Separating Pseudorandom Codes from Quantum Random Oracles

Nico Döttling[1⋆], Anne Müller[1,2], and Mahesh Sreekumar Rajasree[3]

[1] CISPA Helmholtz Center for Information Security, Saarbrücken, Germany
doettling@cispa.de,
[2] Graduate School of Computer Science, Saarland University, Germany
anne.mueller.sb@gmx.de
[3] University of Warsaw
srmahesh1994@gmail.com

Abstract. Pseudorandom codes (PRCs) combine two seemingly conflicting goals: robustness to noise and computational indistinguishability from random strings. Recent work has shown that classical symmetric-key PRCs over binary alphabets cannot be constructed from minimal cryptographic assumptions, such as one-way functions, via black-box reductions. In this work, we revisit this question in the presence of quantum computation and quantum information.
We first consider pseudorandom codes with classical codewords but quantum algorithms. We prove that symmetric-key pseudorandom codes robust to Bernoulli noise admit no black-box construction from quantum random oracles, and hence from quantum-secure one-way functions. This establishes that the classical black-box barrier for PRCs persists even in the quantum computational setting. Motivated by this separation, we investigate the notion of *quantum pseudorandom codes* (QPRCs), in which codewords are quantum states and security is defined via computational indistinguishability from Haar-random states.
We then present several positive results. Assuming quantum-secure one-way functions, we construct QPRCs that are secure against a single encoding query, allow multiple copies of the challenge state, and tolerate an error rate matching that of any classical uniquely decodable binary error-correcting code against a dephasing or bit-flip channel. We also show that the PRFS with *proof of destruction* due to Behera et al. (TCC 2023) can be extended to a 0-bit QPRC that is secure against multiple encoding queries and tolerates unbounded phase errors but only a constant number of bit-flip errors. Finally, assuming the existence of quantum-secure single-query pseudorandom states (1PRS), we construct single-query single-copy QPRC that tolerate arbitrary Pauli noise up to the decoding radius of a quantum error-correcting code with unique decoding.

Keywords. pseudorandom codes, quantum, black-box separation

1 Introduction

Error correction and cryptography traditionally address fundamentally different challenges. Coding theory is concerned with reliable transmission and storage of information in the presence of noise. By introducing structured redundancy, error-correcting codes ensure that messages can be recovered even when a portion of the transmitted data is corrupted. Cryptography, in contrast, is primarily concerned with secrecy and adversarial security. Its central organizing principle is *pseudorandomness*: the ability to efficiently generate objects that are computationally indistinguishable from truly random ones, thereby concealing any exploitable structure.

Pseudorandom Codes (PRCs), introduced by Christ and Gunn [CG24], reconcile these seemingly contradictory goals. A pseudorandom code is an error correcting code whose codewords

⋆ Funded by the European Union (ERC, LACONIC, 101041207). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

are computationally indistinguishable from uniformly random strings to any adversary lacking a secret key, while still allowing efficient decoding by a legitimate receiver who holds the key. PRCs were originally motivated by applications to watermarking machine learning models, where one seeks encodings that are robust to small perturbations yet hidden from unauthorized observers [ZDR19, KJGR21, Aar22, KGW+23, FGJ+25, ZALW24, KTHL24].

More generally, pseudorandom codes can be viewed as a cryptographic strengthening of steganography [HLvA02]: they enable information to be embedded in malleable content in a way that is simultaneously hidden (pseudorandom) and decodable under corruption. In this sense, they unify steganography and coding theory.

Classical limitations. Despite their conceptual appeal, constructing classical pseudorandom codes has proven to be highly nontrivial. In particular, naive approaches that combine standard error-correcting codes with cryptographic masking techniques fail due to a fundamental synchronization problem between encoding and decoding. Known constructions of PRCs therefore rely on strong cryptographic assumptions that are themselves closely tied to coding theory, such as the McEliece assumption, Learning Parity with Noise, planted XOR-type assumptions, the existence of local weak pseudorandom Functions or the permuted codes conjecture [CG24, MG24, GG24, AAC+24, CGG+25]. Recent work also obtains PRCs from LWE, either with subexponential hardness or together with a planted t-SUM assumption [Lee26].

Recent works have shown that this reliance on coding-theoretic assumptions is not merely an artifact of existing constructions. In particular, two independent works established black-box separations demonstrating that symmetric-key pseudorandom codes over binary alphabets cannot be constructed from one-way functions, or even from stronger cryptographic primitives via black-box constructions [DMSR25, GGW25]. The first work [DMSR25] proves impossibility results for constructions based on a broad class of so-called *local* oracles, whose outputs can be efficiently simulated outside a polynomially bounded set of related queries, while the other [GGW25] establishes separations from a class of *cryptographic* oracles. Together, these results identify a fundamental barrier: any black-box construction of classical pseudorandom codes must necessarily begin with a primitive that already exhibits strong non-locality and inherent noise resilience.

Pseudorandom codes in the quantum setting: Given that these barriers are purely *classical*, a natural and intriguing question arises in the quantum setting: can quantum resources overcome this classical barrier? Quantum information exhibits phenomena such as superposition, entanglement, and the no-cloning principle that have no classical analogue and have already been shown to qualitatively change the landscape of cryptography. Notably, cryptographic primitives that are classically separated from one-way functions admit quantum analogues that can be constructed assuming quantum secure one-way functions [LQWY14, BGH+23]. This motivates the following question.

Can pseudorandom codes be constructed from minimal cryptographic assumptions, such as one-way functions, in the presence of quantum computation?

The landscape of quantum cryptography differs in striking ways from its classical counterpart. In the classical setting, one-way functions are the weakest known cryptographic primitive and characterize the so-called *minicrypt* regime. In contrast, the quantum setting admits an even weaker regime, often referred to as *microcrypt*, which is strictly separated from minicrypt. Primitives in microcrypt can be realized using quantum computational power alone, even in hypothetical worlds where $\mathsf{P} = \mathsf{NP}$ [Kre21, KQST23a]. This qualitative difference motivates the study of *quantum pseudorandom codes* (QPRCs).

Informally, a quantum pseudorandom code encodes classical messages into quantum states that satisfy two competing requirements: (i) robustness against a specified class of quantum noise channels, allowing efficient decoding with a secret key, and (ii) computational indistinguishability

from Haar-random quantum states to any efficient adversary who does not possess the key. A further central question addressed in this work is the following:

Can quantum pseudorandom codes be constructed from minimal cryptographic assumptions, such as one-way functions or even pseudorandom states, in a black-box manner?

More broadly, we ask whether quantumness fundamentally alters the relationship between pseudorandomness and error correction, or whether the classical barriers persist even in the quantum setting. Answering this question sheds light on the power and limitations of quantum cryptographic primitives and clarifies the role of noise resilience in the design of pseudorandom objects.

1.1 Our Results

We begin by studying the most basic setting of symmetric-key pseudorandom codes against a Bernoulli noise channel. In a symmetric-key pseudorandom code, both the encoder Encode and the decoder Decode use the same secret key sk. In a Bernoulli channel with *bias parameter* $\rho \in [0,1]$, each transmitted bit is independently flipped with probability $\frac{1}{2} - \frac{1}{2}\rho$. Equivalently, an error vector $\mathbf{e} \in \{0,1\}^n$ is said to be distributed according to Ber_ρ^n if each bit e_i is independently 1 with probability $\frac{1}{2} - \frac{1}{2}\rho$[4].

Classical codewords and quantum adversaries. We first consider pseudorandom codes with *classical codewords*, but where both the construction and the adversary are quantum. In this setting, we establish a black-box separation from quantum-secure one-way functions.

Theorem 1 (Informal). *For any constant ρ, there is no black-box construction of symmetric-key pseudorandom codes with classical keys and classical codewords, secure against quantum adversaries, that can correct Bernoulli noise with parameter ρ, from quantum random oracles.*

In particular, this result implies a black-box separation between symmetric-key pseudorandom codes and *quantum-secure one-way functions*. We emphasize that our attacker only makes classical oracle queries to mount the attack. Our result shows that pseudorandom codes remain fundamentally resistant to such black-box constructions, even in the presence of quantum computation. See Section 4 for details.

Quantum pseudorandom codes. Quantum pseudorandom codes can be informally understood as pseudorandom quantum states that admit reliable decoding in the presence of noise. From this perspective, the problem of constructing error-tolerant pseudorandom states emerges as a natural and largely unexplored research direction.

Given the inherently noisy nature of both near-term and fault-tolerant quantum computation, several applications of pseudorandom states such as quantum money and quantum signature schemes [JLS18, AQY22, MY22] would substantially benefit from built-in noise tolerance.

Positive results: constructions of QPRC. We provide several positive results demonstrating that quantum pseudorandom codes can be constructed under standard cryptographic assumptions, albeit under restricted security notions.

Our first construction shows that quantum-secure one-way functions already suffice to obtain nontrivial quantum pseudorandom codes.

Theorem 2 (Informal). *Assuming quantum-secure one-way functions exist, there exists a symmetric-key quantum pseudorandom code that is secure against a single encoding query and can tolerate an X-error fraction matching that of any classical binary error-correcting code with unique decoding.*

[4] We follow the convention from the analysis of Boolean functions literature [O'D14], where the Bernoulli distribution is parametrized by its bias rather than the bit-flip probability.

The security notion here allows the adversary to obtain multiple copies of the challenge state, which is either a quantum codeword generated by the QPRC or a Haar-random state, but restricts the adversary to a single encoding query. Despite this restriction, the result is meaningful: any QPRC immediately implies a pseudorandom state generator (PRS), and proving security of PRS is known to be highly nontrivial.

The construction can be made robust against Z-type errors by applying an appropriate Hadamard transformation during encoding and decoding. We refer to Section 5.1 for details.

Non-adaptive constructions. Our construction is based on quantum-secure pseudorandom functions, which are known to exist assuming quantum-secure one-way functions, via black-box reductions [Zha12]. A notable feature of our scheme is that it is *non-adaptive* with respect to the pseudorandom function: the encoder makes only a single quantum query to the PRF. Despite this strong restriction, the construction achieves robustness matching the best known classical error-correction guarantees.

This stands in sharp contrast to the setting of pseudorandom states alone. In particular, it is impossible to expand a PRS by even a single bit in the non-adaptive setting [BEHMV25], which precludes achieving any nontrivial error correction using PRS directly.

Achieving multi-query security. While classical pseudorandom codes tolerating a small fraction of errors can be constructed from one-way functions [GM24, DMSR25], obtaining analogous guarantees in the quantum setting appears substantially more challenging. This gap stems from the intrinsic fragility of pseudorandom states, which unlike classical strings, do not admit flexible combination, reuse, truncation [BEM24] or amplification [BEHMV25] across multiple invocations.

Our next result shows that multi-query security can nevertheless be achieved in a restricted sense. Specifically, we observe that the pseudorandom function like state construction with *proof of destruction* due to Behera *et al.* [BBSS23] can be naturally extended to a multi-query *zero-bit* quantum pseudorandom code. This construction tolerates an unbounded number of Z-errors, while remaining robust against a small number of X-errors.

Theorem 3 (Informal). *Assuming quantum-secure one-way functions exist, there exists a symmetric-key 0-bit quantum pseudorandom code that is secure against multiple encoding queries and can tolerate unbounded Z-errors but a constant number of X-errors.*

We refer to Section 5.2 for details. At first glance, one might wonder whether this result subsumes Theorem 2, since it offers multi-query security and robustness to unbounded Z-errors. However, the two constructions achieve fundamentally different guarantees. The QPRC of Theorem 2, although secure only against a single encoding query and tolerant of only a small fraction of Z-errors, supports *multi-bit* encoding: arbitrary-length classical messages can be encoded and fully recovered during decoding. In contrast, the construction underlying Theorem 3 is inherently zero-bit, offering robust detection of a valid encoding. Consequently, the two results are incomparable and capture distinct trade-offs between query security, error tolerance, and message recovery.

Pauli noise and 1QPRC**.** Finally, we consider a stronger noise model allowing arbitrary Pauli errors, but under a more restricted security notion. Specifically, we study quantum pseudorandom codes secure against a single-query, single-copy adversary. This notion parallels the well-studied concept of 1PRS in the literature [MY22, GJMZ23, CCS25], where the non-triviality stems from the fact that the key length is strictly smaller than the number of qubits in the output.

Theorem 4 (Informal). *Assuming the existence of quantum-secure* 1PRS*, there exists a secure* 1QPRC *that can tolerate a Pauli error fraction matching that of any quantum error-correcting code with unique decoding.*

We refer to Section 5.3 for details.

2 Technical Overview

In this section, we provide a high-level overview of the main technical ideas underlying our separation results and constructions. We begin by recalling the classical impossibility result for pseudorandom codes and then explain how these ideas extend to the quantum setting. Throughout this overview, we focus on intuition and omit many technical details, which are developed formally in the main body of the paper.

2.1 Separating Pseudorandom Codes from Quantum One-way Functions

Recap: Classical impossibility of pseudorandom codes. We first recall the black-box impossibility result for classical pseudorandom codes established in prior work [DMSR25, GGW25]. For concreteness, we focus on the framework and proof strategy of [DMSR25]. In both works, it was shown that symmetric-key pseudorandom codes over binary alphabets cannot be constructed from random oracles via black-box reductions.

At a high level, the argument considers a candidate construction $\mathsf{PRC}^{\mathcal{O}} = (\mathsf{KeyGen}, \mathsf{Encode}^{\mathcal{O}}, \mathsf{Decode}^{\mathcal{O}})$ relative to a random oracle $\mathcal{O}$. The key generation algorithm outputs a uniformly random secret key, while the encoder and decoder make only polynomially many oracle queries. Security is defined against an unbounded distinguisher with polynomially many oracle queries, modeling the random oracle as the sole source of computational hardness.

The core idea of the separation is to show that any such decoder can be transformed into an *oracle-free* decoder with only a small loss in correctness. Once oracle access is eliminated, a simple brute-force distinguisher over the secret-key space suffices to break pseudorandomness. The main technical challenge lies in eliminating oracle queries made by the decoder while preserving correctness under Bernoulli noise.

The proof proceeds by analyzing the *query function* induced by the decoder. When the decoder makes the first oracle query on input $\mathbf{x}$, this query can be described by a function $\mathsf{Quer}_{\mathsf{sk}}(\mathbf{x})$ mapping inputs to oracle queries. Correctness under noise forces this query function to be *noise robust* on codewords. This is because the random oracle itself does not correct errors - two distinct inputs yield independent, uniformly random outputs. Using tools from the Fourier analysis of Boolean functions, in particular the Bonami–Beckner hypercontractivity theorem, one can show that such noise robustness is only possible if the query function concentrates most of its mass on a small number of *heavy* queries. An offline learner can identify these heavy queries and their oracle responses, allowing the oracle query to be compiled into an augmented secret key. This yields an oracle-free decoder and enables a distinguishing attack.

Model in the quantum setting. In the present work, we revisit this separation in the quantum setting. While the secret key and codewords remain classical, the encoder and decoder are quantum algorithms with *quantum access* to a random oracle. The adversary is a computationally unbounded[5] quantum algorithm restricted to making a polynomial number of superposition queries to the random oracle. A superposition query to a random function $H : \{0,1\}^n \to \{0,1\}^m$ is modeled as

$$\sum_i \alpha_i |x_i\rangle|y_i\rangle \;\longrightarrow\; \sum_i \alpha_i |x_i\rangle|y_i \oplus H(x_i)\rangle.$$

Here, $|x\rangle$ is the query register, $|y\rangle$ is the response register, and the function H is applied coherently in superposition.

[5] This models that the oracle is the only source of hardness.

Heavy Queries and Exceptionality. For this overview, we focus on the case where the decoder makes a single query of the above form. For generality, we also allow the query to be a mixed state $\sigma \in \mathcal{D}(Q)$, where Q denotes the oracle query space. We can then define a query function parameterized by the secret key, which takes as input a classical codeword and outputs a quantum query:

$$\mathsf{Quer}_{\mathsf{sk}} : \{0,1\}^n \to \mathcal{D}(Q).$$

Due to the error-correcting properties of pseudorandom codes, we expect this query function to be noise robust. The random oracle itself does not correct errors: two distinct inputs yield independent, uniformly random outputs. Moreover, for decoding to succeed, the encoder and decoder must obtain *overlapping information* about the oracle. If they obtain no common information, then oracle replies are indistinguishable from local randomness, and the oracle could be replaced by local coins.

Error robustness up to t errors implies that there are $\sum_{i=1}^{t} \binom{n}{i}$ possible corruptions of a codeword, and querying all such error patterns is infeasible due to the polynomial query bound. Therefore, correctness under noise suggests that the query function of the decoder must make *similar queries* on related inputs of the form $\mathbf{c} + \mathbf{e}_1$ and $\mathbf{c} + \mathbf{e}_2$, where $\mathbf{c}$ is a codeword.

This observation motivates the use of techniques from the analysis of Boolean functions, as in [DMSR25]. It is known that for a Boolean function to be noise robust, it must be close to constant. Translating this intuition, we define a query q to be ϵ-*heavy* if its *global weight* satisfies $\mathsf{w}(q) \geq \epsilon$, where

$$\mathsf{w}(q) \;=\; \mathbb{E}_{\mathbf{u}}\Big[\mathsf{Tr}(|q\rangle\langle q|\,\mathsf{Quer}(\mathbf{u}))\Big],$$

and the expectation is over uniformly random $\mathbf{u} \in \{0,1\}^n$. Intuitively, the global weight measures how likely the query q is to appear when the decoder is run on a uniformly random input.

Analogously, a query q is called δ-*locally heavy* with respect to a codeword $\mathbf{c}$ if its *local weight* satisfies $\ell_{\mathbf{c}}(q) \geq \delta$, where

$$\ell_{\mathbf{c}}(q) = \mathbb{E}_{\mathbf{e}}\Big[\mathrm{Tr}(|q\rangle\langle q|\mathsf{Quer}(\mathbf{c}+\mathbf{e}))\Big]$$

and the expectation is over $\mathbf{e}$ drawn from a Bernoulli noise distribution. The local weight captures how likely the query q is to occur when the decoder is run on noisy perturbations of $\mathbf{c}$.

A query is called (ϵ, δ)-*exceptional* if it is δ-locally heavy but not ϵ-globally heavy. This induces a notion of exceptionality for strings/codewords: a string $\mathbf{c}$ is exceptional if the query function makes an exceptional query on input $\mathbf{c}$. Importantly, although the query function arises from a quantum computation, both the global and local weight functions are classical. Therefore, we can rely on the following lemma from [DMSR25], which shows that exceptional strings are unlikely.

Lemma 1 (Informal).

$$\Pr[\mathbf{u} \text{ is exceptional}] \;\leq\; \frac{\epsilon^c}{\delta^2}.$$

where the probability is over random strings $\mathbf{u}$*. Here, c depends only on the bias of the Bernoulli probability.*

In fact, using the pseudorandomness of the underlying PRC, [DMSR25] showed that this exceptionality property also holds for a codeword generated by the encoder. And, this is also true in our setting.

