Power of Information Loss:
Implications for One-Wayness and Obfuscation

Pouria Fallahpour[1], Alex B. Grilo[1], Garazi Muguruza[2], and Mahshid Riahinia[3]*

[1] Sorbonne Université, CNRS and LIP6, France
[2] QuSoft, Informatics Institute, University of Amsterdam, Netherlands
[3] DIENS, École Normale Supérieure, CNRS, Inria, PSL University, Paris, France

Abstract. We study the class of problems admitting *lossy reductions*—mappings that lose information about their inputs, e.g., via compression, mixing, randomization, or obfuscation. A series of works use such reductions to derive cryptographic and complexity-theoretic results. In particular, Ball et al. (ITCS 2020) use lossy reductions to build one-way functions (OWF). Despite their importance, a unified definition capturing a broad range of intuitively-lossy reductions and enabling analysis of their power remains lacking. We fill this gap by studying the power of lossy reductions, both definition-wise and implication-wise. We introduce *sparse lossiness*, a flexible and fine-grained notion of lossiness that provably captures a broad range of reductions. Our detailed analysis offers a framework that provides substantial flexibility and lays the foundation for the subsequent results:

Positive results: We show that fine-grained OWFs exist if any problem Π has either (i) an AND, OR-reduction that compresses m instances of size n to $mn^{1-\varepsilon}$ bits for $\varepsilon > 0$, or (ii) a worst-case to average-case Karp reduction with constant error and $O(1/\sqrt{n})$ statistical distance from the average distribution, where both run in time slightly better than the best worst-case solver of Π. We extend this to OWFs under a more restricted regime of parameters and discuss generalizations of these statements. We partially extend these findings to the quantum setting, basing the existence of one-way state generators from quantum sparsely lossy reductions.

Impossibility results: Using sparse lossiness, we derive strong impossibility results for the existence of (imperfect) statistical obfuscation (sO). An obfuscation is α-statistical if the distributions of the output over two equivalent circuits have statistical distance at most α. We show that unless the Polynomial Hierarchy collapses, α must be negligibly close to 1. This is improved to inverse-exponentially close to 1 under a variant of Exponential Time Hypothesis. This significantly enlarges the parameter sets for which sO is impossible, leaving only a narrow space for the existence of iO. We also obtain new instance compression and instance randomization impossibilities for kSAT. These impossibilities have various applications in kernelization and parameterized complexity.

Keywords: one-way functions, obfuscation, one-way state generators, fine-grained cryptography, worst-case to average-case reductions .

*Part of this work was done when the author was visiting IRIF, Université Paris Cité, Paris, France.

Table of Contents

1 Introduction

Reductions are among the most fundamental tools in complexity theory and modern cryptography, and they have been the subject of extensive study. They play a central role in proving security of cryptographic schemes, enabling efficiency, and providing separations across the landscape of cryptographic schemes, and establish the limits of our computational models. Consequently, understanding the properties of reductions is the backbone of cryptography and complexity theory. A property of reductions that has recently come to light is *lossiness* that refers to when the reduction loses some information about the input in an irreversible way. Intuitively, such loss of information can occur in various ways—for example, through compression, mixing, randomization, or obfuscation. Take compression as an example; it produces an output string shorter than the input and therefore necessarily discards some information. Moreover, *lossy reductions* naturally arise in several cryptographic settings, including worst-case to average-case reductions, randomized encodings, and homomorphic evaluations. However, despite their broad presence, we still lack a unified definition that captures all—or even a wide range—of such reductions.

A series of works by Harnik and Naor [44], Fortnow and Santhanam [37], Drucker [35], and Ball et al. [10], have brought new insights to the use of lossy reductions in cryptography as well as complexity theory. For example, [35, 37] leverages the lossy behaviour of compressing reductions to prove that the k-Satisfiability problem (kSAT) does not have compressing reductions (within some particular parameters) unless $\mathsf{NP} \subseteq \mathsf{coNP/Poly}$, i.e., the Polynomial Hierarchy PH collapses to its third level. Inspired by this approach, Ball et al. [10] introduced the following definition of lossiness: A reduction R is lossy if $I(X; R(X)) \ll n$ for *all distributions* X on n-bit inputs. They then leverage worst-case hardness of any problem $\Pi \in \mathsf{SZK}$ that admits a lossy reduction (within restricted parameters) with respect to their definition to build—arguably the most fundamental tools in cryptography—*one-way functions* (OWFs). While [10] claims that worst-case to average-case reductions and randomized encodings are lossy as well, to the best of our knowledge, their arguments hold under the condition that the output of the reduction is fully independent from the input. Note that OWFs are central in cryptography and can be viewed as the minimal assumption required for cryptography. Therefore, the result of Ball et al. [10] is an evidence to the importance of studying lossy reductions and highlights their implicatory power for cryptography.

Various questions come to mind in this regard: Is there a mathematical definition of lossiness that captures a broad range of intuitive examples of lossy reductions? What is the minimal amount of lossiness for a reduction of runtime T that can be used to build OWFs? How much a reduction of runtime T can compress kSAT, without contradicting complexity-theory assumptions? In other words:

What are the limits of lossy reductions, both definition-wise and implication-wise?

This perspective can be taken even further by viewing cryptographic schemes themselves as forms of lossy reductions. From lossy trapdoor functions and homomorphic secret sharings to obfuscation and collision-resistant hash functions, these protocols exhibit, in one way or another, a lossy behavior. We can therefore ask: *Is it possbile to interpret any of these schemes as lossy reductions for particular computational problems? And if so, what would this imply about their existence or non-existence in a generic way?*

In praise of generic approaches. Unconditional proofs for the existence of cryptographic primitives almost never happen. Instead, their existence is often based on the *worst-case hardness* of concrete computational problems such as DLOG [34], RSA [69], lattice and code-based problems [3, 5, 62, 68] and more. All of these problems lie in the complexity class $\mathsf{NP} \cap \mathsf{coNP}$ [1, 2, 23, 24], yet it is unknown if they are complete for any class—therefore if one of these problems were efficiently solved in the future, it would likely have limited impact on general complexity theory. In other words, while a polynomial-time algorithm for any of these problems would have considerable implications for cryptography, the widespread belief of their worst-case hardness is based less on concrete *complexity-theoretical* evidence and more on the failed *algorithmic* attempts in directly solving them. In this regard, these computational problems offer only *weak* evidence for the existence of cryptography. In the broader sense, relying on just a few specific problems remains limiting.

A more robust approach for building cryptography is to look for *stronger* evidence, i.e., by basing the existence of cryptographic schemes on the worst-case hardness of a *class* of problems. This can be done for example by assuming the worst-case hardness of a class-complete problem such as an SZK-complete or an NP-complete problem, or by assuming worst-case hardness of a large class of problems

satisfying a generic property. These types of connections provide stronger evidence for the existence of cryptographic protocols in the sense that refuting them, i.e., if cryptography built from them does not exist, would imply unexpected or counterintuitive results on much wider sets of problems. Similarly, refuting the existence of a cryptographic scheme should also ideally be based on strong evidences. Such a generic approach, therefore, provides a more robust method for proving the *non-existence* of cryptographic schemes.

As previously mentioned, lossiness occurs across a wide range of algorithms and cryptographic settings. Studying lossy reductions is therefore aligned with this generic approach and holds the potential to provide strong evidence for cryptography and complexity theory. Formally defining the class of problems that admit lossy reductions, analyzing its magnitude and connections to cryptography is therefore of high importance.

1.1 Our Contribution

In this work, we address the questions raised above in three complementary directions:

1. **A new notion:** We introduce a flexible and fine-grained notion of lossiness, which we term *sparse lossiness*, that provably captures a wide range of worst-case to average-case reductions and randomized encodings, as well as the earlier definition from [10]. More precisely, by carefully analyzing these reductions, we relate their concrete sparse lossiness to properties such as the error of the reduction, the privacy of the randomized encoding, or the distance between the output distribution of a worst-case to average-case reduction and the target average-case distribution. This detailed analysis provides substantial flexibility and lays the foundation for the subsequent results, which support the validity and usefulness of our approach.
2. **Positive result:** We use sparsely lossy reductions to study the connection between fine-grained worst-case hardness and the existence of one-wayness. By leveraging sparse lossiness, we strengthen the results of [10] to build OWFs using a larger class of problems and further extend their result to *fine-grained one-way functions* (FGOWFs). OWFs are undoubtedly the most important cryptographic primitive and their fine-grained variant has recently gained much interest [11,25,29,58]. We also present a possible approach to build FGOWFs from the subexponential hardness of kSAT. Finally, we partially extend these findings to the quantum setting by drawing connections between *one-way state generators* (OWSGs) and quantum sparsely lossy reductions.
3. **Impossibility result:** We show that the powerful machinery of sparsely lossy reductions can be used to derive strong impossibility results on the existence of statistical obfuscation schemes. Roughly speaking, an obfuscation is an algorithm that compiles a circuit into a functionally equivalent one while (statistically or computationally) hiding its white-box aspects, such as its hard-wired secrets or its explicit code. Numerous works are dedicated to study notions of obfuscation (e.g. see [13, 21, 42]), its existence (e.g. see [12, 14, 38, 55, 56, 82]) and its importance in cryptography as well as complexity theory (e.g. see [9, 17, 26, 31, 32, 39, 41, 72]). We significantly enlargen the parameter sets for which statistical obfuscation is impossible under reasonable assumptions. Moreover, we obtain new instance compression and instance randomization impossibilites for kSAT. The instance compression of kSAT has various applications in kernelization and parameterized complexity (e.g. see [33, Ch. 15]).

Setting the context. For a Boolean function $f : \{0,1\}^m \to \{0,1\}$, an f-reduction from a promise problem Π to Π' is a mapping R such that $R(x_1, \cdots, x_m)$ is a YES instance in Π' *iff* $f(\chi_\Pi(x_1), \cdots, \chi_\Pi(x_m)) = 1$, where χ_Π is the characteristic function of Π. In our results, f can be any non-constant function that remains invariant under applying a permutation on its input bits. This notably includes OR, AND, MAJ, PARITY, MOD_k, and THRESHOLD_k.

In this work, we are exclusively interested in the non-uniform setting where algorithms are allowed to have additional advice that only depends on the size of the inputs. In this context, we introduce the following quantity that reflects the exact worst-case hardness of a problem: for a problem Π, let $\tau_\Pi(n)$ be the infimum value of T taken over all non-uniform Turing machines with advice that solve *any* instance of size n of Π in time $O(2^T)$. Note that $\tau_\Pi(n) \leq n$. This is because an algorithm given $\Pi_{YES} \cap \{0,1\}^n$ as advice, which is of size 2^n can solve any instance of size n. [1]

[1] Throughout this work, Π is a promise problem, unless stated otherwise. Formally, Π consists of two disjoint sets $\Pi_Y, \Pi_N \subset \{0,1\}^*$ representing YES and NO instances. The problem is to decide, given an instance promised to lie in $\Pi_Y \cup \Pi_N$, whether it belongs to Π_Y or Π_N.

1.2 Positive Results

Our positive results concern the existence of one-way functions. Informally, a function is one-way if it is easy to compute but hard to invert. We consider the more general case of *fine-grained* one-way functions. Roughly speaking, a function F is a (η, θ)-fine-grained one-way function (FGOWF) if it is computed in time T_{F} but no algorithm can invert it in time T_{F}^{η} with probability better than θ, where η is a constant greater than 1. If F is a (η, θ)-FGOWF for all constants $\eta > 1$, then F is a OWF.

Our first result addresses *compressing* reductions—a special case of lossy reductions—that are reductions which shrink the size of the input, e.g., by mapping n-bit instances to $\ll n$ bits. Our first contribution is the following:

Theorem (Informal) 1 (One-Wayness from Compressing Reductions)**.** *Let Π be a problem and $\lambda(n) \geq 0$ such that $\lambda(n) < \tau_{\Pi}(n)/3$. If Π has a compressing f-reduction that maps m instances of size n to $m \cdot \lambda(n)$ bits, runs in time $o(2^{\tau_{\Pi}(n)/3})$, and has error $\leq 2^{-\lambda(n)-8}$, then FGOWFs exist. When $\lambda(n) = O(1)$, FGOWFs exist if the reduction runs in time $O(2^{\tau_{\Pi}(n)/c})$ for some $c > 1$.*

Moreover, if $\lambda(n) = o(\tau_{\Pi}(n))$ and the runtime is $2^{o(\tau_{\Pi}(n))}$, then OWFs exist.

We also extend this result to *worst-case to average-case reductions* and *randomized encodings*. A worst-case to average-case reduction for Π maps *any* instance of Π to a distribution that is d-close (d is called the distance of the reduction) to an efficiently-samplable distribution which is independent of the given instance. Moreover, a randomized encoding for the characteristic function χ_{Π} of Π is a function E such that $E(x)$ encodes the value of $\chi_{\Pi}(x)$ without revealing information about x. It can be therefore viewed as a reduction for Π. Such an encoding further requires the existence of two efficiently-samplable distributions D_{YES} and D_{NO} for respectively simulating the encoding of YES and NO instances of Π within the statistical distance d (d is called the privacy of the encoding).

Theorem (Informal) 2 (One-Wayness from Instance Randomizations)**.** *Let Π be a problem. Any worst-case to average-case reduction with distance d for Π, or any randomized encoding with privacy d for χ_{Π}, that runs in time $O(2^{\tau_{\Pi}(n)/c})$ for some $c > 1$, with $d \leq 2^{-25.5}/\sqrt{n}$ and error $\leq 2^{-21}$, implies FGOWFs. If the runtime is $2^{o(\tau_{\Pi}(n))}$, then OWFs exist.*

The theorem above extends to worst-case to average-case f-reductions or randomized encodings of $f \circ \chi_{\Pi}$ when d is as small as $2^{-6} \cdot (2^{39}mn)^{-m/2}$ and f is m-arry.

Changing the perspective, Theorem 2 implies that if a worst-case to average-case reduction (or randomized encoding) is known for a problem Π that runs in time T, with error and distance (or privacy) satisfying the conditions of the theorem, then (i) assuming $\tau_{\Pi} < \log T$, FGOWFs exist, and (ii) assuming $\tau_{\Pi} = \omega(\log T)$, OWFs exist. We present the following example to clarify the idea and showcase its utility.

Example (One-wayness from kSAT). Let us consider the well-known NP-complete problem kSAT. According to the Exponential Time Hypothesis (ETH), kSAT cannot be solved in sub-exponential time. More precisely, ETH states that there exists a constant s_k such that no instance of kSAT of N variables can be solved in time $o(2^{s_k N})$. Several variants of ETH—including the non-uniform version—have gained substantial credibility over the past decades and a disproof would yield major consequences in parameterized complexity, derandomization, and circuit lower bounds [30, 51, 52, 61, 79]. On the algorithmic side, decades of extensive research on algorithms for 3SAT [43, 45–47, 49, 54, 66, 67, 70, 74–76] have led to the runtime being improved from $O(1.362^N)$ by [66] to $O(1.306973^N)$ by [74]. [1] Expressing in terms of the instance size n, the best known algorithm for 3SAT has runtime $O(2^{0.06437 \cdot n/\log n})$. It therefore appears that discovering an algorithm with runtime, for instance, $O(2^{0.003 \cdot n/\log n})$ remains far-fetched. Let us therefore assume that $\tau_{3\text{SAT}}(n) = 0.003 \cdot n/\log n$. Under this assumption, according to Theorem 1, finding an f-reduction for 3SAT that compresses m instances of n bits to $0.001 \cdot mn/\log n$ bits and runs in time $O(2^{0.001 \cdot n/\log n})$ would imply FGOWFs. Under the same assumption, finding a worst-case to average-case reduction for 3SAT with distance $\leq 2^{-25.5}/\sqrt{n}$ and error $\leq 2^{-21}$ that runs in time $O(2^{0.001 \cdot n/\log n})$ would imply FGOWFs via Theorem 2. This opens a new path for relating the existence of FGOWFs on the subexponential hardness of NP, and offers a new perspective to the question posed by [11] regarding whether FGOWFs can be constructed under

[1] In fact, [74] achieves $O(1.306973^N)$ for UNIQUE-3SAT which is slightly better than the best current runtime for 3SAT.

ETH.

We further initiate the study of the cryptographic implications of *quantum "lossy" reductions* that output pure states. Similarly to the classical setting, a quantum lossy reduction is roughly defined as a reduction R that satisfies $I_q(X; R(X)) \ll n$ for all distributions X on inputs of size n, where I_q denotes the quantum mutual information.[1] In our work, we consider quantum reductions such that (i) for every instance x the outcome $R(x)$ is a pure quantum state (ii) and there exists a (possibly unbounded) binary quantum measurement that, given $R(x)$, decides x. We show that such reductions imply one-way state generators (OWSGs)—a type of quantum functions that are easy to evaluate on classical inputs but hard to invert given the quantum output. Let τ_{Π}^{Q} be defined similarly as τ_{Π} but with respect to quantum algorithms. We prove the following:

Theorem (Informal) 3 (OWSGs from Quantum Compressing Reductions)**.** *Let $\lambda(n) \geq 0$ such that $\lambda(n) = o(\tau_{\Pi}^{Q}(n))$. If Π has a pure-outcome compressing quantum f-reduction that maps m instances of n bits to $m \cdot \lambda(n)$ qubits, runs in time $2^{o(\tau_{\Pi}^{Q}(n))}$, and has error $\leq 2^{-2\lambda(n)-11}$, then OWSGs exist.*

The proofs of the theorems above relativize, meaning that the theorems hold even when all of the considered algorithms have access to a common arbitrary oracle.

More General Case. The aforementioned classical theorems are both special cases, tuned to their setting and parameters, of a more general theorem that we prove in our work. More precisely, we refine the definition of lossiness and consider a less restrictive notion that we call *sparse lossiness*. At a high level, sparse lossiness requires information loss *only with respect to sparse uniform distributions* over the input. More precisely, we say that an f-reduction R for Π is $(\lambda(n), \gamma(n))$-sparsely lossy if $I(X_1, \cdots, X_m; R(X_1, \cdots, X_m)) \leq \lambda(n)$ for all uniform and independent distributions $\{X_i\}_i$'s with a support size roughly equal to $1/\gamma(n)$[3]. [2] A reduction that is sparsely lossy for every $\gamma \in (0, 1]$ is lossy according to the definition of [10]. We prove the following statement:

Theorem (Informal) 4 (One-Wayness from Sparsely Lossy Reductions)**.** *Let Π be a problem and $\lambda(n) \geq 0$ such that $\lambda(n) < \tau_{\Pi}(n)/3$. There exists a $\Gamma_u \in (0, 1]$ such that if Π has an f-reduction that runs in time $o(2^{\tau_{\Pi}(n)/3})$, is $(\lambda(n), \gamma(n))$-sparsely lossy for some $\gamma(n) \leq \Gamma_u$, and has error $\leq 2^{-\lambda(n)-8}$, then FGOWFs exist. Moreover, there exists a $\Gamma_\ell \leq \Gamma_u$ such that if also $\gamma(n) \geq \Gamma_\ell$, $\lambda(n) = O(1)$, and the reduction runs in time $O(2^{\tau_{\Pi}(n)/c})$ for some $c > 1$, then FGOWFs exist.*[3]

Moreover, if $\lambda(n) = o(\tau_{\Pi}(n))$ and the reduction runs in time $2^{o(\tau_{\Pi}(n))}$, then OWFs exist.

The statement also extends to non-adaptive Turing reductions,[4] under some conditions. Such a reduction from Π to Σ is an algorithm that, given an instance x, outputs oracle queries $y_1, \cdots, y_k$ and a circuit C such that $C(y_1, \mathcal{O}(y_1), \cdots, y_k, \mathcal{O}(y_k)) = \chi_{\Pi}(x)$, where $\mathcal{O}$ is an oracle solver for Σ. We consider a variant of non-adaptive Turing reductions where $(y_1, \ldots, y_k)$ and C does not leak much information about x. This allows to further extend the results to decision-to-search reductions when $\mathcal{O}$ is a search oracle.[5] Unfortunately, in the worst-case to average-case reductions for the well-known problems Permanent, 3Sum, and OV (*e.g.*, see [11, 36, 59]) only the marginal distribution of each y_i follows a distribution that is independent of x and the joint distribution of $(y_1, \ldots, y_k)$ reveals x, preventing applying Theorem 4 to these problems.

A compressing f-reduction that maps m instances of size n to $m \cdot \lambda(n)$ bits is (λ, γ)-sparsely lossy for any choice of γ. Therefore, Theorem 4 immediately implies Theorem 1. For Theorem 2, concerning worst-case to average-case reductions and randomized encodings, we require a finer analysis to compute their sparse lossiness parameters. For the range of parameters presented Theorem 2, we show that such a reduction is (λ, γ)-sparsely lossy for every $\gamma \in [\Gamma_\ell, \Gamma_u]$.

A detailed comparison of our work with prior results relating OWFs to worst-case hardness of large classes of problems is provided in Table 1.

[1] More precisely, we consider our refined notion of lossiness which we discuss later.

[2] The exact support-size matters in the proofs, however, it is out of scope of the introduction. See Def 18 for more details.

[3] In fact, $\Gamma_\ell = m^2 \cdot \omega(2^{-\tau_{\Pi}(n)/c})$, and $\Gamma_u = 2^{-\lambda(n)-4}$.

[4] Truth-table reductions

[5] See Section 4.1 for more details.

— Based on Polynomial-Time Reductions from Π to Complexity Classes —

Work	$\tau_\Pi(n)$[1]	Reduction Parameters		Additional Assumptions	Result
		Runtime	Target Class		
[65]	$\omega(\log n)$	$\mathsf{poly}(n)$	SZK	—	AI-OWF[2]
[8]			SRE	—	OWF
[48]			NP	$\mathsf{NP} \subseteq \mathsf{CZK}$	OWF

— Based on Sparsely Lossy Reductions from a Problem Π —

Work	$\tau_\Pi(n)$	Reduction Parameters					Result
		$\lambda(n)$	$\gamma(n)$[3]	f	Runtime	Error	
[10]	$\omega(\log n)$	$1/100$	$\forall\gamma \in (0,1]$	$\textsc{Or}_m$	$\mathsf{poly}(n)$	$O(1)$	OWF[4]
		$1/100$	$\forall\gamma \in (0,1]$	$\textsc{Maj}_m$	$\mathsf{poly}(n)$	$O(1)$	OWF
		$O(\log n)$	$\forall\gamma \in (0,1]$	$\textsc{Or}_m$	$\mathsf{poly}(n)$	0	OWF
Our Work	Any	$o(\tau_\Pi(n))$	$\leq \Gamma_u$[5]	NC-PI[6]	$2^{o(\tau_\Pi(n))}$	$\leq 2^{-\lambda(n)-8}$	OWF
		$< \tau_\Pi(n)/3$	$\leq \Gamma_u$[5]	NC-PI	$o(2^{\tau_\Pi(n)/3})$	$\leq 2^{-\lambda(n)-8}$	FGOWF
		$O(1)$	$\in [\Gamma_\ell, \Gamma_u]$[5]	NC-PI	$O(2^{\tau_\Pi(n)/c})$[7]	$\leq 2^{-\lambda(n)-8}$	FGOWF

— Based on WC-AVG Reductions and Randomized Encodings for a Problem Π —

Work	$\tau_\Pi(n)$	Reduction Parameters				Result
		d	f	Runtime	Error	
Our Work	Any	$(mn)^{-m/2}$	NC-PI	$2^{o(\tau_\Pi)}$	$\leq 2^{-21}$	OWF
		$(mn)^{-m/2}$	NC-PI	$O(2^{\tau_\Pi/c})$[7]	$\leq 2^{-21}$	FGOWF

[1] $\tau_\Pi(n)$: the infimum value of T taken over all non-uniform Turing machines with advice that solve all instances of size n of Π in time $O(2^T)$. When $\tau_\Pi(n) = \omega(\log n)$, it means Π is worst-case hard for polynomial-time algorithms.
[2] AI-OWF: Auxiliary-Input One-Way Function.
[3] Note that for compressing reductions, γ is irrelevant, as a compressing reduction is sparsely lossy for any $\gamma \in (0,1]$.
[4] Moreover, this result requires the reduction to be from Π to Π.
[5] $\Gamma_\ell = m^2 \cdot \omega(2^{-\tau_\Pi(n)/c})$, and $\Gamma_u = 2^{-\lambda(n)-4}$.
[6] NC-PI: Non-Constant Permutation-Invariant Function.
[7] c is any constant greater than 1.

Table 1: State-of-the-art results building one-way functions from worst-case hardness of generic problems.

1.3 Impossibility Results

Our first impossibility result pertains to the existence of circuit obfuscation. A obfuscation scheme O is an efficient algorithm that compiles a circuit into another one in a way that it only preserves its input-output functionality and leaks only small amount of information about the circuit. We say that the scheme has ε if for every $x \in \{0,1\}^N$ and every circuit over N-bit inputs, it holds that $\Pr[\mathsf{O}(C)(x) \neq C(x)] \leq \varepsilon$ where the probability is taken over the random coins of O. We also say that the scheme is α-correlated if there exists an efficient simulator algorithm Sim such that for any two functionally equivalent circuits C_1 and C_2, the statistical distance between $\mathsf{O}(C_1)$ and $\mathsf{Sim}(C_2)$ is at most α. In this case, the scheme is called α-Statistical Obfuscation (sO). An Indistinguishability Obfuscation (iO) is a statistical obfuscation where the distributions are indistinguishable for any polynomial-time distinguisher and the error is negligible.

Goldwasser and Rothblum [42] show that $\mathsf{negl}(N)$-statistical obfuscation with zero error does not exist unless PH collapses to its second level. In a recent work, Volkovich [80] proves that α-statstical obfuscation with error ε where $\alpha > (1-\varepsilon)^2$ does not exist unless PH collapses to its third level. We show the following result.

Theorem (Informal) 5. *Assuming* $\mathsf{NP} \not\subseteq \mathsf{coNP/poly}$, *i.e.,* PH *does not collapse to its third level, then* $(1 - 1/\mathsf{poly}(N))$*-statistical obfuction with error* $1/2 - 1/\mathsf{poly}(N)$ *does not exist.*

Moreover, assuming nuNETH, then $(1 - 1/\mathsf{subexp}(N))$*-statistical obfuction with error* $1/2 - 1/\mathsf{subexp}(N)$ *does not exist.*

The non-uniform Non-deterministic Exponential Time Hypothesis (nuNETH), which is introduced by [28], states that kSAT does not have a non-uniform reduction to coNP that runs in time subexponential in N, where N is the number of variables in the kSAT instance. The credibility of nuNETH has been studied by [28, 30].

A comparison of our work with the previous ones is presented in Table 2.

Work	Assumption	Impossible Correlation Obfuscation	
		Correlation (α)	Error (ε)
Folklore	—	$\alpha > 1 - 2\varepsilon$	
[42]	$\mathsf{PH} \neq \sum_2^P$ [1]	$\mathsf{negl}(N)$ [2]	0
[80]	$\mathsf{PH} \neq \sum_3^P$ [3]	$\alpha > (1-2\varepsilon)^2$	
Our Work	$\mathsf{PH} \neq \sum_3^P$ [3]	$1 - 1/\mathsf{poly}(N)$	$1/2 - 1/\mathsf{poly}(N)$
	nuNETH	$1 - 1/\mathsf{subexp}(N)$	$1/2 - 1/\mathsf{subexp}(N)$

[1] $\mathsf{PH} = \sum_i^P$ denotes that the Polynomial Hierarchy collapses to its i-th level. More precisely, the assumption is $\mathsf{coNP} \not\subseteq \mathsf{AM}$.

[2] N is the size of the input of the circuit.

[3] More precisely, $\mathsf{NP} \not\subseteq \mathsf{coNP/poly}$.

Table 2: Impossibility of correlation obfuscation under different assumptions.

We note that Theorem 5 leaves a narrow space for the existence of iO; the computational distance of the outputs of iO must be negligible while their statistical distance is exponentially (or negligibly) close to 1 assuming nuNETH (or that PH does not collapse to its third level).

Implications for the complexity of kSAT. We answer a question left open in [35] regarding compressing f-reductions of kSAT. It was previously known that polynomial-time compressing f-reductions of m instances of size n of kSAT into $O(m \log n)$ bits is impossible unless the Polynomial Hierarchy collapses [35, 37]. We extend this result as follows.

