the output of G_x and uniform $m(n)$ bit string for every choice of x of length n . Indeed, by the correctness of A , A must output No with probability at least $3/4 - \text{neg}(n)$ over the uniform distribution. On the other hand, $K^t(G_x(z)) \leq |z| + |x| + O(\log n) \leq \epsilon \cdot m(n)/2 + n + (\log n) \leq \epsilon \cdot m(n)$ for any z, x . Thus, A must output \perp or Yes on any output of G_x .

Finally, to construct G_x we use the PRG construction of HILL from the one-way function f_x to get a function $G': \{0,1\}^{\epsilon \cdot m'(n)/2} \rightarrow \{0,1\}^{m'(n)}$. Let p' be a polynomial that bound the running time of G' , and let $G_x(z_1, \dots, z_{q(n)}) = G'_x(z_1) \parallel \dots \parallel G'_x(z_{q(n)})$. Then G' can be computed in time roughly $q(n) \cdot p'(n)$, and the output length of G' is $m(n) := q(n) \cdot m'(n)$. By taking $q(n) \geq \max\{(p'(n))^{1/\epsilon}, 2n/\epsilon\}$ we get that G_x can be computed in time at most $(m(n))^{1+\epsilon}$. \square

We directly get the following corollaries; the second one using the proof of Theorem 1.1.

Corollary A.2. *Let $\epsilon > 0$ be a constant. Assuming that ai-OWF exist, $(\text{MK}^t\text{P}[(1-\epsilon)n], U_{m(n)}) \notin \text{ioAvgBPP/poly}$ for any $t(n) \geq n^{1+\epsilon}$ and some $m \in \text{poly}$.*

Corollary A.3. *Assume that $\text{MK}^t\text{P}[\epsilon n, (1-\epsilon)n] \in \text{ZKA}$ for some constant ϵ , and that $\text{ZKA} \notin \text{ioP/poly}$. Then one-way functions exist.*