Lemma 2 (Informal). *For random secret key* sk*, bit* b*,*

$$\Pr[\mathbf{c} \text{ is exceptional}] \leq 3\epsilon^c/\delta^2 = p^*.$$

where the probability is over $\mathbf{c} \leftarrow \mathsf{Encode}^{\mathcal{O}}(\mathsf{sk}, b)$*. Here, c depends only on the bias of the Bernoulli probability.*

Oracle-free decoder. To eliminate the decoder's oracle query, we define a learning algorithm $\mathsf{LearnQuery}^{\mathcal{O}}$ that identifies all ϵ-heavy queries. The (unbounded) learner runs the decoder on all inputs, measures the resulting oracle queries, and estimates their distribution. For queries that satisfy the ϵ-heavy criterion, $\mathsf{LearnQuery}^{\mathcal{O}}$ queries the random oracle $\mathcal{O}$ and learns the corresponding responses. Given this set $\mathcal{L}$ of heavy query–response pairs, we define an oracle-free decoder that simulates the oracle query as follows: if the query appears in $\mathcal{L}$, return the stored response; otherwise, return a uniformly random string. This can be simulated using the compressed oracle technique [Zha19].

Correctness of oracle-free decoder. We bound the correctness error of the oracle-free decoder by bounding the trace distance between the post-query states produced by the real oracle and by the simulation. To this end, we assume for the sake of simplicity that the query is a pure state, i.e., $\sum_q \alpha_q |q\rangle|0\rangle$ and analyze three cases:

1. q **is globally heavy:** simulation is perfect because heavy queries are learned directly from the real oracle and stored in $\mathcal{L}$.
2. q **is not globally heavy and does not intersect with queries made by** KeyGen **or** Encode**:** These queries are simulated by a fresh compressed oracle which samples uniformly random strings for previously unseen queries. Recall that the compressed oracle is perfectly indistinguishable from a truly random oracle. Concretely, consider the following scenario. After KeyGen and Encode have made their queries the compressed-oracle database is a state $D_{K,E}$ of query-response pairs whose size is bounded by the number of oracle queries made. Queries outside of $D_{K,E}$ are uniformly random. Equivalently, the codeword produced by the encoder contains information about the oracle only through the set $D_{K,E}$. Using the compressed oracle viewpoint, we conclude that all queries outside $D_{K,E}$ are answered with fresh uniform randomness and are therefore simulated perfectly.
3. q **is not globally heavy and intersects with queries made by** KeyGen **or** Encode

The third case, where q is not a heavy query but belongs to $D_{K,E}$, is potentially problematic. To address this case, we further subdivide it into queries that are locally heavy and those that are not. Queries that are locally heavy but not globally heavy are precisely the exceptional queries, whose probability is bounded by Lemma 2. This forms the first source of error in the correctness analysis.

For queries that are neither globally nor locally heavy, their total amplitude is small. Therefore even though the answer is not simulated correctly their error contribution is small. The relevant overlap is the weight of the joint query-database state on basis pairs in which the query lies outside $\mathcal{L}$ and is present in the database. Conditional independence gives

$$p_{\text{hit}} = \sum_{q \notin \mathcal{L}} \ell_{\mathbf{c}}(q)\, \mathsf{Tr}(\Pi_q D_{K,E}).$$

where (Π_q) projects on queries q that are not in the set $\mathcal{L}$. If $\mathbf{c}$ is non-exceptional, every query outside $\mathcal{L}$ has local weight less than δ. Hence

$$p_{\text{hit}} \leq \delta \sum_{q \notin \mathcal{L}} |D_{K,E}| \leq |D_{K,E}|\delta.$$

The database has at most $|D_{K,E}|$ occupied queries outside the learned set, and the decoder places weight at most δ on each one. The formal argument in (5) expresses this overlap as the expectation of a projector on the joint state.

Finally we argue that the state of the decoder using the simulation is close to the state of the decoder using the real oracle. This can be done by arguing that the simulation (on non-exceptional codewords) is correct outside of the database overlap which we just calculated to be $p_{hit} \leq |D_{K,E}|\delta$. Define the projector Π_S on queries outside of $\mathcal{L}$, Then $\mathsf{Tr}(\Pi_S \rho) \leq |D_{K,E}\delta|$. Since the simulation is perfect on the orthogonal subspace, we can apply the following standard fact (see Lemma 6).

Lemma 3. *Let A, B be finite-dimensional quantum registers, let Π be an orthogonal projector on A, and let $P = \mathbb{I}_A - \Pi$. Let Φ, Ψ be completely positive trace-preserving (CPTP) maps from A to B such that*

$$\Phi(PXP) = \Psi(PXP) \qquad \text{for every operator } X \text{ on } A. \tag{1}$$

Equivalently, the channels agree on every density operator supported in $\operatorname{im} P$*, including superpositions within that subspace. Then, for every density operator ρ_A and $p \in [0,1]$ with $\mathsf{Tr}(\Pi\rho_A) \leq p$,*

$$\mathsf{TD}\big(\Phi(\rho_A), \Psi(\rho_A)\big) \leq 2\sqrt{p}.$$

Applying this lemma with Φ corresponding to the real oracle and Ψ to the simulation, we obtain a trace-distance bound of $2\sqrt{|D_{K,E}|\delta}$. Together with the bound on a codeword being exceptional, the decoding error of the oracle-free decoder satisfies

$$\Pr[\text{Decoding error}] \;\leq\; p^* + 2\sqrt{|D_{K,E}|\delta}.$$

Compiling out multiple queries. To eliminate multiple oracle queries, we apply the above argument iteratively. Care is required to ensure that oracle answers remain consistent and that the decoder cannot detect the transition from simulated to real oracle access. By a union bound, the decoding error of the fully simulated decoder is

$$\Pr[\text{Decoding error}] \;\leq\; d\big(p^* + 2\sqrt{|D_{K,E}|\delta}\big),$$

where d is the number of oracle queries made by the decoder. By choosing parameters appropriately, we obtain an inverse-polynomial error probability while ensuring that the query learner makes only a polynomial number of oracle queries.

2.2 Quantum Pseudorandom Codes: Definitions and Constructions

Quantum pseudorandom codes. Quantum pseudorandom codes (qPRCs) are the quantum analogue of pseudorandom error-correcting codes [CG24]. A qPRC is a keyed family of encodings in which *codewords are quantum states* but decoding is still possible under a prescribed noise process. Formally, a secret-key qPRC consists of QPT algorithms $(\mathsf{KeyGen}, \mathsf{Encode}, \mathsf{Decode})$. Setup outputs a classical secret key sk; encoding maps a message m (and fresh randomness) to an $n(\lambda)$-qubit pure state; and decoding maps a possibly corrupted state to either m or $\perp$. Correctness requires decoding to succeed with probability $1 - \mathrm{negl}(\lambda)$ after the application of a fixed quantum channel $\mathcal{E}$. Security requires that access to the real encoding oracle (fresh randomness per query[6]) is computationally indistinguishable from access to an oracle that outputs independent Haar-random $n(\lambda)$-qubit pure states. The notion of Quantum Pseudorandom Codes has also been studied in the recent work of [HY26], albeit their definitions differ from ours; see Section 2.6 for a comparison to our work.

Variants. We consider several variants tailored to the regimes arising in our results. A *zero-bit qPRC* is the degenerate case where only a single message is encodable. Here decoding becomes a predicate and we additionally require *soundness*, namely that a uniformly keyed decoder accepts an arbitrary state only with negligible probability. We also distinguish *multi-query* pseudorandomness (the default, where the adversary may query the encoding oracle polynomially many times) from *single-query* pseudorandomness (the adversary receives only one encoding but still multiple copies). Finally, we define an even weaker notion, denoted 1QPRC, in which the adversary is restricted to a single query and receives only a single codeword, mirroring the notion of single-copy pseudorandom states (1PRS).

[6] Note that for the same message, there can be multiple quantum codewords. For this reason, we explicitly need randomness in the encoding algorithm to efficiently reproduce the same quantum states

2.3 Construction I: Single-query qPRC for bounded Z-errors

High-level idea. Our first construction achieves *single-query* pseudorandomness (but allows the adversary to obtain multiple copies of the single challenge codeword) and is *perfectly robust* against a bounded fraction of Z-errors. The construction combines (i) a classical uniquely decodable binary error-correcting code ECC and (ii) a quantum-secure PRF PRF with output in $\{0,1\}$. At a high level, we start from the standard *binary-phase* pseudorandom states of the form

$$|\psi_{\mathsf{prf.k}}\rangle = \frac{1}{\sqrt{2^n}} \sum_{y\in\{0,1\}^n} (-1)^{\mathsf{PRF.Eval}(\mathsf{prf.k},y)} |y\rangle,$$

and encode a classical message by applying a Z-mask derived from an error-correcting encoding of the message. This encoding-by-phase-masking approach closely mirrors the blueprint used in Coladangelo's construction of quantum trapdoor functions [Col23], where classical information is embedded into a pseudorandom phase state via Pauli phase operators.

Encoding and decoding. Given $m \in \{0,1\}^d$, compute the classical codeword $x \leftarrow \mathsf{ECC.Encode}(\mathsf{ecc.pp}, m) \in \{0,1\}^n$ and output

$$|\psi_m\rangle = Z^x |\psi_{\mathsf{prf.k}}\rangle = \frac{1}{\sqrt{2^n}} \sum_{y\in\{0,1\}^n} (-1)^{x\cdot y+\mathsf{PRF.Eval}(\mathsf{prf.k},y)} |y\rangle.$$

Under Z-noise, the received state has the effect of replacing x by $x+e$ in the phase. The decoder cancels the PRF-dependent phase using the keyed unitary $G_{\mathsf{prf.k}}$, applies $H^{\otimes n}$, measures to recover the classical string $x+e$, and finally applies ECC.Decode to recover m. Thus correctness reduces to classical decoding: as long as $\mathsf{Ham}(e)$ is within the decoding radius of ECC, recovery is perfect.

Why single-query security holds. Security is proved by a standard hybrid argument. Replacing PRF with a truly random function yields a *random phase state* as in [JLS18], which is known to be single-copy (and hence single-query) indistinguishable from Haar-random states [BS19]. Therefore, a single encoding oracle call (even with many copies of that one output state) is computationally indistinguishable from Haar. We stress that the construction is *not* secure under multi-query access (even on distinct messages), since in the random-function hybrid all message encodings lie in a structured Z-orbit of a single underlying phase state; this additional structure can be exploited information-theoretically once many distinct message states are available.

Handling X-errors. By conjugating the scheme with Hadamards, we obtain an analogous construction robust against bounded X-errors, using $X = HZH$.

2.4 Construction II: Multi-query 0-bit QPRC from pseudorandom states with proof of destruction

High-level idea. Our second construction leverages the pseudorandom states with proof of destruction framework of Behera *et al.* [BBSS23] to obtain a *multi-query* zero-bit quantum pseudorandom code. The key observation is that these states are inherently robust against arbitrary Z-errors, while tolerating a small number of X-errors via classical post-processing. Unlike the previous construction, robustness here is achieved without explicit quantum error correction. Instead, it relies on a structural property of the encoding and a keyed permutation.

Construction. The construction is parameterized by a post-quantum secure pseudorandom permutation (PRP). To encode, the encoder prepares a uniform superposition over strings permuted by the PRP and applies a key-dependent phase determined by the PRP. The resulting quantum state is supported on computational basis states whose locations are also determined by the PRP, embedding a hidden classical structure as follows.

$$|\psi_{k,s}\rangle := \frac{1}{2^{\lambda/2}} \sum_{y\in\{0,1\}^\lambda} \omega_{2^{5\lambda}}^{\mathsf{PRP.Eval}(k,\, y\|s\|1^{3\lambda})}\, |\mathsf{PRP.Eval}(k,\, y\|s\|0^{3\lambda})\rangle.$$

Decoding proceeds by measuring in the computational basis and checking whether the outcome (up to a small number of bit flips) lies in this structured subset, which can be efficiently verified using the secret key by inverting and checking whether the last bits are zero.

Robustness and soundness. Robustness against Z-errors is immediate, as such errors only affect phases and are erased by measurement in the computational basis. A small number of X-errors correspond to bit flips in the measurement outcome and are handled by explicitly checking all nearby candidates. The soundness is established by assuming nonuniform PRP security.

Pseudorandomness (multi-query). Pseudorandomness follows directly from the security guarantees of [BBSS23], which show that polynomially many copies of the encoded states are computationally indistinguishable from Haar-random states. As a result, the construction yields a multi-query zero-bit QPRC that is robust to unbounded Z-errors and a constant number of X-errors, under standard post-quantum assumptions.

Remark 1. It is important to note that Construction 1 remains non-trivial. While the above scheme yields only a zero-bit QPRC, Construction 1 allows for full recovery of the original message.

2.5 Construction III: 1QPRC from 1PRS via quantum error correction

High-level idea. Our third construction targets Pauli noise and uses a quantum error-correcting code (QECC) to protect a *single-copy* pseudorandom state. The security notion is correspondingly weaker: we construct a *zero-bit* 1QPRC, where the adversary receives only a single codeword and makes only a single distinguishing attempt, analogous to 1PRS.

Construction. Let $\{|\psi_k\rangle\}_{k\in\{0,1\}^\lambda}$ be a 1PRS with output length $m(\lambda) > \lambda$, and let V be the Clifford encoder of an $[[N, m+\lambda, d]]$ stabilizer/CSS code. The secret key is $\mathsf{sk} = (k, t, s)$ where $k, t \leftarrow \{0,1\}^\lambda$ and $s \leftarrow \{0,1\}^{N-m-\lambda}$. To encode, prepare $|\psi_k\rangle \otimes |t\rangle \otimes |s\rangle$ and output

$$|\Psi_{k,t,s}\rangle := V\big(|\psi_k\rangle \otimes |t\rangle \otimes |s\rangle\big).$$

To decode, the key (t, s) specifies a Pauli correction that (by Clifford normalisation) can be applied to map the codeword back to an encoding with fixed ancillas; then the QECC decoder recovers $(|\psi_k\rangle, t)$ as long as the Pauli error weight is within the unique-decoding radius. The decoder finally checks whether the recovered tag equals the secret t, outputting 1 on a match and 0 otherwise. Note that the size of the secret key $|\mathsf{sk}| = N - m + \lambda < N$ is strictly less than the size of the codeword (which is N) because $m > \lambda$.

Robustness and soundness. Robustness follows directly from the QECC: after applying the key-dependent Pauli correction, any Pauli error of weight at most $\lfloor (d-1)/2 \rfloor$ is corrected by the decoder. Soundness follows because the tag t is uniformly random and is only compared at the final step; hence for any fixed input state ρ, the acceptance probability is at most $2^{-\lambda}$.

Pseudorandomness (single-copy). Security is shown by a two-step hybrid. First replace $|\psi_k\rangle$ with a Haar-random m-qubit state using the 1PRS assumption. Then observe that averaging over the random classical pads (t, s) yields the maximally mixed state on N qubits, and for a single-copy distinguisher this matches the Haar average, implying indistinguishability. As a consequence, we obtain a zero-bit 1QPRC tolerating Pauli noise up to the decoding radius of the underlying QECC.

2.6 Related Work

Black-box separations. The study of black-box separations plays a central role in understanding the relative power of cryptographic assumptions. The seminal work of Impagliazzo and Rudich [IR89] showed that key-agreement protocols cannot be constructed from one-way permutations via black-box reductions. This result established a fundamental barrier to deriving public-key cryptography from weaker primitives.

Rudich subsequently strengthened this line of work in his thesis [Rud88] by proving a black-box separation between one-way functions (OWFs) and one-way permutations. Together with the Impagliazzo–Rudich separation, this implies that OWFs are insufficient for black-box constructions of key agreement and related public-key primitives. Later, Rudich [Rud91] extended these results by separating trapdoor functions from key agreement in the black-box model.

Further separations refined the cryptographic hierarchy. Simon [Sim98] established a black-box separation between one-way permutations and collision-resistant hash functions. The original Impagliazzo–Rudich separation was revisited by Barak and Mahmoody [BMG09], who obtained a more efficient black-box attack, and by Brakerski et al. [BKSY11], who presented a simpler and more modular proof technique.

Additional impossibility results include the work of Matsuda and Matsuura [MM11], who ruled out fully black-box constructions of one-way permutations from length-increasing injective one-way functions, even when the expansion is only one bit. Mahmoody and Pass [MP12] showed that non-interactive commitment schemes cannot be constructed from OWFs in a fully black-box manner. More recently, Döttling and Mour [DM24] proved that any fully black-box construction of correlation-intractable hash functions for t-wise independent relations from collision-resistant hash functions or one-way permutations must make at least $\Omega(t)$ calls to the underlying primitive.

In the quantum setting, significant effort has been devoted to understanding how these classical separations interact with quantum computation. In particular, several works study the relationship between post-quantum one-way functions and post-quantum key-agreement protocols [ACC+22, BGV+23, LLLL24, LLLL25b, LLLL25a, LLLL26, KNVY26], showing that many classical black-box barriers persist even against quantum adversaries.

Constructions of pseudorandom codes. Christ and Gunn [CG24] introduced pseudorandom codes and defined correctness with respect to random errors. In subsequent work [AAC+25], they showed security against adversarially chosen errors. Both secret-key and public-key variants of PRCs are studied, where in the public-key setting encoding is public while decoding requires a secret key.

Golowich and Moitra [MG24] constructed a zero-bit PRC under the assumption of local weak pseudorandom functions, which are implied by the hardness of learning $\log(n)$-juntas. Relatedly, learning junta distributions has been shown to be computationally equivalent to learning parity functions with noise [Ber25].

Additional constructions appear in [GG24], which derive public-key PRCs from the planted hyperloop assumption and pseudorandom generators, as well as statistically secure PRCs against space-bounded adversaries. They also present a construction based solely on one-way functions, though it tolerates only an $o(1)$ fraction of errors from a non-adaptive noise channel.

Christ et al. [CGG+25] construct binary PRCs with robustness to worst-case edits and to adversaries that know the decoding key, under the permuted codes conjecture. Their construction also

targets subexponential distinguishing security. Huang et al. [HLMZ26] give generic transformations from Hamming-robust PRCs to edit-robust PRCs. Their instantiations attain rates approaching $1/2$ over the binary alphabet and 1 over sufficiently large constant alphabets for sublinear polynomial edit budgets; over polynomial-size alphabets, they obtain insertion–deletion tradeoffs approaching the Singleton bound.

Recent work also broadens the assumptions and security notions available for classical PRCs. Lee [Lee26] constructs PRCs from subexponential LWE, or from standard LWE together with a planted t-SUM assumption, using rounded LWE samples and Fourier analysis to establish robustness. Chen, Mao, and Yi [CMY26] introduce a Dense-Planted LPN construction and achieve CCA1 security in the standard model. Döttling et al. [DJK+26] obtain full CCA-secure public-key PRCs in the standard model with constant rate and tolerance of a constant fraction of adversarial errors, and give a strongly adaptively robust construction under a Noisy McEliece assumption.