Theorem (Informal) 6. *Under nuNETH, kSAT does not have any compressing f-reduction that maps m instances of size n to $mn^{1-\varepsilon}$ bits (for any $\varepsilon > 0$) nor any worst-case to average-case reduction with distance $\leq 2^{-25.5}/\sqrt{n}$ and error $\leq 2^{-21}$, that runs in time $o(2^{n/\log n})$.*

This results shows that under nuNETH, not only solving kSAT requires $2^{\Omega(n/\log n)}$-time algorithms but also any non-constant compression or any loose instance randomization of kSAT runs in time $2^{\Omega(n/\log n)}$. This impossibility can be tightened even more under nuETH and the assumption that FGOWF does not exist via contraposing the statements of Theorem 1 and 2. We summarize these findings in Table 3.

Work	Assumption	s_k	Impossible Compression		Impossible Instance Randomization		
			Compression	Runtime	Max Distance	Runtime	Max Error
	nuETH	$\Omega(1)$	—	—	—	—	—
[37] [35]	$\mathsf{PH} \neq \sum_3^P$ [1]	—	$O(m \log n)$	$\mathsf{poly}(n)$	—	—	—
Our Work	nuETH & ¬FGOWF	$\Omega(1)$	$mn^{1-\varepsilon}$	$2^{s_k n/(6k \log n)}$	$2^{-25.5}/\sqrt{n}$	$2^{s_k n/(2k \log n)}$	2^{-21}
	nuNETH	$\Omega(1)$	$mn^{1-\varepsilon}$	$2^{o(n/\log n)}$	$2^{-25.5}/\sqrt{n}$	$2^{o(n/\log n)}$	2^{-21}

[1] $\mathsf{PH} = \sum_3^P$ denotes that the Polynomial Hierarchy collapses to its third level. More precisely, the assumption is $\mathsf{NP} \not\subseteq \mathsf{coNP/Poly}$.

Table 3: Impossibility of instance compression and randomization for of kSAT under various assumptions.

Limitations of our method for building OWFs from *k*SAT. For building OWFs, the choice of Π in Theorem 1, 2, and 4 is subject to barriers. More precisely, we prove that if a problem Π satisfies the conditions of Theorem 1, 2, or 4, for building OWFs, then Π reduces to SZK in time $2^{o(\tau_\Pi)}$, which contradicts nuNETH.[1] This restriction also applies to the results of [10]. Such barriers, however for black-box constructions, for building OWFs are already known by [4, 18, 19, 36].

1.4 Technical Overview

In this section, we briefly present the core technical tools that we use.

An extended disguising lemma. Let $R : \{0,1\}^* \to \{0,1\}^*$ be a function and consider the problem of finding x given $R(x)$. Fano's inequality gives a lower bound for the amount of information about x that an unbounded algorithm can recover from $R(x)$, for any choice of x. The original variant of the disguising lemma by Drucker [35] is a distinguishing variant of Fano's inequality. The disguising lemma states that for any subset $S \subseteq \{0,1\}^n$, there exists a sparse[2] distribution D_S supported on S, such that for all $y \in S$, the distribution of $R(y)$ is statistically close to that of $R(D_S)$, where the closeness depends on the amount of compression achieved by R—therefore, if R is sufficiently compressing, then R *disguises* y. More precisely, $\mathbb{E}_{D_S}[\|R(y) - R(D_S)\|_1] \leq \delta^*$, where $\delta^* \approx 1 + 2^{-\lambda}$, if R compresses n-bit instance to λ bits. [3]

In order to sketch our improvements of this lemma, we first briefly go over Drucker's proof. Consider the following simultaneous-move two-player game: Given a set S and $d \leq |S|$, Player 1 chooses a d-sparse distribution D supported on S, and Player 2 chooses an element $y \in S$. Define the payoff as $\|R(y) - R(D)\|_1$. Drucker first shows that if R is compressing, then for any strategy $\mathcal{Y}$ for

[1] The idea behind Theorem 6 is to leverage this reduction.

[2] A distribution is d-sparse if d outcomes have non-zero probability.

[3] The ℓ_1 norm $\|X - Y\|_1$ is equal to 2 times of the statistical distance $\Delta(X, Y)$ when variables X, Y are classical distributions, or 2 times of the trace distance $\mathrm{Tr}(X, Y)$ when they are quantum states.

choosing y, there exist a d-sparse distribution D, that is more precisely the uniform distribution over d samples of $\mathcal{Y}$, such that:

$$\mathbb{E}_{\mathcal{Y}^{\otimes d}}[\|R(y) - R(D)\|_1] \leq \delta^* \,. \tag{1}$$

By the minimax theorem, we can switch the quantifiers in above, implying that there exists a distribution $\mathcal{D}_S$ such that, for any choice of $y \in S$, it holds that $\mathbb{E}_{D\sim\mathcal{D}_S}[\|R(y) - R(D)\|_1] \leq \delta^*$. However, note that $\mathcal{D}_S$ is not guaranteed to be sparse. Drucker finally uses a result by Lipton and Young [60, Theorem 2] which roughly states that restricting the strategies of Player 1 to uniform strategies with support size $\ln(\#\{\text{choices of Player 2}\})/\gamma^2$ only changes the optimal expectation payoff with an additive factor γ. [1] This therefore *sparsifies* the support of $\mathcal{D}_S$. Conversely, it incurs a loss of at most an additive factor γ in the expectation bound in Equation (1) and obtains $\delta^* + \gamma$.

We relax the requirement on R and show that *sparse lossiness* of R suffices to obtain a similar result. More precisely, we show that if R loses information on input distributions that are uniform with support size roughly $1/\gamma^3$, then one loses nothing but an(other) additive factor γ in Equation (1). This relies on a double use of the result by Lipton and Young [60, Theorem 2]; we apply it once for Player 2 and once more for Player 1. Before using the minimax theorem, we restrict Player 2's strategies to be uniform distributions with support size roughly $1/\gamma^3$, and we choose $d \approx 1/\gamma$. By showing that Equation (1) remains correct even with this new restriction, we obtain an additive γ-approximation of the value of the game (first use of [60, Theorem 2] for Player 2). Following the minimax theorem, and sparsifying $\mathcal{D}_S$ (via a second use of [60, Theorem 2] for Player 1), we achieve the final upper bound $\delta^* + 2\gamma$. We stress that this step is crucial for our results—otherwise, we could not sufficiently bound the lossiness of worst-case to average-case reductions or randomized encodings.

Extending to multivariate setting. To be more precise, all of the above has been analyzed by Drucker [35] in the setting where R is multivariate, *e.g.*, taking m instances as input; Drucker bounds the distance of $R(D_S, \cdots, y, \cdots, D_S)$—where there are $m-1$ samples of D_S and exactly one y in a random place, from $R(D_S, \cdots, D_S)$—where there are m samples of D_S, for a compression R. Similarly as above, we extend the disguising lemma in the multivariate setting, by showing that the distance of the two aforementioned distributions are bounded by $\delta^* + 2\gamma$, when R is sparsely lossy.[2]

Extending to disjoint sets. It is furthermore proved by [10] that for any choice of $p \in \{0, \ldots, m\}$ and $b \in \{0, 1\}$, there exist two sparse distributions D_{S_0} and D_{S_1} over two disjoint sets S_0 and S_1, such that when R is lossy, for every $y \in S_b$, the two distributions $R(\pi(D_{S_0}, \cdots, y, \cdots, D_{S_1}))$ and $R(\pi(D_{S_0}, \cdots, D_{S_b}, \cdots, D_{S_1}))$ are "close", where π is a uniformly random permutation and the number of D_{S_0} and D_{S_1} samples in the input of the latter are respectively p and $m - p$. Note the constraint that y must have the same support as the distribution that it replaces. Our variant of disguising lemma with sparsely lossy reductions also extends to this setting (we refer to Section 3 for more details). In the rest of this section, let

$$R_p[\star] := R(\pi(D_{S_0}, \cdots, \star, \cdots, D_{S_1}))\,, \tag{2}$$

where the number of D_{S_0} and D_{S_1} samples in the input are $p-1$ and $m-p$, respectively, and $\star$ can posses a fixed quantity or a random variable.

One-wayness from sparse lossiness. Let Π be a decision problem and let R be a sparsely lossy reduction over m instances of Π. Define $S_0 := \Pi_N \cap \{0,1\}^n$ and $S_1 := \Pi_Y \cap \{0,1\}^n$. Due to our disguising lemma, for any $0 \leq p \leq m$, there exist two sparse distributions D_{S_0} and D_{S_1} such that $\mathbb{E}[\|R_p[y] - R_p[D_{S_0}]\|_1] \leq \delta^* + 2\gamma$, for all $y \in S_0$ (recall $R_p[\star]$ as per Equation (2)).

Looking closer, $R_p[D_{S_0}]$ and $R_p[y]$ use an internal randomness to sample from D_{S_0} and D_{S_1}. In other words, they can be viewed as two circuits that, given uniformly random strings, sample elements from D_{S_0} and D_{S_1}, and return the evaluation of R over these samples. Let C_0 and $C_1[y]$ respectively denote these two circuits. We claim that if Π is worst-case hard, then the following is a one-way function:

$$\mathsf{F}(b, r) := \begin{cases} C_0(r) & \text{if } b = 0 \\ C_1[y](r) & \text{if } b = 1 \end{cases}\,, \tag{3}$$

where y is sampled from D_{S_0}.

[1] The same holds for Player 2.

[2] In fact, m possibly changes the upper bound, but by tuning $d \approx m/\gamma$, one can keep the bound the same.

We now sketch the proof of the one-wayness of F. Let $\mathcal{A}$ be an inverter for F. We use $\mathcal{A}$ to decide any instance $\hat{y}$ of Π as follows: Compute $C_0(r)$, and $C_1[\hat{y}](r)$ for a random value r, and then return $C_b(r)$ to $\mathcal{A}$. If $b = 0$, then $\mathcal{A}$ receives an instance of the function F and can therefore invert it. This does not help us with solving Π! Conversely, when $b = 1$, if $\hat{y}$ is a NO instance, then by our disguising lemma, $C_1[\hat{y}]$ would be close to $C_1[y]$.[1] Therefore, $\mathcal{A}$ would succeed to invert it. On the other hand, when $b = 1$, if $\hat{y}$ is a YES instance, then C_0 and $C_1[\hat{y}]$ are far from each other, since R is a reduction. We also know that C_0 and $C_1[y]$ are close. Hence, $C_1[y]$ and $C_1[\hat{y}]$ must be far. Consequently, the image spaces of $C_1[y]$ and $C_1[\hat{y}]$ have small intersection. Therefore, if $b = 1$ and $\hat{y}$ is a YES instance, then there would be no pre-image (except with a small probability) for the value that $\mathcal{A}$ aims to invert, which results in $\mathcal{A}$ failing. By repeating this test several times, we can decide $\hat{y}$ by observing the success rate of $\mathcal{A}$.

We note that our candidate one-way function frequently appears in the SZK literature (*e.g.*, see [16, 71]). In [10], it is shown that when the reduction R is perfect and lossy and Π is worst-case hard (with respect to polynomial-time algorithms), then $C_0(\cdot)$ is a weak one-way function. We also examine the proposal of [10] in detail, however, we find our construction more robust, given that it can be easily extended to OWSGs. When the circuits are quantum, there are only two more technical details to fix: (i) showing that the image spaces of two quantum circuits $C_1[y]$ and $C_1[\hat{y}]$ have small intersection even in the quantum case, (ii) computing the success rate of $\mathcal{A}$. The latter uses SWAP test and requires that for any fixed randomness r, the outputs of $C_0(r)$ and $C_1(r)$ be pure.

Runtime Analysis. Since both distributions $R_p[D_{S_0}]$ and $R_p[y]$ are uniform over m multisets of size $d \approx m/\gamma$ that contain n-bit elements, sampling one of each set requires $O(\log(m/\gamma))$ bits in size and, with an appropriate data structure, $O(m/\gamma)$ in time. Therefore, if T_R is the runtime of R, the total runtime of C_0 or C_1 (and consequently F) is $O(T_R + (m^2/\gamma))$. Note that applying a random permutation takes $O(m)$ steps using the Fisher-Yates' algorithm. Reducing the worst-case hardness of Π to the one-wayness of F as above requires repeating this computation, where the total number of repetition θ depends on the sparse lossiness of R. In total, assuming $\omega((T_R + (m^2/\gamma))\theta^{-1})$-hardness of Π implies fine-grained OWFs. For more details, we refer to Sections 6, 7, and 9.

Reduction to the statistical difference (SD) problem. The statistical difference $\mathrm{SD}_{\alpha,\beta}$ problem asks, given two circuits (C_0, C_1), whether on uniformly random inputs their induced distributions are at least β-far or at most α-far (with respect to the statistical distance) under the promise that one is true. This problem is complete for SZK under polynomial-time reductions when $\beta^2 - \alpha$ is a positive constant.

Recall from the proof of one-wayness above that any instance $\hat{y}$ of Π can be mapped to two circuits $(C_0, C_1[\hat{y}])$ such that

- if $\hat{y}$ is a NO instance of Π, the two circuits have statistical distance at most $\delta^* + 2\gamma$,
- and if $\hat{y}$ is a YES instance of Π, the two distributions are far since R is a reduction.

More precisely, when $\hat{y}$ is a YES instance, the statistical distance of the two distributions is at least $1-\mu^*$, where μ^* is the error of R. By setting $\alpha = (\delta^*+2\gamma)$ and $\beta = 1-\mu^*$, this yields a reduction from Π to $\mathrm{SD}_{\alpha,\beta}$ that runs in time $O(T_R+(m^2/\gamma))$. Note that depending on the quantities α and β, the problem $\mathrm{SD}_{\alpha,\beta}$ is not necessarily inside SZK. However, by using the known polarization tools for SD (e.g. see [71, 81]), we are able to derive a reduction from Π to SZK running in a time that only depends on the sparse lossiness parameters of R. More precisely, for most of the (λ, γ)-sparsely lossy reductions R that we consider in this work, the reduction of Π to SZK runs in time $O(2^{\lambda}T_R)$.[2]

Sparse Lossiness of WC-AVG reductions and randomized encodings. We say that a reduction R from Π is worst-case to average-case if there exist a small $d < 1$ and a distribution $D = \{D_n\}_{n\in\mathbb{N}}$ over $\{0,1\}^*$, such that:

$$\forall x \in \Pi \cap \{0,1\}^n : \ \Delta\left(R(x), D\right) \leq d \ . \tag{4}$$

This definition can be viewed as a generalization of worst-case to average-case reductions in the sense that (i) the reduction is oblivious to the target average-case problem, and (ii) the reduction

[1] As discussed before, this closeness, and therefore the success probability of this reduction depends on the sparse lossiness of R.

[2] The reason that the runtime analysis does not extend to all sparsely lossy reductions is that we require a mild constraint on γ that does not always hold.

maps inputs to a distribution that is *not* necessarily efficiently samplable. The latter does not impose any issues in our setting, since we are only discussing lossiness of the reductions.

In order to prove the sparse lossiness of these reductions, we first translate the mutual information $I(X; R(X))$ in terms of KL-divergence, and then use an inverse Pinsker inequality. It is shown by Sason [73], that for every two random variables X and Y, we have

$$D_{KL}(X\|Y) \leq \log\left(1 + \frac{2 \cdot \Delta(X,Y)^2}{\alpha_X}\right),$$

where $\alpha_X = \min_x \Pr(X = x) > 0$. The term $\Delta(X, R(X))$ is bounded by the worst-case to average-case property, therefore, it suffices to bound α_X. Since the sparse lossiness concerns uniform distributions X with a support of size roughly $1/\gamma^3$, we can bound α_X by $\Omega(\gamma^3)$. By working out the details, we get $\lambda(n) \leq \max\left\{1/m, 9 + 4/m + \log\left(mn/\gamma^3\right) + 2\log d/m\right\}$. We also prove the sparse lossiness of randomized encodings, by similarly calculating their λ as above.

Non-existence of obfuscation. The key idea to derive our impossibility result on obfuscations is to leverage them to build polynomial-time worst-case to average-case reductions for the UniqueSat problem. This problem asks to decide whether a CNF formula over N variables has a satisfiable assignment under the promise that it has at most one satisfiable assignment. Works of [57] and [50] observe that statistical obfuscation gives a worst-case to average-case reduction for UniqueSat. More precisely, if the input circuit C is a YES instance of UniqueSat, i.e., it evaluates to 1 on exactly one input, then a random shift of C, namely the circuit $C_z(x) = C(x \oplus z)$, where $z \xleftarrow{\$} \{0,1\}^N$, has a truth-table whose distribution is identical to that of a random point function. Therefore, the obfuscation of C_z must be statistically close to the obfuscation of a random point function. Moreover, if the obfuscation is perfect, i.e., error $\varepsilon = 0$, then the obfuscation of C_z is a YES instance of UniqueSat.

Recall our earlier discussion that a (λ, γ)-sparsely lossy reduction for a problem Π, with some mild constraints on γ, yields a reduction from Π to SZK in time $O(2^\lambda T)$. As we show earlier, a worst-case to average-case reduction is sparsely lossy. Depending the correctness and security parameters of sO, we modify the aforementioned worst-case to average-case reduction of UniqueSat such that it becomes a sparsely lossy reduction with constant lossiness, i.e., $\lambda = O(1)$. This modification only affects the runtime of the reduction. When sO is α-statistical for some $\alpha = 1 - 1/\mathsf{poly}(N)$, the runtime remains polynomial, while for $\alpha = 1 - 1/\mathsf{subexp}(N)$, the runtime becomes subexponential. Therefore, we obtain two reductions of UniqueSat to SZK in these two different regimes. Finally, we analyze the possibility of these reductions by resorting the complexity of Sat and the Valiant-Vazirani polynomial-time reduction of Sat to UniqueSat.

More general reductions. Our results cover decision-to-decision or decision-to-search non-adaptive Turing reductions. For this purpose, we introduce distinguisher reductions which unify all these variants. For a Boolean function $f : \{0,1\}^m \to \{0,1\}$, we define an f-distinguisher reduction for a problem Π as a mapping $R : \{0,1\}^* \to \{0,1\}^*$ for which there exists an unbounded distinguisher $\mathcal{D}$ that can distinguish between $R(x_1, \ldots, x_m)$ and $R(x'_1, \ldots, x'_m)$, given one $\{x_i\}_i$ at random, if $f(\chi_\Pi(x_1), \cdots, \chi_\Pi(x_m)) \neq f(\chi_\Pi(x'_1), \cdots, \chi_\Pi(x'_m))$. We prove all of our results regarding f-distinguisher reductions and we show that all decision-to-decision or decision-to-search non-adaptive Turing reductions, including Karp reductions, are special cases of f-distinguisher reductions.

2 Preliminaries

In this work, we always consider non-uniform algorithms. All classical algorithms are quantum algorithms, therefore, we mostly use the quantum formalism for generalization and simplification. When the distinction is necessary, we explicitly mention it in the beginning of a section or inside a statement.

Notation. We let n denote the security parameter, and all variables are implicitly parametrized by n. We let MS_n denote the set of all mixed states over n qubits and we define $\mathsf{MS}_* := \cup_{n=1}^{\infty} \mathsf{MS}_n$. For a positive integer n, we let $[n]$ denote $\{1, 2, \cdots, n\}$. The set of all permutations over $[n]$ is $\mathfrak{S}_n$. We abuse the notation and use the same symbol to refer to the uniform distribution over all permutations of $[n]$. The set of natural numbers $\{1, 2, 3, \cdots\}$ is denoted by $\mathbb{N}$. We denote by $\mathbb{R}^+$ the set of positive real numbers.

A function $f(n)$ is $\mathsf{poly}_\ell(n)$ if $f(n) = O(n^\ell)$ (we drop ℓ when the degree is not specified), is $\mathsf{negl}(n)$ if $f(n) = 1/n^{\omega(1)}$, and $\mathsf{subexp}(n)$ if $f(n) = 2^{O(n^c)}$ for some constant $c < 1$.

Uniform and s-Uniform Distributions. For any set S, we let $\mathcal{U}_S$ denote the uniform distribution over S. A distribution is called s-uniform if it is uniform over a multiset of s elements.

Boolean functions. A Boolean function $f : \{0,1\}^m \to \{0,1\}$ is called non-constant if it is not always 0 nor always 1.

Promise Problems. A Promise Problem Π consists of two disjoint sets $\Pi_Y, \Pi_N \subset \{0,1\}^*$, respectively referred to as the set of YES and NO instances. Problem Π asks to decide whether a given instance, which is promised to lie in $\Pi_Y \cup \Pi_N$, belongs Π_Y or Π_N.

Definition 1 (Characteristic Function of a Promise Problem). *For a promise problem Π, the characteristic function of Π is the map $\chi_\Pi(x) : \{0,1\}^* \to \{0,1,\star\}$ given by*

$$\chi_\Pi(x) = \begin{cases} 1 & \text{if } x \in \Pi_Y \\ 0 & \text{if } x \in \Pi_N \\ \star & \text{otherwise} \end{cases} .$$

Search Problems. We recall the definition of a search problem, inspired by that of [15]. We define a *search* problem Π_{search} as a binary relation over $\{0,1\}^* \times \{0,1\}^*$. For any $(x,w) \in \Pi_{\text{search}}$, we call x an *instance* and w a *witness*. For any $x \in \{0,1\}^*$, we define $\Pi_{\text{search}}(x) = \{w \in \{0,1\}^* \mid (x,w) \in \Pi_{\text{search}}\}$. We refer to the sets $\Pi_{\text{search}_{|Y}} = \{x \in \{0,1\}^* \mid \Pi_{\text{search}}(x) \neq \emptyset\}$, and $\Pi_{\text{search}_{|N}} = \{0,1\}^* \setminus \Pi_{\text{search}_{|Y}}$ as the set of YES and NO instances, respectively.

We say that an algorithm $\mathcal{A}$ solves Π_{search}, if for any $x \in \{0,1\}^*$ for which $\Pi_{\text{search}}(x) \neq \emptyset$, $\mathcal{A}$ returns some $w \in \Pi_{\text{search}}(x)$, and otherwise, outputs $\perp$.

We denote the decision language defined by Π_{search} as $\Pi = \{x \in \{0,1\}^* \mid \exists w \in \{0,1\}^*, (x,w) \in \Pi_{\text{search}}\}$. Each decision language Π can have multiple associated *search problems*, one for every relation Π_{search} that defines Π. Given $x \in \Pi$, the Π_{search}-search problem consists on finding $w \in \Pi_{\text{search}}(x)$.

Two-Player games. A two-player, simultaneous-move, zero-sum game is specified by a matrix $\mathbf{M} \in \mathbb{R}^{a \times b}$. Player 1 chooses a row index $i \in [a]$ and Player 2 chooses a column index $j \in [b]$, and Player 2 receives the payoff $\mathbf{M}_{ij}$ from Player 1. The goal of Player 1 is minimizing the expected payoff, while Player 2 opts to maximize it. The row and column indices are called the pure strategies of Player 1 and Player 2, respectively. The mixed strategies are distributions or possible choices of indices. A mixed strategy is s-uniform if it is sampled uniformly from a multiset of at most s pure strategies.

Lemma 2.1 ([77]). *Let $\mathcal{P}$ and $\mathcal{Q}$ be two mixed strategies for Player 1 and 2, respectively. It holds that*

$$\min_{\mathcal{P}} \max_{j} \mathbb{E}_{i \sim \mathcal{P}}[\mathbf{M}_{ij}] = \max_{\mathcal{Q}} \min_{i} \mathbb{E}_{j \sim \mathcal{Q}}[\mathbf{M}_{ij}].$$

The value of the game, which we denote by $\omega(\mathbf{M})$, is the optimal expected value guaranteed by the above lemma. The following lemma shows that each player has nearly-optimal s-uniform strategy when s is chosen to be logarithm of the number of pure strategies of the opponent.

Lemma 2.2 ([60, Theorem 2]). *For any real $\varepsilon > 0$, any $\mathbf{M} \in \mathbb{R}^{a \times b}$, and any integer $s \geq \ln(b)/(2\varepsilon^2)$, it holds that*

$$\min_{\mathcal{P} \in \mathfrak{P}_s} \max_{j} \mathbb{E}_{i \sim \mathcal{P}}[\mathbf{M}_{ij}] \leq \omega(\mathbf{M}) + \varepsilon(\mathbf{M}_{max} - \mathbf{M}_{min}) ,$$

where $\mathfrak{P}_s$ denotes the set of all s-uniform strategies for Player 1. Similar statement holds for Player 2, namely,

$$\max_{\mathcal{Q} \in \mathfrak{Q}_s} \min_{i} \mathbb{E}_{j \sim \mathcal{Q}}[\mathbf{M}_{ij}] \geq \omega(\mathbf{M}) - \varepsilon(\mathbf{M}_{max} - \mathbf{M}_{min}) ,$$

where $\mathfrak{Q}_s$ denotes the set of all s-uniform strategies for Player 2.

Classical information. Given two probability distributions X and Y over Σ, their statistical distance, also called total variation distance, is defined as

$$\Delta(X,Y) := \frac{1}{2}\sum_{x\in\Sigma} |\Pr(X=x) - \Pr(Y=x)| \ .$$

The Kullback–Leibler divergence or classical relative entropy of X with respect to Y is defined as

$$D_{KL}(X||Y) := \sum_{x\in\Sigma} \Pr(X=x)\log\left(\frac{\Pr(X=x)}{\Pr(Y=x)}\right) \ .$$

Lemma 2.3 (Chernoff-Hoeffding Bound). *Let $X_1, X_2, \cdots, X_k$ be mutually independent random variables in $[0,1]$ and $\mu := \mathbb{E}\left[\sum_{i=1}^k X_i\right]$. For every $t > 0$, it holds that $\Pr\left[|\sum_{i=1}^k X_i - \mu| > t\right] \leq 2e^{-2t^2/k}$.*

Quantum information. For a mixed state ρ, we let $\|\rho\|_1$ denote its 1-norm. We denote by $\mathrm{Tr}(\rho,\sigma)$ the the trace distance between any two states ρ and σ, with $\mathrm{Tr}(\rho,\sigma) := \|\rho-\sigma\|_1/2$. For an operator Φ, we let $\|\Phi\|_{op}$ denote its operator norm. Let $R : \{0,1\}^n \to \mathsf{MS}_m$ be any quantum mapping and X a random variable supported over $\{0,1\}^n$. We let

$$\rho_{X,R(X)} := \sum_{x\in\{0,1\}^n} \Pr_X(x)\,|x\rangle\langle x| \otimes R(x) \ . \tag{5}$$

For a mixed state ρ, we let $S(\rho) := \mathrm{Tr}(\rho\log_2\rho)$ denote the Von Neumann entropy of ρ. The quantum mutual information of two subsystems A and B is defined as follows. Let ρ_{AB} be their joint state, then

$$I_q(A;B)_\rho := S(\rho_A) + S(\rho_B) - S(\rho_{AB}) \ ,$$

where $\rho_A = \mathrm{Tr}_B(\rho_A B)$ and $\rho_B = \mathrm{Tr}_A(\rho_A B)$. For the sake of simplicity, we sometimes drop the subscripts q and ρ in I_q. When working with quantum systems A, B, the notation $I(A;B)$ implicitly refers to $I_q(A;B)$.

The following lemma states that if the outcome of a measurement is close to deterministic, then it must not alter much the state.

Lemma 2.4 (Gentle Measurement Lemma [83]). *Let ρ be a mixed state and $\{\Lambda, I-\Lambda\}$ a two-outcome POVM with $\mathrm{Tr}(\Lambda\rho) \geq 1-\varepsilon$, then $\|\rho-\rho'\|_1 \leq \sqrt{\varepsilon}$, where $\rho' = \frac{\sqrt{\Lambda}\rho\sqrt{\Lambda}}{\mathrm{Tr}(\Lambda\rho)}$.*

For two quantum states σ, ρ stored in two different registers A, B, the swap test is executed on the registers A, B and a control register C initialized to $|1\rangle\langle 1|$. It applies Hadamard on C, swaps A and B conditioned on C, and measures B on the Hadamard basis.

Lemma 2.5 (SWAP Test [27]). *The SWAP test on input (σ,ρ) outputs 1 with probability $(1+\mathrm{Tr}(\rho\sigma))/2$, in which case we say that it passes the test. For pure states $|\sigma\rangle, |\rho\rangle$, it equals to $(1+|\langle\rho|\sigma\rangle|^2)/2$.*

Given that the trace distance of two pure states $|\sigma\rangle, |\rho\rangle$ can be expressed in terms of the inner product uniquely as $\sqrt{1-|\langle\rho|\sigma\rangle|^2}$, the SWAP test can also be used to calculate their trace distance.

Definition 2 (ℓ_1 distance for classical distributions and quantum states). *We use the notation $\|X-Y\|_1$ to refer to 2 times of the statistical distance $\Delta(X,Y)$ when variables X, Y are classical distributions, or 2 times of the trace distance $\mathrm{Tr}(X,Y)$ when they are quantum states.*

Worst-case hardness. In this work, we consider fine-grained worst-case hardness, as introduced below.

Definition 3. *For a function $T : \mathbb{N} \to \mathbb{R}^+$, a promise problem Π is said to be $T(n)$-hard, if for any non-uniform classical-advice algorithm $\mathcal{A}$ with runtime at most $T(n)$ over n-bit inputs, and any sufficiently large $n \in \mathbb{N}$, there exists an input $x \in (\Pi_Y \cup \Pi_N) \cap \{0,1\}^n$ such that $\Pr[\mathcal{A}(x) = \chi_\Pi(x)] < 2/3$.*

One can without loss of generality assume that the size of the advice is not larger than the runtime.

Polynomial Hierarchy. We let $\Sigma_0^p := \Pi_0^p := \mathsf{P}$ be the class of problems solvable in polynomial time. The levels of the Polynomial Hierarchy PH are defined as follows: for every $i \in \mathbb{N}$, we let

$$\Sigma_{i+1}^p := \mathsf{NP}^{\Sigma_i^p} \text{ and } \Pi_{i+1}^p := \mathsf{coNP}^{\Sigma_i^p} .$$

It is a common belief that PH does not collapse to any of its levels, namely, it is unlikely to have $\mathsf{PH} = \Sigma_i^p = \Pi_i^p$ for any $i \geq 2$. The following result will be used in our theorems.

Lemma 2.6 ([85]). *If* $\mathsf{NP} \subseteq \mathsf{coNP/poly}$*, then* $\mathsf{PH} = \Sigma_3^p = \Pi_3^p$.

Complexity class (Q)SZK. We recall the quantum state distinguishability problem below. We refer to [81] for more details.

Definition 4 (Quantum State Distinguishability). *Let* $\alpha, \beta \in [0,1]$ *such that* $\alpha < \beta$*. Given two quantum circuits* $\mathcal{C}_0$ *and* $\mathcal{C}_1$*, let* ρ_0 *and* ρ_1 *be the (mixed) quantum states that they produce by running on all-zero states with the promise that either* $\|\rho_0 - \rho_1\|_1 \geq \beta$ *(corresponds to no instances) or* $\|\rho_0 - \rho_1\|_1 \leq \alpha$ *(corresponds to yes instances). The* $\mathrm{QSD}_{\alpha,\beta}$ *problem is to decide which one is the case.*

The above problem enjoys a polarization property. The lemma below is adapted from [71, 81].

Lemma 2.7. *Let* n *be a positive integer. Let* $\alpha, \beta : \mathbb{N} \to [0,1]$*, and* $\theta : \mathbb{R} \to (1, +\infty)$ *be functions of* n *such that* $\theta := \beta^2/\alpha$*. There exists a deterministic classical algorithm* Polarize *that given a pair of (quantum) circuits* (C_0, C_1) *as well as a unary parameter* 1^n*, outputs a pair of (quantum) circuits* (P_0, P_1) *such that*

$$\|C_0 |0\rangle - C_1 |0\rangle\|_1 \leq \alpha \Rightarrow \|P_0 |0\rangle - P_1 |0\rangle\|_1 \leq 2^{-n} ,$$
$$\|C_0 |0\rangle - C_1 |0\rangle\|_1 \geq \beta \Rightarrow \|P_0 |0\rangle - P_1 |0\rangle\|_1 \geq 1 - 2^{-n} .$$

Moreover, the runtime and output size of Polarize *are of* $O(n \log(8n)(|C_0| + |C_1|)/\log(\theta))$ *when* $n \to +\infty$.

There are various equivalent definitions of the complexity class QSZK. The following definition suffices for our purposes.

Definition 5 (QSZK). *The class* QSZK *is consisted of all promise problems that have many-to-one polynomial-time reductions to* $\mathrm{QSD}_{1/4,3/4}$.

All definitions and lemmas above can be restricted to classical algorithms. In this case, we let SZK denote the corresponding classical complexity class and SD denote the statistical difference problem (classical variant of QSD). We also have the following lemma about SZK that is derived from [35, Thm.4.11] and [21, Thm. 4].

Lemma 2.8. *It holds that* $\mathsf{SZK/poly} \subseteq \mathsf{NP/poly} \cap \mathsf{coNP/poly}$.

Cryptographic primitives. One-way functions are defined as follows:

Definition 6 (Non-Uniform One-Way Functions). *Let* $T : \mathbb{N} \to \mathbb{R}^+$ *and* $\theta : \mathbb{N} \to [0,1]$*. A family of non-uniform PPT algorithms* $\mathsf{F} := \{\mathsf{F}_n\}_{n\in\mathbb{N}}$ *is said to be a* (T, θ)*-one-way function (OWF) if for all sufficiently large* n *and any* $T(n)$*-time algorithm* $\mathcal{A}$*, it holds that*

$$\Pr_{x \sim \mathcal{U}_{\{0,1\}^n}} [\mathsf{F}(\mathcal{A}(\mathsf{F}(x))) = \mathsf{F}(x)] \leq \theta(n) .$$

Furthermore, we say that F *is a* θ*-OWF for an algorithm* $\mathcal{A}$ *if the above inequality holds without imposing any bound on the runtime of* $\mathcal{A}$.

When $T = \mathsf{poly}(n)$ and $\theta = \mathsf{negl}(n)$, the above definition corresponds to the common definition of one-way functions. If θ is $1 - 1/n^c$ for some constant c, this corresponds to weak one-way functions. It is shown by [84] that weak one-way functions imply one-way functions.

Below, we define efficiently samplable statistically far but computationally indistinguishable quantum states (EFI).

Definition 7 (Non-Uniform EFI). *Let $T : \mathbb{N} \to \mathbb{R}^+$ and $d, D : \mathbb{N} \to [0,1]$ be functions. A non-uniform (T, D, d)-EFI scheme is a QPT algorithm $\mathsf{EFI}_h(1^n, b)$ that is given a classical $\mathsf{poly}(n)$-size advice h and a bit b, outputs a quantum state ρ_b, such that for any sufficiently large $n \in \mathbb{N}$ has the following specifications:*

1. ***Computational indistinguishability.*** *For all non-uniform (possibly quantum) $T(n)$-time algorithms $\mathcal{A}$:*
$$\left|\Pr\left[\mathcal{A}(\rho_0) = 1\right] - \Pr\left[\mathcal{A}(\rho_1) = 1\right]\right| \leq d(n).$$
2. ***Statistical Distance.*** $\|\rho_0 - \rho_1\|_1 \geq D(n)$.

Furthermore, we say that EFI is a (D, d)-EFI for an algorithm $\mathcal{A}$, if the computational indistinguishability holds for $\mathcal{A}$ without requiring any bound of the runtime of $\mathcal{A}$.

Remark 1. When restricted to classical algorithms, EFI pairs with $D - d \geq 1/\mathsf{poly}(n)$ and $T = \mathsf{poly}(n)$ imply the existence of one-way functions (*e.g.*, see [16, 40, 64]). The state of the art for the quantum EFI pairs is more restricted. More precisely, an EFI pair with mixed states and $D^2 - \sqrt{d} \geq O(1)$ implies quantum bit commitment (see [20, Corollary 8.8] for EFI polarization and [22] for the generic transformation to construct quantum bit commitments from EFI pairs).

In this work, we consider the inefficient-verifier one-way state generators.

Definition 8 (Non-Uniform One-Way State Generators). *Let $T : \mathbb{N} \to \mathbb{R}^+$ and $\theta : \mathbb{N} \to [0,1]$. A (T, θ)-one-way state generator (OWSG) is a tuple of algorithms $\mathsf{G} := (\mathsf{KeyGen}, \mathsf{StateGen}, \mathsf{Ver})$ with the following specification:*

- $\mathsf{KeyGen}_h(1^n) \to k$*: is a QPT algorithm that given the security parameter 1^n and a $\mathsf{poly}(n)$-size classical advice h, outputs a classical string $k \in \{0,1\}^n$;*
- $\mathsf{StateGen}(k) \to \rho_k$*: is a QPT algorithm that given a classical string k, outputs an m-qubit quantum state;*
- $\mathsf{Ver}(k, \rho) \in \{0,1\}$*: is a (possibly unbounded) algorithm that given a classical string k and a quantum state ρ outputs either 0 or 1.*

Further, they satisfy the following properties:

1. ***Correctness.*** *Outputs of the samplers $(\mathsf{KeyGen}, \mathsf{StateGen})$ pass the verification with overwhelming probability, i.e.,*
$$\Pr_{\substack{k \leftarrow \mathsf{KeyGen}_h \\ \rho_k \leftarrow \mathsf{StateGen}(k)}}\left[\mathsf{Ver}(k, \rho_k) = 1\right] \geq 1 - \mathsf{negl}(n)\,.$$
2. ***Security.*** *For every non-uniform $T(n)$-time adversary $\mathcal{A}$, and any polynomial $t(n)$*
$$\Pr_{\substack{k \leftarrow \mathsf{KeyGen}_h \\ \rho_k \leftarrow \mathsf{StateGen}(k) \\ k' \leftarrow \mathcal{A}(\rho_k^{\otimes t}; h)}}\left[\mathsf{Ver}(k', \rho_k) = 1\right] \leq \theta(n)\,.$$

Furthermore, we say that G is a θ-OWSG for an algorithm $\mathcal{A}$ if the inequality concerning security (Property 2) holds for $\mathcal{A}$ without requiring any bound on the runtime of $\mathcal{A}$.

A weak OWSG can be recovered by the above definition for $T = \mathsf{poly}(n)$ and $\theta = 1 - 1/n^c$ for some constant c. It is shown in [63] that weak OWSGs imply OWSGs.

Fine-Grained primitives. In fine-grained one-way functions, there is at most a polynomial gap between the runtime of the function and runtime of the adversary.

Definition 9 (Fine-grained OWF). *Let $\eta > 1$ be a real number and $\theta : \mathbb{N} \to [0,1]$. A family of non-uniform algorithms $\mathsf{F} := \{\mathsf{F}_n\}_{n \in \mathbb{N}}$ is said to be an (η, θ)-fine-grained one-way function (FGOWF) if for any $O(T_{\mathsf{F}}^{\eta})$-time algorithm $\mathcal{A}$, for all sufficiently large n, it holds that*
$$\Pr_{x \sim \mathcal{U}_{\{0,1\}^n}}\left[\mathsf{F}(\mathcal{A}(\mathsf{F}(x))) = \mathsf{F}(x)\right] \leq \theta(n)\,,$$
where T_{F} is the runtime of F. If θ is constant, we simply say that F is a weak η-FGOWF.

Using Yao's direct-product construction [84], under some conditions on η and θ, an (η, θ)-FGOWF can be transformed into a weak (η')-FGOWG.

3 Lossy Mappings and Disguising Lemma

[35] derives a quantitative approach (called disguising distribution lemma) to measure how much information can be recovered from the ouput of a compressing mapping about its input, based on the compression size; a distinguishing variant of Fano's inequality. Such mappings are indeed a special type of lossy mappings, an observation upon which Ball et al. [10] develop their work.

In this section, we focus on variants of lossy mappings and their properties, and extend the disguising lemma. In our analysis, we consider both randomized functions and quantum mappings. All the statements hold with respect to both cases. For simplicity and generality, we only refer to quantum mappings. We explicitly highlight the distinction when the analysis requires to distinguish between the two cases.

Classically, a randomized function $R : \{0,1\}^* \to \{0,1\}^*$ is said to be ℓ-lossy for a distribution X if $I(X; R(X)) \leq \ell$. Below, we also consider general mappings with classical input and quantum output.

Definition 10 (Lossy Mapping). *Let $\ell, m \geq 0$. Let $R : \{0,1\}^* \to S$ be a mapping, where $S = \{0,1\}^*$ (classical mapping) or $S = \mathsf{MS}_*$ (quantum mapping). We say that R is ℓ-lossy for an m-tuple distribution $X = (X_1, X_2, \ldots, X_m)$ over $\{0,1\}^*$, if it holds that*

$$I(X; R(X)) \leq m\ell \ .$$

For the sake of simplicity, we say that R is ℓ-lossy, if it is ℓ-lossy for all m-tuple distributions.

The results by [10, 35] rely on the lossiness of the mapping for all distributions. Such a condition seems quite strong. We simplify this condition in two different directions. First, we consider lossy mappings over a particular class of distributions as follows:

Definition 11 (Splitting Distribution). *Let $S_0, S_1 \subseteq \{0,1\}^*$ be two disjoint sets. We say that a distribution $X = (X_1, \ldots, X_m)$ splits over the pair (S_0, S_1) if for each $i \in [m]$, either* $\mathrm{Supp}(X_i) \subseteq S_0$ *or* $\mathrm{Supp}(X_i) \subseteq S_1$.

Later, for the lossy reductions of a problem Π, we choose S_0 and S_1 as the sets Π_{N} and Π_{Y}. Splitting the distribution in such a way allows us to precisely calculate the lossiness of randomized encodings.

In the disguising distribution lemma in [35] and its improvement by [10], the lossiness (compression in the former and lossiness in the latter) is considered as in Definition 10 with respect to all possible input distributions. Instead, we show that the lemma remains almost intact for lossy maps over all splitting uniform distributions with sparse support. This is obtained by a more refined analysis but yet very similar to those of [10, 35]. Below, we have the main lemma of this section.

Lemma 3.1 (Extended Disguising Lemma). *Let n, m, m_0, m_1 be positive integers such that $m = m_0 + m_1 + 1$, and $R : \{0,1\}^* \to \mathsf{MS}_*$ be any quantum mapping. Further, let $S_0, S_1 \subseteq \{0,1\}^n$ be two disjoint sets, d be a positive integer, $\varepsilon > 0$ be real, and $s := \lceil n \ln 2/(2\varepsilon^2) \rceil$.*

For any choice of positive real ℓ, if R is ℓ-lossy for all ds-uniform distributions that split over the pair (S_0, S_1), then there exist two collections $K_1, \ldots, K_s$ and $T_1, \ldots, T_s$ of multisets of d elements respectively contained in S_0 and S_1, such that

- *for any $y \in S_0$, it holds that*

$$\mathop{\mathbb{E}}_{\substack{a \sim \mathcal{U}_{[s]} \\ \pi \sim \mathfrak{S}_m}} \left[\left\| R\left(\pi\left(\mathcal{U}_{K_a}^{\otimes m_0}, y, \mathcal{U}_{T_a}^{\otimes m_1}\right)\right) - R\left(\pi\left(\mathcal{U}_{K_a}^{\otimes (m_0+1)}, \mathcal{U}_{T_a}^{\otimes m_1}\right)\right) \right\|_1 \right] \leq \delta + \frac{2(m+1)}{d+1} + 2\varepsilon \ ;$$

- *and for any $y \in S_1$, it holds that*

$$\mathop{\mathbb{E}}_{\substack{a \sim \mathcal{U}_{[s]} \\ \pi \sim \mathfrak{S}_m}} \left[\left\| R\left(\pi\left(\mathcal{U}_{K_a}^{\otimes m_0}, y, \mathcal{U}_{T_a}^{\otimes m_1}\right)\right) - R\left(\pi\left(\mathcal{U}_{K_a}^{\otimes m_0}, \mathcal{U}_{T_a}^{\otimes (m_1+1)}\right)\right) \right\|_1 \right] \leq \delta + \frac{2(m+1)}{d+1} + 2\varepsilon \ ,$$

where

$$\delta := \min\left\{\sqrt{\frac{\ell \ln 2}{2m}}, 1 - 2^{-\frac{\ell}{m}-2}\right\}.$$

Note that the states inside the trace distance are mixed states since the inputs of R are randomized classical distributions.

The proof requires some background definitions and lemmas. Similar to [10,35], we define distributional stability as follows.

Definition 12. *Let n, m, m_0, m_1 be positive integers such that $m = m_0 + m_1 + 1$. For a real $\delta \in [0,1]$, a quantum mapping $R : \{0,1\}^{mn} \to \mathsf{MS}_*$ is said to be δ-quantumly-distributionally stable (δ-QDS) with respect to two distributions $(\mathcal{D}_0, \mathcal{D}_1)$ over $\{0,1\}^n$ if the following holds:*

$$\mathop{\mathbb{E}}_{\substack{y\sim\mathcal{D}_0\\ \pi\sim\mathfrak{S}_m}}\left[\left\|R\left(\pi\left(\mathcal{D}_0^{\otimes m_0}, y, \mathcal{D}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\mathcal{D}_0^{\otimes(m_0+1)}, \mathcal{D}_1^{\otimes m_1}\right)\right)\right\|_1\right] \leq \delta\ .$$

Note that the order of the pair $(\mathcal{D}_0, \mathcal{D}_1)$ matters. Furthermore, when $m_1 = 0$, we simply say that the mapping is δ-QDS with respect to $\mathcal{D}_0$.

Below, we recall an adaptation of [35, Lemma 8.10].

Lemma 3.2. *Assume that $R : \{0,1\}^{m\cdot n} \to \mathsf{MS}_*$ satisfies the properties in Lemma 3.1 for $m_1 = 0$. Then R is δ-QDS with respect to any ds-uniform distribution $\mathcal{D}_0$ supported on either S_0 or S_1.*

In the original lemma from [35], compression is used to bound the entropy of the mutual information. However, note that this can be argued directly from splitting lossiness, and that any restriction on the input distributions will give a result for the same restricted case.

The following lemma is the generalization of the above one.

Lemma 3.3. *Assume that $R : \{0,1\}^{mn} \to \mathsf{MS}_*$ satisfies the properties in Lemma 3.1. Then R is δ-QDS with respect to any ds-uniform independent distributions $(\mathcal{D}_0, \mathcal{D}_1)$ each supported on either S_0 or S_1.*

Proof. The proof is similar to that of [10, Proposition B.1]. Let $\pi \in \mathfrak{S}_m$ be a fixed permutation. One can rewrite it as the composition of two partial permutations π_0 and π_1, i.e., $\pi = \pi_0 \circ \pi_1$, such that π_1 only acts on the last m_1 arguments of the input. Let $\rho_\pi(y)$ be as follows

$$\rho_\pi(y) := R\left(\pi\left(\mathcal{D}_0^{\otimes m_0}, y, \mathcal{D}_1^{\otimes m_1}\right)\right)\ .$$

For $y, y' \sim \mathcal{D}_0$, two independent random variables, and $\pi \sim \mathfrak{S}_m$, we want to prove that

$$\mathbb{E}_{y,\pi}\left[\left\|\rho_\pi(y) - \rho_\pi(y')\right\|_1\right] \leq \delta\,.$$

Note that it is enough to bound the conditional distributions since

$$\mathbb{E}_{y,\pi}\left[\left\|\rho_\pi(y) - \rho_\pi(y')\right\|_1\right] = \mathbb{E}_\pi\left[\mathbb{E}_{y,\pi|\pi_1}\left[\left\|\rho_\pi(y) - \rho_\pi(y')\right\|_1\right]\right],$$

by the law of total probability. Let $R'(x_1, x_2, \ldots, x_{m_0+1})$ be the mapping that first samples π then evaluates $R\left(\pi_1\left(x_1, x_2, \ldots, x_{m_0+1}, \mathcal{D}_1^{\otimes m_1}\right)\right)$. For any fixed π_1, we show that R' is ℓ-lossy for all *ds*-uniform distributions that split over (S_0, S_1). Indeed, let $(\mathcal{X}_1, \ldots, \mathcal{X}_{m_0+1})$ be independent *ds*-uniform random variables with $\mathrm{Supp}(\mathcal{X}_i) \subseteq S_0$ or $\mathrm{Supp}(\mathcal{X}_i) \subseteq S_1$ for each $i \in [m_0+1]$, and $(\mathcal{Z}_1, \ldots, \mathcal{Z}_{m_1}) \sim \mathcal{D}_1^{\otimes m_1}$, thus $\mathrm{Supp}(\mathcal{Z}_i) \subseteq \mathrm{Supp}(\mathcal{D}_1) \subseteq S_j$ for all $i \in [m_1]$ and some $j \in \{0,1\}$. By the splitting lossiness of R for any *ds*-uniform distribution, we can bound the loss of R':

$$\begin{aligned}
\ell &\geq I_q(\pi_1(\mathcal{X}_1, \ldots, \mathcal{X}_{m_0+1}, \mathcal{Z}_1, \ldots, \mathcal{Z}_{m_1});\\
&\qquad\qquad R(\pi_1(\mathcal{X}_1, \ldots, \mathcal{X}_{m_0+1}, \mathcal{Z}_1, \ldots, \mathcal{Z}_{m_1})))\\
&= I_q(\mathcal{X}_1, \ldots, \mathcal{X}_{m_0+1}, \mathcal{Z}_1, \ldots, \mathcal{Z}_{m_1}; R(\pi_1(\mathcal{X}_1, \ldots, \mathcal{X}_{m_0+1}, \mathcal{Z}_1, \ldots, \mathcal{Z}_{m_1})))\\
&\geq I_q(\mathcal{X}_1, \ldots, \mathcal{X}_{m_0+1}; R(\pi_1(\mathcal{X}_1, \ldots, \mathcal{X}_{m_0+1}, \mathcal{Z}_1, \ldots, \mathcal{Z}_{m_1}))).
\end{aligned}$$

Finally, by Lemma 3.2 a splitting lossy map must also be δ-QSD, thus

$$\begin{aligned}
&\mathbb{E}_{y,\pi|\pi_1}\left[\left\|\rho_\pi(y)-\rho_\pi(y')\right\|_1\right]\\
&=\mathbb{E}_{y,\pi|\pi_1}\left[\left\|R\left(\pi\left(\mathcal{D}_0^{\otimes m_0},y,\mathcal{D}_1^{\otimes m_1}\right)\right)-R\left(\pi\left(\mathcal{D}_0^{\otimes m_0},y',\mathcal{D}_1^{\otimes m_1}\right)\right)\right\|_1\right]\\
&=\mathbb{E}_{y,\pi_0}\left[\left\|R'\left(\mathcal{D}_0^{\otimes m_0},y\right)-R'\left(\mathcal{D}_0^{\otimes m_0},y'\right)\right\|_1\right]\\
&\leq\delta\,.
\end{aligned}$$

□

If a mapping is distributionally stable with respect to a pair of distributions, then one can "sparsify" the distributions while nearly keeping the stability.

Lemma 3.4. *Let $n,m,m_0,m_1,\ell,S_0,S_1,R$ and δ be as in Lemma 3.1. Let $\mathcal{D}_0$ and $\mathcal{D}_1$ be two independent distributions with supports over S_0 and S_1, respectively. Let $\{x_i^{(0)}\}_{i\in[d+1]}$ and $\{x_i^{(1)}\}_{i\in[d+1]}$ be independent samples from $\mathcal{D}_0$ and $\mathcal{D}_1$, respectively. For each $j\in\{0,1\}$, let $y_j^*:=x_{i^*}^{(j)}$ be uniformly chosen from $\{x_i^{(j)}\}_{i\in[d+1]}$ and let $\widehat{\mathcal{D}}_j$ be the uniform distribution over the multiset $\{x_i^{(j)}\}_{i\in[d+1]\setminus\{i^*\}}$. Then it holds that*

$$\begin{aligned}
\mathbb{E}_{\pi\sim\mathfrak{S}_m}\left[\left\|R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes m_0},y_0^*,\widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right)-R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes(m_0+1)},\widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right)\right\|_1\right]\\
\leq\delta+\frac{2m_0+1}{d+1}\,,
\end{aligned}$$

and

$$\begin{aligned}
\mathbb{E}_{\pi\sim\mathfrak{S}_m}\left[\left\|R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes m_0},y_1^*,\widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right)-R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes m_0},\widehat{\mathcal{D}}_1^{\otimes(m_1+1)}\right)\right)\right\|_1\right]\\
\leq\delta+\frac{2m_1+1}{d+1}\,.
\end{aligned}$$

Proof. We prove the first statement. The other one is implied similarly. Let $\widetilde{\mathcal{D}}_0$ denote the uniform distribution over $\{x_i^{(0)}\}_{i\in[d+1]}$. For any fixed set of of multisets as above and any choice of permutation π and quantum mapping R, we have

$$\begin{aligned}
&\left\|R\left(\pi\left(\widetilde{\mathcal{D}}_0^{\otimes(m_0+1)},\widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right)-R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes(m_0+1)},\widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right)\right\|_1\\
&\leq\left\|\widetilde{\mathcal{D}}_0^{\otimes(m_0+1)}\otimes\widehat{\mathcal{D}}_1^{\otimes m_1}-\widehat{\mathcal{D}}_0^{\otimes(m_0+1)}\otimes\widehat{\mathcal{D}}_1^{\otimes m_1}\right\|_1\\
&\leq\left\|\widetilde{\mathcal{D}}_0^{\otimes(m_0+1)}-\widehat{\mathcal{D}}_0^{\otimes(m_0+1)}\right\|_1\\
&\leq(m_0+1)\left\|\widetilde{\mathcal{D}}_0-\widehat{\mathcal{D}}_0\right\|_1\,,
\end{aligned}$$

where we used the quantum data processing inequality for the first two upper bounds, and the property of tensor product for the last one. Since both $\widehat{\mathcal{D}}_0$ and $\widetilde{\mathcal{D}}_0$ are classical, their trace distance coincides with their statistical distance. Therefore, we have

$$\begin{aligned}
\left\|\widetilde{\mathcal{D}}_0-\widehat{\mathcal{D}}_0\right\|_1&=\frac{1}{2}\sum_{x\in\{x_i^{(0)}\}_{i\in[d+1]}}\left|\Pr_{\widetilde{\mathcal{D}}_0}(x)-\Pr_{\widehat{\mathcal{D}}_0}(x)\right|\\
&=\frac{1}{2(d+1)}+\frac{1}{2}\sum_{x\in\{x_i^{(0)}\}_{i\in[d+1]\setminus\{i^*\}}}\left|\frac{1}{d+1}-\frac{1}{d}\right|\\
&=\frac{1}{d+1}\,.
\end{aligned}$$

Similarly, it holds that

$$\left\|R\left(\pi\left(\widetilde{\mathcal{D}}_0^{\otimes m_0},y_0^*,\widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right)-R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes m_0},y_0^*,\widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right)\right\|_1\leq\frac{m_0}{d+1}\,.$$

From the triangle inequality, it follows that

$$\begin{aligned}
&\left\| R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes m_0}, y_0^*, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes (m_0+1)}, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right)\right\|_1 \\
&\leq \left\| R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes m_0}, y_0^*, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\widetilde{\mathcal{D}}_0^{\otimes m_0}, y_0^*, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right)\right\|_1 \\
&\quad + \left\| R\left(\pi\left(\widetilde{\mathcal{D}}_0^{\otimes m_0}, y_0^*, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\widetilde{\mathcal{D}}_0^{\otimes (m_0+1)}, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right)\right\|_1 \\
&\quad\quad + \left\| R\left(\pi\left(\widetilde{\mathcal{D}}_0^{\otimes (m_0+1)}, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes (m_0+1)}, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right)\right\|_1 \\
&< \left\| R\left(\pi\left(\widetilde{\mathcal{D}}_0^{\otimes m_0}, y_0^*, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\widetilde{\mathcal{D}}_0^{\otimes (m_0+1)}, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right)\right\|_1 \\
&\quad + \frac{2m_0+1}{d+1} \ .
\end{aligned}$$

Recall that R is ℓ-lossy with respect to all ds-uniform distributions that split over (S_0, S_1). Therefore, by Lemma 3.3 it is δ-QSD with respect to all ds-uniform pair of distributions each supported on either S_0 or S_1, including $(\widetilde{\mathcal{D}}_0, \widehat{\mathcal{D}}_1)$. Finally, by taking expectation from both sides above with respect to π, and using the fact that R is δ-QSD with respect to $(\widetilde{\mathcal{D}}_0, \widehat{\mathcal{D}}_1)$, one obtains the claimed upper bound.