On concrete security, Wang et al. [WWC+26] analyze practical LDPC-based PRCs, giving distinguishing and robustness attacks together with parameter and key-generation countermeasures. Their results concern concrete instantiations and parameter choices.

Earlier work by Katz et al. [KKRT16] studied a weaker notion of pseudorandom codes that does not require efficient decoding or full pseudorandomness. In their setting, codewords need only be far apart, and such constructions can be obtained from one-way functions.

Two recent independent works [DMSR25, GGW25] establish black-box impossibility results for pseudorandom codes in the random oracle model. Both rely on analytic techniques such as hypercontractivity. One work proves impossibility for a broad class of *local* oracles, while the other establishes separation from *cryptographic oracles* [LMW25], which are stateless efficient algorithms with access to a random oracle.

For a survey of the overall landscape of PRC definitions, transformations, assumptions, and applications, see Christ [Chr26].

Microcrypt. The quantum cryptographic landscape admits primitives that have no classical analogue and may exist even in a world without one-way functions. This regime, often referred to as *microcrypt*, is strictly weaker than the classical *minicrypt* regime and arises uniquely in the presence of quantum information.

Pseudorandom states (PRS), introduced by Ji, Liu, and Song [JLS18], are a canonical example of microcrypt primitives. PRS are efficiently generable quantum states that are computationally indistinguishable from Haar-random states, even given polynomially many copies. Kretschmer [Kre21] exhibited a quantum oracle relative to which PRS exist but quantum-secure classical one-way functions do not, ruling out relativizing constructions of the latter from PRS. Subsequent works developed additional constructions and variants of PRS [BS19, BS20, BBSS23] and studied extensions such as pseudorandom unitaries and channels [AGQY22, ABGL25, GLMY25].

More recently, it was shown that PRS can be separated from one-way functions even relative to a *classical* oracle [KQST23b], strengthening the evidence that PRS inhabit a cryptographic regime incomparable with minicrypt. Several works have explored the limitations of PRS in black-box constructions, showing that they are difficult to shrink [BEM24] or expand [BEHMV25].

The notion of *single-copy* or *single-query* pseudorandom states (1PRS) was introduced by Morimae and Yamakawa [MY22] and subsequently studied in greater depth in [GJMZ23, CCS25]. More recently, Coladangelo and Mutreja [CM24] established a black-box separation between digital signature schemes and pseudorandom states.

The work of [BMB+25] investigates noise tolerant EFI pairs and noise tolerant PRS. Their notion of error tolerance is different from our notion since they do not require decodability under noise but only indistinguishability of noisy PRS from noiseless Haar random states.

Concurrent Work on Quantum Pseudorandom Codes. The concurrent and independent work of Hsieh and Yamada [HY26] studies quantum pseudorandom error-correcting codes. Their construction is

more general than ours in several aspects. They allow for quantum messages to be encoded while we define QPRC to encode classical messages and they achieve strong robustness guarantees against local error channels. Assuming LPN hardness against subexponential quantum algorithms, they construct two variants of QPRCs. First, they construct a constant-rate pseudorandom isometric error-correcting code, whose encoding oracle is indistinguishable from a fixed Haar-random isometry, correcting arbitrary noise supported on $o(N \log\log N/\log N)$ physical qubits, where N is the code length. Under the same assumption, they also construct codes where pseudorandomness is defined relative to the maximally mixed state. The code corrects noise on a constant fraction of physical qubits. Their robustness is formulated in diamond distance after averaging over the secret key. Our definitions instead allow for requesting multiple copies of an encoding of a classical message which means the oracle returns multiple copies of the same Haar-random states. Thus the ideal objects, message spaces, and access models differ. Their results obtain broader noise correction under a structured LPN assumption, whereas our constructions explore what quantum-secure one-way functions achieve under restricted query and copy guarantees. The results in our paper were previously published as part of a PhD thesis [Mü26].

3 Preliminaries

This section collects notation, general definitions, and technical facts that will be used throughout the paper.

3.1 Notation and Basic Conventions

The security parameter is denoted by $\lambda \in \mathbb{N}$. All Hilbert spaces are finite-dimensional. For $t \geq 1$, let $\mathcal{H}_t := (\mathbb{C}^2)^{\otimes t}$ and $\mathcal{D}(\mathcal{H}_t)$ be the set of density operators on $\mathcal{H}_t$. For a pure state $|\psi\rangle$, we write $\psi := |\psi\rangle\langle\psi|$ for its density matrix. The maximally mixed state on a d-dimensional Hilbert space is denoted by $\frac{\mathbb{I}_d}{d}$. We denote by $\|\sigma\|_1 = \mathsf{Tr}[\sqrt{\sigma^\dagger\sigma}]$ the trace norm.

We denote $\mathsf{Ham}(\mathbf{x})$ as the Hamming weight of $\mathbf{x}$, i.e., the number of '1' in the binary string $\mathbf{x}$. $\mathsf{Ham}(\mathbf{x},\mathbf{y}) := \mathsf{Ham}(\mathbf{x}-\mathbf{y})$ denotes Hamming distance. Addition of binary words is coordinatewise modulo two.

Definition 1 (Bernoulli distribution). *A binary random variable $x \in \{0,1\}$ is said to follow a **Bernoulli distribution** with bias parameter $\gamma \in [0,1]$, denoted as Ber_γ, if $\Pr[x=1] = \frac{1-\gamma}{2}$. We say that $\mathbf{x} \in \{0,1\}^n$ follows Ber_γ^n, if each component x_i of $\mathbf{x}$ independently follows Ber_γ.*

Claim 1. Let $r, n \in \mathbb{N}$ such that $n \geq 1$ and $r \leq n/2$. Then, for $\gamma = 1 - \frac{r}{n}$,

$$\Pr_{\mathbf{x}\leftarrow\mathsf{Ber}_\gamma^n}[\mathsf{Ham}(\mathbf{x}) \geq r] \leq e^{-r/6}$$

Proof. We have $\Pr_{\mathbf{x}\leftarrow\mathsf{Ber}_\gamma^n}[x_i = 1] = \frac{r}{2n}$. Thus $X := \mathsf{Ham}(\mathbf{x})$ has mean $\mu = r/2$. For $r > 0$, the multiplicative Chernoff bound gives $\Pr[X \geq 2\mu] \leq \exp(-\mu/3) = \exp(-r/6)$.

3.2 Trace Distance

For an operator X, we write $\|X\|_1$ for its trace norm. For two quantum states ρ, σ, the *trace distance* is defined as

$$\mathsf{TD}(\rho,\sigma) := \tfrac{1}{2}\|\rho-\sigma\|_1 .$$

For pure states, the trace distance can be defined in terms of the inner product $\mathsf{TD}(|\phi\rangle,|\psi\rangle) = \sqrt{1-|\langle\phi|\psi\rangle|^2}$.

Lemma 4 (Basic trace-distance inequalities). *Let ρ, σ be quantum states on the same Hilbert space, and let τ be any quantum state. Then, tensoring with a fixed state does not increase distance, i.e.,*

$$\|\rho \otimes \tau - \sigma \otimes \tau\|_1 \;=\; \|\rho - \sigma\|_1 .$$

This follows from $\|A \otimes B\|_1 = \|A\|_1 \|B\|_1$ and $\|\tau\|_1 = 1$.

Lemma 5 (Data processing / Contractivity). *For any CPTP map Φ and any density operators ρ, σ,*

$$\mathsf{TD}(\Phi(\rho), \Phi(\sigma)) \leq \mathsf{TD}(\rho, \sigma).$$

Lemma 6 (Channel agreement outside a low-weight subspace). *Let A, B be finite-dimensional quantum registers, let Π be an orthogonal projector on A, and let $P = \mathbb{I}_A - \Pi$. Let Φ, Ψ be completely positive trace-preserving (CPTP) maps from A to B such that*

$$\Phi(PXP) = \Psi(PXP) \qquad \text{for every operator } X \text{ on } A. \tag{2}$$

Equivalently, the channels agree on every density operator supported in $\operatorname{im} P$, including superpositions within that subspace. Then, for every density operator ρ_A and $p \in [0,1]$ with $\mathsf{Tr}(\Pi\rho_A) \leq p$,

$$\mathsf{TD}\big(\Phi(\rho_A), \Psi(\rho_A)\big) \leq 2\sqrt{p}.$$

More generally, for any reference register R and joint density operator ρ_{AR} with $\mathsf{Tr}[(\Pi \otimes \mathbb{I}_R)\rho_{AR}] \leq p$,

$$\mathsf{TD}\big((\Phi \otimes \mathrm{id}_R)(\rho_{AR}), (\Psi \otimes \mathrm{id}_R)(\rho_{AR})\big) \leq 2\sqrt{p},$$

where id_R is the identity channel on R.

Proof. We prove the statement with a reference register. Set $q := \mathsf{Tr}[(\Pi \otimes \mathbb{I}_R)\rho_{AR}] \leq p$. If $q = 1$, the result follows from $\mathsf{TD} \leq 1$. For $q < 1$, define the normalized projected state

$$\rho^0_{AR} := \frac{(P \otimes \mathbb{I}_R)\rho_{AR}(P \otimes \mathbb{I}_R)}{1 - q}.$$

Let $|\phi\rangle_{ARE}$ purify ρ_{AR}, and let

$$|\phi_0\rangle_{ARE} := \frac{(P \otimes \mathbb{I}_{RE})|\phi\rangle_{ARE}}{\sqrt{1-q}}.$$

This vector purifies ρ^0_{AR} and satisfies $|\langle\phi|\phi_0\rangle| = \sqrt{1-q}$. The pure-state formula for trace distance and contractivity under the partial trace therefore give

$$\mathsf{TD}(\rho_{AR}, \rho^0_{AR}) \leq \mathsf{TD}(|\phi\rangle, |\phi_0\rangle) = \sqrt{q}.$$

Write $\widehat{\Phi} := \Phi \otimes \mathrm{id}_R$ and $\widehat{\Psi} := \Psi \otimes \mathrm{id}_R$. Expanding ρ^0_{AR} in matrix blocks on R and applying (2) to each block shows that $\widehat{\Phi}(\rho^0_{AR}) = \widehat{\Psi}(\rho^0_{AR})$. The triangle inequality and Lemma 5 now imply

$$\begin{aligned}\mathsf{TD}\big(\widehat{\Phi}(\rho_{AR}), \widehat{\Psi}(\rho_{AR})\big) &\leq \mathsf{TD}\big(\widehat{\Phi}(\rho_{AR}), \widehat{\Phi}(\rho^0_{AR})\big) + \mathsf{TD}\big(\widehat{\Psi}(\rho^0_{AR}), \widehat{\Psi}(\rho_{AR})\big) \\ &\leq 2\mathsf{TD}(\rho_{AR}, \rho^0_{AR}) \leq 2\sqrt{q} \leq 2\sqrt{p}.\end{aligned}$$

3.3 Pseudorandom Codes

Definition 2 (Black-box Construction of Pseudorandom Codes).
A construction of a Pseudorandom Code $\mathsf{PRC} = (\mathsf{KeyGen}, \mathsf{Encode}, \mathsf{Decode})$ *from a family of sets of oracles* $\{\overline{\mathcal{O}}_\lambda\}$ *is* black-box *if the following properties hold for any* $\mathcal{O}$ *from the family* $\overline{\mathcal{O}}_\lambda$*:*

- ***Construction Efficiency****: The algorithms* $\mathsf{KeyGen}^{\mathcal{O}}$, $\mathsf{Encode}^{\mathcal{O}}$, *and* $\mathsf{Decode}^{\mathcal{O}}$ *run in time polynomial in* $\lambda \in \mathbb{N}$ *(and input lengths) with inputs and outputs being classical and make at most* $\mathsf{poly}(\lambda)$ *superposition queries to* $\mathcal{O}$*.*
- ***Correctness****: For any* $\lambda \in \mathbb{N}$*, message* $\mathsf{m} \in \{0,1\}^{d(\lambda)}$*,*

$$\Pr\left[\mathsf{Decode}^{\mathcal{O}}(\mathsf{sk}, \mathbf{c}+\mathbf{e}) \neq \mathsf{m} \;\middle|\; \begin{array}{c} \mathsf{sk} \leftarrow \mathsf{KeyGen}^{\mathcal{O}}(1^\lambda) \\ \mathbf{c} \leftarrow \mathsf{Encode}^{\mathcal{O}}(\mathsf{sk}, \mathsf{m}) \\ \mathbf{e} \leftarrow \mathsf{Ber}_\gamma^n \end{array}\right] \leq \kappa(\lambda).$$

  where the probability includes all internal randomness and measurement outcomes of KeyGen, Encode, *and* Decode*, as well as the independent error vector* $\mathbf{e}$*.*
- ***Security****: For any (unbounded) quantum adversary* $\mathcal{A}$ *that makes polynomially many queries, there exists a negligible function* $\mathrm{negl}(\cdot)$ *such that for all* $\lambda \in \mathbb{N}$*,*

$$\left| \Pr_{\substack{\mathcal{O}\leftarrow\overline{\mathcal{O}}_\lambda \\ \mathsf{sk}\leftarrow\mathsf{KeyGen}^{\mathcal{O}}(1^\lambda)}} \left[\mathcal{A}^{\mathcal{O}(\cdot),\mathsf{O}_{\mathsf{sk}}(\cdot)}(1^\lambda)=1\right] - \Pr_{\substack{\mathcal{O}\leftarrow\overline{\mathcal{O}}_\lambda \\ U}} \left[\mathcal{A}^{\mathcal{O}(\cdot),U(1^\lambda,\cdot)}(1^\lambda)=1\right] \right| \leq \mathrm{negl}(\lambda),$$

  where $\mathcal{O}$ *is quantum accessible but* O_{sk} *and* U *are classically accessible oracles* [7] *with the following description.* O_{sk} *on any (even previously queried) input* m*, samples a random* $r \in \{0,1\}^{r'(\lambda)}$ *and outputs* $\mathsf{Encode}^{\mathcal{O}}(\mathsf{sk}, \mathsf{m}; r)$ *and* U *denotes an oracle that on any (even previously queried) input returns a fresh uniform string in* $\{0,1\}^{n(\lambda)}$ *.*

Definition 3 (Black-box Separation). *We say a PRC is black-box separated from a family of oracle sets if for any black-box construction of a PRC from* $\{\overline{\mathcal{O}}_\lambda\}$ *there exists an (unbounded) adversary* $\mathcal{A}$ *which makes only polynomially many oracle queries, and has non-negligible distinguishing advantage in the pseudorandomness experiment.*

Remark 2. Note that robustness is defined with respect to the Bernoulli distribution. In the context of error-correcting codes, one typically specifies a decoding radius α such that any word of the form $\mathbf{c}+\mathbf{e}$, where $\mathsf{Ham}(\mathbf{e}) < \alpha n$, can be successfully decoded. By setting $\gamma = 1-\alpha$, we ensure that an error vector $\mathbf{e} \leftarrow \mathsf{Ber}_\gamma^n$ satisfies $\mathsf{Ham}(\mathbf{e}) < \alpha n$ with high probability, provided α is a constant. This follows from Claim 1 and the fact that for pseudorandomness, $n = \omega(\log \lambda)$.

3.4 Pseudorandom Function

A pseudorandom function $\mathsf{PRF} = (\mathsf{KeyGen}, \mathsf{Eval})$ consists of PPT algorithms
$\mathsf{KeyGen} : \{0,1\}^\lambda \to \{0,1\}^{s(\lambda)}$ and $\mathsf{Eval} : \{0,1\}^{s(\lambda)} \times \{0,1\}^{n(\lambda)} \to \{0,1\}^{m(\lambda)}$ where s, n, m are polynomials. Also, for any PPT (QPT) adversary $\mathcal{A}$, there exists a negligible function $\mathrm{negl}(\cdot)$ such that for any $\lambda \in \mathbb{N}$, the following advantage function $\mathrm{Adv}^{\mathrm{PRF}}(\lambda, \mathcal{A})$ is at most $\mathrm{negl}(\lambda)$:

[7] We emphasize that restricting these oracles to classical (rather than quantum) access does not weaken our result. On the contrary, it strengthens the separation, since classical access is a special case of quantum access and any distinguisher with quantum access can simulate classical queries.

$$\mathrm{Adv}^{\mathrm{PRF}}(\lambda,\mathcal{A}) = \left| \Pr_{k\leftarrow \mathsf{KeyGen}(1^\lambda)}\left[\mathcal{A}^{\mathsf{Eval}(k,\cdot)}(1^\lambda)=1\right] - \Pr_{f\leftarrow\mathcal{F}}\left[\mathcal{A}^{f(\cdot)}(1^\lambda)=1\right]\right|$$

where $\mathcal{F}$ is the set of all functions $f:\{0,1\}^{n(\lambda)}\to\{0,1\}^{m(\lambda)}$.

3.5 Pseudorandom Permutations

A pseudorandom permutation $\mathsf{PRP}=(\mathsf{KeyGen},\mathsf{Eval},\mathsf{Inv})$ consists of PPT algorithms $\mathsf{KeyGen}:\{0,1\}^\lambda\to\{0,1\}^{s(\lambda)}$, $\mathsf{Eval}:\{0,1\}^{s(\lambda)}\times\{0,1\}^{n(\lambda)}\to\{0,1\}^{n(\lambda)}$ and $\mathsf{Inv}:\{0,1\}^{s(\lambda)}\times\{0,1\}^{n(\lambda)}\to\{0,1\}^{n(\lambda)}$ such that for any $\mathsf{sk}\leftarrow\mathsf{KeyGen}(1^\lambda)$, $\mathsf{Eval}(\mathsf{sk},\cdot):\{0,1\}^n\to\{0,1\}^n$ is a permutation, $\mathsf{Inv}(\mathsf{sk},\cdot)$ is the inverse of $\mathsf{Eval}(\mathsf{sk},\cdot)$ and s,n are polynomials. Also, for any PPT (QPT) adversary $\mathcal{A}$, there exists a negligible function $\mathrm{negl}(\cdot)$ such that for any $\lambda\in\mathbb{N}$, the following advantage function $\mathrm{Adv}^{\mathrm{PRP}}(\lambda,\mathcal{A})$ is at most $\mathrm{negl}(\lambda)$:

$$\mathrm{Adv}^{\mathrm{PRP}}(\lambda,\mathcal{A}) = \left| \Pr_{k\leftarrow \mathsf{KeyGen}(1^\lambda)}\left[\mathcal{A}^{\mathsf{Eval}(k,\cdot),\mathsf{Inv}(k,\cdot)}(1^\lambda)=1\right] - \Pr_{p\leftarrow\mathcal{P}}\left[\mathcal{A}^{p(\cdot),p^{-1}(\cdot)}(1^\lambda)=1\right]\right|$$

where $\mathcal{P}$ is the set of all permutations $p:\{0,1\}^{n(\lambda)}\to\{0,1\}^{n(\lambda)}$.

3.6 Pseudorandom States

We use $|\psi\rangle\leftarrow\mu_d$ to denote a Haar-random pure state in $\mathcal{H}_d$.