□

Proof of Lemma 3.1. Consider the following two-player, simultaneous-move, zero-sum game:

- Player 1: chooses a pair of multisets $K \subseteq S_0$ and $T \subseteq S_1$, each of size d.
- Player 2: chooses an element $y \in S_0 \cup S_1$
- Payoff: if $y \in S_0$, Player 2 gains

$$\mathbb{E}_{\pi \sim \mathfrak{S}_m}\left[\left\| R\left(\pi\left(\mathcal{U}_K^{\otimes m_0}, y, \mathcal{U}_T^{\otimes m_1}\right)\right) - R\left(\pi\left(\mathcal{U}_K^{\otimes (m_0+1)}, \mathcal{U}_T^{\otimes m_1}\right)\right)\right\|_1\right] \ ,$$

otherwise, Player 2 gains

$$\mathbb{E}_{\pi \sim \mathfrak{S}_m}\left[\left\| R\left(\pi\left(\mathcal{U}_K^{\otimes m_0}, y, \mathcal{U}_T^{\otimes m_1}\right)\right) - R\left(\pi\left(\mathcal{U}_K^{\otimes m_0}, \mathcal{U}_T^{\otimes (m_1+1)}\right)\right)\right\|_1\right] \ .$$

Consider a ds-uniform strategy for Player 2, i.e. a distribution $\mathcal{Y}$ of y that is uniform over a multiset of pure strategies of size ds. We explain a strategy $(\mathcal{K}, \mathcal{T})$ for Player 1 that bounds the expected payoff. Player 1 chooses K by sampling d independent instances of the restriction of $\mathcal{Y}$ to S_0, and chooses T by sampling d independent instances of the restriction of $\mathcal{Y}$ to S_1. The expected payoff is

$$\begin{aligned}
E := &\Pr_{y \sim \mathcal{Y}}(y \in S_0)\ \mathbb{E}_{\pi,K,T}\left[\left\| R\left(\pi\left(\mathcal{U}_K^{\otimes m_0}, y, \mathcal{U}_T^{\otimes m_1}\right)\right)\right.\right. \\
&\qquad\qquad \left.\left. - R\left(\pi\left(\mathcal{U}_K^{\otimes (m_0+1)}, \mathcal{U}_T^{\otimes m_1}\right)\right)\right\|_1 \middle| y \in S_0\right] \\
&+ \Pr_{y \sim \mathcal{Y}}(y \in S_1)\ \mathbb{E}_{\pi,K,T}\left[\left\| R\left(\pi\left(\mathcal{U}_K^{\otimes m_0}, y, \mathcal{U}_T^{\otimes m_1}\right)\right)\right.\right. \\
&\qquad\qquad \left.\left. - R\left(\pi\left(\mathcal{U}_K^{\otimes m_0}, \mathcal{U}_T^{\otimes (m_1+1)}\right)\right)\right\|_1 \middle| y \in S_1\right] \ .
\end{aligned}$$

Let $x_1^{(0)}, x_2^{(0)}, \ldots, x_{d+1}^{(0)}$ and $x_1^{(1)}, x_2^{(1)}, \ldots, x_{d+1}^{(1)}$ be $d+1$ independent samples from $\mathcal{Y}|_{S_0}$ and $\mathcal{Y}|_{S_1}$, respectively. Sample $i^* \xleftarrow{\$} [d+1]$ and for $j \in \{0,1\}$, let $y_j^* := x_{i^*}^{(j)}$. Let $\widehat{\mathcal{Y}}_0$ and $\widehat{\mathcal{Y}}_1$ be the uniform distributions over the multisets $\{x_i^{(0)}\}_{i \in [d+1] \setminus \{i^*\}}$ and $\{x_i^{(1)}\}_{i \in [d+1] \setminus \{i^*\}}$, respectively. For $j \in \{0,1\}$, we have that $(y_j^*, \widehat{\mathcal{Y}}_0, \widehat{\mathcal{Y}}_1) \sim (\mathcal{Y}|_{S_j}, \mathcal{K}, \mathcal{T})$. Then, by Lemma 3.4, we have

$$\begin{aligned}
&\mathbb{E}_\pi\left[\left\| R\left(\pi\left(\widehat{\mathcal{Y}}_0^{\otimes m_0}, y, \widehat{\mathcal{Y}}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\widehat{\mathcal{Y}}_0^{\otimes (m_0+1)}, \widehat{\mathcal{Y}}_1^{\otimes m_1}\right)\right)\right\|_1 \ \middle|\ y \in S_0\right] \\
&\qquad\qquad \leq \delta + \frac{2m_0+1}{d+1} \ ,
\end{aligned}$$

and

$$\mathbb{E}_{\pi}\left[\left\|R\left(\pi\left(\widehat{\mathcal{Y}}_0^{\otimes m_0}, y, \widehat{\mathcal{Y}}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\widehat{\mathcal{Y}}_0^{\otimes m_0}, \widehat{\mathcal{Y}}_1^{\otimes (m_1+1)}\right)\right)\right\|_1 \;\middle|\; y \in S_1\right] \leq \delta + \frac{2m_1+1}{d+1}\,.$$

Therefore, we obtain $E \leq \delta + 2(m+1)/(d+1)$.

Above, we showed that for every ds-uniform strategy for Player 2, there exists a strategy for Player 1 that bounds the expected payoff by $\delta + 2(m+1)/(d+1)$. Let $\mathbf{M} := [\mathbf{M}_{ij}]_{i,j}$ be the matrix such that $\mathbf{M}_{ij}$ corresponds to the payoff when Player 1 outputs i and Player 2 outputs j. By Lemma 2.2, we have

$$\begin{aligned}\delta + 2(m+1)/(d+1) &\geq \max_{\mathcal{Q}\in\mathfrak{Q}_{ds}} \min_i \mathbb{E}_{j\sim\mathcal{Q}}[\mathbf{M}_{ij}] \geq \omega(\mathbf{M}) - \varepsilon(\mathbf{M}_{\max} - \mathbf{M}_{\min}) \\ &\geq \omega(\mathbf{M}) - \varepsilon\,,\end{aligned}$$

where $\mathfrak{Q}_{ds}$ is the set of all ds-uniform strategies for Player 2. It follows that $\omega(\mathbf{M}) \leq \delta + 2(m+1)/(d+1) + \varepsilon$.

Now we use Lemma 2.2 in other way around. In fact, the number of possible choices for Player 1 is $|S_0 \cup S_1| \leq 2^n$. Therefore, Lemma 2.2 asserts that there exists a s-uniform strategy for Player 2 such that for any possibly mixed strategy for Player 1, the expected payoff is at most ε-far from the value of the game $\omega(\mathbf{M})$. In other words, for this particular strategy of Player 1, the expected payoff is always at most

$$\omega(\mathbf{M}) + \varepsilon \leq \delta + 2(m+1)/(d+1) + 2\varepsilon\,.$$

Recall that a s-uniform strategy is, by definition, a uniformly sampled element from a size-s multiset of choices of the player. Note that Player 1 chooses a pair (K, T). Therefore, this strategy is essentially a uniform distribution over some multiset $\{(K_1, T_1), \ldots, (K_s, T_s)\}$, which concludes the proof.

□

4 Sparsely Lossy Problems

In this section, we first put forward a new abstraction, called f-distinguisher reduction, that is suitable for our analysis and implies definitions of f-reductions (adapted from Drucker [35]) as well as Karp and non-adaptive Turing reductions. Then, by considering the lossiness property(as defined in Section 3), we introduce sparsely lossy problems which will be the core of our analysis in the subsequent sections. Our analysis applies to both classical and quantum reductions. For the sake of simplicity and generality, we only refer to quantum reductions and we explicitly highlight the distinction when necessary.

4.1 f-Distinguisher Reductions

A Karp decision-to-decision reduction R from Π to Σ has the following property: $\chi_\Pi(x) = 1$ if and only if $\chi_\Sigma(R(x))$ (up to some error). In our work, the target problem Σ is not restricted and does not play any roles. Therefore, we consider the following more general notion: a mapping R is a reduction if there exists a (possibly unbounded) distinguisher $\mathcal{D}$ that can tell $R(x)$ and $R(x')$ apart, when $\chi_\Pi(x) \neq \chi_\Pi(x')$ (up to some error). A reduction is therefore a mapping that preserves the distinguishing power of the unbounded algorithm. [1] In other words, it preserves some information about the inputs. When the reduction is to a search problem, there must also exist an inverting algorithm such that given x and the solution (or witness) of $R(x)$, outputs $\chi_\Pi(x)$. To include such reductions, we generalize this definition once more by allowing the distinguisher to have one and only one of the instances x or x'. To see how this helps, we give an example: the reduction from ParamSat to MaxSat. In ParamSat, an instance $x := (\varphi, k)$, with φ a CNF formula and k an integer, is a

[1] Note that and unbounded algorithm can always distinguish YES and NO instances of a problem by simply solving them.

YES instance if and only if at least k clauses of φ are satisfiable. The MaxSat problem asks to find an assignment that satisfies the maximum number of clauses. Consdier the decision-to-search reduction as follows: given an instance $x := (\varphi, k)$ of ParamSat, the outputs of the reduction is φ. By having k and an assigment w_φ satisfying the maximum number of clauses of φ (solution of φ as a MaxSat instance), it computes $\chi_{\text{ParamSat}}(x)$ by comparing k and the number of satified clauses by w_φ. Note that it is necessary for the inverting algorithm to know k. In this subsection, we show that such reductions can be captured by the generalized distinguisher reductions:

Definition 13 (f-Distinguisher Reduction). *Let n, m be positive integers, and $\mu : \mathbb{N} \to [0,1]$ be a function of n. Let $f : \{0,1\}^m \to \{0,1\}$, and Π be a promise problem. A (μ, f^m)-distinguisher reduction for Π is a mapping $R : \{0,1\}^* \to S$, where $S = \{0,1\}^*$ (classical) or $S = \mathsf{MS}_*$ (quantum), for which there exists an unbounded distinguisher $\mathcal{D}$, such that for all $(x_1, \ldots, x_m)$ and $(x'_1, \ldots, x'_m)$ in $((\Pi_Y \cup \Pi_N) \cap \{0,1\}^n)^m$ where $f(\chi_\Pi(x_1), \ldots, \chi_\Pi(x_m)) \neq f(\chi_\Pi(x'_1), \ldots, \chi_\Pi(x'_m))$, we have*

$$\mathop{\mathbb{E}}_{i \sim \mathcal{U}_{[m]}} \left| \Pr[1 \leftarrow \mathcal{D}(h_i, R(x_1, \ldots, x_m))] - \Pr[1 \leftarrow \mathcal{D}(h_i, R(x'_1, \ldots, x'_m))] \right| \geq 1 - 2\mu(n)\,,$$

where $h_i := (x_i, \{\chi_\Pi(x_j)\}_j, \{\chi_\Pi(x'_j)\}_j)$. We call μ the error of the reduction.

f-Reductions

Drucker [35, Definition 8.2] defines an f-compression reduction for a promise problem Π in a somewhat similar fashion that we define f-distinguisher reductions: as a mapping that sends an instances $x_1, \ldots, x_m$ of size n to a quantum state ρ, such that there exists a binary measurement $\mathcal{M}$ (not necessarily efficient) that outputs $f(\chi_\Pi(x_1), \ldots, \chi_\Pi(x_m))$ with probability more than $1 - \mu$. We adapt this definition as below.

Definition 14 (f-Reduction). *Let n, m be positive integers, and $\mu : \mathbb{N} \to [0,1]$ be a function of n. Let $f : \{0,1\}^m \to \{0,1\}$, and Π be a promise problem. A (μ, f^m)-reduction for Π is a mapping $R : \{0,1\}^{mn} \to S$, where $S = \{0,1\}^*$ (classical) or $S = \mathsf{MS}_*$ (quantum), for which there exists a family of unbounded algorithms $\{\mathcal{M}_k\}_{k \in \mathbb{N}}$, such that for all $(x_1, \ldots, x_m) \in ((\Pi_Y \cup \Pi_N) \cap \{0,1\}^n)^m$,*

$$\Pr\left[\mathcal{M}(R(x_1, \ldots, x_m)) = f(\chi_\Pi(x_1), \ldots, \chi_\Pi(x_m))\right] \geq 1 - \mu(n)\,,$$

where the probability is taken over the randomness of R and $\mathcal{M}$. We call μ the error of the reduction.[1]

In the following, we show that f-reductions are special cases of f-distinguisher reductions (per Definition 13) when the hint h_i is set to be empty.

Lemma 4.1. *Let $f : \{0,1\}^m \to \{0,1\}$, and Π be a promise problem. If R is a (μ, f^m)-reduction for Π, then R is also a (μ, f^m)-distinguisher reduction for Π.*

Proof. Recall that for an f-reduction there exists an algorithm $\mathcal{M}$ such that

$$\Pr[\mathcal{M}(R(x_1, \ldots, x_m)) = f(\chi_\Pi(x_1), \ldots, \chi_\Pi(x_m))] \geq 1 - \mu(n)\,,$$

which implies that $\mathcal{M}$ can distinguish $R(x_1, \ldots, x_m)$ from $R(x'_1, \ldots, x'_m)$ with probability at least $1 - 2\mu$. Therefore, there exists an unbounded distinguisher $\mathcal{D}$ such that for h_i per Definition 13, we have

$$\begin{aligned}\mathop{\mathbb{E}}_{i \sim \mathcal{U}_{[m]}} \left| \Pr[1 \leftarrow \mathcal{D}(h_i, R(x_1, \ldots, x_m))] - \Pr[1 \leftarrow \mathcal{D}(h_i, R(x'_1, \ldots, x'_m))] \right| \\ \geq \left| \Pr[1 \leftarrow \mathcal{M}(R(x_1, \ldots, x_m))] - \Pr[1 \leftarrow \mathcal{M}(R(x'_1, \ldots, x'_m))] \right| \\ \geq 1 - 2\mu\,,\end{aligned}$$

where for the first inequality we used the fact that revealing more information to the distinguisher does not decrease its advantage. □

[1] When considering quantum mappings, $\mathcal{M}$ can be a binary quantum measurement.

Turing and Karp Reductions

In this part, we focus on (non-adaptive) Turing and Karp reductions, demonstrating that they are f-distinguisher reductions. This supports the generality of Definition 13 and will be used in Section 8.

In the following, we first recall the definition of Karp and (non-adaptive) Turing reductions in Definitions 15 and 16, and prove in Lemmas 4.2 and 4.3 that the two are f-distinguisher reductions.

Definition 15 (Non-Adaptive Turing f-Reduction). *Let n be a positive integer and $\mu : \mathbb{N} \to [0,1]$ be a function of n. Let $f : \{0,1\}^m \to \{0,1\}$, Π be a promise problem, and Σ be a promise or search problem. A non-adaptive (μ, f^m)-Turing reduction from Π to Σ consists of an algorithm R_{Turing} that on input $(x_1, \ldots, x_m)$, where $x_i \in \{0,1\}^n$ for $i \in [m]$, outputs $(y_1, \ldots, y_k) \in \{0,1\}^*$ and a circuit C such that*

- *if Σ is a promise problem:*

$$\Pr\left[C(y_1, \chi_\Sigma(y_1), \ldots, y_k, \chi_\Sigma(y_k)) = f\left(\chi_\Pi(x_1), \ldots, \chi_\Pi(x_m)\right)\right] \geq 1 - \mu(n) \ .$$

- *if Σ is a search problem:*

$$\Pr\left[C(y_1, w_{y_1}, \ldots, y_k, w_{y_k}) = f\left(\chi_\Pi(x_1), \ldots, \chi_\Pi(x_m)\right)\right] \geq 1 - \mu(n) \ ,$$

where w_{y_i} is the witness of y_i in Σ for all $i \in [k]$.

The definition above can be generalized in the following manner: y_i's can be instances of different problems Σ_i's instead of one single problem Σ. All our results also hold in this setting.

Definition 16 (Karp f-Reduction). *Let n be a positive integer and $\mu : \mathbb{N} \to [0,1]$ be a function of n. Let $f : \{0,1\}^m \to \{0,1\}$ and Π be a promise problem and Σ be a promise or search problem. A (μ, f^m)-Karp reduction from Π to Σ consists of an algorithm R_{Karp} and a circuit C, where R_{Karp} on input $(x_1, \ldots, x_m)$, where $x_i \in \{0,1\}^n$ for $i \in [m]$, outputs $y \in \{0,1\}^*$ such that*

- *if Σ is a promise problem:*

$$\Pr\left[C(y, \chi_\Sigma(y)) = f\left(\chi_\Pi(x_1), \ldots, \chi_\Pi(x_m)\right)\right] \geq 1 - \mu(n) \ .$$

- *if Σ is a search problem:*

$$\Pr\left[C(y, w_y) = f\left(\chi_\Pi(x_1), \ldots, \chi_\Pi(x_m)\right)\right] \geq 1 - \mu(n) \ ,$$

where w_y is the witness of y in Σ.

Note that in a Karp reduction, the circuit C does not depend on the instance x. In fact, in a standard definition of a Karp reduction to a promise problem, C simply outputs $\chi_\Pi(x)$.

In the following lemma, we show that all non-adaptive Turing reductions are f-distinguisher reduction.

Lemma 4.2 (Turing f-Reduction is f-Distinguisher Reduction). *Let $\mu : \mathbb{N} \to [0,1]$. Let Π be a promise problem and Σ be a promise or search problem. If R_{Turing} is a non-adaptive (μ, f^m)-Turing reduction (Definition 15) from Π to Σ, then it is (μ, f^m)-distinguisher reduction for Π.*

Proof. The distinguisher $\mathcal{D}$ in Figure 1 satisfies the definition of (μ, f^m)-distinguisher reductions (Definition 13). This is because if $B = ((y_1, \ldots, y_k), C)$ is an output of $R_{\text{Turing}}(x_1, \ldots, x_m)$, then by the correctness of the reduction, it holds with high probability that $C(y_1, \chi_\Sigma(y_1), \ldots, y_k, \chi_\Sigma(y_k)) = f\left(\chi_\Pi(x_1), \ldots, \chi_\Pi(x_m)\right)$, if Σ is a promise problem, and similarly $C(y_1, w_{y_1}, \ldots, y_k, w_{y_k})$ $f\left(\chi_\Pi(x_1), \ldots, \chi_\Pi(x_m)\right)$, if Σ is a search problem.

□

Lemma 4.3 (R_{Karp} is f-Distinguisher Reduction). *Let $\mu : \mathbb{N} \to [0,1]$. Let Π be a promise problem and Σ be a promise or search problem. If R_{Karp} is a (μ, f^m)-Karp reduction (Definition 16) from Π to Σ, then it is (μ, f^m)-distinguisher reduction for Π.*

Proof. Since any Karp reduction is a Turing reduction, the statement holds due to Lemma 4.2. □

Algorithm 1 Distinguisher $\mathcal{D}$ for non-adaptive Turing reductions.

Parameters: n, m, f, Π, Σ
Input: A pair (h_i, B), where $h_i := (x_i, \{\chi_\Pi(x_j)\}_j, \{\chi_\Pi(x'_j)\}_j)$ for a uniformly random $i \in [m]$ and $B = ((y_1, \ldots, y_k), C)$.
Promise: Either $B \leftarrow R(x_1, \ldots, x_m)$ or $B \leftarrow R(x'_1, \ldots, x'_m)$ for some $(x'_1, \ldots, x'_m) \in (\{0,1\}^n)^m$ such that $f(\chi_\Pi(x_1), \ldots, \chi_\Pi(x_m)) \neq f(\chi_\Pi(x'_1), \ldots, \chi_\Pi(x'_m))$.
Output: A bit b.

Parse $h_i := (x_i, \{\chi_\Pi(x_j)\}_j, \{\chi_\Pi(x'_j)\}_j)$ and $B = ((y_1, \ldots, y_k), C)$.
if Σ is a promise problem: **then**
    Compute $\chi_\Sigma(y_1), \ldots, \chi_\Sigma(y_k)$.
    Compute $\widehat{b} \leftarrow C(y_1, \chi_\Sigma(y_1), \ldots, y_k, \chi_\Sigma(y_k))$.
else
    Compute the witnesses $w_{y_1}, \ldots, w_{y_k}$ in Σ.
    Compute $\widehat{b} \leftarrow C(y_1, w_{y_1}, \ldots, y_k, w_{y_k})$.
if $\hat{b} = f(\chi_\Pi(x_1), \ldots, \chi_\Pi(x_m))$ **then**
    Return 1.
else
    Return 0.

4.2 Sparsely Lossy Problems

To analyze the lossiness of f-distinguisher reductions, we fix the set of functions f to those ones that are invariant under permuting their inputs.

Definition 17 (Permutation-Invariant Boolean Function). *We call a Boolean function $f : \{0,1\}^m \to \{0,1\}$ permutation-invariant if for every $\pi \in \mathfrak{S}_m$, it holds that $f(\pi(b_1, b_2, \ldots, b_m)) = f(b_1, b_2, \ldots, b_m)$.*

This set of functions is of great interest. The functions And, Or, and Maj that were considered in [10,35] are all non-constant permutation-invarant. Moreover, the (non-monotone) functions Parity and Mod_k are of this type as well as Threshold_k.

We use the following technical lemma about non-constant permutation-invariant functions.

Lemma 4.4. *Let $f : \{0,1\}^m \to \{0,1\}$ be a non-constant permutation-invariant function. Then there exists an integer $1 \leq p \leq m$ such that*

$$f(\underbrace{1, 1, \ldots, 1}_{p-1}, 0, 0, \ldots, 0) = 0\,, \quad \text{and} \quad f(\underbrace{1, 1, \ldots, 1}_{p}, 0, 0, \ldots, 0) = 1\,.$$

We let $p(f)$ denote the minimum choice of such an integer.

Proof. The set $\{0,1\}^m$ can be partitioned into $m+1$ equivalence classes where each class consists of strings with the same number of 1's. We note that the result of a permutation on an input falls in the same equivalence class. Therefore, since the function is permutation-invariant, then the evaluation of f over each input is determined by its class. Because the function is non-constant, there must exist two consecutive classes (the classes can be ordered by the number of 1's that they represent) with different evaluation under f. This completes the proof. □

Finally, we introduce the notion of *sparsely lossy problems* which are promise problems that admit sparsely lossy f-distinguisher reductions.

Definition 18 (Sparsely Lossy Problems). *Let n, m be positive integers, λ, T be positive reals, $\gamma \in (0,1]$, and $\mu \in [0, 1/2)$. A promise problem Π is said to be $(T, \mu, f^m, \lambda, \gamma)$-sparsely lossy if there exists a non-uniform (μ, f^m)-distinguisher reduction R (per Definition 13) for Π with the following properties:*

1. *f is some non-constant permutation-invariant function $f : \{0,1\}^m \to \{0,1\}$, and*
2. *the reduction R runs in time T, and*
3. *R is λ-lossy (per Definition 10) for all pairwise independent B-uniform distributions over n-bit strings that split over (Π_Y, Π_N) (per Definition 11), where $B = \lceil 4(m+1)/\gamma \rceil \cdot \lceil 8n \ln 2/\gamma^2 \rceil$.*

We explicitly mention the type of the reduction R (classical or quantum) when the distinction is necessary. Also, we interchangeably say that the reduction R as above is sparsely lossy.

We recall δ from the upper bound for splitting lossy functions in Lemma 3.1 for clarity, as it will be frequently used in the following sections.

Definition 19. *We let $\delta : \mathbb{R}^+ \to \mathbb{R}^+$ to be the following function*

$$\delta(\lambda) := \min\left\{\sqrt{\frac{\lambda \ln 2}{2}}, 1 - 2^{-\lambda-2}\right\}.$$

5 Zero-Knowledgeness from Sparsely Lossy Problems

In this section, we show that lossy problems admit Karp reductions to the statistical difference problem or the quantum state distinguishability problem, depending on the type of the lossy reduction. We provide a fine-grained analaysis. When restricted to polynomial-time AND-compression reductions, this recreates the result of Drucker [35, Theorem 8.14]: roughly, if a promise problem Π has a (quantum) polynomial-time AND-compression reduction, then Π must belong to SZK (resp., QSZK). Similar statement holds for the AND- or MAJ-lossy reductions (see [10]). We note that our result holds for any non-constant permutation-invarinat function, requires a less restricted notion of lossiness, and allows a fine-grained runtime analysis

Theorem 5.1. *Let Π be $(T, \mu, f^m, \lambda, \gamma)$-sparsely lossy. Let us assume that $\theta_{\mathsf{szk}} := (1-2\mu)^2/(\delta(\lambda)+\gamma) > 1$, with $\delta(\lambda)$ as in Definition 19. Then Π reduces to a problem in* QSZK*, with zero error, in time $O((T+m^2\gamma^{-1})/\log\theta_{\mathsf{szk}})$, and with a classical advice of size $4mn/\gamma$, as described in Algorithm 2. Moreover, the reduction is deterministic (but non-uniform) and Π reduces to* SZK *if Π is lossy with respect to a classical reduction.*

Algorithm 2 Reduction from Π to $\mathrm{QSD}_{1/4,3/4}$.

Parameters: $n, m, \mu, f, \lambda, \gamma, R, \Pi$ as in Definition 18. Further

$$S_0 := \Pi_N \cap \{0,1\}^n, \quad S_1 := \Pi_Y \cap \{0,1\}^n, \quad \varepsilon := \frac{\gamma}{4}, \quad d := \left\lceil \frac{m+1}{\varepsilon} \right\rceil, \; s := \left\lceil \frac{n \ln 2}{2\varepsilon^2} \right\rceil,$$

and $K_1, \dots K_s, T_1, \dots, T_s$ as in Lemma 3.1.

Input: An instance $y \in \{0,1\}^n$.

Advice: $p := p(f)$ as in Lemma 4.4, $b_Y, b_N \in \{0,1\}$ respectively representing whether $\Pi_Y \cap \{0,1\}^n$ and $\Pi_N \cap \{0,1\}^n$ are empty. K_a, T_a, π for some uniformly chosen $a \in [s]$ and $\pi \in \mathfrak{S}_m$.

Output: A pair of circuits (C_0, C_1).

1: If $b_N = 1$, return (Y_0, Y_1) where $\|Y_0 - Y_1\|_1 \le 1/4$.
2: If $b_Y = 1$, return (N_0, N_1) where $\|N_0 - N_1\|_1 \ge 3/4$.
3: Let $\widehat{C}_0$ be the following circuit: it samples $\widetilde{x} \sim \left(\mathcal{U}_{K_a}^{\otimes m-p+1}, \mathcal{U}_{T_a}^{\otimes p-1}\right)$, then it outputs $R(\pi(\widetilde{x}))$.
4: Let $\widehat{C}_1$ be the following circuit: it samples $\widetilde{x} \sim \left(\mathcal{U}_{K_a}^{\otimes m-p}, y, \mathcal{U}_{T_a}^{\otimes p-1}\right)$, then it outputs $R(\pi(\widetilde{x}))$.
5: Compute $(C_0, C_1) \leftarrow \mathsf{Polarize}(\widehat{C}_0, \widehat{C}_1, 1^2)$.
6: Return (C_0, C_1).

Remark 2 (Input-output type of the circuits). Consider $(\widehat{C}_0, \widehat{C}_1)$ to be the circuit pair in Algorithm 2, Lines 3 and 4. When R is a randomized reduction, the two circuits are also randomized. Part of their randomness input is used to sample $\widetilde{x}$ and the other part is fed to R. Let κ be the size of the total randomness. For $r \in \{0,1\}^\kappa$ and any $b \in \{0,1\}$, we let $\widehat{C}_b(r)$ denote the outcome of $\widehat{C}_b$ given the randomness r. On the other hand, when R is quantum, the circuits will be mixed algorithms; classical randomness is required for sampling $\widetilde{x}$. Let κ' be the size of total randomness.[1] For any $r \in \{0,1\}^{\kappa'}$

[1]Note that κ and κ' are possibly different depending on how much classical randomness R requires.

and any $b \in \{0,1\}$, we let the mixed outcome of $\widehat{C}_b$ be $\widehat{C}_b\,|r,\mathbf{0}\rangle$ where $|\mathbf{0}\rangle$ is some appropriate-size ancilla, emphasizing its mixed classical-quantum nature. When it is not relevent, we drop the dependency on r for simplification.

Proof of Theorem 5.1. In the following, we assume that R is quantum. The classical case is similar with the only difference being the type of the inputs and outputs of $(\widehat{C}_0, \widehat{C}_1)$.

Consider the case $y \in \mathit{\Pi}_Y$. We bound the ℓ_1 distance (per Definition 2) of the outcomes of $\widehat{C}_0$ and $\widehat{C}_1$ from below. Sample a uniform coin $b \sim U_{\{0,1\}}$, and let $z \leftarrow \widehat{C}_b\,|r,\mathbf{0}\rangle$ where r follows the uniform distribution. We drop the dependency on r for simplification. Let $\mathcal{A}$ be a (possibly unbounded) distinguisher that takes z as input and guesses which circuit ($\widehat{C}_0$ or $\widehat{C}_1$) is used to compute z. Let $\mathcal{A}$ be the quantum distinguisher of the (μ, f^m)-distinguisher reduction (that comes from Definition 18) for $\mathit{\Pi}$. On the one hand, if z is computed by $\widehat{C}_0$, we have that $\widetilde{x} := (x_1, \ldots, x_m) \sim \left(\mathcal{U}_{K_a}^{\otimes m-p+1}, \mathcal{U}_{T_a}^{\otimes p-1}\right)$ with $K_a \subseteq \mathit{\Pi}_N \cap \{0,1\}^n$ and $T_a \subseteq \mathit{\Pi}_Y \cap \{0,1\}^n$. Then, since $\widetilde{x}$ contains $p-1$ YES instances by Lemma 4.4, for any $\pi \in \mathfrak{S}_m$, we have

$$f(\pi(\chi_{\mathit{\Pi}}(x_1), \ldots, \chi_{\mathit{\Pi}}(x_m))) = 0\ .$$

On the other hand, if z is computed by $\widehat{C}_1$, we have that $\widetilde{x}$ contains one more YES instance $y \in \mathit{\Pi}_Y \cap \{0,1\}^n$, therefore,

$$f(\pi(\chi_{\mathit{\Pi}}(x_1), \ldots, \chi_{\mathit{\Pi}}(x_m))) = 1\ .$$

Moreover, revealing π with the description of the circuits does not decrease the success probability of the distinguisher, thus by the quantum f-distinguishability of the reduction, we have

$$\begin{aligned}
&\|\widehat{C}_0\,|\mathbf{0}\rangle - \widehat{C}_1\,|\mathbf{0}\rangle\,\|_1 \\
&\geq \mathbb{E}_{i\sim\mathcal{U}_{[m]}}\left|\Pr\left[1 \leftarrow \mathcal{D}(x_i, \widehat{C}_0\,|\mathbf{0}\rangle)\right] - \Pr\left[1 \leftarrow \mathcal{D}(x_i, \widehat{C}_1\,|\mathbf{0}\rangle)\right]\right| \\
&\geq 1 - 2\mu(n)\,.
\end{aligned}$$

Now, we discuss the case of $y \in \mathit{\Pi}_N$. We consider a modification of the distinguishing game where the random variables a and π are also given to the distinguisher. Revealing a, π along with z does not decrease the success probability of the distinguisher, thus we can bound the original distinguishing probability by the distinguishing probability of the new task. It holds that

$$\begin{aligned}
&\|\widehat{C}_0\,|\mathbf{0}\rangle - \widehat{C}_1\,|\mathbf{0}\rangle\,\|_1 \\
&\qquad\qquad \leq \left\| R\left(\pi\left(\mathcal{U}_{K_a}^{\otimes m-p+1}, \mathcal{U}_{T_a}^{\otimes p-1}\right)\right) - R\left(\pi\left(\mathcal{U}_{K_a}^{\otimes m-p}, y, \mathcal{U}_{T_a}^{\otimes p-1}\right)\right)\right\|_1,
\end{aligned}$$

By taking the expectation over a and π, we have

$$\begin{aligned}
&\|\widehat{C}_0\,|\mathbf{0}\rangle - \widehat{C}_1\,|\mathbf{0}\rangle\,\|_1 \\
&\qquad \leq \mathbb{E}_{a\sim\mathcal{U}_{[s]}, \pi\sim\mathfrak{S}_m}\Big[\Big\| R\left(\pi\left(\mathcal{U}_{K_a}^{\otimes m-p+1}, \mathcal{U}_{T_a}^{\otimes p-1}\right)\right) \\
&\qquad\qquad\qquad - R\left(\pi\left(\mathcal{U}_{K_a}^{\otimes m-p}, y, \mathcal{U}_{T_a}^{\otimes p-1}\right)\right)\Big\|_1\Big]\,. \qquad (6)
\end{aligned}$$

By our choice of ε, d, s, $K_1, \ldots, K_s, T_1, \ldots, T_s$ and Lemma 3.1, we conclude that

$$\|\widehat{C}_0\,|\mathbf{0}\rangle - \widehat{C}_1\,|\mathbf{0}\rangle\,\|_1 \leq \delta + \frac{2(m-p+1)}{d+1} + 2\varepsilon \leq \delta + \gamma\,.$$

Let $\alpha := (\delta+\gamma)$ and $\beta := (1-2\mu)$. Above, we proved that $(\widehat{C}_0, \widehat{C}_1)$ is an instance of $\mathrm{QSD}_{\alpha,\beta}$. The runtime of each circuit is $T + m^2/\gamma$ since R runs in T, each sampling of $\mathcal{U}_{K_a}$ (or $\mathcal{U}_{T_a}$) takes $O(d) = O(m/\gamma)$, and applying the permutation takes $O(m)$ according to the Fisher-Yates' algorithm. Recall that $\theta_{\mathsf{szk}} = \beta^2/\alpha$. Therefore, the runtime of the algorithm $\mathsf{Polarize}(\widehat{C}_0, \widehat{C}_1, 1^2)$ and its output size are both of $O((T+m^2)/(\gamma \log \theta_{\mathsf{szk}}))$ according to Lemma 2.7.

□

6 One-Way Functions from Sparsely Lossy Problems

In this section and in Section 7, we discuss how sparsely lossy problems can be used to build cryptographic primitives. In Theorem 6.1, we construct EFI schemes. The statement allows both classical reductions and quantum reductions. We immediately obtain one-way functions (or quantum bit commitments if the reduction is quantum), by taking into account the known transforms from EFI schemes (see Remark 1). However, the required condition on the lossiness is highly restrictive. More precisely, λ must be a small constant. In Theorem 6.2 and 7.1, we explain how one can tackle this issue using different constructions. The construction in Theorem 6.2 is inspired by [10], and resist adaptations to the quantum settings. On the other hand, the construction in Theorem 7.1 is quite flexible and allows obtaining one-way state generators. Finally, we note that the latter does imply one-way functions, too, but for simplicity, we only discuss one-way state generators.

Theorem 6.1. *Let Π be $(T, \mu, f^m, \lambda, \gamma)$-sparsely lossy. Let us assume that $\theta_{\mathsf{efi}} := (1-2\mu) - 3(\delta(\lambda) + \gamma) > 0$, with $\delta(\lambda)$ as in Definition 19. Then there exists an algorithm EFI that runs in $O(T + m^2\gamma^{-1})$ and an oracle algorithm $\mathcal{C}$, such that for any algorithm $\mathcal{A}$ one and only one of the following statements holds:*

I. $\mathcal{C}^{\mathcal{A}}$ solves $\Pi \cap \{0,1\}^n$ in time $O((T + m^2\gamma^{-1})\theta_{\mathsf{efi}}^{-2})$ with $O(\theta_{\mathsf{efi}}^{-2})$ queries to $\mathcal{A}$,
II. EFI is $(1-2\mu, 1-2\mu-\theta_{\mathsf{efi}}/2)$-EFI for $\mathcal{A}$.

Moreover, if the sparsely lossy reduction of Π is classical, EFI would also be classical.

Remark 3. From the conditions of Theorem 6.1, it must hold that $\delta < 1/3$, therefore, λ must be small. Most notably, the statement does not include perfect 1-sparsely lossy reductions. However, this can be overcome as follows: Let R be 1-sparsely lossy and perfect. Consider the new reduction R' that with probability 0.35 randomly outputs a YES or a NO instance of the target language (note that instance can be given as advice). Otherwise, it applies R. The new reduction is 0.35-sparsely lossy with error 0.375 which satisfies the condition $(1-2\mu) - 3(\delta(\lambda) + \gamma) > 0$.

Proof. We prove the case where R is quantum. The classical case can be done similarly. Let Π be the promised problem in the statement. Let $\mathcal{F}$ denote Algorithm 2 that returns the two circuits in Lines 3 and 4, and h be its advice as follows: $h := (K_a, T_a, p, b_Y, b_N)$. The construction of the non-uniform EFI is the following:

- $\mathsf{EFI}_h(1^n, b)$: Sample $y \sim \mathcal{U}_{T_a}$. Compute $(\widehat{C}_0, \widehat{C}_1) \leftarrow \mathcal{F}(y)$. Return the state $\widehat{C}_b |\mathbf{0}\rangle$.

Note that T_a has only YES instances.

The two output states are statistically far. By Theorem 5.1, the pair of circuits $(\widehat{C}_0, \widehat{C}_1) \leftarrow \mathcal{F}(y)$ is a $\mathrm{QSD}_{1-2\mu,\delta+\gamma}$ instance. Since $y \in \Pi_Y$, then $\|\widehat{C}_0 |\mathbf{0}\rangle - \widehat{C}_1 |\mathbf{0}\rangle\|_1 \geq 1 - 2\mu$. This concludes the statistical distinguishability.

On the computational indistinguishability, we will argue by contradiction. Assume there exists an adversary $\mathcal{A}$ that distinguishes the EFI states $\widehat{C}_b |\mathbf{0}\rangle$ with advantage ν that is to be determined later. Let us consider an algorithm $\mathcal{B}$ targetting Π as follows: given an instance $z \in \{0,1\}^n$, it first computes $(C'_0, C'_1) \leftarrow \mathcal{F}(z)$, then it samples a uniform coin $b \sim \mathcal{U}_{\{0,1\}}$ and relays $C'_b |\mathbf{0}\rangle$ to the distinguisher $\mathcal{A}$. Finally, $\mathcal{B}$ will return 1 if $\mathcal{A}$ returns b, and 0 otherwise.

Case $z \in \Pi_Y$: Suppose that z has been sampled from $\mathcal{U}_{T_a}$. Then, the (mixed) state $C'_b |\mathbf{0}\rangle$ that we deliver to the adversary $\mathcal{A}$ would be identical to the EFI state $\widehat{C}_b |\mathbf{0}\rangle$. Therefore, from the ν-distinguishability of EFI states for $\mathcal{A}$, we would have

$$\Pr(\mathcal{B}(z) = 1) = \Pr(\mathcal{A}(P_b |0\rangle) = b) \geq \frac{1}{2} + \frac{\nu}{2} .$$

We know that z does not necessarily follow the distribution $\mathcal{U}_{T_a}$. However, one can argue that $\widehat{C}_b$ is not far from C'_b by leveraging the disguising lemma. We have that

$$
\begin{aligned}
&\|\widehat{C}_0 \otimes \widehat{C}_1 \,|\mathbf{0},\mathbf{0}\rangle - C'_0 \otimes C'_1 \,|\mathbf{0},\mathbf{0}\rangle\,\|_1 \\
&\leq \|\widehat{C}_1 \,|\mathbf{0}\rangle - C'_1 \,|\mathbf{0}\rangle\,\|_1 \\
&\leq \mathbb{E}_{a\sim\mathcal{U}_{[s]},\pi\sim\mathfrak{S}_m}\Big[\Big\|R\left(\pi\left(\mathcal{U}_{K_a}^{\otimes m-p},\mathcal{U}_{T_a}^{\otimes p}\right)\right) \\
&\qquad\qquad - R\left(\pi\left(\mathcal{U}_{K_a}^{\otimes m-p},y,\mathcal{U}_{T_a}^{\otimes p-1}\right)\right)\Big\|_1\Big] \\
&\leq \delta + \frac{2(m+1)}{d+1} + \varepsilon \\
&\leq \delta + \gamma\ ,
\end{aligned}
$$

where we used the fact that $\widehat{C}_0 = C'_0$, properties of trace distance, and Lemma 3.1. Using the fact that the trace distance is decreasing under partial trace, for any $b \in \{0,1\}$, we obtain

$$
\|\widehat{C}_b \,|\mathbf{0}\rangle - C'_b \,|\mathbf{0}\rangle\,\|_1 \leq \delta + \gamma\ .
$$

The adversary $\mathcal{A}$ can thus distinguish the general C'_b with probability

$$
\begin{aligned}
\Pr(\mathcal{B}(z)=1) &= \Pr(\mathcal{A}(C'_b\,|\mathbf{0}\rangle)=b) \\
&= \frac{1}{2} + \frac{1}{2}\left|\Pr_{x\leftarrow C'_0}(\mathcal{A}(x)=1) - \Pr_{x\leftarrow C'_1}(\mathcal{A}(x)=1)\right| \\
&\geq \frac{1}{2} + \frac{1}{2}\Bigg(\left|\Pr_{x\leftarrow \widehat{C}_0}(\mathcal{A}(x)=1) - \Pr_{x\leftarrow \widehat{C}_1}(\mathcal{A}(x)=1)\right| \\
&\qquad - \left|\Pr_{x\leftarrow \widehat{C}_0}(\mathcal{A}(x)=1) - \Pr_{x\leftarrow C'_0}(\mathcal{A}(x)=1)\right| \\
&\qquad - \left|\Pr_{x\leftarrow \widehat{C}_1}(\mathcal{A}(x)=1) - \Pr_{x\leftarrow C'_1}(\mathcal{A}(x)=1)\right|\Bigg) \\
&\geq \frac{1}{2} + \frac{\nu}{2} - \delta - \gamma\ .
\end{aligned}
\tag{7}
$$

Case $z \in \Pi_N$: By Theorem 5.1, the two circuits $(C'_0, C'_1) \leftarrow \mathcal{F}(z)$ are close in trace distance, namely,

$$
\|C'_0\,|\mathbf{0}\rangle - C'_1\,|\mathbf{0}\rangle\,\|_1 \leq \delta + \gamma\ .
$$

Recall that the trace distance provides the maximum distinguishability advantage for *any* distinguisher, including $\mathcal{A}$, therefore

$$
\begin{aligned}
\Pr(\mathcal{B}(z)=1) &= \Pr(\mathcal{A}(C'_b\,|\mathbf{0}\rangle)=b) \\
&\leq \frac{1}{2}(1 + \|C'_0\,|\mathbf{0}\rangle - C'_1\,|\mathbf{0}\rangle\,\|_1) \\
&\leq \frac{1}{2}(1+\delta+\gamma)\ .
\end{aligned}
\tag{8}
$$

Conclusion: We need one more algorithm that will leverage the capacity of $\mathcal{B}$ to decide Π. Let $k \in \mathbb{N}$, and $\mathcal{C}$ be an algorithm that on instance $z \in \{0,1\}^n$, runs $\mathcal{B}(z)$ for k times independently. Let $b_1, \ldots, b_k$ be k corresponding independent outputs of $\mathcal{B}(z)$. Then $\mathcal{C}$ returns as follows:

$$
\begin{cases}
0 & \text{if } \left|\frac{1}{k}\sum_i b_i - \frac{1}{2}\right| \geq \tau, \\
1 & \text{otherwise,}
\end{cases}
$$

where $\tau(n)$ is chosen such that

$$
\tau := \frac{\nu}{4} - \frac{3(\delta+\gamma)}{4}\ .
\tag{9}
$$

Then, we have

$$\begin{aligned}
&\Pr(\mathcal{C}(z)=0|z\in \Pi_Y)\\
&=\Pr\left(\left|\frac{1}{k}\sum_i b_i-\frac{1}{2}\right|\geq\tau \;\middle|\; b_1,\dots,b_k\leftarrow\mathcal{B}(z), z\in\Pi_Y\right)\\
&\geq\Pr\left(\frac{1}{k}\sum_i b_i\geq\frac{1}{2}+\tau \;\middle|\; b_1,\dots,b_k\leftarrow\mathcal{B}(z), z\in\Pi_Y\right)\\
&\geq\Pr\left(\frac{1}{k}\left(\sum_i b_i-\mathbb{E}(\mathcal{B}_i(z))\right)\geq-\tau \;\middle|\; b_1,\dots,b_k\leftarrow\mathcal{B}(z), z\in\Pi_Y\right)\\
&\geq 1-\exp\bigl(-2k\tau^2\bigr)\ ,
\end{aligned}$$

where we used $\mathbb{E}(\mathcal{B}_i(z))-\tau\geq\frac{1}{2}+\tau$ for $z\in\Pi_Y$ by Equation (7) in the second inequality, and Hoeffding's lemma in the last inequality. On the other hand, we have

$$\begin{aligned}
&\Pr(\mathcal{C}(z)=1|z\in \Pi_N)\\
&=\Pr\left(\left|\frac{1}{k}\sum_i b_i-\frac{1}{2}\right|<\tau \;\middle|\; b_1,\dots,b_k\leftarrow\mathcal{B}(z), z\in\Pi_N\right)\\
&=\Pr\left(\left|\frac{1}{k}\sum_i b_i-\frac{1}{k}\sum_i\mathbb{E}(\mathcal{B}_i(z))\right|<\tau \;\middle|\; b_1,\dots,b_k\leftarrow\mathcal{B}(z), z\in\Pi_N\right)\\
&\geq 1-\exp\bigl(-2k\tau^2\bigr)\ ,
\end{aligned}$$

where we once again used Hoeffding's lemma and Equation (8). For $k:=1/\tau^2$, any sufficiently large $n\in\mathbb{N}$, and any $z\in(\Pi_Y\cup\Pi_N)\cap\{0,1\}^n$, it holds that

$$\Pr(\mathcal{C}(z)=\chi_\Pi(z))\geq 1-\exp\left(-2k\tau^2\right)\geq\frac{2}{3}\ ,$$

This breaks the worst-case hardness of Π.

Since $\theta_{\mathsf{efi}}:=(1-2\mu)-3(\delta+\gamma)$, we can set $\nu:=(1-2\mu)-\theta_{\mathsf{efi}}/2$, and the number of repetitions in the last step becomes

$$1/\tau^2=\frac{4^2}{(\nu-3(\delta+\gamma))^2}=\frac{4^3}{\theta_{\mathsf{efi}}^2}\,.$$

Runtime: We compute the runtime of EFI as follows. It first samples $2m$ instances from $\mathcal{U}_{K_a}$ (or $\mathcal{U}_{T_a}$), applies the permutations π twice to each half of the samplings, and computes R on each half. One single sampling from $\mathcal{U}_{K_a}$ (or $\mathcal{U}_{T_a}$) takes time $O(d)$, where $d\leq(m+1)/\gamma$ is the size of $\mathcal{U}_{K_a}$. The permutations can be applied in time $O(m)$ using the Fisher-Yates's algorithm. Therefore, the total runtime of EFI is $O(T+m^2/\gamma)$.

Note that $\mathcal{C}$ runs $\mathcal{B}$ for $O(1/{\theta_{\mathsf{efi}}}^2)$ times. Each execution of $\mathcal{B}$ evaluates $\widehat{C}_b$, queries $\mathcal{A}$, and performs an equality check. All of this takes $O((T+m^2/\gamma)/{\theta_{\mathsf{efi}}}^2)$ with $O(1/{\theta_{\mathsf{efi}}}^2)$ queries to $\mathcal{A}$.

□

Next, we consider larger values of λ, for instance when $\lambda\geq 2$. The following concerns only *classical* one-way functions.

Theorem 6.2. *Let Π be $(T,\mu,f^m,\lambda,\gamma)$-sparsely lossy with a classical reduction. Assume that $\theta_{\mathsf{owf}}:=(1-10\mu)-(\delta(\lambda)+\gamma)>0$, with $\delta(\lambda)$ as in Definition 19. Then there exists an algorithm F that runs in time$O(T+m^2\gamma^{-1})$ and an oracle algorithm $\mathcal{C}$, such that for any algorithm $\mathcal{A}$ one and only one of the following holds:*

I. $\mathcal{C}^{\mathcal{A}}$ solves $\Pi\cap\{0,1\}^n$ in time $O((T+m^2\gamma^{-1})\theta_{\mathsf{owf}}^{-2})$ with $O(\theta_{\mathsf{owf}}^{-2})$ queries to $\mathcal{A}$,
II. F is a $(1-\theta_{\mathsf{owf}}/2)$-OWF for $\mathcal{A}$.

Proof. Consider the circuit $\widehat{C}_0$ in Line 3 of Algorithm 2. This circuit is independent of the input of Algorithm 2 and is randomized. Part of its randomness is used to sample $\widetilde{x}$ and the other part is fed to R. Let κ be the size of the total randomness. For $r \in \{0,1\}^\kappa$, we let $\widehat{C}_0(r)$ be the outcome of the circuit when it is given r as the randomness. We show that F, defined by $\widehat{C}_0(\cdot) : \{0,1\}^\kappa \to \{0,1\}^*$, is a $(\theta_{\mathsf{owf}}/2)$-weak one-way function. This suffices for the proof since weak one-way functions imply one-way functions.

The proof works by a reduction to the worst-case hardness of Π. Assume that we are given a to-be-decided instance y of Π. Apply Algorithm 2 up to Line 4 to obtain $(\widehat{C}_0, \widehat{C}_1)$. Assume that there exists an adversary $\mathcal{A}$ that inverts $\widehat{C}_0(\cdot)$ with probability more than $1-\theta_{\mathsf{owf}}/2$. Consider the following oracle algorithm $\mathcal{B}^{\mathcal{A}}$:

- $\mathcal{B}^{\mathcal{A}}(\widehat{C}_0, \widehat{C}_1, y)$: samples a uniform $r \in \{0,1\}^\kappa$ and a uniform $b \in \{0,1\}$, and computes $z := \widehat{C}_b(r)$. Runs the adversary $r' \leftarrow \mathcal{A}(z)$, and computes $z' = \widehat{C}_0(r')$. If $z = z'$ it outputs 1, otherwise it outputs 0.

We show that $\mathcal{B}$ can distinguish between the YES and NO instances of Π by analysing the probability of outputting 1. More precisely, we study the following random variable:

$$X(\widehat{C}_0, \widehat{C}_1, y) := \left|\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widehat{C}_0, \widehat{C}_1, y) = 1 | b = 0\Big) - \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widehat{C}_0, \widehat{C}_1, y) = 1 | b = 1\Big)\right| .$$

Case $y \in \Pi_Y$: We show the following bound for every $y \in \Pi_Y$:

$$X(\widehat{C}_0, \widehat{C}_1, y) > 1 - \theta_{\mathsf{owf}}/2 - 10\mu \ .$$

Instead of proving the inequality directly for the circuits $(\widehat{C}_0, \widehat{C}_1)$, we will show it for two similar circuits $(\widetilde{C}_0, \widetilde{C}_1)$ with disjoint image sets. Let $\widehat{D}_0$ and $\widehat{D}_1$ be respectively the outcome distributions of $\widehat{C}_0$ and $\widehat{C}_1$ when given uniform input, and A be the following set

$$A := \{a \mid \Pr_{\widehat{D}_0}(a) \geq \Pr_{\widehat{D}_1}(a)\} .$$

Let $\widetilde{C}_0$ be the restriction of $\widehat{C}_0$ to A and $\widetilde{C}_1$ the restriction of $\widehat{C}_1$ to A^c. We will show that

$$X(\widetilde{C}_0, \widetilde{C}_1, y) \leq X(\widehat{C}_0, \widehat{C}_1, y) + 8\mu.$$

Indeed in Theorem 5.1, we showed that for every $y \in \Pi_Y$, the statistical distance between the outcome distributions of $\widehat{C}_0$ and $\widehat{C}_1$ when given uniform input is at least $1 - 2\mu$. Moreover, we have

$$\begin{aligned}
\|\widehat{D}_0 - \widehat{D}_1\|_1 &= \frac{1}{2}\sum_a |\Pr_{\widehat{D}_0}(a) - \Pr_{\widehat{D}_1}(a)| \\
&= \frac{1}{2}\sum_{a\in A} \Pr_{\widehat{D}_0}(a) - \Pr_{\widehat{D}_1}(a) + \frac{1}{2}\sum_{a\in A^c} \Pr_{\widehat{D}_1}(a) - \Pr_{\widehat{D}_0}(a) \\
&= \frac{1}{2}(\Pr_{\widehat{D}_0}(A) - \Pr_{\widehat{D}_0}(A^c) + \Pr_{\widehat{D}_1}(A^c) - \Pr_{\widehat{D}_1}(A)) \\
&= \frac{1}{2}(\Pr_{\widehat{D}_0}(A) - (1 - \Pr_{\widehat{D}_0}(A)) + \Pr_{\widehat{D}_1}(A^c) - (1 - \Pr_{\widehat{D}_1}(A^c))) \\
&= \Pr_{\widehat{D}_0}(A) + \Pr_{\widehat{D}_1}(A^c) - 1 .
\end{aligned}$$

It follows that $\Pr_{\widehat{D}_0}(A) + \Pr_{\widehat{D}_1}(A^c) \geq 2 - 2\mu$. Therefore, we have

$$(\Pr_{\widehat{D}_0}(A) \geq 1 - \mu) \wedge (\Pr_{\widehat{D}_1}(A^c) \geq 1 - 2\mu) ,$$

or

$$(\Pr_{\widehat{D}_0}(A) \geq 1 - 2\mu) \wedge (\Pr_{\widehat{D}_1}(A^c) \geq 1 - \mu) .$$

Then for either of cases above, we have

$$\|\widehat{D}_0 - \widetilde{D}_0\|_1 \leq 2\mu\,, \quad \text{and} \quad \|\widehat{D}_1 - \widetilde{D}_1\|_1 \leq 2\mu\,, \tag{10}$$

where $\widetilde{D}_0$ and $\widetilde{D}_1$ are respectively the outcome distributions of $\widetilde{C}_0$ and $\widetilde{C}_1$. Pretend that not only does $\mathcal{A}$ invert F, but also tries to distinguish between $\widehat{C}_b$ and $\widetilde{C}_b$ for $b \in \{0,1\}$. Consider the following sequence of games that modifies $\mathcal{B}^{\mathcal{A}}$:

Game $\mathcal{G}_1$: In this game $\mathcal{B}$ behaves originally as above.

Game $\mathcal{G}_2$: In this game $\mathcal{B}$ replaces $\widehat{C}_0$ with $\widetilde{C}_0$. Note that $\mathcal{A}$ can distinguish this modification with probability at most 2μ according to Equation (10). It follows that

$$\begin{aligned}
&X(\widetilde{C}_0, \widehat{C}_1, y) \\
&= \left|\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widehat{C}_1, y) = 1 | b = 0\Big) - \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widehat{C}_1, y) = 1 | b = 1\Big)\right| \\
&\leq \left|\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widehat{C}_1, y) = 1 | b = 0\Big) - \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widehat{C}_0, \widehat{C}_1, y) = 1 | b = 0\Big)\right| \\
&\quad + \left|\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widehat{C}_0, \widehat{C}_1, y) = 1 | b = 0\Big) - \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widehat{C}_0, \widehat{C}_1, y) = 1 | b = 1\Big)\right| \\
&\quad\quad + \left|\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widehat{C}_0, \widehat{C}_1, y) = 1 | b = 1\Big) - \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widehat{C}_1, y) = 1 | b = 1\Big)\right| \\
&= X(\widehat{C}_0, \widehat{C}_1, y) + 4\mu\,.
\end{aligned}$$

Game $\mathcal{G}_3$: In this game, $\mathcal{B}$ replaces $\widehat{C}_1$ with $\widetilde{C}_1$. Note that $\mathcal{A}$ can identify this modification with probability at most 2μ. We obtain

$$\begin{aligned}
&X(\widetilde{C}_0, \widetilde{C}_1, y) \\
&= \left|\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widetilde{C}_1, y) = 1 | b = 0\Big) - \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widetilde{C}_1, y) = 1 | b = 1\Big)\right| \\
&\leq \left|\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widetilde{C}_1, y) = 1 | b = 0\Big) - \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widehat{C}_1, y) = 1 | b = 0\Big)\right| \\
&\quad + \left|\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widehat{C}_1, y) = 1 | b = 0\Big) - \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widehat{C}_1, y) = 1 | b = 1\Big)\right| \\
&\quad\quad + \left|\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widehat{C}_1, y) = 1 | b = 1\Big) - \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widetilde{C}_1, y) = 1 | b = 1\Big)\right| \\
&= X(\widetilde{C}_0, \widehat{C}_1, y) + 4\mu \\
&\leq X(\widehat{C}_0, \widehat{C}_1, y) + 8\mu\,.
\end{aligned}$$

To prove the inequality for the YES instances, it suffices to show that $X(\widetilde{C}_0, \widetilde{C}_1, y) > 1 - \theta_{\mathsf{owf}}/2 - 2\mu$. Recall that

$$\begin{aligned}
&X(\widetilde{C}_0, \widetilde{C}_1, y) \\
&\qquad := \left|\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widetilde{C}_1, y) = 1 | b = 0\Big) - \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widetilde{C}_1, y) = 1 | b = 1\Big)\right|\,.
\end{aligned}$$

First, when $b = 0$ and hence $z = \widetilde{C}_0(r)$ with $\|\widehat{D}_0 - \widetilde{D}_0\|_1 \leq 2\mu$, the adversary $\mathcal{A}$ succeeds with probability at least $1 - \theta_{\mathsf{owf}}/2 - 2\mu$ to invert $\widetilde{C}_0$, which is equal to the probability that $\mathcal{B}^{\mathcal{A}}$ outputs 1. Second, when $b = 1$ and hence $z = \widetilde{C}_1(r)$, since the supports of $\widetilde{C}_0$ and $\widetilde{C}_1$ are distinct, $\mathcal{A}$ never succeeds to find an r' such that $\widetilde{C}_0(r') = \widetilde{C}_1(r)$, i.e., the probability of $\mathcal{B}$ outputting one is zero. This completes the first part.