Definition 4 (Pseudorandom States (PRS) [JLS18]). *A* pseudorandom state generator *is a quantum polynomial-time (QPT) algorithm*

$$\mathsf{PRS}:\{0,1\}^\lambda\to\mathcal{H}_{m(\lambda)}$$

that on input a classical seed $k\in\{0,1\}^\lambda$ outputs an $m(\lambda)$-qubit pure state $|\phi_k\rangle$.

The family $\{|\phi_k\rangle\}_{k\in\{0,1\}^\lambda}$ is called a family of pseudorandom states *if for every quantum polynomial-time distinguisher $\mathcal{A}$ and every polynomial $t(\lambda)$, there exists a negligible function $\mathrm{negl}(\lambda)$ such that*

$$\left|\Pr_{k\leftarrow\{0,1\}^\lambda}\left[\mathcal{A}\left((|\phi_k\rangle\langle\phi_k|)^{\otimes t(\lambda)}\right)=1\right] - \Pr_{|\psi\rangle\leftarrow\mu_{m(\lambda)}}\left[\mathcal{A}\left((|\psi\rangle\langle\psi|)^{\otimes t(\lambda)}\right)=1\right]\right|\le\mathrm{negl}(\lambda).$$

If the above indistinguishability holds even for computationally unbounded distinguishers, then the PRS family is said to be statistically secure.

Definition 5 (Single-copy Pseudorandom States (1PRS) [MY22]). *A single-copy pseudorandom state generator (1PRS) is defined as in Definition 4, with the following two modifications:*

1. (***Single-copy security***) *The indistinguishability condition is required to hold only for a single copy of the output state, i.e., $t(\lambda)=1$.*
2. (***Stretch***) *The output length satisfies $m(\lambda)>\lambda$ for all sufficiently large λ.*

Formally, for every quantum polynomial-time distinguisher $\mathcal{A}$, there exists a negligible function $\mathrm{negl}(\lambda)$ such that

$$\left|\Pr_{k\leftarrow\{0,1\}^\lambda}[\mathcal{A}(|\phi_k\rangle\langle\phi_k|)=1] - \Pr_{|\psi\rangle\leftarrow\mu_{m(\lambda)}}[\mathcal{A}(|\psi\rangle\langle\psi|)=1]\right|\le\mathrm{negl}(\lambda).$$

If the above indistinguishability holds against computationally unbounded distinguishers, the 1PRS is said to be statistically secure.

The next theorem states that 1PRS exists if PRS exists in a blackbox sense. This is non-trivial because the length of the output $m(\lambda)$ of the PRS might be smaller than λ.

Theorem 5 ([GJMZ23]). *If there exists secure* PRS*, then there exists secure* 1PRS*.*

3.7 Linear Error-Correcting Codes

Definition 6 (Linear Error-Correcting Code). *Let q be a power of 2 and $\mathbb{F}_q$ be the finite field of order q. A* linear $[n, k, d]$ code *C over $\mathbb{F}_q$ is a k-dimensional subspace of $\mathbb{F}_q^n$ where n is the codeword length, k is the message length and d is the minimum distance. The minimum distance d is defined as*

$$d = \min_{\substack{\mathbf{u},\mathbf{v}\in C \\ \mathbf{u}\neq\mathbf{v}}} \mathsf{Ham}(\mathbf{u},\mathbf{v})$$

The code is defined by two efficient algorithms $\mathsf{Encode}, \mathsf{Decode}$ *such that the following properties hold:*

- ***Correctness:*** *For every* $\mathsf{m} \in \mathbb{F}_q^k$, $\mathsf{Decode}(\mathsf{Encode}(\mathsf{m})) = \mathsf{m}$
- ***Robustness:*** *The code corrects up to t errors if, for every* $\mathsf{m} \in \mathbb{F}_q^k$ *and* $\mathbf{e} \in \mathbb{F}_q^n$ *with* $\mathsf{Ham}(\mathbf{e}) \leq t$, *we have* $\mathsf{Decode}(\mathsf{Encode}(\mathsf{m}) + \mathbf{e}) = \mathsf{m}$

Quantum Error-Correcting Codes. A *quantum error-correcting code* with parameters $[[N, m, d]]$ is a 2^m-dimensional subspace

$$\mathcal{Q} \subseteq \mathcal{H}_N$$

called the *codespace*, together with an encoding isometry

$$V : \mathcal{H}_m \to \mathcal{H}_N,$$

whose image is $\mathcal{Q}$. The parameter d, called the *distance* of the code, is the largest integer such that for every Pauli operator P of weight strictly less than d and for all $|\phi\rangle, |\psi\rangle \in \mathcal{H}_m$,

$$\langle\phi|V^\dagger \, \mathrm{P} \, V|\psi\rangle = c_{\mathrm{P}} \, \langle\phi||\psi\rangle$$

for some scalar $c_{\mathrm{P}} \in \mathbb{C}$ depending only on P. Consequently, the code can correct all Pauli errors acting nontrivially on at most $\lfloor (d-1)/2 \rfloor$ qubits, and quantum channels whose Kraus operators lie in their linear span.

A quantum code is called a *stabilizer code* if the codespace $\mathcal{Q}$ is the common $+1$ eigenspace of an abelian subgroup of the N-qubit Pauli group that does not contain $-\mathrm{I}$. A *CSS code* is a special class of stabilizer codes constructed from pairs of classical linear codes satisfying suitable dual-containment conditions.

The Quantum Random Oracle (QROM) Throughout this paper, we use the compressed-oracle technique [Zha19].Further preliminaries and auxiliary lemmas appear in Appendix A.

4 Black-box Impossibility of Pseudorandom Codes in the Quantum Random Oracle Model

In this section, we investigate the black-box separation for binary pseudorandom codes correcting a constant positive rate of independent bit flips in the quantum random oracle model (QROM), against computationally unbounded adversaries making polynomially many oracle queries[8]. Our model allows the PRC to be constructed using fully quantum polynomial-time algorithms, while all underlying objects including secret keys, messages, and codewords remain classical. All algorithms are permitted to make superposition queries to a single global quantum random oracle.

The proof follows the same high-level structure as the classical impossibility result of [DMSR25]. The main technical challenge is to ensure that the arguments, which rely on properties of query distributions and Boolean analysis, continue to hold when the decoding algorithms are quantum. We will see that some of the machinery can be reused, since the relevant properties extracted from the decoder are purely classical distributions, even though they arise from quantum computations.

[8] See Definition 3

4.1 Exceptional Queries

We begin by introducing terminology related to oracle queries made by decoding algorithms. This subsection closely follows [DMSR25], with the key difference that the algorithms considered here may be quantum and may issue oracle queries in superposition.

Quantum Query Distribution. Let Q be a finite query domain, and let $\mathcal{A}$ be an arbitrary quantum oracle algorithm. For an input word $\mathbf{x} \in \{0,1\}^n$, let $\rho_{\mathbf{x}}^{\mathcal{A}}$ denote the joint quantum state of all registers *immediately before* $\mathcal{A}$ makes its first oracle query on input $\mathbf{x}$. Let Q also denote the designated query register, and define the reduced density matrix on the query register by

$$\rho_{\mathbf{x}}^{Q} := \mathsf{Tr}_{\text{rest}}\left[\rho_{\mathbf{x}}^{\mathcal{A}}\right],$$

where the trace is taken over all registers except the query register. We associate with this state a classical probability distribution, called the *quantum query distribution*,

$$\mathsf{Quer}(\mathbf{x}) \in [0,1]^{|Q|},$$

defined by

$$\mathsf{Quer}(\mathbf{x})[q] := \mathsf{Tr}\Big(|q\rangle\langle q|\,\rho_{\mathbf{x}}^{Q}\Big) \qquad \text{for all } q \in Q.$$

Equivalently, $\mathsf{Quer}(\mathbf{x})[q]$ is the probability of observing q when measuring the query register in the computational basis immediately before the first oracle call. This measurement is purely conceptual and is introduced only for analysis; no measurement is actually performed during the execution of the algorithm. By definition, $\mathsf{Quer}(\mathbf{x})$ is a valid probability distribution:

$$\mathsf{Quer}(\mathbf{x})[q] \geq 0 \quad \text{and} \quad \sum_{q \in Q} \mathsf{Quer}(\mathbf{x})[q] = 1.$$

Unless specified otherwise, all probabilities and expectations below are taken over independent random choices $\mathbf{u} \leftarrow \{0,1\}^n$ and $\mathbf{e} \leftarrow \mathsf{Ber}_\gamma^n$, where Ber_γ denotes the Bernoulli distribution with bias parameter γ.

We now define two weight functions associated with the *first* oracle query of $\mathcal{A}$. These definitions apply to quantum algorithms, but the resulting objects are entirely classical.

- The *global weight function* $\mathsf{w}(q)$ measures how likely $\mathcal{A}$ is to query the basis element $q \in Q$ on its first oracle query when the input is uniformly random. It is defined as

$$\mathsf{w}(q) := \mathsf{E}_{\mathbf{u}}\big[\mathsf{Quer}(\mathbf{u})[q]\big].$$

- The *local weight function* $\ell_{\mathbf{x}}(q)$ measures how likely $\mathcal{A}$ is to query q on its first oracle query when the input is drawn from a noisy neighborhood of a fixed word $\mathbf{x}$. It is defined as

$$\ell_{\mathbf{x}}(q) := \mathsf{E}_{\mathbf{e}}\big[\mathsf{Quer}(\mathbf{x}+\mathbf{e})[q]\big],$$

Although these weights arise from a quantum computation, both w and $\ell_{\mathbf{x}}$ are classical functions with classical inputs and outputs. This observation will be crucial for lifting classical arguments to the quantum setting.

Definition 7 (L_p-Norm and Smoothing of Functions on the Boolean Hypercube). *Let $f : \{0,1\}^n \to \mathbb{R}$ be a real-valued function. For $1 \leq p < \infty$, we define the normalized L_p-norm of f by*

$$\|f\|_{(p)} = \left(\mathsf{E}_{\mathbf{u}}[|f(\mathbf{u})|^p]\right)^{1/p},$$

and

$$\|f\|_\infty = \max_{\mathbf{x}\in\{0,1\}^n} |f(\mathbf{x})|.$$

For a noise parameter $\gamma \in [0,1]$, the noise operator $\mathbf{T}_\gamma$ *maps f to the function $\mathbf{T}_\gamma f : \{0,1\}^n \to \mathbb{R}$ defined by*

$$(\mathbf{T}_\gamma f)(\mathbf{x}) = \mathsf{E}_{\mathbf{e}\leftarrow \mathsf{Ber}_\gamma^n}[f(\mathbf{x}+\mathbf{e})].$$

Note that $\|f\|_{(p)}$ differs from the standard L_p norm of the truth table of f by a normalization factor of $2^{-n/p}$; that is,

$$\|f\|_{(p)} = 2^{-n/p} \cdot \|f\|_p.$$

As a consequence of this normalization, the usual relationship between L_p and L_q norms is reversed: for $1 \le p < q$, we have $\|f\|_{(p)} \le \|f\|_{(q)}$. To avoid confusion, we consistently use the notation $\|\cdot\|_{(p)}$ throughout.

Theorem 6 (Hypercontractivity Theorem [Bon70, Bec75, O'D14]). *Let $f : \{0,1\}^n \to \mathbb{R}$, let $0 \le s \le t < \infty$, with $t > 0$ and let $0 \le \gamma \le \sqrt{s/t}$. Then*

$$\|\mathbf{T}_\gamma f\|_{(1+t)} \le \|f\|_{(1+s)}.$$

We now recall a key consequence of hypercontractivity established in [DMSR25]. Importantly, although the query distribution Quer may arise from a quantum algorithm, the lemma below depends only on the classical function $\mathbf{x} \mapsto \mathsf{Quer}(\mathbf{x})[q]$ and therefore applies unchanged in the quantum setting.

Lemma 7 ([DMSR25]). *Let $0 < \gamma < 1$ and let $t \ge 0$. Fix a query $q \in Q$. Then*

$$\mathsf{E}_{\mathbf{u}}[\ell_{\mathbf{u}}(q)^{1+t}] \le (\mathsf{w}(q))^{(1+t)/(1+\gamma^2 t)}.$$

Furthermore, for any $0 < \delta \le 1$,

$$\Pr_{\mathbf{u}}[\ell_{\mathbf{u}}(q) \ge \delta] \le \frac{(\mathsf{w}(q))^{(1+t)/(1+\gamma^2 t)}}{\delta^{1+t}}.$$

This lemma shows that a query which is globally unlikely cannot, for most inputs, become locally likely under noise.

Definition 8 (Exceptionality). *Fix $0 < \epsilon < \delta < 1$. A word $\mathbf{x} \in \{0,1\}^n$ is said to be (ϵ,δ)-*exceptional *with respect to a query function $\mathsf{Quer} : \{0,1\}^n \to [0,1]^Q$ if there exists a query $q \in Q$ such that*

$$\mathsf{w}(q) < \epsilon \quad \text{and} \quad \ell_{\mathbf{x}}(q) \ge \delta.$$

We denote this event by $\mathsf{EXC}_{\mathsf{Quer},\epsilon,\delta}(\mathbf{x})$.

Intuitively, exceptionality captures the situation in which the decoder's behavior becomes highly non-uniform in a noisy neighborhood of $\mathbf{x}$, despite the particular query having small global weight. Such behavior is expected for structured inputs like codewords when decoding succeeds reliably in their vicinity. The following theorem shows that this phenomenon is unlikely for uniformly random inputs.

Theorem 7 ([DMSR25]). *Fix $n > 0$, $\gamma \in (0,1)$, any finite set Q, and any query function $\mathsf{Quer} : \{0,1\}^n \to [0,1]^Q$. Let w and ℓ be defined as above. Then for any $t \ge 0$ and any $0 < \epsilon < \delta < 1$,*

$$\Pr_{\mathbf{u}}[\mathsf{EXC}_{\mathsf{Quer},\epsilon,\delta}(\mathbf{u})] \le \frac{\epsilon^{(1+t)/(1+\gamma^2 t)-1}}{\delta^{1+t}}.$$

The exponent $(1+t)/(1+\gamma^2 t) - 1 = \frac{t(1-\gamma^2)}{1+\gamma^2 t}$ is strictly positive for all $\gamma \in (0,1)$ and $t > 0$. For concreteness, we fix $t = 1$ for the remainder of this section.

The above bound is independent of the input length and the query domain size. We will use it to argue that pseudorandom codewords cannot be exceptional with noticeable probability. Before doing so, we describe how to eliminate oracle queries from the decoding algorithm by means of augmented decoding.

Definition 9 (Augmented Decoding).
A secret-key PRC $\mathsf{PRC}^{\mathcal{O}} = (\mathsf{KeyGen}^{\mathcal{O}}, \mathsf{Encode}^{\mathcal{O}}, \mathsf{Decode}^{\mathcal{O}})$ *has* κ-***augmented correctness*** *with respect to augmented decoding algorithms* $(\mathsf{LearnQuery}^{\mathcal{O}}, \mathsf{ADecode}^{\mathcal{O}})$ *where the learner* $\mathsf{LearnQuery}^{\mathcal{O}}$ *receives only the classical secret key and the accuracy parameter* ϵ*; it has no access to the message or codeword and produces a classical string* aux *of polynomially bounded length. Its computation may be unbounded, but its oracle-query count must be polynomially bounded.* $\mathsf{ADecode}^{\mathcal{O}}$ *is a quantum polynomial-time oracle algorithm and for all* $\lambda \in \mathbb{N}$*,*

$$\Pr\left[\mathsf{ADecode}^{\mathcal{O}}(\mathsf{sk}, \mathsf{aux}, \mathbf{c}+\mathbf{e}) \neq m \;\middle|\; \begin{array}{l} \mathsf{sk} \leftarrow \mathsf{KeyGen}^{\mathcal{O}}(1^\lambda) \\ m \leftarrow \{0,1\}^d \\ \mathbf{c} \leftarrow \mathsf{Encode}^{\mathcal{O}}(\mathsf{sk}, m) \\ \mathbf{e} \leftarrow \mathsf{Ber}_\gamma^n \\ \mathsf{aux} \leftarrow \mathsf{LearnQuery}^{\mathcal{O}}(\mathsf{sk}, \epsilon) \end{array}\right] \leq \kappa.$$

The following lemma, proven in [DMSR25], bounds the probability that codewords are exceptional. We use $\mathsf{EXC}_{\mathsf{sk}',\epsilon,\delta}(\mathbf{x})$ as shorthand for $\mathsf{EXC}_{\mathsf{Quer}_{\mathsf{sk}'},\epsilon,\delta}(\mathbf{x})$.

Lemma 8 ([DMSR25]). *Let* $\mathsf{PRC}^{\mathcal{O}} = (\mathsf{KeyGen}^{\mathcal{O}}, \mathsf{Encode}^{\mathcal{O}}, \mathsf{Decode}^{\mathcal{O}}, \mathsf{LearnQuery}^{\mathcal{O}}, \mathsf{ADecode}^{\mathcal{O}})$ *be an augmented pseudorandom code. Let* $\mathsf{Quer}_{\mathsf{sk}'} : \{0,1\}^n \to [0,1]^Q$ *denote the first-query distribution induced by the augmented key* $\mathsf{sk}' = (\mathsf{sk}, \mathsf{aux})$*. Let* $p^*(\epsilon,\delta) = \frac{\epsilon^{(1-\gamma^2)/(1+\gamma^2)}}{\delta^2}$*. If*

$$\Pr[\mathsf{EXC}_{\mathsf{sk}',\epsilon,\delta}(\mathbf{c})] > 3p^*$$

over the joint randomness of sk*,* m*,* $\mathbf{c}$*, and* aux*, then there exists a distinguisher* $\mathcal{D}$ *which distinguishes codewords of* $\mathsf{PRC}^{\mathcal{O}}$ *from the uniform distribution on* $\{0,1\}^n$ *with advantage at least* $\frac{1}{2}p^* - 2^{-\lambda}$ *using* $N = (k+\lambda)/p^*$ *samples, without access to* sk *or to* $\mathcal{O}$*. In particular, if* k *and* $1/p^*$ *are polynomially bounded and* $\mathsf{PRC}^{\mathcal{O}}$ *is pseudorandom against unbounded polynomial-query distinguishers, then, for all sufficiently large* λ*,*

$$\Pr[\mathsf{EXC}_{\mathsf{sk}',\epsilon,\delta}(\mathbf{c})] \leq 3p^*.$$

Although Theorem 7 and Lemma 8 were originally proven in the classical setting, its proof relies exclusively on the classical distributions Quer, w, and ℓ induced by the decoder's first query. It does not depend on any internal structure of the decoding algorithm and therefore applies verbatim when the decoder is quantum.

We now present the main theorem of this section.