Case $y \in \mathit{\Pi}_N$: In Theorem 5.1, we also proved that for every $y \in \mathit{\Pi}_N$, the outcomes of the two circuits $(\widehat{C}_0, \widehat{C}_1)$ is at most $\delta + \gamma$. Therefore, the adversary $\mathcal{A}$ cannot distinguish them with a probability larger than $\delta + \gamma$. The information processing inequality then implies that

$$X(\widehat{C}_0, \widehat{C}_1, y) \leq \delta + \gamma\,.$$

Conclusion: The quantity $X(\widehat{C}_0, \widehat{C}_1, y)$ diverges for YES and NO instances of y. For our choice of parameters, we know that

$$1 - \theta_{\mathsf{owf}}/2 - 10\mu - (\delta + \gamma) = \theta_{\mathsf{owf}}/2\,.$$

We denote by $\mathcal{C}^{\mathcal{A}}$ an algorithm that runs $\mathcal{B}$ for $O(1/\theta_{\mathsf{owf}}^2)$ many times, and approximates the quantity above within error less than $\theta_{\mathsf{owf}}/4$. If this value is more than $\delta+\gamma+\theta_{\mathsf{owf}}/4$, then y must be a YES instance, otherwise it is a NO instance. Therefore, we finally obtain a algorithm that solves Π.

Runtime: The runtime of F can be computed as follows. It samples m instances from $\mathcal{U}_{K_a}$ (or $\mathcal{U}_{T_a}$), applies a permutation π, and computes R on top of it. Each time, sampling from $\mathcal{U}_{K_a}$ (or $\mathcal{U}_{T_a}$) takes time $O(d)$. The permutation takes $O(m)$ staps by using the Fisher-Yates's shuffle. Then the total runtime of F is $O(T+m^2/\gamma)$.

For the runtime of $\mathcal{C}^{\mathcal{A}}$, note that $\mathcal{C}$ runs $\mathcal{B}$ for $O(1/{\theta_{\mathsf{owf}}}^2)$ times. Each execution of $\mathcal{B}$ evaluates $\widehat{C}_b$, queries $\mathcal{A}$, and performs an equality check. All of this takes $O((T+m^2/\gamma)/{\theta_{\mathsf{owf}}}^2)$ with $O(1/{\theta_{\mathsf{owf}}}^2)$ queries to $\mathcal{A}$.

□

7 One-Way State Generators from Sparsely Lossy Problems

In the next theorem, we discuss the adaptation to the quantum settings, when λ is relatively large.

Theorem 7.1. *Let Π be $(T,\mu,f^m,\lambda,\gamma)$-sparsely lossy where the outcome of the reduction is a pure state. Also assume that $\theta_{\mathsf{ows}} := 1-(\delta(\lambda)+\gamma+4\sqrt{2\mu})>0$ and $\tau_{\mathsf{ows}} := 1-2\mu-(\delta(\lambda)+\gamma)>0$, with $\delta(\lambda)$ as in Definition 19. Then there exists an algorithm $\mathsf{G}=(\mathsf{StateGen},\mathsf{Ver})$ such that $\mathsf{StateGen}$ runs in time $O(T+m^2\gamma^{-1})$, and an oracle algorithm $\mathcal{C}$, such that for every algorithm $\mathcal{A}$ one and only one of the following statements holds:*

I. $\mathcal{C}^{\mathcal{A}}$ solves $\Pi\cap\{0,1\}^n$ in time $O((cT+cm^2\gamma^{-1}+\tau_{\mathsf{ows}}^{-2})\theta_{\mathsf{ows}}^{-2})$ with $O(\theta_{\mathsf{ows}}^{-2})$ classical queries to $\mathcal{A}$,
II. G is a c-copy $(1-\theta_{\mathsf{ows}}/4)$-OWSG for $\mathcal{A}$.

Proof. Sample $z\sim\mathcal{U}_{K_a}$ and apply Algorithm 2 up to Line 4 on input z to obtain the two circuits (C_0^*,C_1^*). Note that the two circuits are mixed; a classical randomness is used to sample $\widetilde{x}$ but the algorihm R is a pure quantum circuit. Let κ be the size of the randomness of these circuits. For any $r\in\{0,1\}^\kappa$ and $b\in\{0,1\}$, let $C_b^*\left|r,\mathbf{0}\right\rangle$ be the pure state obtained by sampling $\widetilde{x}$ using r and applying R to $\pi(\widetilde{x})$ and a possibly ancilla $\left|\mathbf{0}\right\rangle$ with an appropriate size. We show that G, defined as follows:

- $\mathsf{StateGen}(r,b)$: output $C_b^*\left|r,\mathbf{0}\right\rangle$.
- $\mathsf{Ver}((r,b),\rho)$: If $\|C_b^*\left|r,\mathbf{0}\right\rangle-\rho\|_1\leq\delta+\gamma$ output 1, otherwise output 0.

is a $(1-\theta_{\mathsf{ows}}/4)$-weak one-way state generator.

Assume that there exists an adversary $\mathcal{A}$ that breaks the scheme above with probability more than $1-\theta_{\mathsf{ows}}/4$. We use $\mathcal{A}$ to construct an algorithm for Π. Consider the following oracle algorithm $\mathcal{B}^{\mathcal{A}}$:

- $\mathcal{B}^{\mathcal{A}}(\widehat{C}_0,\widehat{C}_1,y)$: computes $(\widehat{C}_0,\widehat{C}_1(y))$ as in Algorithm 2 up to Line 4 on input y. Samples a uniform $r\in\{0,1\}^\kappa$ and a uniform $b\in\{0,1\}$, and computes $\rho:=\widehat{C}_b\left|r,\mathbf{0}\right\rangle$. Runs the adversary $(r',b')\leftarrow\mathcal{A}(\rho^{\otimes n})$, and computes $\rho'=\widehat{C}_{b'}\left|r',\mathbf{0}\right\rangle$. If $\|\rho-\rho'\|_1\leq\delta+\gamma$ it outputs 1, otherwise it outputs 0.

We compute the advantage of $\mathcal{B}$ in distinguishing between YES and NO instances of Π by analyzing the probability $\Pr\left(\mathcal{B}^{\mathcal{A}}(\widehat{C}_0,\widehat{C}_1,y)=1\right)$.

Case $y\in\Pi_Y$: We show that for every $y\in\Pi_Y$, we have:

$$\Pr\left(\mathcal{B}^{\mathcal{A}}(\widehat{C}_0,\widehat{C}_1,y)=1\right)\leq\frac{1}{2}+2\sqrt{2\mu}\ .$$

Instead of proving the inequality directly for the circuits $(\widehat{C}_0,\widehat{C}_1)$, we will show it for two similar circuits $(\widetilde{C}_0,\widetilde{C}_1)$ with disjoint images. Let $\widehat{\rho}_0$ and $\widehat{\rho}_1$ be respectively the mixed states $\widehat{C}_0\left|r,\mathbf{0}\right\rangle$ and $\widehat{C}_1\left|r,\mathbf{0}\right\rangle$ when r follows the uniform distribution. For any POVM $\mathcal{M}=\{M_i\}_i$, let us define by $A_{\mathcal{M}}$ the following set:

$$A_{\mathcal{M}}:=\{i\mid\mathrm{Tr}(M_i\widehat{\rho}_0)\geq\mathrm{Tr}(M_i\widehat{\rho}_1)\}\ .$$

In Theorem 5.1, we showed that for every $y \in \Pi_Y$, the statistical distance between $\widehat{\rho}_0$ and $\widehat{\rho}_1$ is at least $1 - 2\mu$. Moreover, we can rewrite the trace distance in terms of the POVMs as

$$\begin{aligned}
\|\widehat{\rho}_0 - \widehat{\rho}_1\|_1 &= \max_{\{M_i\}_i} \frac{1}{2} \sum_i |\mathrm{Tr}(M_i\widehat{\rho}_0) - \mathrm{Tr}(M_i\widehat{\rho}_1)| \\
&= \max_{\{M_i\}_i} \frac{1}{2} \left[\sum_{i \in A_{\mathcal{M}}} (\mathrm{Tr}(M_i\widehat{\rho}_0) - \mathrm{Tr}(M_i\widehat{\rho}_1)) \right. \\
&\qquad \left. + \sum_{i \in A^c_{\mathcal{M}}} (\mathrm{Tr}(M_i\widehat{\rho}_1) - \mathrm{Tr}(M_i\widehat{\rho}_0)) \right] \\
&= \max_{\{M_i\}_i} \left\{ \sum_{i \in A_{\mathcal{M}}} \mathrm{Tr}(M_i\widehat{\rho}_0) + \sum_{i \in A^c_{\mathcal{M}}} \mathrm{Tr}(M_i\widehat{\rho}_1) - 1 \right\}.
\end{aligned}$$

It follows that there exists a particular POVM $\mathcal{M}$, such that if we define the projections of $\widehat{C}_0$ and $\widehat{C}_1$ onto $A_{\mathcal{M}}$ and $A^c_{\mathcal{M}}$ by $\widetilde{C}_0$ and $\widetilde{C}_1$ respectively, i.e.,

$$\widetilde{C}_0 = \sum_{i \in A_{\mathcal{M}}} M_i \widehat{C}_0, \quad \text{and} \quad \widetilde{C}_1 = \sum_{i \in A^c_{\mathcal{M}}} M_i \widehat{C}_1 \,,$$

we have $\mathrm{Tr}(\widetilde{\rho}_0) + \mathrm{Tr}(\widetilde{\rho}_1) \geq 2 - 2\mu$, where $\widetilde{\rho}_b$ is the mixed state $\widetilde{C}_b\,|r, \mathbf{0}\rangle$ and r is uniform. Therefore

$$(\mathrm{Tr}(\widetilde{\rho}_0) \geq 1 - \mu) \wedge (\mathrm{Tr}(\widetilde{\rho}_1) \geq 1 - 2\mu)$$

or

$$(\mathrm{Tr}(\widetilde{\rho}_0) \geq 1 - 2\mu) \wedge (\mathrm{Tr}(\widetilde{\rho}_1) \geq 1 - \mu)\,.$$

By the Gentle Measurement Lemma 2.4, for either of cases above, we have

$$\|\widehat{\rho}_0 - \widetilde{\rho}_0\|_1 \leq \sqrt{2\mu}\,, \quad \text{and} \quad \|\widehat{\rho}_1 - \widetilde{\rho}_1\|_1 \leq \sqrt{2\mu}\,. \tag{11}$$

Pretend that $\mathcal{A}$ also tried to distinguish between for $\widehat{C}_b$ and $\widetilde{C}_b$ for $b \in \{0, 1\}$, and consider the following sequence of games that modifies $\mathcal{B}^{\mathcal{A}}$.

Game $\mathcal{G}_1$: In this game $\mathcal{B}$ behaves originally as above.

Game $\mathcal{G}_2$: In this game $\mathcal{B}$ replaces $\widehat{C}_0$ with $\widetilde{C}_0$. Note that $\mathcal{A}$ can distinguish this modification with probability at most $\sqrt{2\mu}$ according to Equation (11). It follows that

$$\begin{aligned}
&\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widehat{C}_0, \widehat{C}_1, y) = 1\Big) \\
&\leq \Big|\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widehat{C}_0, \widehat{C}_1, y) = 1\Big) - \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widehat{C}_1, y) = 1\Big)\Big| \\
&\qquad + \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widehat{C}_1, y) = 1\Big) \\
&\leq \sqrt{2\mu} + \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widehat{C}_1, y) = 1\Big)\,.
\end{aligned}$$

Game $\mathcal{G}_3$: In this game, $\mathcal{B}$ replaces $\widehat{C}_1$ with $\widetilde{C}_1$. Note that $\mathcal{A}$ can identify this modification with probability at most $\sqrt{2\mu}$. We obtain

$$\begin{aligned}
&\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widehat{C}_0, \widehat{C}_1, y) = 1\Big) \\
&\leq \Big|\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widehat{C}_0, \widehat{C}_1, y) = 1\Big) - \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widehat{C}_1, y) = 1\Big)\Big| \\
&\qquad + \Big|\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widehat{C}_1, y) = 1\Big) - \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widetilde{C}_1, y) = 1\Big)\Big| \\
&\qquad\qquad + \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widetilde{C}_1, y) = 1\Big) \\
&\leq 2\sqrt{2\mu} + \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widetilde{C}_1, y) = 1\Big)
\end{aligned}$$

Now, note that the projection onto the supports of $\widetilde{C}_0$ and $\widetilde{C}_1$ are orthogonal to each other. Therefore, the adversary never succeeds when the bit b (chosen by $\mathcal{B}$) is equal to 1; there exists no r' such that $\|\widetilde{C}_0 |r, |\mathbf{0}\rangle\rangle - \widetilde{C}_1 |r', \mathbf{0}\rangle \|_1 \leq \delta + \gamma$. So

$$\begin{aligned}
&\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widetilde{C}_1, y) = 1\Big) \\
&= \frac{1}{2}\left(\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widetilde{C}_1, y) = 1 | b = 0\Big) + \Pr\Big(\mathcal{B}^{\mathcal{A}}(\widetilde{C}_0, \widetilde{C}_1, y) = 1 | b = 1\Big)\right) \\
&\leq \frac{1}{2}\,.
\end{aligned}$$

Case $y \in \Pi_N$: By Lemma 3.1, the trace distance of the outcomes of $\widehat{C}_1$ and C_1^* is at most $\delta + \gamma$. Moreover, $\widehat{C}_0$ is exactly the same as C_0^*. Therefore, if the bit b, chosen by $\mathcal{B}$ is equal to 0, then $\mathcal{A}$ succeeds with probability at least $1 - \theta_{\mathsf{ows}}/4$, and if $b = 1$, it succeeds with probability $1 - \theta_{\mathsf{ows}}/4 - (\delta + \gamma)$. In total, we obtain

$$\begin{aligned}
&\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widehat{C}_0, \widehat{C}_1, y) = 1\Big) \\
&\geq \frac{1}{2}(1 - \frac{\theta_{\mathsf{ows}}}{4}) + \frac{1}{2}(1 - \frac{\theta_{\mathsf{ows}}}{4} - (\delta + \gamma)) \\
&= 1 - \frac{\theta_{\mathsf{ows}}}{4} - \frac{(\delta + \gamma)}{2}\,.
\end{aligned}$$

Conclusion: We showed that the quantity of $\Pr\Big(\mathcal{B}^{\mathcal{A}}(\widehat{C}_0, \widehat{C}_1, y) = 1\Big)$ diverges for YES and NO instances of y. For our choice of parameters, we have

$$\begin{aligned}
1 - \frac{\theta_{\mathsf{ows}}}{4} - \frac{(\delta + \gamma)}{2} - \left(\frac{1}{2} + 2\sqrt{2\mu}\right) &= \frac{1 - (\delta + \gamma + 4\sqrt{2\mu})}{2} - \frac{\theta_{\mathsf{ows}}}{4} \\
&= \frac{\theta_{\mathsf{ows}}}{4}\,.
\end{aligned}$$

Let $\mathcal{C}$ be an algorithm that runs $\mathcal{B}$ for $O(1/\theta_{\mathsf{ows}}^2)$ many times, and approximates the quantity above within error less than $\theta_{\mathsf{ows}}/4$. If this value is more than $1 - \theta_{\mathsf{ows}}/4 - (\delta + \gamma)/2$, then y must be a NO instance, otherwise it is a YES instance. Therefore, we finally obtain a algorithm that solves Π. Note that $\mathcal{B}$ verifies whether $\|\widehat{C}_b |r, \mathbf{0}\rangle - \widehat{C}_{b'} |r', \mathbf{0}\rangle \|_1$ is smaller than $\delta + \gamma$. Since the reduction R is pure and r, r' are fixed, these states are pure, therefore $\mathcal{B}$ can perform a SWAP test for $O(1/\tau_{\mathsf{ows}}^2)$ number of times on them to approximate their ℓ_1 distance.

□

8 Sparse Lossiness and Instance Randomization

In Section 4 we introduced sparsely lossy problems, promise problems that admit reductions that *lose* some information about the input, and in Section 6 and 7 we constructed cryptography primitives from these. In this section we show that sparsely lossy problems are not uncommon by proving that both worst-case to average-case reductions and randomized encodings imply sparse lossiness, given a classical reduction.

8.1 Worst-Case to Average-Case Reductions

In this section we analyse the sparse lossiness of worst-case to average-case reductions. Since we discuss sparse lossiness of such reductions, as motivated in Section 4, we focus on worst-case to average-case *f-distinguisher* reductions (Definition 13). In Definition 20, we put forward the definition of *worst-case to distribution f-distinguisher reduction* which can be viewed as a generalization of worst-case to average-case reductions in the sense that (i) the reduction is oblivious to the target average-case problem (inherited from being f-distinguisher), and (ii) the reduction maps inputs to a distribution that is *not* necessarily efficiently samplable. The latter does not impose any issues in our setting, since we are only discussing sparse lossiness of the reductions, and not the hardness of the problems. We then prove, in Lemma 8.1, that such reductions are lossy and specify the sparse lossiness parameters.

Definition 20 (Worst-Case to Distribution f-Distinguisher Reduction). *Let Π be a promise problem, $n \in \mathbb{N}$, and $d \in [0,1]$. We say that a reduction R is a (T, μ, f^m, d)-worst-case to distribution, denoted WC-DIST, reduction for Π if*

1. *R is a (μ, f^m)-distinguisher reduction for Π (Definition 13), and*
2. *for all $x \in \Pi \cap \{0,1\}^n$, $R(x)$ runs in time $T(n)$, and*
3. *there exists a distribution $D = \{D_n\}_{n\in\mathbb{N}}$ over $\{0,1\}^*$, such that*

$$\forall (x_1, \cdots, x_m) \in (\Pi \cap \{0,1\}^n)^m : \ \Delta(R(x_1, \cdots, x_m), D) \leq d \ .$$

The upper bound d is called the distance of the reduction.

If there exist two distributions D_Y and D_N over $\{0,1\}^*$ such that for inputs $x \in \Pi_Y$ the distribution D_Y approximates $R(x)$ up to error d, and for inputs $x \in \Pi_N$ the distribution D_N approximates $R(x)$ up to error d, we say that the reduction R is a (T, μ, f^m, d)-worst-case to distribution splitting-reduction for Π.

Lemma 8.1 (Sparse Lossiness of WC-DIST f-Distinguisher Classical Reductions). *Let $\Pi = \Pi_Y \cup \Pi_N$ for two disjoint sets $\Pi_Y, \Pi_N \subset \{0,1\}^*$. Assume there exists a (T, μ, f^m, d)-WC-DIST classical splitting-reduction R for Π (see Definition 20), such that f is a non-constant permutation-invariant function. Then for any $\gamma > 0$, Π is $(T, \mu, f^m, \lambda, \gamma)$-sparsely lossy, where*

$$\lambda = \max\left\{\frac{1}{m}, 9 + \frac{4}{m} + \log\left(\frac{mn}{\gamma^3}\right) + \frac{2\log d}{m}\right\} .$$

Proof. The proof consists of showing that the reduction R satisfies Definition 18. Let $\gamma > 0$ and $X = (X_1, \cdots, X_m)$ such that X_i's are pairwise independent B-uniform distributions over n-bit strings where $B = \lceil 4(m+1)/\gamma \rceil \cdot \lceil 8n \ln 2/\gamma^2 \rceil$ as in Definition 18. We show that

$$I(X; R(X)) \leq \max\left\{1, 4 + 9m + m\log\left(\frac{mn}{\gamma^3}\right) + 2\log d\right\} \ .$$

Letting $p_X(y) := \Pr(X = y)$, we first rewrite the mutual information in terms of Kullback-Leibler divergence.

$$I(X; R(X)) = \sum_{y \in \mathrm{Supp}(R)} p_{R(X)}(y) \cdot D_{KL}\left(p_{X|R(X)=y} \parallel p_X\right) \ . \tag{12}$$

From a reverse Pinsker inequality due to [73], the KL divergence of two distributions decreases as their trace distance does, in particular

$$D_{KL}\left(p_{X|R(X)=y} \parallel p_X\right) \leq \log\left(1 + \frac{2 \cdot \Delta(X_{|R(X)=y}, X)^2}{\alpha_X}\right)$$

where $\alpha_X = \min_x p_X(x) > 0$. If $\Delta(X_{|R(X)=y}, X) = 0$, then $I(X; R(X)) = 0$.[1] Otherwise, since for any value $a \in (0,1]$, we have that $\log(1+a) \leq \max\{1, 1 + \log(a)\}$, we can write

$$D_{KL}\left(p_{X|R(X)=y} \parallel p_X\right) \leq \max\{1, 2 + 2\log\big(\Delta(X_{|R(X)=y}, X)\big) - \log(\alpha_X)\} \ ,$$

Substituting above in Equation (12), we obtain:

$$I(X; R(X)) \leq \max\{1, \ 2 - \log(\alpha_X) + 2 \sum_{y \in \mathrm{Supp}(R)} p_{R(X)}(y) \cdot \log\big(\Delta(X_{|R(X)=y}, X)\big)\} \ . \tag{13}$$

We split the bound on the right-hand side of the Inequality (13) into two terms.

Bounding $\mathsf{term}_1 = -\log(\alpha_X)$: Since X_i's are pairwise independent B-uniform distributions, we have $\alpha_X = 1/B \leq (\gamma^3/2^9 mn)^m$. Therefore $-\log(\alpha_X) \leq m(9 + \log\big(mn/\gamma^3\big))$.

[1] However, this is very unlikely!

Bounding $\mathsf{term}_2 = \sum_{y \in \mathrm{Supp}(R)} p_{R(X)}(y) \cdot \log\big(\Delta(X_{|R(X)=y}, X)\big)$**:** Firstly, for any $y \in \mathrm{Supp}(R)$, we have

$$\begin{aligned}
&\Delta(X_{|R(X)=y}, X) && (14)\\
&= \frac{1}{2} \sum_x \left|\Pr\left(X = x | R(X) = y\right) - \Pr(X = x)\right| \\
&= \frac{1}{2} \sum_x \left| \frac{\Pr\left(X = x \wedge R(X) = y\right)}{\Pr(R(X) = y)} - \Pr(X = x) \right| \\
&= \frac{1}{2} \sum_x \frac{1}{\Pr(R(X) = y)} \left|\Pr\left(X = x \wedge R(X) = y\right)\right. && (15)\\
&\qquad\qquad \left. - \Pr(X = x) \cdot \Pr(R(X) = y)\right| \\
&= \frac{1}{\Pr(R(X) = y)} \cdot \Delta((X, R(X) = y), X \cdot (R(X) = y)) \ . && (16)
\end{aligned}$$

Rewriting $\mathsf{term}_2 = \mathbb{E}_{R(X)}\left[\log\big(\Delta(X_{|R(X)=y}, X)\big)\right]$, we now have to bound

$$\begin{aligned}
\mathsf{term}_2 &= \mathbb{E}_{R(X)}\left[\log \Delta(X_{|R(X)=y}, X)\right] \\
&\leq \log \mathbb{E}_{R(X)}\left[\Delta(X_{|R(X)=y}, X)\right] \quad \text{(by Jensen's inequality)} \\
&= \log\left(\sum_{y \in \mathrm{Supp}(R)} \Pr(R(X) = y) \cdot \Delta(X_{|R(X)=y}, X)\right) \\
&= \log\left(\sum_{y \in \mathrm{Supp}(R)} \Delta((X, R(X) = y), X \cdot (R(X) = y))\right) \ . && (17)
\end{aligned}$$

where the last equality holds by Equation 16. Analysing the term inside the logarithm above, we have

$$\begin{aligned}
&\sum_{y \in \mathrm{Supp}(R)} \Delta((X, R(X) = y), X \cdot (R(X) = y)) \\
&= \frac{1}{2} \sum_{y \in \mathrm{Supp}(R)} \sum_{x \in X} \left|\Pr(R(X) = y | X = x) \cdot \Pr(X = x)\right. && (18)\\
&\qquad\qquad \left. - \Pr(R(X) = y) \cdot \Pr(X = x)\right| \\
&= \frac{1}{2} \sum_{y \in \mathrm{Supp}(R)} \sum_x \Pr(X = x) \cdot \left|\Pr(R(x) = y) - \Pr(R(X) = y)\right| \\
&= \sum_x \Pr(X = x) \cdot \Delta(R(x), R(X)) \\
&\leq \max_x \Delta(R(x), R(X)) \ .
\end{aligned}$$

We therefore have that $\mathsf{term}_2 \leq \max_x \log(\Delta(R(x), R(X)))$. Finally, note that since R is a (T, μ, f^m, d)-WC-DIST reduction, for any $x \in \Pi_Y \cap \{0,1\}^n$, it holds that $\Delta(R(x), D_{n,Y}) \leq d$. Therefore $\Delta(R(X), D_{n,Y}) \leq d$ for any distribution X over $\Pi_Y \cap \{0,1\}^n$. We conclude that for any $x \in \Pi_Y \cap \{0,1\}^n$, $\Delta(R(x), R(X)) \leq 2d$ for any distribution X over $\Pi_Y \cap \{0,1\}^n$, which yields $\mathsf{term}_2 \leq 1 + \log(d)$. Note that the same argument holds for $x \in \Pi_N \cap \{0,1\}^n$ and distributions $D_{n,N}$.

Combining upper bounds on term_1 and term_2, we finish by proving that

$$I(X; R(X)) \leq \max\left\{1, 4 + 9m + m \log\left(\frac{mn}{\gamma^3}\right) + 2 \log d\right\} \ ,$$

for splitting lossy distributions X.

□

WC-DIST Turing Reductions

All reductions in the rest of the work until Section 9.5 are classical. In this part, we give an adapted version of the worst-case to distribution reduction (Definition 20) to the case of non-adaptive randomized Turing reductions.

Definition 20 covers the notion of worst-case to average-case *Karp* reductions, that is the type of most cryptographic reductions. However, in order to discuss the sparse lossiness of WC-DIST Turing reductions, we have to slightly refine this definition; Recall from Section 4 that a non-adaptive randomized Turing reduction from Π to Σ, maps an input x to $(y_1, \ldots, y_k)$, where each y_i is an instance of Σ, as well as a Boolean circuit C. Since C depends on x, it can carry some information about the input and affect the sparse lossiness. On the other hand, the requirement of Definition 20 requires analysing the joint distribution of $((y_1, \ldots, y_k), C)$ that might be tedious. We therefore relax the above definition to this case and discuss the sparse lossiness of randomized Turing reductions in this relaxed setting.