Theorem 8. *Let* $\mathsf{PRC}^{\mathcal{O}} = (\mathsf{KeyGen}^{\mathcal{O}}, \mathsf{Encode}^{\mathcal{O}}, \mathsf{Decode}^{\mathcal{O}})$ *be a pseudorandom code with* κ*-correctness and* γ*-robustness relative to a random oracle* $\mathcal{O}$*. Then there exists an (unbounded) oracle algorithm* $\mathsf{LearnQuery}^{\mathcal{O}}$ *and an efficient decoder* Decode' *such that for all* $\lambda \in \mathbb{N}$ *the scheme* $(\mathsf{KeyGen}^{\mathcal{O}}, \mathsf{Encode}^{\mathcal{O}}, \mathsf{Decode}')$ *is a PRC with* $(\kappa + 1/\mathsf{poly}(\lambda))$*-augmented correctness and* γ*-robustness.*

We prove Theorem 8 by iteratively eliminating oracle queries made by the decoder. Each iteration removes one oracle query at the cost of a inverse-polynomial increase in the augmented correctness error. The key step is the following lemma.

Lemma 9. *Let* $\mathsf{PRC}^{\mathcal{O}} = (\mathsf{KeyGen}^{\mathcal{O}}, \mathsf{Encode}^{\mathcal{O}}, \mathsf{Decode}^{\mathcal{O}})$ *be a PRC relative to a random oracle* $\mathcal{O}$*, and suppose it admits an augmented decoder* $(\mathsf{LearnQuery}^{\mathcal{O}}, \mathsf{ADecode}^{\mathcal{O}})$ *with correctness error* κ *and decoding radius* γ*. Assume that* $\mathsf{ADecode}^{\mathcal{O}}$ *makes at most* ℓ *oracle queries. Then for every* $\kappa' = 1/\mathsf{poly}(\lambda)$*, there exists an augmented decoder* $(\mathsf{LearnQuery}'^{\mathcal{O}}, \mathsf{ADecode}'^{\mathcal{O}})$ *for* $\mathsf{PRC}^{\mathcal{O}}$ *with the following properties:*

1. $\mathsf{ADecode}'^{\mathcal{O}}$ *makes at most* $(\ell - 1)$ *oracle queries.*
2. *For all sufficiently large* λ*,* $(\mathsf{KeyGen}^{\mathcal{O}}, \mathsf{LearnQuery}'^{\mathcal{O}}, \mathsf{Encode}^{\mathcal{O}}, \mathsf{ADecode}'^{\mathcal{O}})$ *has* $(\kappa + \kappa')$*-augmented correctness and decoding radius* γ*.*
3. *The number of oracle queries made by* $\mathsf{LearnQuery}'^{\mathcal{O}}$ *is at most the number made by* $\mathsf{LearnQuery}^{\mathcal{O}}$ *plus* $1/\epsilon$*, where* ϵ *depends only on* κ' *and* γ*.*
4. *The size of the output of* $\mathsf{LearnQuery}'^{\mathcal{O}}$ *is at most* $\mathsf{poly}(\lambda)/\epsilon$*.*

Proof. We describe the construction of $\mathsf{LearnQuery}'^{\mathcal{O}}$ and $\mathsf{ADecode}'^{\mathcal{O}}$.

Description of $\mathsf{LearnQuery}'^{\mathcal{O}}$*.* On input (sk, ϵ), the learner $\mathsf{LearnQuery}'^{\mathcal{O}}$ proceeds as follows:

1. Run $\mathsf{aux} \leftarrow \mathsf{LearnQuery}^{\mathcal{O}}(\mathsf{sk}, \epsilon)$ and set $\mathsf{sk}' := (\mathsf{sk}, \mathsf{aux})$.
2. Let $\mathsf{Quer}_{\mathsf{sk}'}$ denote the quantum query distribution induced by the *first oracle query* of $\mathsf{ADecode}^{\mathcal{O}}(\mathsf{sk}', \cdot)$. Since $\mathsf{LearnQuery}'^{\mathcal{O}}$ is allowed to be unbounded, it computes the global weight function $\mathsf{w}(\cdot)$ associated with $\mathsf{Quer}_{\mathsf{sk}'}$ by simulating $\mathsf{ADecode}^{\mathcal{O}}$ up to its first oracle call and extracting the induced distribution.
3. Let
$$\tilde{Q} := \{q \in Q : \mathsf{w}(q) \geq \epsilon\}$$
be the set of globally ϵ-heavy queries. Query the oracle $\mathcal{O}$ on all $q \in \tilde{Q}$ using basis queries, measure the responses, and store the resulting query–response pairs in a classical list $\mathcal{L}$.
4. Output $\mathsf{aux}' := \mathsf{aux} \,\|\, \mathcal{L}$.

Note that by definition of $\tilde{Q}$, there can be at most $1/\epsilon$ globally ϵ-heavy queries, and hence $\mathsf{LearnQuery}'^{\mathcal{O}}$ makes at most $1/\epsilon$ classical oracle queries beyond those made by $\mathsf{LearnQuery}^{\mathcal{O}}$.

Description of $\mathsf{ADecode}'^{\mathcal{O}}$*.* On input $(\mathsf{sk}, \mathsf{aux}', \mathbf{c})$, where $\mathsf{aux}' = \mathsf{aux} \| \mathcal{L}$, the decoder $\mathsf{ADecode}'^{\mathcal{O}}$ proceeds as follows:

1. Run $\mathsf{ADecode}^{\mathcal{O}}(\mathsf{sk}, \mathsf{aux}, \mathbf{c})$ until it reaches the point immediately before applying its *first* oracle query.
2. Replace this first oracle query by a query to a *local initialized compressed oracle* as follows. $\mathsf{ADecode}'$ maintains a private database register D, initialized to the database state $(|\emptyset\rangle_D \otimes |\mathcal{L}\rangle_L)$ (see Appendix A.1).
Writing the joint pre-query state as
$$\Big(\sum_{q,z} \alpha_{q,z} |q\rangle_Q |z\rangle_Z\Big) \otimes |\emptyset\rangle_D |\mathcal{L}\rangle_L,$$
the first oracle query is answered by applying the compressed-oracle query unitary on (Q, Z, D), corresponding to an initialized compressed oracle with initial database $\mathcal{L}$. Concretely, on each basis state $|q\rangle_Q |z\rangle_Z |\mathbf{d}\rangle_D |\mathcal{L}\rangle_L$, this unitary acts as follows:
   (a) If $q \in \mathrm{Dom}(\mathcal{L})$, then XOR the stored value $\mathcal{L}[q]$ into the answer register contained in Z, and leave the database register unchanged.
   (b) If $q \notin \mathrm{Dom}(\mathcal{L})$, then coherently apply the compressed-oracle query rule: $\mathsf{CStO}(|q\rangle_Q |z\rangle_Z |\mathbf{d}\rangle_D)$. This means queries are answered with fresh uniformly random values.

This implements exactly one query to an initialised compressed oracle with initial database $\mathcal{L}$. Note that after this step, every database basis state in the support of D has at most one entry.

3. For every *subsequent* oracle query made by ADecode, the decoder proceeds as follows. On a query q, where the database register D may be in a superposition of classical databases, the following rule is applied coherently:
   - If $q \in \mathrm{Dom}(\mathcal{L})$, return the stored value $\mathcal{L}[q]$ by XORing it into the answer register.
   - If $q \in \mathrm{Dom}(\mathbf{d})$, apply the compressed oracle unitary.
   - Otherwise, forward the query coherently to the real oracle $\mathcal{O}$, where the response $\mathcal{O}(q)$ is XORed into the answer register.

   That is, after the first query, $\mathsf{ADecode}'$ uses the local database D and list $\mathcal{L}$ whenever possible, and otherwise defers to the real oracle $\mathcal{O}$.
4. Output the final output of ADecode.

Bounding the augmented correctness error. Fix $b \in \{0,1\}$.

Consider the standard decoding experiment where $\mathsf{sk} \leftarrow \mathsf{KeyGen}^{\mathcal{O}}(1^\lambda)$, $\mathbf{c} \leftarrow \mathsf{Encode}^{\mathcal{O}}(\mathsf{sk}, b)$, and $\mathbf{e} \leftarrow \mathsf{Ber}_\gamma^n$. Define the following hybrids:

Hyb_0**:** Run $\mathsf{aux} \leftarrow \mathsf{LearnQuery}^{\mathcal{O}}(\mathsf{sk}, \epsilon)$ and then run $\mathsf{ADecode}^{\mathcal{O}}(\mathsf{sk}, \mathsf{aux}, \mathbf{c}+\mathbf{e})$ where $\mathbf{e} \leftarrow \mathsf{Ber}_\gamma^n$. Let

$$\mathsf{Fail} := \{\mathsf{ADecode}^{\mathcal{O}}(\mathsf{sk}, \mathsf{aux}, \mathbf{c}+\mathbf{e}) \neq b\}.$$

Hyb_1**:** Run $\mathsf{aux}' \leftarrow \mathsf{LearnQuery}'^{\mathcal{O}}(\mathsf{sk}, \epsilon)$ but still run $\mathsf{ADecode}^{\mathcal{O}}(\mathsf{sk}, \mathsf{aux}, \mathbf{c}+\mathbf{e})$ where $\mathbf{e} \leftarrow \mathsf{Ber}_\gamma^n$. Let

$$\mathsf{Fail}^* := \{\mathsf{ADecode}^{\mathcal{O}}(\mathsf{sk}, \mathsf{aux}, \mathbf{c}+\mathbf{e}) \neq b\}.$$

Hyb_2**:** Run $\mathsf{aux}' \leftarrow \mathsf{LearnQuery}'^{\mathcal{O}}(\mathsf{sk}, \epsilon)$ and then run $\mathsf{ADecode}'^{\mathcal{O}}(\mathsf{sk}, \mathsf{aux}', \mathbf{c}+\mathbf{e})$ where $\mathbf{e} \leftarrow \mathsf{Ber}_\gamma^n$. Let

$$\mathsf{Fail}' := \{\mathsf{ADecode}'^{\mathcal{O}}(\mathsf{sk}, \mathsf{aux}', \mathbf{c}+\mathbf{e}) \neq b\}.$$

By assumption, $\Pr[\mathsf{Fail}] \leq \kappa(\lambda)$ in Hyb_0.

From Hyb_0 *to* Hyb_1. The only difference between Hyb_0 and Hyb_1 is that LearnQuery is replaced by $\mathsf{LearnQuery}'$. The additional oracle queries made by $\mathsf{LearnQuery}'$ beyond those of LearnQuery are basis queries whose outcomes are measured and stored in $\mathcal{L}$, but these outcomes are not accessed by ADecode in Hyb_1. This transformation is perfectly indistinguishable. Hence,

$$\Pr_{\mathsf{Hyb}_0}[\mathsf{Fail}] = \Pr_{\mathsf{Hyb}_1}[\mathsf{Fail}^*].$$

We now analyze Hyb_2. As mentioned earlier, we work in the compressed-oracle representation to model the random oracle. After running KeyGen and Encode they have made at most M queries to the oracle. Then the learner $\mathsf{LearnQuery}'$ makes its oracle queries.

Exceptional codewords.

Let $\mathsf{Quer}_{\mathsf{sk}}$ be the query distribution induced by the first oracle query of $\mathsf{ADecode}^{\mathcal{O}}(\mathsf{sk}, \cdot)$, and let $\mathsf{w}, \ell_{\mathbf{c}}$ be the associated global and local weight functions. Let EXC denote the event that $\mathbf{c}$ is (ϵ, δ)-exceptional. Then

$$\Pr[\mathsf{Fail}'] \leq \Pr[\mathsf{Fail}' \mid \neg\mathsf{EXC}] + \Pr[\mathsf{EXC}].$$

By Lemma 8 and Theorem 7, $\Pr[\mathsf{EXC}] \leq 3p^*(\epsilon, \delta)$.

Bounding $\Pr[\mathsf{Fail}' \mid \neg\mathsf{EXC}]$**.**

Assume $\mathbf{c}$ is not exceptional. Recall that $\mathcal{L}$ is the set of learned queries and let $D_{K,E,\mathcal{L}}$ denote the state of the compressed oracle database after KeyGen, Decode and $\mathsf{LearnQuery}'$ have made their queries. Let $D_{K,E} = D_{K,E,\mathcal{L}} \setminus \mathcal{L}$.

If the query q produced due to $\mathbf{c}$ overlaps with $D_{K,E}$, i.e. is not in $\mathcal{L}$ we can get a decoding error, so we want to bound this possibility.

Occupancy outside the learned set. For each oracle address q, define the occupancy projector on D by

$$E_q = \sum_{\mathbf{d}:\, q \in \mathrm{Dom}(\mathbf{d})} |\mathbf{d}\rangle\langle\mathbf{d}|, \qquad N_{\overline{\mathcal{L}}} = \sum_{q \notin \mathcal{L}} E_q.$$

The operator $N_{\overline{\mathcal{L}}}$ counts occupied addresses outside $\mathcal{L}$:

$$N_{\overline{\mathcal{L}}}|\mathbf{d}\rangle = |\mathrm{Dom}(\mathbf{d}) \setminus \mathcal{L}|\,|\mathbf{d}\rangle.$$

Before the additional basis queries, the database is supported on databases of size at most M. Those additional queries act only at addresses in $\mathcal{L}$, so they cannot increase overlap outside $\mathcal{L}$. Conditioning on the classical outputs and learned values preserves this support bound. Consequently $\mathsf{Tr}(N_{\overline{\mathcal{L}}}\rho_D^{a,\mathbf{c}}) \le M$.

Let Q be the decoder's query register and W its other private registers. The decoder receives the classical input $(\mathsf{sk}, \mathsf{aux}, \mathbf{c} + \mathbf{e})$, where $\mathbf{e} \leftarrow \mathrm{Ber}_\gamma^n$ is independent of the preceding computation, and initialises fresh private workspace. Its computation up to the first oracle query is a channel acting only on these inputs and workspace. Consequently, conditioned on $a = (\mathsf{sk}, \mathsf{aux}, \mathcal{L})$ and $\mathbf{c}$ and averaged over $\mathbf{e}$ and the decoder's private randomness, the joint state at that query is

$$\omega_{QWD}^{a,\mathbf{c}} = \sigma_{QW}^{a,\mathbf{c}} \otimes \rho_D^{a,\mathbf{c}}, \qquad \mathsf{Tr}\big[(|q\rangle\langle q|_Q \otimes \mathbb{I}_W)\sigma_{QW}^{a,\mathbf{c}}\big] = \ell_{\mathsf{sk},\mathbf{c}}(q), \tag{3}$$

per definition of the local weight function and where $\sigma_{QW}^{a,\mathbf{c}}$ denotes the decoder's state.

To bound the overlap let $\Pi_{\mathrm{hit}}^{\mathcal{L}}$ project onto pairs consisting of a query outside $\mathcal{L}$ and a database containing that query:

$$\Pi_{\mathrm{hit}}^{\mathcal{L}} = \sum_{q \notin \mathcal{L}} |q\rangle\langle q|_Q \otimes \mathbb{I}_W \otimes E_q.$$

The summands are mutually orthogonal projectors because their query registers are orthogonal. Define $p_{\mathrm{hit}}(a, \mathbf{c}) = \mathsf{Tr}(\Pi_{\mathrm{hit}}^{\mathcal{L}}\omega^{a,\mathbf{c}})$. Using (3),

$$p_{\mathrm{hit}}(a, \mathbf{c}) = \sum_{q \notin \mathcal{L}} \ell_{a,\mathbf{c}}(q)\,\mathsf{Tr}(E_q\rho_D^{a,\mathbf{c}}). \tag{4}$$

Suppose $\mathbf{c}$ is not (ϵ, δ)-exceptional for $\mathsf{Quer}_{\mathsf{sk}}$. Since $\mathcal{L}$ contains every globally ϵ-heavy address, every $q \notin \mathcal{L}$ satisfies $\ell_{a,\mathbf{c}}(q) < \delta$. Thus

$$\begin{aligned} p_{\mathrm{hit}}(a, \mathbf{c}) &\le \delta \sum_{q \notin \mathcal{L}} \mathsf{Tr}(E_q\rho_D^{a,\mathbf{c}}) \\ &= \delta\,\mathsf{Tr}(N_{\overline{\mathcal{L}}}\rho_D^{a,\mathbf{c}}) \le M\delta. \end{aligned} \tag{5}$$

This establishes that the overlap with $D_{K,E} \setminus \mathcal{L}$ is bounded by $M\delta$.

Finally, we have to argue that the queries outside of $D_{K,E} \cup \mathcal{L}$ are perfectly simulated. Conditioned on $(\mathsf{sk}, \mathbf{c}, D)$ the value $O(q)$ has not been fixed yet by any prior oracle interaction and is therefore uniformly random. In this case the augmented decoder $\mathsf{ADecode}'$ also samples a fresh uniform value for q by leveraging the compressed oracle mechanism. Hence the answers are identically distributed.

Therefore, conditioned on $(\mathsf{sk}, \mathbf{c}, \mathcal{D})$, the real oracle $\mathcal{O}$ and the simulated oracle used by $\mathsf{ADecode}'$ agree on all basis queries except those in the set $S = D_{\mathsf{K},\mathsf{E}} \setminus \mathcal{L}$.

Let $P_{\bar{S}} := \mathbb{I} - \Pi_{\mathrm{hit}}^{\mathcal{L}}$. Suppose that the real and simulated continuations are implemented as CPTP maps $\Phi_{a,\mathbf{c}}$ and $\Psi_{a,\mathbf{c}}$ on a common input space QWD, with their output consisting only of the decoder's final registers. Fresh local simulator workspace is appended in its prescribed initial state

as part of $\Psi_{a,\mathbf{c}}$; private oracle and simulator memory is traced out at the end of each continuation. Since $\Phi_{a,\mathbf{c}}$ and $\Psi_{a,\mathbf{c}}$ agree on $P_{\bar{S}}$ they fulfill the precondition of Lemma 6

$$\Phi_{a,\mathbf{c}}(P_{\bar{S}}XP_{\bar{S}}) = \Psi_{a,\mathbf{c}}(P_{\bar{S}}XP_{\bar{S}}) \quad \text{for every operator } X \text{ on } QWD. \tag{6}$$

Then we can apply Lemma 6 to $\omega_{QWD}^{a,\mathbf{c}}$ and use (5) to obtain

$$\mathsf{TD}\big(\Phi_{a,\mathbf{c}}(\omega_{QWD}^{a,\mathbf{c}}), \Psi_{a,\mathbf{c}}(\omega_{QWD}^{a,\mathbf{c}})\big) \leq 2\sqrt{p_{\text{hit}}(a,\mathbf{c})} \leq 2\sqrt{M\delta}.$$

If $M\delta > 1$, the same inequality follows directly from $\mathsf{TD} \leq 1$. The final output measurement cannot increase this distance.

Hence,

$$\Pr_{\mathsf{Hyb}_2}[\mathsf{Fail}' \mid a, \mathbf{c}] \leq \Pr_{\mathsf{Hyb}_1}[\mathsf{Fail}^* \mid a, \mathbf{c}] + 2\sqrt{M\delta}.$$

The classical data $(a, \mathbf{c}, b)$ have the same distribution in these two hybrids.

Averaging over $(\mathsf{sk}, \mathbf{c}, \mathcal{D})$ and using $\Pr[\mathsf{Fail}] \leq \kappa(\lambda)$ gives

$$\begin{aligned}\Pr_{\mathsf{Hyb}_2}[\mathsf{Fail}'] &\leq \Pr_{\mathsf{Hyb}_1}[\mathsf{Fail}^*] + 2\sqrt{M\delta} + \Pr[\mathsf{EXC}] \\ &\leq \kappa(\lambda) + 2\sqrt{M\delta} + 3p^*(\epsilon, \delta).\end{aligned}$$

Choosing $\delta = \kappa'^2/(16M)$ and $\epsilon = (\frac{\kappa'\delta^2}{6})^{(1+\gamma^2)/(1-\gamma^2)}$ yields $\Pr[\mathsf{Fail}'] \leq \kappa(\lambda) + \kappa'$, completing the proof.

Proof (*Proof of Theorem 8*). Let ℓ be an upper bound on the number of oracle queries made by $\mathsf{Decode}^{\mathcal{O}}$. Applying Lemma 9 iteratively ℓ times removes all oracle queries, yielding an oracle-free decoder Decode'. The total auxiliary information size is at most $\ell \cdot \mathsf{poly}(\lambda)/\epsilon$, and the total correctness loss is additive. Choosing parameters appropriately completes the proof.

Finally, we can invoke the following theorem from [DMSR25].

Theorem 9. *Let* $\mathsf{PRC}^{\mathcal{O}} = (\mathsf{KeyGen}^{\mathcal{O}}, \mathsf{Encode}^{\mathcal{O}}, \mathsf{Decode}^{\mathcal{O}})$ *be a PRC relative to a random oracle* $\mathcal{O}$*. Let* $\mathsf{LearnQuery}^{\mathcal{O}}$ *be an unbounded algorithm and* ADecode *be a decoder such that the correctness error with respect to the augmented scheme* $(\mathsf{KeyGen}, \mathsf{LearnQuery}^{\mathcal{O}}, \mathsf{Encode}^{\mathcal{O}}, \mathsf{ADecode})$ *is* κ *where* $\kappa < 1/8$*. Then, there exists an unbounded distinguisher* $\mathcal{D}$ *that breaks the pseudorandomness property of* $\mathsf{PRC}^{\mathcal{O}}$ *with advantage* $\frac{1}{16} - 2^{-\lambda}$*.*

The unbounded adversary can distinguish codewords produced by the PRC from uniformly random words by brute-forcing over all secret keys. This concludes the proof.

5 Quantum Pseudorandom Codes

We now define quantum pseudorandom codes (QPRCs), which are the quantum analogue of pseudorandom error-correcting codes introduced in [CG24]. As in the classical setting, a QPRC can be viewed as a family of error-correcting codes indexed by encoding and decoding keys.

Formally, a QPRC is specified by three quantum polynomial-time algorithms: a key generation algorithm, an encoding algorithm, and a decoding algorithm. The encoding algorithm outputs quantum codewords, while the decoding algorithm takes as input a (possibly corrupted) quantum state and outputs either the original message or the failure symbol $\perp$.

Definition 10 (Secret-key quantum pseudorandom code (QPRC)). *A* secret-key quantum pseudorandom code *with robustness to a quantum channel $\mathcal{E}$ is a triple of quantum polynomial-time (QPT) algorithms* $(\mathsf{KeyGen}, \mathsf{Encode}, \mathsf{Decode})$ *satisfying the following properties.*

- ***Syntax.*** *There exist polynomially bounded functions ℓ, n, k, r such that for all $\lambda \in \mathbb{N}$:*
  - $\mathsf{KeyGen}(1^\lambda, 1^{k(\lambda)})$*, on input a security parameter λ and message length $k(\lambda)$, outputs a secret key* $\mathsf{sk} \in \{0,1\}^{\ell(\lambda)}$.
  - $\mathsf{Encode}(\mathsf{sk}, m; s)$*, on input a message $m \in \{0,1\}^{k(\lambda)}$ and a seed $s \in \{0,1\}^{r(\lambda)}$*[9]*, outputs an $n(\lambda)$-qubit pure quantum state in $\mathcal{H}_{n(\lambda)}$.*
  - $\mathsf{Decode}(\mathsf{sk}, \rho)$*, on input an $n(\lambda)$-qubit quantum state ρ, outputs either a message in $\{0,1\}^{k(\lambda)}$ or the symbol $\perp$.*
- ***Error correction (robustness).*** *There is a negligible function bounding the error uniformly over every λ, every message $m \in \{0,1\}^{k(\lambda)}$, such that*

$$\Pr\left[\mathsf{Decode}(\mathsf{sk}, \mathcal{E}(|\psi_{\mathsf{sk},m,s}\rangle)) = m \;\middle|\; \begin{array}{r} \mathsf{sk} \leftarrow \mathsf{KeyGen}(1^\lambda), \\ s \leftarrow \{0,1\}^{r(\lambda)}, \\ |\psi_{\mathsf{sk},m,s}\rangle \leftarrow \mathsf{Encode}(\mathsf{sk}, m; s) \end{array}\right] \geq 1 - \mathrm{negl}(\lambda).$$

  Here $\mathcal{E}(|\psi_{\mathsf{sk},m,s}\rangle)$ abbreviates $\mathcal{E}(|\psi_{\mathsf{sk},m,s}\rangle\langle\psi_{\mathsf{sk},m,s}|)$, where $\mathcal{E}$ is a completely positive trace-preserving map on the codeword register. The probability is taken over the randomness of KeyGen*, the seed s, the channel $\mathcal{E}$, and any measurements performed by* Decode*. Each construction specifies its allowed channel family, including the identity channel.*
- ***t-Pseudorandomness.*** *Fix a polynomially bounded copy count $t = t(\lambda) \geq 1$. For every quantum polynomial-time adversary $\mathcal{A}$ making polynomially many (possibly adaptive)* classical[10] *oracle queries, there exists a negligible function* $\mathrm{negl}(\lambda)$ *such that*

$$\left|\Pr\left[\mathcal{A}^{\mathcal{O}_{\mathsf{Encode}(\mathsf{sk},\cdot)}}(1^\lambda) = 1 \;\middle|\; \mathsf{sk} \leftarrow \mathsf{KeyGen}(1^\lambda)\right] - \Pr\left[\mathcal{A}^{\mathcal{O}_{\mathrm{Haar}}}(1^\lambda) = 1\right]\right| \leq \mathrm{negl}(\lambda),$$

  where $\mathcal{O}_{\mathsf{Encode}(\mathsf{sk},\cdot)}$ is an oracle that, on each query message m, samples an independent fresh seed $s \leftarrow \{0,1\}^{r(\lambda)}$ for that query and outputs t copies of $\mathsf{Encode}(\mathsf{sk}, m; s)$*, and $\mathcal{O}_{\mathrm{Haar}}$ is an oracle that, on each query, outputs t copies of an independent Haar-random $n(\lambda)$-qubit pure state.*[11]

Variants of QPRC. We consider several variants of quantum pseudorandom codes.

- **Zero-bit QPRC.** If the message space consists of a single message (which we denote by 1), then the QPRC is called a *zero-bit QPRC*. In such QPRCs, the decoder outputs 1 upon successful decoding and outputs 0 upon failure. As in the classical definition of PRCs [CG24], an additional property called soundness is required for zero-bit QPRCs, in order to ensure that decoding rejects quantum states that are unrelated to the code.
  Soundness. For any $\lambda \in \mathbb{N}$ and any quantum state ρ chosen independently of the sampled secret key,

$$\Pr_{\mathsf{sk} \leftarrow \mathsf{KeyGen}(1^\lambda)}[\mathsf{Decode}(\mathsf{sk}, \rho) = 1] \leq \mathrm{negl}(\lambda).$$

[9] The explicit seed s allows the same message to admit multiple pure quantum encodings and enables efficient regeneration of a fixed encoding. Relying on measurement-induced randomness would not suffice, as it does not permit reproducible preparation of the same quantum state.

[10] One can also define a quantum analogue of this notion. In this work, we focus on classical oracle access.

[11] When the security guarantee holds for all $t = \mathsf{poly}(\lambda)$, we omit the parameter t from the notation.

– **Single-query vs. multi-query pseudorandomness.** By default, pseudorandomness refers to indistinguishability even when the adversary is given oracle access to polynomially many independent encodings. One may also define a weaker notion, called a *single-query QPRC*, where the adversary makes a single encoding-oracle call and receives t copies of that call's codeword.
  An even weaker notion, which we call 1QPRC, is a single-query QPRC in which the adversary is restricted to a single oracle query and receives only one encoding, mirroring the notion of single-copy pseudorandom states (1PRS). To match the nontrivial stretching notion of 1PRS used below, we additionally require deterministic encoding ($r = 0$) and $n(\lambda) > \ell(\lambda)$ for all sufficiently large λ.
– **Computational vs. statistical security.** A QPRC is said to be *statistically secure* if the pseudorandomness requirement holds even against computationally unbounded adversaries. Otherwise, security is understood to be computational with respect to quantum polynomial-time adversaries.

In the quantum setting, an error-correcting code needs to correct bit flip (X) or phase flip (Z) errors or a combination thereof. We specify in each section which errors the error channel can apply.

5.1 Single Query QPRC for Bounded Z-Errors

We present a construction of a quantum pseudorandom code (QPRC) that is perfectly robust against a bounded fraction of Z-errors. The construction is single-query secure and its security follows from the pseudorandomness of the underlying quantum states. The construction makes use of the following two primitives:

1. a classical error-correcting code $\mathsf{ECC} = (\mathsf{ECC.Setup}, \mathsf{ECC.Encode}, \mathsf{ECC.Decode})$ that encodes d-bit strings into $n = n(d)$-bit codewords, and
2. a pseudorandom function (PRF) $\mathsf{PRF} = (\mathsf{PRF.Setup}, \mathsf{PRF.Eval})$ with range $\{0,1\}$.

$\mathsf{KeyGen}(1^\lambda, 1^d)$**:** Run $\mathsf{ecc.pp} \leftarrow \mathsf{ECC.Setup}(1^\lambda, 1^d)$ to generate parameters for the error-correcting code, and $\mathsf{prf.k} \leftarrow \mathsf{PRF.Setup}(1^\lambda, 1^n)$ to generate a PRF key. Output the secret key $\mathsf{sk} = (\mathsf{ecc.pp}, \mathsf{prf.k})$.

$\mathsf{Encode}(\mathsf{sk}, m \in \{0,1\}^d)$**:** Compute the classical codeword $x \leftarrow \mathsf{ECC.Encode}(\mathsf{ecc.pp}, m)$. First prepare the uniform superposition

$$|\psi_1\rangle = \frac{1}{\sqrt{2^n}} \sum_{y \in \{0,1\}^n} |y\rangle.$$

Apply the PRF-dependent phase operation to obtain

$$|\psi_2\rangle = \frac{1}{\sqrt{2^n}} \sum_{y \in \{0,1\}^n} (-1)^{\mathsf{PRF.Eval}(\mathsf{prf.k}, y)} |y\rangle.$$

Finally, apply the operator Z^x and output

$$|\psi\rangle = \frac{1}{\sqrt{2^n}} \sum_{y \in \{0,1\}^n} (-1)^{x \cdot y + \mathsf{PRF.Eval}(\mathsf{prf.k}, y)} |y\rangle.$$

$\mathsf{Decode}(\mathsf{sk}, \rho')$**:** Apply the unitary

$$\mathrm{G}_{\mathsf{prf.k}} : |z\rangle \mapsto (-1)^{\mathsf{PRF.Eval}(\mathsf{prf.k}, z)} |z\rangle, \qquad z \in \{0,1\}^n,$$

to obtain $\rho = \mathrm{G}_{\mathsf{prf.k}}\, \rho'\, \mathrm{G}_{\mathsf{prf.k}}^\dagger$. Then apply the Hadamard transform on all qubits, $\sigma = \mathrm{H}^{\otimes n}\, \rho\, \mathrm{H}^{\otimes n}$. Measure σ in the computational basis to obtain x', and output $m' \leftarrow \mathsf{ECC.Decode}(\mathsf{ecc.pp}, x')$.

Perfect Robustness. Let α denote the fraction of errors decodable by the ECC. Suppose the encoding algorithm outputs $|\psi\rangle$ and the received state is $|\psi'\rangle = Z^e |\psi\rangle$, where $e \in \{0,1\}^n$ indicates the locations of Z-errors. Then

$$|\psi'\rangle = \frac{1}{\sqrt{2^n}} \sum_{y\in\{0,1\}^n} (-1)^{(x+e)\cdot y + \mathsf{PRF.Eval}(\mathsf{prf.k},y)} |y\rangle.$$

Applying $\mathsf{G}_{\mathsf{prf.k}}$ cancels the PRF-dependent phase and yields

$$|\phi\rangle = \frac{1}{\sqrt{2^n}} \sum_{y\in\{0,1\}^n} (-1)^{(x+e)\cdot y} |y\rangle.$$

Applying $\mathsf{H}^{\otimes n}$ gives $|x+e\rangle$. Measuring in the computational basis yields $x' = x + e$, which can be decoded correctly as long as $\mathsf{Ham}(e) \leq \alpha n$. Hence the scheme is perfectly robust against an α-fraction of Z-errors.

Security. We prove single-query pseudorandomness via a hybrid argument.

Hybrid Hyb_0**:** This is the real construction, where the phase is determined by $\mathsf{PRF.Eval}(\mathsf{prf.k}, \cdot)$.
Hybrid Hyb_1**:** Replace $\mathsf{PRF.Eval}(\mathsf{prf.k}, \cdot)$ with a truly random function $g : \{0,1\}^n \to \{0,1\}$. By the quantum security of the PRF, no quantum polynomial-time adversary can distinguish Hybrid 0 from Hybrid 1 except with negligible probability.
Hybrid Hyb_2**:** Replace the resulting quantum codewords with Haar-random n-qubit pure states.

Conditioned on Hybrid 1, the encoding outputs states of the form

$$|\psi\rangle = \frac{1}{\sqrt{2^n}} \sum_{y\in\{0,1\}^n} (-1)^{f(y)} |y\rangle,$$

where $f(y) = g(y) + x \cdot y \pmod 2$ remains a uniformly random Boolean function for each fixed message and public code. The multi-copy binary-phase theorem of [BS19] bounds the trace distance between the averaged t-copy state and the Haar t-copy average by $4t^2/2^n$. This is negligible for every polynomial t when $n = \omega(\log \lambda)$. For $t = 1$, the averages are exactly equal to $\mathbb{I}_{2^n}/2^n$.

Combining the above, the construction is a secure single-query QPRC.

Remark 3 (*Deterministic encoding and single-query security*). In contrast to the classical definition of pseudorandom codes [CG24], our encoding algorithm is deterministic for a fixed secret key and message, and thus always outputs the same quantum codeword for the same message. This does not violate our security notion, since we only require *single-query* pseudorandomness. In particular, the adversary is restricted to making a single encoding query, and hence never observes two encodings of the same message. Under this restriction, the output state is computationally indistinguishable from a Haar-random state, which suffices for security in our setting.

Remark 4 (*Failure of multi-query security for distinct messages*). Even if the adversary is restricted to making multiple queries on *distinct* messages, the above construction does not satisfy multi-query pseudorandomness. To see this, consider Hybrid Hyb_1, in which the PRF is replaced by a truly random function $f : \{0,1\}^n \to \{0,1\}$. In this hybrid, the encoding of a message m produces the state

$$|\psi_m\rangle = \frac{1}{\sqrt{2^n}} \sum_{y\in\{0,1\}^n} (-1)^{\mathsf{ECC.Encode}(\mathsf{ecc.pp},m)\cdot y + f(y)} |y\rangle = Z^{x_m} |\psi_f\rangle,$$

where $x_m := \mathsf{ECC.Encode}(\mathsf{ecc.pp}, m)$ and $|\psi_f\rangle$ is a fixed pseudorandom phase state determined solely by f. Thus, encodings of different messages correspond to applying Z^{x_m} to the *same* underlying state $|\psi_f\rangle$. Crucially, the strings x_m range over the image of the error-correcting code ECC,

which is a highly structured subset of $\{0,1\}^n$ with only d bits of entropy. As a result, the family of states $\{|\psi_m\rangle\}$ lies in a highly constrained orbit of $|\psi_f\rangle$ under Z-operators. An information-theoretic adversary, given polynomially many such states corresponding to distinct messages, can exploit this structure to distinguish the resulting joint state from a collection of independently sampled Haar-random states.

Remark 5. By conjugating the scheme with Hadamard gates, we obtain an analogous construction that corrects X-errors. Specifically, applying $\mathrm{H}^{\otimes n}$ to each codeword converts Z-errors into X-errors, using the identity $\mathrm{X} = \mathrm{H}\,\mathrm{Z}\,\mathrm{H}$. The modified decoder first applies $\mathrm{H}^{\otimes n}$ to the received state and then runs the original decoder. Security is preserved by unitary invariance of the Haar distribution.

5.2 Multi-Query QPRC for Unbounded Z-Errors and Few X-Errors

In this section we observe that the *pseudorandom function-like state generator with proof of destruction* construction of Behera *et al.* [BBSS23] can be extended to obtain a *multi-query* zero-bit quantum pseudorandom code (0-bit QPRC) that is robust against *unbounded* Z-errors and against a *constant* number of X-errors (for simplicity, we present the case of a single X-error). The construction uses a strong quantum-secure pseudorandom permutation $\mathsf{PRP} = (\mathsf{PRP.Setup}, \mathsf{PRP.Eval}, \mathsf{PRP.Eval}^{-1})$ over $\{0,1\}^{5\lambda}$, with efficient evaluation and inversion and security under superposition access to both directions, assuming security against nonuniform polynomial-size adversaries.

$\mathsf{KeyGen}(1^\lambda)$**:** Sample $k \leftarrow \mathsf{PRP.Setup}(1^\lambda)$ and output k.
$\mathsf{Encode}(k; s)$**:** Output the (5λ)-qubit state

$$|\psi_{k,s}\rangle := \frac{1}{2^{\lambda/2}} \sum_{y\in\{0,1\}^\lambda} \omega_{2^{5\lambda}}^{\mathsf{PRP.Eval}(k,\, y\|s\|1^{3\lambda})}\, |\mathsf{PRP.Eval}(k,\, y\|s\|0^{3\lambda})\rangle,$$

where $s \in \{0,1\}^\lambda$ is freshly sampled for each encoding and $\omega_N = \exp(2\pi i/N)$.

$\mathsf{Decode}(k, \rho)$**:** Measure ρ in the computational basis to obtain $z \in \{0,1\}^{5\lambda}$. For each $i \in \{0, 1, \ldots, 5\lambda\}$, where $e_0 = 0^{5\lambda}$ and e_i is the ith unit vector for $i > 0$, compute

$$x_i \;:=\; \mathsf{PRP.Eval}^{-1}(k,\, z \oplus e_i),$$

and check whether x_i is of the form $x_i = y\|s\|0^{3\lambda}$ for some $y, s \in \{0,1\}^\lambda$. If the check succeeds for some i, output 1. Otherwise output 0.