Definition 21 (WC-DIST Non-Adaptive Randomized Turing f-Reductions). *Let Π be a promise problem. We say that R_{Turing} is a (T, μ, f^m, d, h)-worst-case to distribution (WC-DIST) non-adaptive randomized Turing reduction for Π, if*

1. *R_{Turing} is a non-adaptive (f^m, μ)-Turing reduction from Π to some promise or search problem Σ (per Definition 15), and*
2. *for all $x \in \Pi \cap \{0,1\}^n$, $R_{Turing}(x)$ runs in time $T(n)$, and*
3. *there exists a distribution $D = \{D_n\}_{n \in \mathbb{N}}$ over $\{0,1\}^*$, such that:*

$$\forall x \in \Pi \cap \{0,1\}^n : \ \Delta((y_1, \ldots, y_k), D_n) \leq d \ ,$$

   where $((y_1, \ldots, y_k), C) \leftarrow R_{Turing}(x)$, and
4. *for all distributions X over n-bit strings:*

$$I\left(X; C | (Y_1, \ldots, Y_k)\right) \leq h,$$

   where $((Y_1, \ldots, Y_k), C) \leftarrow R_{Turing}(X)$.

We now state the following lemma, on the sparse lossiness of worst-case to distribution Turing reductions.

Lemma 8.2 (Sparse Lossiness of WC-DIST Non-Adaptive Randomized Turing Reductions). *Let Π be a promise problem. If there exists a (T, μ, f^m, d, h)-WC-DIST non-adaptive randomized Turing reduction R_{Turing} for Π (per Definition 21), then for any $\gamma > 0$, Π is $(T, \mu, f^m, \lambda, \gamma)$-sparsely lossy, where*

$$\lambda = \frac{h}{m} + \max\left\{\frac{1}{m}, 9 + \frac{4}{m} + \log\left(\frac{mn}{\gamma^3}\right) + \frac{2\log d}{m}\right\} .$$

Proof. Similarly to the proof of Lemma 8.1, we show that for any $\gamma > 0$, the reduction R_{Turing} is λ-lossy for all distributions $X = (X_1, \cdots, X_m)$ where X_i's are pairwise independent B-uniform distributions over n-bit inputs, where $B = \lceil 4(m+1)/\gamma \rceil \cdot \lceil 8n \ln 2/\gamma^2 \rceil$ and $\lambda = h/m + \max\{1/m, 9/m + 4 + \log(mn/\gamma^3) + 2\log d/m\}$. In other words,

$$I(X; R_{\text{Turing}}(X)) \leq h + \max\left\{1, 9 + 4m + m\log\left(\frac{mn}{\gamma^3} + 2\log d\right)\right\} .$$

For any distribution X let $((Y_1, \ldots, Y_k), C)$ denote the distribution of $R_{\text{Turing}}(X)$. By the chain rule for the mutual information, we have

$$\begin{aligned} I\left(X; \ ((Y_1, \ldots, Y_k), C)\right) &= I(X; \ (Y_1, \ldots, Y_k)) + I\left(X; C | (Y_1, \ldots, Y_k)\right) \\ &\leq I(X; \ (Y_1, \ldots, Y_k)) + h, \end{aligned}$$

where we used the inequality $I\left(X; C | (Y_1, \ldots, Y_k)\right) \leq h$ imposed by the conditions. The rest of the proof is similar to that of Lemma 8.1 and consists of using the condition $\Delta((y_1, \ldots, y_k), D_n) \leq d$ to derive $I(X; \ (Y_1, \ldots, Y_k)) \leq \max\{1, 9 + 4m + m\log(mn/\gamma^3) + 2\log d\}$. It therefore concludes the proof.

□

8.2 Randomized Encodings

We now discuss the sparse lossiness of *randomized encodings* [6,7,53]. In Lemma 8.3, we show that a randomized encoding of a Boolean function is in fact a worst-case to distribution reductions (Definition 20). Hence, we conclude the sparse lossiness of randomized encodings and their utility in building one-way functions in Corollary 8.1.

We first recall the definition of randomized encodings.

Definition 22 (Randomized Encoding (Adapted from [7])). *Let $\mu, d \in [0,1]$ and let $F : \{0,1\}^* \to \{0,1\}^*$ be a function. We say that a function $E : \{0,1\}^* \to \{0,1\}^*$ is a (T, μ, d)-randomized encoding of F, if*

1. *for all $x \in \{0,1\}^n$, $E(x)$ can be computed in time $T(n)$, and*
2. *(μ-**correctness**) there exists an algorithm Dec such that for all $x \in \{0,1\}^n$:*

$$\Pr\left[\mathsf{Dec}(E(x)) \neq F(x)\right] \leq \mu \ ,$$

and

3. *(d-**privacy**) there exists an algorithm Sim such that for all $x \in \{0,1\}^n$:*

$$\Delta(\mathsf{Sim}(F(x)), E(x)) \leq d \ .$$

When F is the characteristic function of a promise problem Π, We say that E is a randomized encoding for Π.

Lemma 8.3. *Let $E : \{0,1\}^* \to \{0,1\}^*$ be a (T, μ, d)-randomized encoding for a Boolean function $F : \{0,1\}^* \to \{0,1\}$. Then E is a (T, μ, id, d)-worst-case to distribution splitting-reduction for Π, where $\Pi = \Pi_Y \cup \Pi_N$ is a promise problem defined as $\Pi_Y = \{x \mid F(x) = 1\}$, and $\Pi_N = \{x \mid F(x) = 0\}$, and $id : x \mapsto x$ is the identity function.*

Proof. We start by showing that $E(\cdot, \mathcal{U}_m)$ is a (μ, id)-reduction for Π as in Definition 14, which by definition implies that it is a (μ, id)-distinguisher reduction. Let $x, x' \in \Pi \cap \{0,1\}^*$ such that $\chi_\Pi(x) \neq \chi_\Pi(x')$, i.e. without loss of generality we can assume that $F(x) = 1$ and $F(x') = 0$. By μ-correctness of the randomized encoding E, there is a distinguisher Dec such that

$$\begin{aligned}
&|\Pr(\mathsf{Dec}(E(x)) = 1) - \Pr(\mathsf{Dec}(E(x')) = 1)| \\
&\quad = |\Pr(\mathsf{Dec}(E(x)) = F(x)) - \Pr(\mathsf{Dec}(E(x')) \neq F(x'))| \\
&\quad \geq (1 - \mu) - \mu.
\end{aligned}$$

For $x \in \Pi_Y \cap \{0,1\}^*$, we have $F(x) = 1$, thus $\mathsf{Sim}(1) = \mathsf{Sim}(F(x))$ is a distribution over the YES instances, by a similar argument $\mathsf{Sim}(0)$ is a distribution over the NO instances. By d-secrery of the randomized encoding, for every $x \in \Pi_Y \cap \{0,1\}^*$, we have that

$$\Delta(E(x) - \mathsf{Sim}(1)) \leq d \ ,$$

and the same approximation holds for $E(x)$ with instances $x \in \Pi_N \cap \{0,1\}^*$ with respect to $\mathsf{Sim}(0)$, leading to the desired result. □

The function F above can be chosen as χ_Π or $f^m \circ \chi_\Pi$ for a m-bit input Boolean function f^m. For these choices, we have the following statement.

Corollary 8.1 (Sparse Lossiness of Randomized Encodings). *If there exists a (T, μ, d)-randomized encoding* E *for $f^m \circ \chi_\Pi$, then for any $\gamma > 0$, Π is $(T, \mu, f^m, \lambda, \gamma)$-sparsely lossy, where*

$$\lambda = \max\left\{\frac{1}{m}, 9 + \frac{4}{m} + \log\left(\frac{mn}{\gamma^3}\right) + \frac{2\log d}{m}\right\} \ .$$

9 Applications

In the previous sections, we analysed the conditions under which a sparsely lossy reduction or a WC-DIST reduction of Π implies one-way functions under the hardness of Π. In this section, we discuss the concrete parameters. Except in Section 9.5, all statements are subject to classical algorithms.

9.1 Hardness vs One-Wayness

Let us discuss the implications of generic sparsely lossy reductions. We will explicit some particular conditions under which (fine-grained) one-way functions exist. Before that, we require the following quantitative measure of hardness.

Definition 23 (τ_Π ; Exact Hardness of Problems). *For a problem Π, let $\tau_\Pi(n) := \inf_{\tau_i(n)\in\Upsilon}\{\tau_i\}$ (the limit is taken point-wise), where Υ is the set of family of functions τ_i such that $\Pi \cap \{0,1\}^n$ can be solved with $O(2^{\tau_i(n)})$-time Turing machines with advice on all instances with probability $\geq 2/3$.*

Note that always $\tau_\Pi(n) \leq n$. This is because algorithms with an advice of size 2^n (maximum size of the truth table of χ_Π) can solve any instance of size n.

We also need following lemma.

Lemma 9.1. *For a non-constant permutation-invariant function f^m, if an f^m-distinguisher reduction has an error μ that is within a constant distance from $1/2$, then it must have runtime $\Omega(m)$.*

Proof. Assume that the reduction has runtime $o(m)$. Supposing that reading each input of the reduction takes constant time, the assumption implies that the circuit evaluating the reduction ignores $m - o(m)$ number of inputs. Let $\mathcal{I}$ be the indices of the discarded inputs, and let $p(f)$ be as in Lemma 4.4. As shown in the same lemma, function f only depends on the number of 1's in its inputs. On each input with $p(f)-1$ number of 1's (which evaluates to 0), one can flip one of the 0's to 1 and obtain an input that evaluates to 1. However, if the index of this input is in $\mathcal{I}$, it will be discarded by the reduction. Therefore, on $|\mathcal{I}| = m - o(m)$ number of bit-flips, the reduction errs. Consequently, the error must be at least $(m - o(m))/(2m)$. □

The following theorem provides the explicit conditions on sparsely lossy problems that allow to build (fine-grained) one-way functions.

Theorem 9.1 (One-Wayness from Sparsely Lossy Reductions). *Let f^m be a non-constant permutation-invariant function and Π be a promise problem. Let $n \in \mathbb{N}$ and $\gamma \in (0,1], \lambda, c \in \mathbb{R}^+$ be functions of n.*

1. *If $c \geq 3$, $\lambda < \tau_\Pi/c$, and Π is $(T, \mu, f^m, \lambda, \gamma)$-sparsely lossy such that $T = o(2^{\tau_\Pi/c})$ and $10\mu + \gamma \leq 2^{-\lambda-3}$, then there exists a (c,θ)-fine-grained one-way function where θ is the following function: $x \mapsto 1 - 1/(16x)$.*
2. *If $c > 1$, $\lambda = O(1)$, and Π is $(T, \mu, f^m, \lambda, \gamma)$-sparsely lossy such that $T = O(2^{\tau_\Pi/c})$, $10\mu + \gamma \leq 2^{-\lambda-3}$, and $\gamma = m^2 \cdot \omega(2^{-\tau_\Pi/c})$, then $(c, O(1))$-fine-grained one-way functions exist.*
3. *If, in addition to the conditions in Item 1 or Item 2, it also holds that $T = 2^{o(\tau_\Pi)}$ and $\lambda = o(\tau_\Pi)$, then one-way functions exist.*

Proof. We first prove Item 1. For these parameters, we have $\theta_{\mathsf{owf}} := (1 - 10\mu) - (\delta(\lambda) + \gamma) \geq 2^{-\lambda-3}$. Then, in Theorem 6.2, F has runtime $O(T + m^2 2^\lambda)$ and the runtime of the Π-solver is $O(2^{2\lambda}(T + T_\mathcal{A}) + m^2 2^{3\lambda})$, for all sufficiently large n. Let $\kappa := T + m^2 2^\lambda$. Then it holds that

$$\kappa^3 = T^3 + 3T^2 m^2 2^\lambda + 3T m^4 2^{2\lambda} + m^6 2^{3\lambda} \geq T 2^{2\lambda} + m^2 2^{3\lambda}\,.$$

Assume that $\mathcal{A}$ runs in time $\mathsf{poly}_\ell(T, m^2)$ for some degree ℓ. We have

$$O\left(\kappa^{\ell+2}\right) = O\left(\sum_{i=0}^{\ell+2} \binom{\ell+2}{i} T^{\ell+2-i} (m^2 2^\lambda)^i\right) = 2^{2\lambda}\mathsf{poly}_\ell(T, m^2) = 2^{2\lambda} T_\mathcal{A}\,.$$

Therefore, we have $O(2^{2\lambda}(T + T_\mathcal{A}) + m^2 2^{3\lambda}) = O(\kappa^3 + \kappa^{\ell+2})$. For Π being $O(\kappa^{\max\{3,\ell+2\}})$-hard we must have $O(\kappa^{\max\{3,\ell+2\}}) = o(2^{\tau_\Pi})$. In fact, by the change of parameters $c = \max\{3, \ell+2\}$, the $O(\kappa^c)$-hardness of Π holds since $T = o(2^{\tau_\Pi/c})$ and $\lambda < \tau_\Pi/c$. Then by Theorem 6.2, no algorithm $\mathcal{A}$ of runtime $O(\kappa^c)$ can invert F with probability better than $1 - \theta_{\mathsf{owf}}/2$. On the other hand, we have $1 - \theta_{\mathsf{owf}}/2 \leq 1 - 2^{-\lambda-4}$. By Lemma 9.1, we obtain

$$\kappa = T + m^2 2^\lambda \geq cm + m^2 2^\lambda \geq m(c + 2^\lambda)\,.$$

Therefore $2^\lambda \leq \kappa/m - c$. Hence $1 - \theta_{\mathsf{owf}}/2 \leq 1 - 2^{-\lambda-4} \leq 1 - m/(16(\kappa - cm)) \leq 1 - 1/(16\kappa)$. Finally, this means that no algorithm of runtime $O(\kappa^c)$ can invert F (whose runtime is $O(\kappa)$) with advantage more than $\approx 1 - 1/(16\kappa)$. This implies a $(c, 1 - 1/16\kappa)$-fine-grained one-way function.

For Item 2, by differently setting $\kappa := T + m^2\gamma^{-1}$ and a similar argument, the statement follows. Item 3 implies that Item 1 or 2 hold for all constants c which immediately yields one-way functions. □

We first consider the application of theorem above to the compressing reductions. We obtain the following corollary.

Theorem 9.2 (One-Wayness from Compressing Reductions).
Let f^m be a non-constant permutation-invariant function and Π be a promise problem. Let $n \in \mathbb{N}$ and $\lambda, c \in \mathbb{R}^+$ be functions of n. Assume that Π has an f^m-compression reductions that compresses mn bits to $m\lambda$ bits. Then the following statements hold:

1. *If $c \geq 3$, $\lambda < \tau_\Pi/c$, $T = o(2^{\tau_\Pi/c})$ and $\mu \leq 2^{-\lambda-4}/10$, then there exists a (c, θ)-fine-grained one-way function where θ is the following function: $x \mapsto 1 - 1/(16x)$.*
2. *If $c > 1$, $\lambda = O(1)$, $T = O(2^{\tau_\Pi/c})$ and $\mu \leq 2^{-\lambda-4}/10$, then $(c, O(1))$-fine-grained one-way functions exist.*
3. *If, in addition to the conditions in Item 1 or Item 2, it also holds that $T = 2^{o(\tau_\Pi)}$ and $\lambda = o(\tau_\Pi)$, then one-way functions exist.*

Proof. The statement follows by noting that reductions that compress mn bits to $m\lambda$ bits are (λ, γ)-cute for any choice of $\gamma \in (0, 1]$. In fact, the runtime T is independent from the choice of γ. We then set γ to be in $[m^2/T, 2^{-\lambda-4}]$. □

For WC-DIST reductions, we achieve stronger results.

Theorem 9.3 (One-Wayness from WC-DIST Reductions).
Let f^m be a non-constant permutation-invariant function and Π be a promise problem. Assume that Π has a (T, μ, f^m, d)-WC-DIST splitting reduction (as per Definition 20) such that

$$\mu \leq \frac{2^{-13-4/m}}{10}, \quad d \leq 2^{-6} \cdot (2^{39}mn)^{-m/2},$$

then the following statements hold:

1. *If $T = O(2^{\tau_\Pi/c})$ for some $c > 1$, then $(c, O(1))$-fine-grained one-way functions exist.*
2. *If $T = 2^{o(\tau_\Pi)}$, then one-way functions exist.*

Proof. Let $\gamma = 2^{-13-4/m}$. By Lemma 8.1 (or Corollary 8.1), the reduction have lossiness $\lambda \leq 9 + 4/m + \log\left(mn/\gamma^3\right) + 2\log d/m \leq 9 + 4/m$. It also holds that $10\mu + \gamma \leq 2^{-\lambda-3}$. Therefore, the statement follows by using Item 2 of Theorem 9.1.

□

Remark 4 (Relativization). We note that all the statements above relativize; one can assume that all the algorithms have accesss to an arbitrary oracle $\mathcal{O}$. This is reminiscent of the following facts. Firstly, the diguising lemma relativizes. In fact, there is no restriction on the mapping R in Lemma 3.1 and it can particularly set as $R^{\mathcal{O}}$. Secondly, both Theorems 5.1,6.2 which were used for the proofs of the results above, relativize. We also note that all results in the following sections relativize.

9.2 Sparsely Lossy Problems Reduce to SZK

The specific sparse lossiness parameters that are used in Section 9.1 restrict the choice of Π. A problem that admits such sparsely lossy reductions cannot be arbitrary. In this section, we show that if Π has such sparsely lossy reductions, then it reduces to SZK within a runtime that only depends on λ and T. The first theorem concerns the generic sparsely lossy reductions of Theorem 9.1.

Theorem 9.4. *In Theorem 9.1, if Π satisfies the conditions of Item 2 or 3, then Π reduces to SZK in time $O(2^\lambda T)$ with zero error.*

Proof. Let θ_{szk} be as in Theorem 5.1, namely, $\theta_{\mathsf{szk}} = (1-2\mu)^2/(\delta(\lambda)+\gamma)$. By the conditions of the statement, we have $(1-2^{-\lambda-3})^2 < (1-2\mu)^2 \leq 1$, $0 < \gamma \leq 2^{-\lambda-3}$ (for sufficiently large n), and $\delta(\lambda) = 1-2^{-\lambda-2}$. Therefore, we obtain

$$\frac{1-2^{-\lambda-2}}{1-2^{-\lambda-3}} \approx \frac{(1-2^{-\lambda-3})^2}{1-2^{-\lambda-2}+2^{-\lambda-3}} \leq \theta_{\mathsf{szk}} = \frac{(1-2\mu)^2}{\delta(\lambda)+\gamma} \leq \frac{1}{1-2^{-\lambda-2}}\,.$$

By the approximation $\log(1+x) \approx x$, we have $\log \theta_{\mathsf{szk}} = \Theta(2^{-\lambda})$. Then by Theorem 5.1, Π reduces to SZK in time $O(2^{\lambda}(T+m^2/\gamma))$ with a classical advice of size $4mn/\gamma$. Note that since $\gamma = m^2 \cdot \omega(2^{-\tau_\Pi/c})$ and $T = O(2^{\tau_\Pi/c})$, it holds that $\gamma = \Omega(m^2/T)$. Therefore, this runtime can be simplified as $O(2^{\lambda}T)$. □

Similar statement holds for compressing reductions.

Theorem 9.5. *Under the conditions of Theorem 9.2, Π reduces to SZK in time $O(2^{\lambda}T)$ with zero error.*

Proof. Without loss of generality, we assume that the compressing reduction is $(\lambda, \gamma = \Omega(m^2/T))$-sparsely lossy . The rest of the proof is exactly similar to that of Theorem 9.4. □

We also have the following result for WC-DIST reductions.

Theorem 9.6. *Under the conditions of Theorem 9.3, Π reduces to SZK in time $O(T)$ with zero error.*

Proof. As shown in the proof of Theorem 9.3, the reduction is $(\lambda = 9+4/m, \gamma = 2^{-13-4/m})$-sparsely lossy. So without loss of generality, we assume that $\lambda, \gamma = O(1)$. The rest of the proof is exactly similar to that of Theorem 9.4. □

9.3 On the Existence of Fine-Grained One-Way Functions from kSAT

Recall that the kSAT problem asks to decide whether a CNF formula of N variables and M clauses, where each clause has k variables, has a satisfiable assignment. The exact hardness of kSAT has been formulated by Impagliazzo and Paturi [52] as below.

Assumption 1 (non-uniform Exponential Time Hypothesis). *Let $s_k := \inf\{c \in \mathbb{R} \mid k\text{SAT} \in \mathsf{nuTIME}(2^{cN})\}$. Then $s_3 > 0$.*

The above assumption is sometimes denoted simply by nuETH. Impagliazzo and Paturi [52] show that the nuETH assumption is equivalent to $\forall k \geq 3 : \ s_k > 0$. One can also use the number of clauses M instead of the number of variables N in the above definition and still obtain the same quantity for s_k (e.g. see the thesis of Zeijlemaker [86, Cor. 4.8.2]). We also have:

Lemma 9.2. *Under nuETH, we have $\tau_{k\text{SAT}}(n) \in [(s_k/2k)\cdot n/\log n, (4s_k)\cdot n/\log n]$ for sufficiently large n, where n is the instance size.*

Proof. For any fixed k, we have $\lceil N/k \rceil \leq M \leq (2N)^k$. On the other hand, the bit-size of an instance is $n := kM\lceil \log 2N \rceil$. It follows that $M \log M \leq n \leq kM(\log M + \log(2k))$ where the right hand side is smaller than $2kM \log M$ for sufficiently large instances. For sufficiently large instances, we have

$$(M\log M)/\log(M\log M) \leq M \leq 2(M\log M)/\log(M\log M)\,,$$

and

$$(M\log M)/(2\log(M\log M)) \leq n/\log n \leq 2k(M\log M)/\log(M\log M)\,.$$

Hence, $n/(2k\log n) \leq M \leq 4n/\log n$. By rewriting $2^{s_k M}$ in terms of the instance size n, we obtain $2^{s_k M} \in [2^{(s_k/2k)n/\log}, 2^{(4s_k)n/\log n}]$.

□

We obtain the following corollary.

Corollary 9.1 (FGWOF from nuETH). *Let f^m be a non-constant permutation-invariant Boolean function. Then the following statements hold under nuETH:*

1. *Let $\lambda \geq 0$ such that $\lambda < s_k n/(6k \cdot \log n)$. If kSAT has an f^m-reduction that runs in time $O(2^{s_k n/(6k \cdot \log n)})$, compresses m instances of n bits to $m\lambda$ bits and has error $\leq 2^{-\lambda-4}/10$, then there exists a fine-grained one-way function. Furthermore, when $\lambda = O(1)$ fine-grained one-way functions exist if the reduction runs in time $O(2^{s_k n/(2kc \cdot \log n)})$ for some $c > 1$.*
2. *If there exists a worst-case to average-case Karp f^m-reduction with distance d for kSAT or a randomized-encoding with privacy d for $f^m \circ \chi_{k\text{SAT}}$ with the parameters below, then fine-grained one-way functions exist:*

$$T = O(2^{s_k n/(2kc \cdot \log n)}), \quad \mu \leq 2^{-21}, \quad d \leq 2^{-6} \cdot (2^{39} mn)^{-m/2},$$

for some $c > 1$.

Proof. The statements follow by Theorem 9.2 and 9.3. □

Barriers for One-Way Functions from kSAT

The Non-deterministic Strong Exponential Time Hypothesis (NETH) is a hypothesis that has been formalized by [28]. Refuting NETH would imply breakthough results in proof complexity and circuit lower bounds [28, 30], which provides evidence for legitimacy of this assumption. In our non-uniform setting, we are more interested in the non-uniform variant of NETH. More precisely, the assumption is as follows:

Assumption 2 (non-uniform NETH (nuNETH)). *Let kTAUT be the language of k-DNF tautologies. Moreover, let $\sigma_k := \inf\{c \in \mathbb{R} \mid k\text{TAUT} \in \mathsf{nuNTIME}(2^{cN})\}$ where N is the number of variables. Then $\sigma_k > 0$ for every $k \geq 3$.*

The quantity σ_k remains invariant under replacing the number of clauses M with the number of variables N in the above definition. This is reminiscent of the proof of [86, Cor. 4.8.2].

nuNETH implies that kSAT cannot be solved in non-uniform co-nondeterministic subexponential-time.[1] Therefore, it implies the impossibility of building one-way functions based on kSAT using our methods.

Corollary 9.2. *Under nuNETH, kSAT does not satisfy the conditions of Item 3 of Theorem 9.1 or Item 3 of Theorem 9.2 or Item 2 of Theorem 9.3.*

Proof. Theorem 9.4, 9.5, and 9.6 respectively show that under the conditions of Item 3 of Theorem 9.1 or Item 3 of Theorem 9.2 or Item 2 of Theorem 9.3, kSAT reduces to SZK in time $2^{o(\tau_{k\text{SAT}}(n))}$ where n is the instance size of kSAT. Therefore, by Lemma 2.8, kSAT reduces to coNP in time $2^{o(\tau_{k\text{SAT}}(n))}$. On the other hand, nuNETH implies that $\tau_{k\text{SAT}} = \Omega(M)$ (as a function of M). Taking into account the fact that $M = \Theta(n/\log n)$, we obtain a $2^{o(M)}$-time reduction of kSAT to coNP. This is a violatoin of nuNETH. □

However, this argument alone cannot refute building *fine-grained* one-way functions based on kSAT and invalidate Corollary 9.1. For this purpose, we expect that one requires stronger assumptions to establish whether s_k is larger (or smaller) than σ_k. At this moment, we are not aware of any concrete comparison that is implied by nuNETH.

9.4 On the Non-Existence of Statistical Obfuscation

We start this section by defining the obfuscation functionality as below.

Definition 24 (Statistical Obfuscation (sO), [42]). *Let N be a positive integer, and $\alpha, \varepsilon : \mathbb{N} \to [0, 1)$ be functions. An α-statistical obfuscation with error ε is an algorithm sO that, upon receiving a security parameter 1^N and a circuit C of size $\mathsf{poly}(N)$ as inputs, runs in time $\mathsf{poly}(N)$ and outputs a circuit with the following specifications:*

[1] The runtime is parameterized in M.

1. ***Correctness.*** *For any circuit C over inputs of size N and any $x \in \{0,1\}^N$, it holds that*

$$\Pr\left[\mathsf{sO}(1^N, C)(x) \neq C(x)\right] \leq \varepsilon(N)\,,$$

where the probability is taken over the randomness of sO.
2. ***Statistical Distance.*** *For all circuits C_1 and C_2 that are functionally equivalent over inputs of size N such that $|C_1| = |C_2|$, it holds that*

$$\Delta(\mathsf{sO}(1^N, C_1), \mathsf{sO}(1^N, C_2)) \leq \alpha(N)\,.$$

Remark 5. The statistical distance property is defined in a more general way in [42]. It requires the existence of an *efficient* simulator Sim such that for all functionally equivalent circuits C_1 and C_2 with $|C_1| = |C_2|$ over inputs of size N, it holds that $\Delta(\mathsf{sO}(1^N, C_1), \mathsf{Sim}(1^N, C_2)) \leq \alpha(N)$. We note that our impossibility result stated in Theorem 9.7 also applies to this variant even if the simulator is allowed to be unbounded.

Remark 6. In the correctness definition that is proposed by [42], it is required that $\Pr\left[\exists x \in \{0,1\}^N \;:\; \mathsf{sO}(1^N, C)(x) \neq C($ ε. This guarantees the functional equivalency of the obfuscation, while the above definition only guarantees point-wise equivalency of the obfuscation. In the scenarioes where $\varepsilon(N) = \mathsf{negl}(N)$ and only $\mathsf{poly}(N)$ number of evaluations is needed, using the weaker definition incurs only $\mathsf{negl}(N)$ error compared to that of [42].

The high-level idea behind the main result of this section is to apply the obfuscation scheme over instances of a variant of Sat. Recall that the problem Sat asks to decide whether a Circuit (a CNF formula) over N variables has a satisfiable assignment. The variant that we are interested in is UniqueSat that is the same problem under the promise that the input circuit has at most one satisfiable assignment. An obfuscation scheme provides a worst-case to average-case reduction for UniqueSat. This remarkable consequence is observed by [57] and [50]. They observe that if the input circuit C is a YES instance of UniqueSat, i.e., it evaluates to 1 on exactly one input, then a random shift of C, namely the circuit $C_z(x) = C(x \oplus z)$, where $z \xleftarrow{\$} \{0,1\}^N$, has a truth-table whose distribution is identical to that of a random point function. Therefore, the obfuscation of C_z must be statistically close to the obfuscation of a random point function. Moreover, if the obfuscation is perfect, i.e., error $\varepsilon = 0$, then the obfuscation of C_z is a YES instance of UniqueSat.