Security and robustness (overview). Behera *et al.* [BBSS23] show that the multi-query ensemble $\{\mathsf{Encode}(k; s_j)^{\otimes \mathsf{poly}(\lambda)}\}_{j\in[\mathsf{poly}(\lambda)]}$ is computationally indistinguishable from $\{|\Psi_j\rangle^{\otimes \mathsf{poly}(\lambda)}\}_{j\in[\mathsf{poly}(\lambda)]}$ where $|\Psi_j\rangle \leftarrow \mathsf{Haar}_{2^{5\lambda}}$, and hence the above scheme satisfies multi-query pseudorandomness. Robustness against Z-errors is immediate since Z-errors only introduce phases and are erased by the computational-basis measurement in Decode. Robustness against a single X-error follows because X on the computational basis flips exactly one bit of the measured string, which is corrected by checking all $z \oplus e_i$.

Soundness.

Proof. Let $S = \{0,1\}^\lambda\|\{0,1\}^\lambda\|0^{3\lambda}$ and let $E = \{e_0, e_1, \ldots, e_{5\lambda}\}$. For a permutation π, define

$$a_\pi(z) = \mathbf{1}\{\exists e \in E : \pi^{-1}(z \oplus e) \in S\}.$$

For uniform π and every fixed z, a union bound gives $\mathsf{E}_\pi a_\pi(z) \le \mu := (5\lambda+1)2^{-3\lambda}$. A polynomial-time distinguisher, given z as classical advice, computes this predicate with $5\lambda+1$ inverse queries. Nonuniform PRP security therefore implies

$$\max_z \mathsf{E}_k a_{\mathsf{PRP.Eval}(k,\cdot)}(z) \le \mu + \mathrm{negl}(\lambda).$$

Indeed, a sequence of maximizing strings would otherwise be polynomial-length advice for a PRP distinguisher. For any state ρ independent of k, the measurement-first decoder satisfies

$$\Pr_k[\mathsf{Decode}(k,\rho)=1] = \sum_z \langle z|\rho|z\rangle\, \mathsf{E}_k a_{\mathsf{PRP.Eval}(k,\cdot)}(z) \le \mu + \mathrm{negl}(\lambda).$$

This proves the required soundness even when ρ is not efficiently preparable. Under only uniform-adversary PRP security, the same reduction immediately covers efficiently preparable states, but the uniform bound over all states needs the additional assumption above. The proof extends to constant c with $\mu = B_c(5\lambda)2^{-3\lambda}$ where B_c is the volume of a Hamming ball of radius c.

5.3 1QPRC from 1PRS

In this section, we present a construction of 1QPRC from 1PRS. The construction requires the following primitives.

1. a single-copy pseudorandom state generator 1PRS that takes as input a λ-bit seed and outputs an $m = m(\lambda) > \lambda$-qubit pure state.
2. a quantum $[[N, m+\lambda, d]]$ stabilizer/CSS code, implemented as an N-qubit Clifford encoder V, where $N = N(\lambda) \ge m+\lambda$. The first $m+\lambda$ qubits are data and the remaining $N-m-\lambda$ are ancillas.

$\mathsf{KeyGen}(1^\lambda)$**:** Sample $k, t \leftarrow \{0,1\}^\lambda$ and $s \leftarrow \{0,1\}^{(N-m-\lambda)}$ uniformly and independently. Output $\mathsf{sk} = (k,t,s)$.

$\mathsf{Encode}(\mathsf{sk}=(k,t,s))$**:** Prepare the m-qubit 1PRS state $|\psi_k\rangle$. Prepare $|t\rangle$ and ancillas $|s\rangle$. Apply the stabilizer encoder unitary V to all N qubits. Output the N-qubit codeword state

$$|\Psi_{k,t,s}\rangle \;:=\; V\Big(|\psi_k\rangle \otimes |t\rangle \otimes |s\rangle\Big) \;\in\; \mathcal{H}_N.$$

$\mathsf{Decode}(\mathsf{sk},\rho)$**:** Let $\mathrm{P}^{s'}$ be the N-qubit Pauli operator satisfying

$$V(\mathrm{I}_{m+\lambda}\otimes \mathrm{X}^s) = \mathrm{P}^{s'} V \qquad \text{equivalently} \qquad \mathrm{P}^{s'} = V(\mathrm{I}_{m+\lambda}\otimes \mathrm{X}^s)V^\dagger.$$

Compute $\rho' := (\mathrm{P}^{s'})^\dagger \rho\, \mathrm{P}^{s'}$ and run the decoder of the $[[N, m+\lambda, d]]$ code on ρ' to obtain an $(m+\lambda)$-qubit state $\tilde{\rho} \in \mathcal{D}(\mathcal{H}_{m+\lambda})$. Measure the last λ qubits in the computational basis to obtain $\tilde{t}$ and output 1 if $t = \tilde{t}$. Otherwise output 0.

Parameters. The length of the secret key sk is $(N-m-\lambda)+2\lambda = N-m+\lambda$ and the length of the codeword is N. Therefore, for the construction to be non-trivial (i.e., $N > |\mathsf{sk}|$), we require $m > \lambda$, i.e., the 1PRS must have expansion at least 1.

Perfect Robustness. Let

$$|\Psi_{k,t,s}\rangle = V\big(|\psi_k\rangle \otimes |t\rangle \otimes \mathrm{X}^s\,|0\rangle^{N-m-\lambda}\big).$$

Since V is Clifford, there exists a Pauli operator $\mathrm{P}^{s'}$ such that

$$|\Psi_{k,t,s}\rangle = \mathrm{P}^{s'}\; V\big(|\psi_k\rangle \otimes |t\rangle \otimes |0\rangle^{N-m-\lambda}\big).$$

After a Pauli error P^e, the state becomes

$$\mathrm{P}^e\,\mathrm{P}^{s'}\,V\big(|\psi_k\rangle\otimes|t\rangle\otimes|0\rangle^{N-m-\lambda}\big)=\pm\,\mathrm{P}^{s'}\,\mathrm{P}^e\,V\big(|\psi_k\rangle\otimes|t\rangle\otimes|0\rangle^{N-m-\lambda}\big).$$

Therefore, applying $(\mathrm{P}^{s'})^\dagger$ yields

$$\pm\,\mathrm{P}^e\,V\big(|\psi_k\rangle\otimes|t\rangle\otimes|0\rangle^{N-m-\lambda}\big).$$

If P^e has Pauli weight at most $\lfloor(d-1)/2\rfloor$, then the decoder recovers $|\psi_k\rangle\otimes|t\rangle$ and hence $\tilde{t}=t$.

Soundness. We show that for every security parameter λ and every N-qubit state ρ independent of the sampled secret key,

$$\Pr_{\mathsf{sk}\leftarrow\mathsf{KeyGen}(1^\lambda)}[\mathsf{Decode}(\mathsf{sk},\rho)=1]\le 2^{-\lambda},$$

which is negligible in λ.

Fix any $\lambda\in\mathbb{N}$ and any N-qubit state $\rho\in\mathcal{D}(\mathcal{H}_N)$. The secret key is sampled as

$$\mathsf{sk}=(k,t,s)\leftarrow\{0,1\}^\lambda\times\{0,1\}^\lambda\times\{0,1\}^{N-m-\lambda}$$

uniformly and independently. Consider the execution of the decoder $\mathsf{Decode}(\mathsf{sk},\rho)$.

The decoder first applies the Pauli correction $(\mathrm{P}^{s'})^\dagger$ (where s' is determined by (V,s)), then runs the fixed stabilizer decoding algorithm of the $[[N,m+\lambda,d]]$ code, and finally measures the last λ qubits of the resulting $(m+\lambda)$-qubit state in the computational basis to obtain an outcome $\tilde{t}\in\{0,1\}^\lambda$. The decoder outputs 1 if and only if $\tilde{t}=t$.

Crucially, the value t is not used anywhere in the decoding procedure prior to this final comparison. Therefore, for any fixed choice of (k,s), the random variable $\tilde{t}$ (which may depend on ρ, k, s, and internal measurement randomness) is statistically independent of the uniformly random string t.

We now compute the acceptance probability:

$$\begin{aligned}\Pr_{\mathsf{sk}}\big[\mathsf{Decode}(\mathsf{sk},\rho)=1\big]&=\Pr_{k,s,t}\big[t=\tilde{t}\big]\\&=\mathsf{E}_{k,s}\Big[\Pr_t\big[t=\tilde{t}\mid k,s\big]\Big].\end{aligned}$$

Fix any k,s. Since t is uniformly distributed over $\{0,1\}^\lambda$ and independent of $\tilde{t}$, we have

$$\begin{aligned}\Pr_t\big[t=\tilde{t}\mid k,s\big]&=\sum_{u\in\{0,1\}^\lambda}\Pr[\tilde{t}=u\mid k,s]\cdot\Pr[t=u]\\&=\sum_{u\in\{0,1\}^\lambda}\Pr[\tilde{t}=u\mid k,s]\cdot 2^{-\lambda}\\&\le 2^{-\lambda}.\end{aligned}$$

Taking expectation over (k,s) gives

$$\Pr_{\mathsf{sk}}\big[\mathsf{Decode}(\mathsf{sk},\rho)=1\big]\le 2^{-\lambda}.$$

Since $2^{-\lambda}$ is negligible in λ, the scheme satisfies soundness.

Theorem 10. *Assume $\{|\psi_k\rangle\}_{k\in\{0,1\}^\lambda}$ is a* 1PRS *with output length $m=m(\lambda)$ in the sense of Definition 5. Let $N(\lambda)\ge m(\lambda)+\lambda$ be any polynomially bounded function and let V be any N-qubit Clifford circuit implementing an $[[N,m+\lambda,d]]$ quantum error-correcting code. Then the scheme* $(\mathsf{KeyGen},\mathsf{Encode},\mathsf{Decode})$ *defined above is a zero-bit* 1QPRC *of length N with distinguishing advantage at most* $\mathrm{negl}(\lambda)$.

Proof. We prove single-query, single-copy pseudorandomness via a hybrid argument.

Hybrid Hyb_0**:** This is the real construction.
Hybrid Hyb_1**:** Replace $|\psi_k\rangle$ with a Haar-random m-qubit state. By the quantum security of the 1PRS, no quantum polynomial-time adversary can distinguish Hybrid 0 from Hybrid 1 except with negligible probability.
Hybrid Hyb_2**:** Replace the resulting quantum codeword with a Haar-random N-qubit pure state.

To prove that Hyb_1 and Hyb_2 are indistinguishable, we use that the adversary receives only a single copy. In this case, the acceptance probability of any (possibly unbounded) distinguisher depends only on the average output density matrix in each hybrid.

In Hybrid Hyb_1, the average (mixed) output state is

$$\begin{aligned}\mathsf{E}_{|\psi\rangle\leftarrow\text{Haar}}\,\mathsf{E}_{t,s}\left[V\Big(|\psi\rangle\langle\psi|\otimes|t,s\rangle\langle t,s|\Big)V^\dagger\right] &= V\left(\mathsf{E}_{|\psi\rangle\leftarrow\text{Haar}}\left[|\psi\rangle\langle\psi|\right]\otimes\mathsf{E}_{t,s}\left[|t,s\rangle\langle t,s|\right]\right)V^\dagger\\ &= V\left(\frac{\mathbb{I}_{2^m}}{2^m}\otimes\frac{\mathbb{I}_{2^{N-m}}}{2^{N-m}}\right)V^\dagger\\ &= \frac{\mathbb{I}_{2^N}}{2^N}.\end{aligned}$$

In Hybrid Hyb_2, averaging a Haar-random N-qubit pure state also yields $\mathbb{I}_{2^N}/2^N$. Hence, for a single-copy adversary, Hybrids Hyb_1 and Hyb_2 are perfectly indistinguishable.

Using Theorem 5, we have the following corollary.

Corollary 1. *If a secure* PRS *family with sufficiently long output and an efficiently encodable and recoverable stabilizer code with the parameters above are available, then a secure zero-bit* 1QPRC *with the corresponding error tolerance exists.*

References

AAC+24. Omar Alrabiah, Prabhanjan Ananth, Miranda Christ, Yevgeniy Dodis, and Sam Gunn. Ideal pseudorandom codes. *arXiv preprint arXiv:2411.05947*, 2024. 2

AAC+25. Omar Alrabiah, Prabhanjan Ananth, Miranda Christ, Yevgeniy Dodis, and Sam Gunn. Ideal pseudorandom codes. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, STOC '25, page 1638–1647, New York, NY, USA, 2025. Association for Computing Machinery. 11

Aar22. Scott Aaronson. My ai safety lecture for ut effective altruism, 2022. *URL https://scottaaronson.blog*, 2022. 2

ABGL25. Prabhanjan Ananth, John Bostanci, Aditya Gulati, and Yao-Ting Lin. Pseudorandomness in the (inverseless) haar random oracle model. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 138–166. Springer, 2025. 12

ACC+22. Per Austrin, Hao Chung, Kai-Min Chung, Shiuan Fu, Yao-Ting Lin, and Mohammad Mahmoody. On the impossibility of key agreements from quantum random oracles. In *Annual International Cryptology Conference*, pages 165–194. Springer, 2022. 11

AGQY22. Prabhanjan Ananth, Aditya Gulati, Luowen Qian, and Henry Yuen. Pseudorandom (function-like) quantum state generators: New definitions and applications. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022: 20th Theory of Cryptography Conference, Part I*, volume 13747 of *Lecture Notes in Computer Science*, pages 237–265, Chicago, IL, USA, November 7–10, 2022. Springer, Cham, Switzerland. 12

AQY22. Prabhanjan Ananth, Luowen Qian, and Henry Yuen. Cryptography from pseudorandom quantum states. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022, Part I*, volume 13507 of *Lecture Notes in Computer Science*, pages 208–236, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Cham, Switzerland. 3

BBSS23. Amit Behera, Zvika Brakerski, Or Sattath, and Omri Shmueli. Pseudorandomness with proof of destruction and applications. In Guy N. Rothblum and Hoeteck Wee, editors, *TCC 2023: 21st Theory of Cryptography Conference, Part IV*, volume 14372 of *Lecture Notes in Computer Science*, pages 125–154, Taipei, Taiwan, November 29 – December 2, 2023. Springer, Cham, Switzerland. 4, 9, 10, 12, 28

Bec75. William Beckner. Inequalities in fourier analysis. *Annals of Mathematics*, 102(1):159–182, 1975. 19

BEHMV25. Samuel Bouaziz-Ermann, Minki Hhan, Garazi Muguruza, and Quoc-Huy Vu. On limits on the provable consequences of quantum pseudorandomness, 2025. 4, 12

BEM24. Samuel Bouaziz-Ermann and Garazi Muguruza. Quantum pseudorandomness cannot be shrunk in a black-box way. Cryptology ePrint Archive, Report 2024/291, 2024. 4, 12

Ber25. Lorenzo Beretta. New statistical and computational results for learning junta distributions, 2025. 11

BGH+23. Khashayar Barooti, Alex B. Grilo, Loïs Huguenin-Dumittan, Giulio Malavolta, Or Sattath, Quoc-Huy Vu, and Michael Walter. Public-key encryption with quantum keys. In Guy N. Rothblum and Hoeteck Wee, editors, *TCC 2023: 21st Theory of Cryptography Conference, Part IV*, volume 14372 of *Lecture Notes in Computer Science*, pages 198–227, Taipei, Taiwan, November 29 – December 2, 2023. Springer, Cham, Switzerland. 2

BGV+23. Samuel Bouaziz, Alex B Grilo, Damien Vergnaud, Quoc-Huy Vu, et al. Towards the impossibility of quantum public key encryption with classical keys from one-way functions. *Cryptology ePrint Archive*, 2023. 11

BKSY11. Zvika Brakerski, Jonathan Katz, Gil Segev, and Arkady Yerukhimovich. Limits on the power of zero-knowledge proofs in cryptographic constructions. In *Theory of Cryptography: 8th Theory of Cryptography Conference, TCC 2011, Providence, RI, USA, March 28-30, 2011. Proceedings 8*, pages 559–578. Springer, 2011. 11

BMB+25. Nikhil Bansal, Wai-Keong Mok, Kishor Bharti, Dax Enshan Koh, and Tobias Haug. Pseudorandom density matrices. *PRX Quantum*, 6:020322, May 2025. 12

BMG09. Boaz Barak and Mohammad Mahmoody-Ghidary. Merkle puzzles are optimal—an o (n 2)-query attack on any key exchange from a random oracle. In *Annual International Cryptology Conference*, pages 374–390. Springer, 2009. 11

Bon70. Aline Bonami. Etude des coefficients de fourier des fonctions de $l^p(g)$. In *Annales de l'institut Fourier*, volume 20, pages 335–402, 1970. 19

BS19. Zvika Brakerski and Omri Shmueli. (Pseudo) random quantum states with binary phase. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019: 17th Theory of Cryptography Conference, Part I*, volume 11891 of *Lecture Notes in Computer Science*, pages 229–250, Nuremberg, Germany, December 1–5, 2019. Springer, Cham, Switzerland. 9, 12, 27

BS20. Zvika Brakerski and Omri Shmueli. Scalable pseudorandom quantum states. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology – CRYPTO 2020, Part II*, volume 12171 of *Lecture Notes in Computer Science*, pages 417–440, Santa Barbara, CA, USA, August 17–21, 2020. Springer, Cham, Switzerland. 12

CCS25. Boyang Chen, Andrea Coladangelo, and Or Sattath. The power of a single haar random state: Constructing and separating quantum pseudorandomness. In Serge Fehr and Pierre-Alain Fouque, editors, *Advances in Cryptology – EUROCRYPT 2025, Part VII*, volume 15607 of *Lecture Notes in Computer Science*, pages 108–137, Madrid, Spain, May 4–8, 2025. Springer, Cham, Switzerland. 4, 12

CG24. Miranda Christ and Sam Gunn. Pseudorandom error-correcting codes. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology – CRYPTO 2024, Part VI*, volume 14925 of *Lecture Notes in Computer Science*, pages 325–347, Santa Barbara, CA, USA, August 18–22, 2024. Springer, Cham, Switzerland. 1, 2, 8, 11, 24, 25, 27

CGG+25. Miranda Christ, Noah Golowich, Sam Gunn, Ankur Moitra, and Daniel Wichs. Improved pseudorandom codes from permuted puzzles. *arXiv preprint arXiv:2512.08918*, 2025. 2, 11

Chr26. Miranda Christ. A survey of pseudorandom codes. Cryptology ePrint Archive, Paper 2026/2162, 2026. 12

CM24. Andrea Coladangelo and Saachi Mutreja. On black-box separations of quantum digital signatures from pseudorandom states. In Elette Boyle and Mohammad Mahmoody, editors, *TCC 2024: 22nd Theory of Cryptography Conference, Part III*, volume 15366 of *Lecture Notes in Computer Science*, pages 289–317, Milan, Italy, December 2–6, 2024. Springer, Cham, Switzerland. 12