We present a modified version of their reduction in Figure 3.

Algorithm 3 WC-DIST Reduction for UniqueSat.

Parameters: Positive integer N and a positive odd number k.
Input: A circuit C over N-bit size inputs.
Output: A tuple of circuits $(C_1, \cdots, C_k)$.

1: Sample $z \xleftarrow{\$} \{0,1\}^N$.
2: Define $C_z : \{0,1\}^N \to \{0,1\}$ s.t. $C_z(\cdot) := C(\cdot \oplus z)$.
3: For $i \in [k]$, sample $r_i \xleftarrow{\$} \{0,1\}^{\mathsf{poly}(N)}$.
4: For $i \in [k]$, let $C_i := \mathsf{sO}(C_z; r_i)$.
5: Return $(C_1, \cdots, C_k)$.

Below, we show that Algorithm 3 is a WC-DIST splitting reduction for UniqueSat.

Lemma 9.3. *Assume that there exists an α-statistical obfuscation sO with error ε. There exists $k = \mathsf{poly}(N)$ such that Algorithm 3 with parameters N and k is a $(T, \mu \leq 2^{-21}, \mathsf{id}, d \leq 2^{-25.5}/\sqrt{n})$-WC-DIST splitting reduction (as per Definition 20) for* UniqueSat*, where*

1. *$T = \mathsf{poly}(n)$, if $\alpha = 1 - 1/\mathsf{poly}(N)$, and $\varepsilon = 1/2 - 1/\mathsf{poly}(N)$, and*
2. *$T = \mathsf{subexp}(n)$, if $\alpha = 1 - 1/\mathsf{subexp}(N)$, and $\varepsilon = 1/2 - 1/\mathsf{subexp}(N)$.*

In above, $\mathsf{id} : x \mapsto x$ is the identity function, and $n := |C|$ is the instance size.

Proof. We analyze the distance d, error μ and runtime T of the reduction.
Distance of the reduction: If C is not satisfiable, then the truth table of the all-zero function, which we denote by $\mathbb{1}_{\varnothing}$, is equivalent to that of C_z. On the other hand, if C is satisfiable, then the truth table of the point function $\mathbb{1}_{x=x^*}$, where $x^* \xleftarrow{\$} \{0,1\}^N$, is statistically equivalent to that of C_z. Therefore, by the correlation property of sO, for every $i \in [k]$, if C is not satisfiable then $\Delta(C_i, \mathsf{sO}(1^N, \mathbb{1}_{\varnothing})) \leq \alpha(N)$, and if C is satisfiable, then $\Delta(C_i, \mathsf{sO}(1^N, \mathbb{1}_{x=x^*})) \leq \alpha(N)$. Finally, By the direct product lemma, we get the following statements:

1. If C is not satisfiable, then
$$\Delta((C_1, \cdots, C_k), \mathsf{sO}(1^N, \mathbb{1}_{\varnothing})^{\otimes k}) \leq \alpha(N)^k \,.$$
2. If C is satisfiable, then
$$\Delta((C_1, \cdots, C_k), \mathsf{sO}(1^N, \mathbb{1}_{x=x^*})^{\otimes k}) \leq \alpha(N)^k \,.$$

Therefore, Algorithm 3 maps the input C to a tuple $(C_1, \cdots, C_k)$ whose statistical distance to an input-independent distribution is $d = \alpha(N)^k$. Note that the size of the instance $n = |C|$ is $\mathsf{poly}(N)$. Thus, if $\alpha(N) = 1 - 1/\mathsf{poly}(N)$, there exists $k = \mathsf{poly}(N)$ such that $\alpha(N)^k \leq 2^{-25.5}/\sqrt{n}$. Moreover, if $\alpha(N) = 1 - 1/\mathsf{subexp}(N)$, then it suffices to take $k = \mathsf{subexp}(N)$. This concludes the distance analysis.

Error of the reduction: Looking closer, Algorithm 3 actually reduces UniqueSat to the $\textsc{Maj}_k \circ$ UniqueSat problem that is defined as follows: A tuple $(C'_1, \cdots, C'_k)$ of k circuits over N-bit inputs is a YES instance of $\textsc{Maj}_k \circ$ UniqueSat if there exists $x \in \{0,1\}^N$ such that the majority of C'_i's evaluate to 1 on x, and the tuple is a NO instance if such an input x does not exist. Note that every C'_i is fed with the same input x. By the correctness of sO, for every $x \in \{0,1\}^N$ and every $i \in [k]$, we have that $\Pr[C_z(x) \neq C_i(x)] \leq \varepsilon(N)$. For every $i \in [k]$, let X_i be a random variable that outputs 1 if $C_z(x) \neq C_i(x)$ and 0 otherwise. We have $\mu := \mathbb{E}[\sum X_i] \leq \varepsilon k$. Lemma 2.3 implies that
$$\Pr\left[\left|\sum_{i=1}^{k} X_i - \mu\right| > (\frac{1}{4} + \frac{\varepsilon}{2})k\right] \leq 2e^{-2(1/4+\varepsilon/2)^2 k} \,.$$
If $\varepsilon = 1/2 - 1/\mathsf{poly}(N)$, then $(1/4 + \varepsilon/2)k = (1/2 - 1/\mathsf{poly}(N))k$. Therefore, by taking a sufficiently large $k = \mathsf{poly}(N)$, we obtain that
$$\begin{aligned}\Pr\left[\sum_{i=1}^{k} X_i \geq k/2\right] &\leq \Pr\left[\left|\sum_{i=1}^{k} X_i - \mu\right| > (\frac{1}{4} + \frac{\varepsilon}{2})k\right] \\ &\leq 2e^{-2(1/2-1/\mathsf{poly}(N))^2 k} = e^{-\Omega(k)} \,.\end{aligned}$$

Therefore, it holds that $\Pr[\, C_z(x) \neq \textsc{Maj}\{C_i(x)\}_i] \leq e^{-\Omega(k)}$. For the case where $\varepsilon = 1/2 - 1/\mathsf{subexp}(N)$, a similar argument holds by taking a sufficiently large $k = \mathsf{subexp}(N)$. We then use the union bound over all 2^N possible inputs $x \in \{0,1\}^N$ to obtain the following statements:

1. If C is not satisfiable, then $(C_1, \cdots, C_k)$ is a NO instance of $\textsc{Maj}_k \circ$ UniqueSat with probability at least $1 - 2^N e^{-\Omega(k)}$.
2. If C is satisfiable, then $(C_1, \cdots, C_k)$ is a YES instance of $\textsc{Maj}_k \circ$ UniqueSat with probability at least $1 - 2^N e^{-\Omega(k)}$.

By taking a sufficiently large $k = \mathsf{poly}(N)$, we obtain that $1 - 2^N e^{-\Omega(k)} > 1 - 2^{-21}$. Therefore, the reduction in Algorithm 3 has an error $\leq 2^{-21}$ for either of choices of α in the statement.

Runtime of the reduction: The runtime of the reduction is $k \cdot \mathsf{poly}(N)$. Recall that $|C| = \mathsf{poly}(n)$. Without loss of generality, we can assume that $|C| = \Omega(N)$; else, we can pad C with enough garbage bits to add to its length, and throw out these extra bits in the reduction. For the case where $\alpha(N) = 1 - 1/\mathsf{poly}(N)$, we saw that k should be set as $k = \mathsf{poly}(N)$. Therefore, it holds that $k \cdot \mathsf{poly}(N) = \mathsf{poly}(|C|)$. For the case where $\alpha(N) = 1 - 1/\mathsf{subexp}(N)$, we have $k = \mathsf{subexp}(N)$. Therefore, since $|C| = \Omega(N)$, we obtain $k \cdot \mathsf{poly}(N) = \mathsf{subexp}(|C|)$.

□

Valiant and Vazirani [78] show that there exists a reduction from SAT to UNIQUESAT such that given an instance C (a circuit over inputs $x \in \{0,1\}^N$), runs in time $\mathsf{poly}(N)$ and has error at most $1 - 1/\mathsf{poly}(N)$. The following lemma is obtained using their reduction.

Lemma 9.4. *There exists a (non-uniform) reduction from* SAT *to* MAJ $\circ$ UNIQUESAT *that runs in time* $\mathsf{poly}(N)$ *and has zero error.*

Proof. Let C be an instance of SAT with size at most $q(N)$. Let R be the Valiant-Vazirani reduction and p be such that R has error at most $1 - 1/p(N)$. By repeating the reduction for a sufficiently large $k = \mathsf{poly}(p, N)$ times and taking the majority, the error reduces to $e^{-\Omega(k)}$ according to Lemma 2.3. Moreover, the union bound implies that R errs over at least one circuit C of size $q(N)$ with probability at most $2^{q(N)} \cdot e^{-\Omega(k)}$. By carefully tuning k to be a polynomial larger than q, this error quantity becomes strictly smaller than 1. Therefore, there exists at least one random string such that the reduction correctly maps all the instances C of size q. We set this random string as the non-uniform advice. Therefore, this reduction runs in $\mathsf{poly}(N)$ time and has zero error.

□

Finally, we have the main theorem of this section.

Theorem 9.7. *Assuming that* $\mathsf{NP} \not\subseteq \mathsf{coNP/poly}$*, then* $(1 - 1/\mathsf{poly}(N))$*-statistical obfuscation with error* $1/2 - 1/\mathsf{poly}(N)$ *does not exist.*

Moreover, assuming nuNETH, then $(1 - 1/\mathsf{subexp}(N))$*-statistical obfuscation with error* $1/2 - 1/\mathsf{subexp}(N)$ *does not exist.*

Proof. We prove the statements separately:

Proving the first statement: Assume $(1-1/\mathsf{poly}(N))$-statistical obfuscation with error $1/2-1/\mathsf{poly}(N)$ exists. Lemma 9.3 presents a reduction for UNIQUESAT that is $(T = \mathsf{poly}(n), \mu \leq 2^{-21}, \mathsf{id}, d \leq 2^{-25.5}/\sqrt{n})$-WC-DIST splitting. Therefore, due to Theorem 9.6, UNIQUESAT reduces to SZK in time $T = \mathsf{poly}(n)$ with zero error. On the other hand, Lemma 9.4 provides a zero-erro reduction from SAT to MAJ $\circ$ UNIQUESAT. Combining these two reductions, we obtain a zero-erro reduction from SAT to MAJ $\circ$ SZK that runs in polynomial time. The class SZK is closed under majority [71, Cor. 4.14], therefore, the whole chain of reductions imply that SAT $\in \mathsf{SZK/poly} \subseteq \mathsf{coNP/poly}$. In other words, we obtain $\mathsf{NP} \subseteq \mathsf{coNP/poly}$ which contradicts the assumption.

Proving the second statement: Assume $(1 - 1/\mathsf{subexp}(N))$-statistical obfuscation with error $1/2 - 1/\mathsf{subexp}(N)$ exists. Lemma 9.3 presents a reduction for UNIQUESAT that is $(T = \mathsf{subexp}(n), \mu \leq 2^{-21}, \mathsf{id}, d \leq 2^{-25.5}/\sqrt{n})$-WC-DIST splitting. Therefore, by Theorem 9.6, UNIQUESAT reduces to SZK in time $T = \mathsf{subexp}(n)$ and zero error. By combining this with the zero-error reduction of SAT presented in Lemma 9.4, we obtain that SAT reduces to $\mathsf{SZK} \subseteq \mathsf{coNP/poly}$ in time $\mathsf{subexp}(N)$. However, this contradicts nuNETH.

□

Remark 7. We note that this impossibility result should not be confused with the positive result of Brakerski, Brzuska, and Fleischhacker [21]. They construct an efficient obfuscator when $2\varepsilon+\alpha > 1$ [21, Appendix A]. But their correctness definition of obfuscation is quite different. They require that the obfuscator correctly evaluates the original circuit only over a uniformly random input, which is a relatively weaker guarantee.

9.5 Quantum Sparsely Lossy Reductions

In this section, we use quantum sparsely lossy reduction to derive one-way state generators. We first define a measure of quantum hardness as follows:

Definition 25 (Exact Quantum Hardness of Problems). *For a problem* Π*, let* $\tau^Q_\Pi(n) := \inf_{\tau_i(n) \in \Upsilon}\{\tau_i\}$ *(the limit is taken point-wise), where* Υ *is the set of family of functions* τ_i *such that* $\Pi \cap \{0,1\}^n$ *can be solved by quantum algorithms with classical advice in time* $O(2^{\tau_i(n)})$ *on all instances with probability* $\geq 2/3$*.*

Theorem 9.8. *Let f^m be a non-constant permutation-invariant function and Π be a promise problem. Let $n \in \mathbb{N}$, and $\lambda \geq 0$ be a function of n. If Π is $(T, \mu, f^m, \lambda, \gamma)$-sparsely lossy with a pure-outcome quantum reduction such that $\lambda = o(\tau^Q_\Pi)$, $T = 2^{o(\tau^Q_\Pi)}$ and $4\sqrt{2\mu} + \gamma \leq 2^{-\lambda-3}$, then there exists a one-way state generator.*

Proof. We compute θ_{owf} and τ_{owsg} that are required in Theorem 7.1. For the given parameters, we have $\theta_{\mathsf{ows}} = 1-(\delta(\lambda)+\gamma+4\sqrt{2\mu}) \geq 2^{-\lambda-3}$ and $\tau_{\mathsf{ows}} \geq \theta_{\mathsf{ows}} \geq 2^{-\lambda-3}$. The runtime of the construction G in Theorem 7.1 is $O(T+m^2 2^\lambda)$ and the runtime of the Π-solver is $O(2^{2\lambda}(cT+T_{\mathcal{A}}+2^{2\lambda})+cm^2 2^{3\lambda})$, for all sufficiently large n. By letting $\kappa := T + m^2 2^\lambda$ and following a similar argument as in Theorem 9.1, the runtime of G becomes $O(\kappa)$ while the runtime of the Π-solver becomes $O(c\kappa^3+\kappa^{\ell+2})$ where ℓ is the degree of the runtime of $\mathcal{A}$ as a polynomial in κ. As long as $T = 2^{o(\tau_\Pi)}$, $\lambda = o(\tau_\Pi)$, and $c = \mathsf{poly}(\kappa)$ an algorithm that runs in $O(c\kappa^3+\kappa^{\ell+2})$ can not solve Π. Therefore, G must be a c-copy (or $\mathsf{poly}(\kappa)$-copy) $(1-\theta_{\mathsf{ows}}/4)$-one-way state generator. Similar to the proof of Theorem 9.1, we have $(1-\theta_{\mathsf{ows}}/4) \leq 1-1/(16\kappa)$. Finally, one can conclude by noting that weak one-way state generators imply one-way state generators [63, Theorem 3.7]. □

Remark 8. The above result immediately applies to quantum compressing f^m-reductions with error $\mu \leq 2^{-2\lambda-13}$ since any quantum f^m-reduction that compresses m instance of size n to $m\lambda$ qubits is quantum (λ, γ)-sparsely lossy for any $\gamma \in (0, 1]$.

Acknowledgments. The authors thank Damien Vergnaud for helpful discussions. This work is part of HQI initiative[1] and is supported by France 2030 under the French National Research Agency award number ANR-22-PNCQ-0002.

[1] www.hqi.fr

References

1. Manindra Agrawal, Neeraj Kayal, and Nitin Saxena. Primes is in p. *Annals of Mathematics*, 2004.
2. Dorit Aharonov and Oded Regev. Lattice problems in NP ∩ coNP. *J. ACM*, 2005.
3. Miklós Ajtai. The shortest vector problem in ℓ_2 is NP-hard for randomized reductions (extended abstract). In *STOC*, 1998.
4. Adi Akavia, Oded Goldreich, Shafi Goldwasser, and Dana Moshkovitz. On basing one-way functions on NP-hardness. STOC '06, 2006.
5. Michael Alekhnovich. More on average case vs approximation complexity. FOCS '03, 2003.
6. Benny Applebaum. Garbled circuits as randomized encodings of functions: a primer. Cryptology ePrint Archive, Paper 2017/385, 2017.
7. Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. Cryptography in NC^0. *SIAM Journal on Computing*, 2006.
8. Benny Applebaum and Pavel Raykov. On the relationship between statistical zero-knowledge and statistical randomized encodings. In *Advances in Cryptology – CRYPTO 2016*, 2016.
9. Saikrishna Badrinarayanan, Vipul Goyal, Aayush Jain, and Amit Sahai. Verifiable functional encryption. In *Advances in Cryptology – ASIACRYPT 2016*, 2016.
10. Marshall Ball, Elette Boyle, Akshay Degwekar, Apoorvaa Deshpande, Alon Rosen, Vinod Vaikuntanathan, and Prashant Nalini Vasudevan. Cryptography from information loss. In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, 2020.
11. Marshall Ball, Alon Rosen, Manuel Sabin, and Prashant Nalini Vasudevan. Average-case fine-grained hardness. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2017, 2017.
12. Boaz Barak, Zvika Brakerski, Ilan Komargodski, and Pravesh K. Kothari. Limits on low-degree pseudorandom generators (or: Sum-of-squares meets program obfuscation). In *Advances in Cryptology – EUROCRYPT 2018*, 2018.
13. Boaz Barak, Oded Goldreich, Rusell Impagliazzo, Steven Rudich, Amit Sahai, Salil Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In *Advances in Cryptology – CRYPTO 2001*, 2001.
14. Boaz Barak, Samuel B. Hopkins, Aayush Jain, Pravesh Kothari, and Amit Sahai. Sum-of-squares meets program obfuscation, revisited. In *Advances in Cryptology – EUROCRYPT 2019*, 2019.
15. Mihir Bellare and Shafi Goldwasser. The complexity of decision versus search. *SIAM J. Comput.*, 1994.
16. Itay Berman, Akshay Degwekar, Ron D. Rothblum, and Prashant Nalini Vasudevan. Statistical difference beyond the polarizing regime. In *Theory of Cryptography: 17th International Conference, TCC 2019, Nuremberg, Germany, December 1–5, 2019, Proceedings, Part II*, 2019.
17. Nir Bitansky, Omer Paneth, and Alon Rosen. On the cryptographic hardness of finding a nash equilibrium. FOCS '15, 2015.
18. Andrej Bogdanov and Christina Brzuska. On basing size-verifiable one-way functions on NP-Hardness. In *Theory of Cryptography Conference*, 2015.
19. Andrej Bogdanov and Luca Trevisan. On worst-case to average-case reductions for NP problems. *SIAM Journal on Computing*, 2006.
20. John Bostanci, Luowen Qian, Nicholas Spooner, and Henry Yuen. An efficient quantum parallel repetition theorem and applications. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, STOC 2024, 2024.
21. Zvika Brakerski, Christina Brzuska, and Nils Fleischhacker. On statistically secure obfuscation with approximate correctness. In *Advances in Cryptology – CRYPTO 2016*, 2016.
22. Zvika Brakerski, Ran Canetti, and Luowen Qian. On the computational hardness needed for quantum cryptography. In *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, Leibniz International Proceedings in Informatics (LIPIcs), 2023.
23. Zvika Brakerski, Vadim Lyubashevsky, Vinod Vaikuntanathan, and Daniel Wichs. Worst-case hardness for lpn and cryptographic hashing via code smoothing. In *Advances in Cryptology – EUROCRYPT 2019*, 2019.
24. G. Brassard. A note on the complexity of cryptography (corresp.). *IEEE Transactions on Information Theory*, 1979.
25. Chris Brzuska and Geoffroy Couteau. On building fine-grained one-way functions from strong average-case hardness. *J. Cryptol.*, 2024.
26. Christina Brzuska, Pooya Farshim, and Arno Mittelbach. Indistinguishability obfuscation and uces: The case of computationally unpredictable sources. In *Advances in Cryptology – CRYPTO 2014*, 2014.
27. Harry Buhrman, Richard Cleve, John Watrous, and Ronald de Wolf. Quantum fingerprinting. *Phys. Rev. Lett.*, 2001.

28. Marco L. Carmosino, Jiawei Gao, Russell Impagliazzo, Ivan Mihajlin, Ramamohan Paturi, and Stefan Schneider. Nondeterministic extensions of the strong exponential time hypothesis and consequences for non-reducibility. ITCS '16, 2016.
29. Lijie Chen, Shuichi Hirahara, and Neekon Vafa. Average-Case Hardness of NP and PH from Worst-Case Fine-Grained Assumptions. In *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*, 2022.
30. Lijie Chen, Ron D. Rothblum, Roei Tell, and Eylon Yogev. On exponential-time hypotheses, derandomization, and circuit lower bounds. *J. ACM*, 2023.
31. Kai-Min Chung, Huijia Lin, and Rafael Pass. Constant-round concurrent zero-knowledge from indistinguishability obfuscation. In *Advances in Cryptology – CRYPTO 2015*, 2015.
32. Aloni Cohen, Justin Holmgren, Ryo Nishimaki, Vinod Vaikuntanathan, and Daniel Wichs. Watermarking cryptographic capabilities. STOC '16, 2016.
33. Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Daniel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer Publishing Company, Incorporated, 1st edition, 2015.
34. W. Diffie and M. Hellman. New directions in cryptography. *IEEE Trans. Inf. Theory*, 1976.
35. Andrew Drucker. New limits to classical and quantum instance compression. *SIAM Journal on Computing*, 2015.
36. Joan Feigenbaum and Lance Fortnow. Random-self-reducibility of complete sets. *SIAM Journal on Computing*, 1993.
37. Lance Fortnow and Rahul Santhanam. Infeasibility of instance compression and succinct pcps for np. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC '08, 2008.
38. Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. *SIAM Journal on Computing*, 2016.
39. Sanjam Garg, Omkant Pandey, and Akshayaram Srinivasan. Revisiting the cryptographic hardness of finding a nash equilibrium. In *Advances in Cryptology – CRYPTO 2016*, 2016.
40. Oded Goldreich. A note on computational indistinguishability. *Inf. Process. Lett.*, 1990.
41. Shafi Goldwasser, S. Dov Gordon, Vipul Goyal, Abhishek Jain, Jonathan Katz, Feng-Hao Liu, Amit Sahai, Elaine Shi, and Hong-Sheng Zhou. Multi-input functional encryption. In *Advances in Cryptology – EUROCRYPT 2014*, 2014.
42. Shafi Goldwasser and Guy N. Rothblum. On best-possible obfuscation. In *Theory of Cryptography*, 2007.
43. Thomas Dueholm Hansen, Haim Kaplan, Or Zamir, and Uri Zwick. Faster k-sat algorithms using biased-ppsz. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019, 2019.
44. Danny Harnik and Moni Naor. On the compressibility of np instances and cryptographic applications. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*, 2006.
45. Timon Hertli. 3-sat faster and simpler - unique-sat bounds for ppsz hold in general. In *Proceedings of the 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, FOCS '11, 2011.
46. Timon Hertli. Breaking the ppsz barrier for unique 3-sat. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *Automata, Languages, and Programming*, 2014.
47. Timon Hertli, Robin Moser, and Dominik Scheder. Improving ppsz for 3-sat using critical variables. *Symposium on Theoretical Aspects of Computer Science (STACS2011)*, 2010.
48. Shuichi Hirahara and Mikito Nanashima. One-way functions and zero knowledge. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, STOC 2024, 2024.
49. Thomas Hofmeister, Uwe Schöning, Rainer Schuler, and Osamu Watanabe. A probabilistic 3sat algorithm further improved. In *STACS 2002*, 2002.
50. Rahul Ilango and Alex Lombardi. Cryptography meets worst-case complexity: Optimal security and more from iO and worst-case assumptions, 2025.
51. R. Impagliazzo, R. Paturi, and F. Zane. Which problems have strongly exponential complexity? FOCS '98, 1998.
52. Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-sat. *J. Comput. Syst. Sci.*, 2001.
53. Y. Ishai and E. Kushilevitz. Randomizing polynomials: A new representation with applications to round-efficient secure computation. In *Proceedings 41st Annual Symposium on Foundations of Computer Science*, 2000.
54. Kazuo Iwama, Kazuhisa Seto, Tadashi Takai, and Suguru Tamaki. Improved randomized algorithms for 3-sat. In *Algorithms and Computation*, 2010.
55. Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from well-founded assumptions. STOC 2021, 2021.
56. Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from lpn over fp, dlin, and prgs in nc0. In *Advances in Cryptology – EUROCRYPT 2022*, 2022.
57. Ilan Komargodski, Tal Moran, Moni Naor, Rafael Pass, Alon Rosen, and Eylon Yogev. One-way functions and (im)perfect obfuscation. FOCS '14, 2014.

58. Rio LaVigne, Andrea Lincoln, and Virginia Vassilevska Williams. Public-key cryptography in the fine-grained setting. In *Advances in Cryptology – CRYPTO 2019*, 2019.
59. Richard J. Lipton. New directions in testing. In *Distributed Computing And Cryptography*, 1989.
60. Richard J. Lipton and Neal E. Young. Simple strategies for large zero-sum games with applications to complexity theory. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing*, STOC '94, 1994.
61. Daniel Lokshtanov, Dániel Marx, and Saket Saurabh. Slightly superexponential parameterized problems. *SIAM Journal on Computing*, 2018.
62. Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM Journal on Computing*, 2007.
63. Tomoyuki Morimae and Takashi Yamakawa. One-wayness in quantum cryptography. In *19th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2024*, 2024.
64. Moni Naor and Guy N. Rothblum. Learning to impersonate. In *Proceedings of the 23rd International Conference on Machine Learning*, ICML '06, 2006.
65. R. Ostrovsky and A. Wigderson. One-way functions are essential for non-trivial zero-knowledge. In *The 2nd Israel Symposium on Theory and Computing Systems*, 1993.
66. R. Paturi, P. Pudlák, M. E. Saks, and F. Zane. An improved exponential-time algorithm for k-sat. In *Proceedings of the 39th Annual Symposium on Foundations of Computer Science*, FOCS '98, 1998.
67. R. Paturi, P. Pudlak, and F. Zane. Satisfiability coding lemma. In *Proceedings 38th Annual Symposium on Foundations of Computer Science*, 1997.
68. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 2009.
69. R. L. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM*, 1978.
70. Daniel Rolf. Improved bound for the ppsz/schöning-algorithm for 3-sat. *Electronic Colloquium on Computational Complexity (ECCC)*, 2005.
71. Amit Sahai and Salil P. Vadhan. A complete problem for statistical zero knowledge. *J. ACM*, 2003.
72. Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: Deniable encryption, and more. *SIAM Journal on Computing*, 2021.
73. Igal Sason. On reverse pinsker inequalities. *CoRR*, 2015.
74. Dominik Scheder. Ppsz is better than you think. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, 2022.
75. Dominik Scheder and John Steinberger. Ppsz for general k-sat and csp—making hertli's analysis simpler and 3-sat faster. *computational complexity*, 2024.
76. T. Schoning. A probabilistic algorithm for k-sat and constraint satisfaction problems. In *40th Annual Symposium on Foundations of Computer Science (Cat. No.99CB37039)*, 1999.
77. J. v. Neumann. Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, 1928.
78. L.G. Valiant and V.V. Vazirani. NP is as easy as detecting unique solutions. *Theoretical Computer Science*, 1986.
79. Virginia Vassilevska Williams. Hardness of Easy Problems: Basing Hardness on Popular Conjectures such as the Strong Exponential Time Hypothesis. In *10th International Symposium on Parameterized and Exact Computation (IPEC 2015)*, 2015.
80. Ilya Volkovich. The final nail in the coffin of statistically-secure obfuscator. *Inf. Process. Lett.*, 2023.
81. J. Watrous. Limits on the power of quantum statistical zero-knowledge. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.*, 2002.
82. Hoeteck Wee and Daniel Wichs. Candidate obfuscation via oblivious lwe sampling. In *Advances in Cryptology – EUROCRYPT 2021*, 2021.
83. A. Winter. Coding theorem and strong converse for quantum channels. *IEEE Transactions on Information Theory*, 1999.
84. Andrew C. Yao. Theory and application of trapdoor functions. In *23rd Annual Symposium on Foundations of Computer Science (sfcs 1982)*, 1982.
85. Chee K. Yap. Some consequences of non-uniform conditions on uniform classes. *Theoretical Computer Science*, 1983.
86. S. Zeijlemaker. The sparsification lemma for CNF satisfiability. Master's thesis, Eindhoven University of Technology, 7 July 2020.