CMY26. Yu Chen, Xinyu Mao, and Hongxu Yi. Advancing pseudorandom codes: Beyond parity checks and standard-model CCA1 security. Cryptology ePrint Archive, Paper 2026/1228, 2026. 12

Col23. Andrea Coladangelo. Quantum trapdoor functions from classical one-way functions. Cryptology ePrint Archive, Report 2023/282, 2023. 9

DJK+26. Nico Döttling, Antoine Joux, Venkata Koppula, Mahesh Sreekumar Rajasree, and Hendrik Waldner. Chosen ciphertext secure pseudorandom codes in the standard model. Cryptology ePrint Archive, Paper 2026/1227, 2026. 12

DM24. Nico Döttling and Tamer Mour. On the black-box complexity of correlation intractability. In *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*, pages 40–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2024. 11

DMSR25. Nico Döttling, Anne Müller, and Mahesh Sreekumar Rajasree. Separating pseudorandom codes from local oracles. In *Theory of Cryptography Conference*, pages 225–257. Springer, 2025. 2, 4, 5, 6, 12, 17, 18, 19, 20, 24

FGJ+25. Jaiden Fairoze, Sanjam Garg, Somesh Jha, Saeed Mahloujifar, Mohammad Mahmoody, and Mingyuan Wang. Publicly-detectable watermarking for language models. *IACR Communications in Cryptology*, 1(4), 2025. 2

GG24. Surendra Ghentiyala and Venkatesan Guruswami. New constructions of pseudorandom codes. *arXiv preprint arXiv:2409.07580*, 2024. 2, 11

GGW25. Sanjam Garg, Sam Gunn, and Mingyuan Wang. Black-box crypto is useless for pseudorandom codes. In *Theory of Cryptography Conference*, pages 205–224. Springer, 2025. 2, 5, 12

GJMZ23. Sam Gunn, Nathan Ju, Fermi Ma, and Mark Zhandry. Commitments to quantum states. In Barna Saha and Rocco A. Servedio, editors, *55th Annual ACM Symposium on Theory of Computing*, pages 1579–1588, Orlando, FL, USA, June 20–23, 2023. ACM Press. 4, 12, 16

GLMY25. Aditya Gulati, Yao-Ting Lin, Tomoyuki Morimae, and Shogo Yamada. Black-box separation between pseudorandom unitaries, pseudorandom isometries, and pseudorandom function-like states. *arXiv preprint arXiv:2510.04486*, 2025. 12

GM24. Noah Golowich and Ankur Moitra. Edit distance robust watermarks for language models. *arXiv preprint arXiv:2406.02633*, 2024. 4

HLMZ26. Shengtang Huang, Xin Li, Songtao Mao, and Zhaienhe Zhou. High-rate public-key pseudorandom codes for edit errors. arXiv:2605.19402, 2026. 12

HLvA02. Nicholas J. Hopper, John Langford, and Luis von Ahn. Provably secure steganography. In Moti Yung, editor, *Advances in Cryptology – CRYPTO 2002*, volume 2442 of *Lecture Notes in Computer Science*, pages 77–92, Santa Barbara, CA, USA, August 18–22, 2002. Springer Berlin Heidelberg, Germany. 2

HY26. Min-Hsiu Hsieh and Shogo Yamada. Quantum pseudorandom error-correcting codes. arXiv:2609.38986, 2026. 8, 12

IR89. Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *Proceedings of the twenty-first annual ACM symposium on Theory of computing*, pages 44–61, 1989. 11

JLS18. Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018, Part III*, volume 10993 of *Lecture Notes in Computer Science*, pages 126–152, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Cham, Switzerland. 3, 9, 12, 16

KGW+23. John Kirchenbauer, Jonas Geiping, Yuxin Wen, Jonathan Katz, Ian Miers, and Tom Goldstein. A watermark for large language models. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 17061–17084. PMLR, 23–29 Jul 2023. 2

KJGR21. Gabriel Kaptchuk, Tushar M Jois, Matthew Green, and Aviel D Rubin. Meteor: Cryptographically secure steganography for realistic distributions. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, pages 1529–1548, 2021. 2

KKRT16. Vladimir Kolesnikov, Ranjit Kumaresan, Mike Rosulek, and Ni Trieu. Efficient batched oblivious prf with applications to private set intersection. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, CCS '16, page 818–829, New York, NY, USA, 2016. Association for Computing Machinery. 12

KNVY26. Fuyuki Kitagawa, Ryo Nishimaki, Ágí Villanyi, and Takashi Yamakawa. Towards the impossibility of imperfectly complete key agreement in the QROM. arXiv:2608.17610, 2026. 11

KQST23a. William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal. Quantum cryptography in algorithmica. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1589–1602, 2023. 2

KQST23b. William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal. Quantum cryptography in algorithmica. In Barna Saha and Rocco A. Servedio, editors, *55th Annual ACM Symposium on Theory of Computing*, pages 1589–1602, Orlando, FL, USA, June 20–23, 2023. ACM Press. 12

Kre21. William Kretschmer. Quantum pseudorandomness and classical complexity. *arXiv preprint arXiv:2103.09320*, 2021. 2, 12

KTHL24. Rohith Kuditipudi, John Thickstun, Tatsunori Hashimoto, and Percy Liang. Robust distortion-free watermarks for language models. *Transactions on Machine Learning Research*, 2024. 2

Lee26. Keewoo Lee. Pseudorandom codes from LWE. Cryptology ePrint Archive, Paper 2026/2339, 2026. 2, 12

LLLL24. Longcheng Li, Qian Li, Xingjian Li, and Qipeng Liu. How (not) to build quantum pke in minicrypt. In *Annual International Cryptology Conference*, pages 152–183. Springer, 2024. 11

LLLL25a. Longcheng Li, Qian Li, Xingjian Li, and Qipeng Liu. Cryptomania vs minicrypt in a quantum world. *arXiv preprint arXiv:2504.05710*, 2025. 11

LLLL25b. Longcheng Li, Qian Li, Xingjian Li, and Qipeng Liu. Toward the impossibility of perfect complete quantum pke from owfs. In *16th Innovations in Theoretical Computer Science Conference (ITCS 2025)*, pages 71–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2025. 11

LLLL26. Longcheng Li, Qian Li, Xingjian Li, and Qipeng Liu. Impossibility of perfectly complete many-round key agreement in the QROM. arXiv:2608.03824, 2026. 11

LMW25. Wei-Kai Lin, Ethan Mook, and Daniel Wichs. Black box crypto is useless for doubly efficient PIR. In Serge Fehr and Pierre-Alain Fouque, editors, *Advances in Cryptology – EUROCRYPT 2025, Part VI*, volume 15606 of *Lecture Notes in Computer Science*, pages 65–93, Madrid, Spain, May 4–8, 2025. Springer, Cham, Switzerland. 12

LQWY14. Dongdai Lin, Yujuan Quan, Jian Weng, and Jun Yan. Quantum bit commitment with application in quantum zero-knowledge proof. Cryptology ePrint Archive, Report 2014/791, 2014. 2

MG24. Ankur Moitra and Noah Golowich. Edit distance robust watermarks for language models. In *Proceedings of the 38th International Conference on Neural Information Processing Systems*, NIPS '24, Red Hook, NY, USA, 2024. Curran Associates Inc. 2, 11

MM11. Takahiro Matsuda and Kanta Matsuura. On black-box separations among injective one-way functions. In *Theory of Cryptography: 8th Theory of Cryptography Conference, TCC 2011, Providence, RI, USA, March 28-30, 2011. Proceedings 8*, pages 597–614. Springer, 2011. 11

MP12. Mohammad Mahmoody and Rafael Pass. The curious case of non-interactive commitments–on the power of black-box vs. non-black-box use of primitives. In *Annual Cryptology Conference*, pages 701–718. Springer, 2012. 11

MY22. Tomoyuki Morimae and Takashi Yamakawa. Quantum commitments and signatures without one-way functions. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022, Part I*, volume 13507 of *Lecture Notes in Computer Science*, pages 269–295, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Cham, Switzerland. 3, 4, 12, 16

Mü26. Anne Müller. Lower bounds for advanced cryptographic primitives. Universität des Saarlandes, 2026. 13

O'D14. Ryan O'Donnell. *Analysis of boolean functions*. Cambridge University Press, 2014. 3, 19

Rud88. Steven Rudich. *Limits on the provable consequences of one-way functions*. University of California at Berkeley, 1988. 11

Rud91. Steven Rudich. The use of interaction in public cryptosystems. extended abstract. In *Annual International Cryptology Conference*, pages 242–251. Springer, 1991. 11

Sim98. Daniel R Simon. Finding collisions on a one-way street: Can secure hash functions be based on general assumptions? In *Advances in Cryptology—EUROCRYPT'98: International Conference on the Theory and Application of Cryptographic Techniques Espoo, Finland, May 31–June 4, 1998 Proceedings 17*, pages 334–345. Springer, 1998. 11

WWC+26. Tianrui Wang, Anyu Wang, Tianshuo Cong, Delong Ran, Jinyuan Liu, and Xiaoyun Wang. Cryptanalysis of LDPC-based pseudorandom error-correcting codes. In *35th USENIX Security Symposium (USENIX Security 26)*, pages 4287–4306. USENIX Association, 2026. 12

ZALW24. Xuandong Zhao, Prabhanjan Vijendra Ananth, Lei Li, and Yu-Xiang Wang. Provable robust watermarking for ai-generated text. In *The Twelfth International Conference on Learning Representations*, 2024. 2

ZDR19. Zachary M Ziegler, Yuntian Deng, and Alexander M Rush. Neural linguistic steganography. *arXiv preprint arXiv:1909.01496*, 2019. 2

Zha12. Mark Zhandry. How to construct quantum random functions. In *53rd Annual Symposium on Foundations of Computer Science*, pages 679–687, New Brunswick, NJ, USA, October 20–23, 2012. IEEE Computer Society Press. 4

Zha19. Mark Zhandry. How to record quantum queries, and applications to quantum indifferentiability. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part II*, volume 11693 of *Lecture Notes in Computer Science*, pages 239–268, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Cham, Switzerland. 7, 17, 35, 36, 37

A Quantum Random Oracles and Compressed Oracles

We work in the quantum random oracle model (QROM), where the adversary may query the oracle in superposition. A (classical) random oracle is a uniformly random function

$$H : \{0,1\}^m \to \{0,1\}^n.$$

Fixed and standard oracles. In the *fixed standard oracle* (FStO), a function h is sampled uniformly at random at the beginning of the experiment and the oracle deterministically implements the unitary

$$|x, y\rangle \mapsto |x, y \oplus h(x)\rangle.$$

The *standard oracle* (StO) may be viewed as a purification of the fixed standard oracle. Concretely, the oracle maintains a hidden register that is initialized in a uniform superposition over all functions

$$|H\rangle \;=\; \frac{1}{\sqrt{|\mathcal{H}|}} \sum_{h:\{0,1\}^m \to \{0,1\}^n} |h\rangle,$$

and answers queries by coherently applying the unitary

$$|x, y\rangle \otimes |h\rangle \;\longmapsto\; |x, y \oplus h(x)\rangle \otimes |h\rangle.$$

The oracle register is never revealed to the adversary and serves only as a purification of the classical randomness used to sample h. Zhandry [Zha19] showed that these two models are perfectly indistinguishable, even for quantum adversaries.

Lemma 10 ([Zha19], Lemma 2). *For any (possibly unbounded) quantum algorithm $\mathcal{A}$,*

$$\Pr[\mathcal{A}^{\mathsf{FStO}} = 1] = \Pr[\mathcal{A}^{\mathsf{StO}} = 1].$$

Compressed standard oracle. Zhandry [Zha19] further introduced the *compressed standard oracle* (CStO), which simulates the standard oracle with a hidden *database register* D. After q queries, its state is a superposition of databases with at most q entries.

Description of the compressed oracle (unitary view). We recall the compressed-oracle formalism of Zhandry [Zha19], which provides a unitary simulation of a uniformly random function $H : \{0,1\}^m \to \{0,1\}^n$. We emphasize that this is only a high-level description and avoids many low-level details so as to be informative for the reader and easy to read. We highly recommend the reader to refer to [Zha19] for the exact description.

The compressed oracle maintains a private *database register* D, initialized to the empty database $|\emptyset\rangle_D$. A basis state of D corresponds to a finite partial function

$$\mathbf{d} \subseteq \{0,1\}^m \times \{0,1\}^n$$

with distinct inputs. During the computation, D may be in a coherent superposition of such classical databases. For exposition, we will write "$\mathbf{d}(x) = \perp$" to indicate that x is *not* defined in the partial function $\mathbf{d}$ (i.e., x does not appear among the first components of pairs in $\mathbf{d}$); the symbol $\perp$ is only a notational convenience and is not an additional element of $\mathbf{d}$.

A compressed-oracle query acts on a query register X holding $x \in \{0,1\}^m$, a response register Y holding $y \in \{0,1\}^n$, and the private database register D. The oracle unitary is defined as

$$\mathsf{CStO} \;=\; \mathsf{StdDecomp}^\dagger \circ \mathsf{CStO}' \circ \mathsf{StdDecomp} \circ \mathsf{Increase},$$

where Increase is an isometry and the remaining components are unitaries on the enlarged workspace. With all workspace allocated in advance for a fixed query bound, the logical query is unitary on a common Hilbert space.

- **The isometry** Increase**.** The isometry Increase prepares additional private workspace so that the database representation can be coherently extended if needed. Concretely, Increase enlarges the oracle's private register space (e.g., by appending fresh private ancilla initialized to $|0\rangle$) without changing the logical contents represented by the database. One may view this as allocating extra "empty slots" in the internal encoding of D so that new entries can be introduced reversibly during decompression. This step is purely technical and has no observable effect on the logical database contents.
- **The decompression unitary** StdDecomp**.** The unitary StdDecomp coherently expands the oracle databases corresponding to the query input while maintaining necessary consistencies. Informally, if $\mathbf{d}(x) = \perp$ then StdDecomp coherently samples a fresh uniform value and records it in the database. The following identities specify the logical local action when there is sufficient database capacity; Increase supplies the required workspace in the sparse representation:
  - For $\mathbf{d}$ such that $\mathbf{d}(x) = \perp$, StdDecomp coherently samples a fresh uniform value and records it in the database:

$$\mathsf{StdDecomp}_x |\mathbf{d}\rangle_D = \frac{1}{\sqrt{2^n}} \sum_{y' \in \{0,1\}^n} |\mathbf{d} \cup \{(x, y')\}\rangle_D.$$

  - Conversely, for $\mathbf{d}'$ such that $\mathbf{d}'(x) = \perp$, applying $\mathsf{StdDecomp}_x$ to the uniform superposition over all extensions reverses the above transformation:

$$\mathsf{StdDecomp}_x \Big(\frac{1}{\sqrt{2^n}} \sum_{y \in \{0,1\}^n} |\mathbf{d}' \cup \{(x, y)\}\rangle_D \Big) = |\mathbf{d}'\rangle.$$

  - For every $z \in \{0,1\}^n \setminus \{0^n\}$, the corresponding Fourier superposition is invariant under $\mathsf{StdDecomp}_x$:

$$\mathsf{StdDecomp}_x \Big(\frac{1}{\sqrt{2^n}} \sum_{y \in \{0,1\}^n} (-1)^{z \cdot y} |\mathbf{d}' \cup \{(x, y)\}\rangle_D \Big) = \frac{1}{\sqrt{2^n}} \sum_{y \in \{0,1\}^n} (-1)^{z \cdot y} |\mathbf{d}' \cup \{(x, y)\}\rangle_D.$$

These identities define a self-adjoint local unitary, $\mathsf{StdDecomp}_x^\dagger = \mathsf{StdDecomp}_x$, which leaves all entries at other addresses unchanged. The sparse implementation using the additional workspace prepared by Increase is given in [Zha19].

– **The local oracle unitary** CStO'**.** The unitary CStO' performs the actual random-oracle response update on the query registers by XORing the exposed value into the response register. That is, in the decompressed representation where the value $\mathbf{d}(x)$ is available, CStO' implements

$$|x\rangle_X|y\rangle_Y \longmapsto |x\rangle_X|y \oplus \mathbf{d}(x)\rangle_Y,$$

without modifying the oracle's private workspace.
To define a unitary on the entire extended space, let CStO' act as the identity on the response register when the exposed cell is $\perp$.

– **Unitary** $\mathsf{StdDecomp}^\dagger$**.** The unitary $\mathsf{StdDecomp}^\dagger$ recompresses the oracle after the response update. It can change or erase database entries through interference.

The resulting oracle simulates the standard random oracle from the algorithm's perspective. Its joint evolution includes the private database register D. The equivalence is formalized by the following lemma.

Lemma 11 ([Zha19], Lemma 4). *For any (possibly unbounded) quantum algorithm $\mathcal{A}$,*

$$\Pr[\mathcal{A}^{\mathsf{StO}} = 1] = \Pr[\mathcal{A}^{\mathsf{CStO}} = 1].$$

The register D is not accessible to the adversary and should be viewed as part of the oracle's private workspace. By the compressed-oracle invariant in [Zha19], after a total of q oracle queries, the state of D has support only on databases containing at most q input/output pairs. Importantly, CStO is operationally equivalent to StO from the adversary's perspective, and we treat it as a black box with oracle-query access. We refer to [Zha19] for the full construction and proof of equivalence.

A.1 Initializing the Compressed Oracle

Let $\mathcal{L} = \{(x_i, y_i)\}_{i\in[t]}$ be a prescribed classical list of input–output pairs with distinct inputs. An oracle initialized with $\mathcal{L}$ represents a uniformly random function $H : \{0,1\}^m \to \{0,1\}^n$ conditioned on $H(x_i) = y_i$ for every $i \in [t]$. We store $\mathcal{L}$ in a fixed classical register L and initialize a residual compressed database D to $|\emptyset\rangle_D$. The register L is read-only, and every database in the support of D has domain disjoint from $\mathrm{Dom}(\mathcal{L})$.

The initialized oracle is implemented by a controlled unitary U_{init} acting on (X, Y, D, L). On a basis state $|x\rangle_X|y\rangle_Y|\mathbf{d}\rangle_D|\mathcal{L}\rangle_L$, its action is as follows:

1. **Predefined input.** If $x = x_i$ for some $(x_i, y_i) \in \mathcal{L}$, XOR the stored value y_i into the response register:
$$U_{\mathsf{init}}\big(|x_i\rangle_X|y\rangle_Y|\mathbf{d}\rangle_D|\mathcal{L}\rangle_L\big) = |x_i\rangle_X|y \oplus y_i\rangle_Y|\mathbf{d}\rangle_D|\mathcal{L}\rangle_L.$$
The residual database is unchanged.
2. **Unlisted input.** If $x \notin \mathrm{Dom}(\mathcal{L})$, apply the full compressed-oracle query to the residual database, including decompression, the response update, and recompression:
$$U_{\mathsf{init}}\big(|x\rangle_X|y\rangle_Y|\mathbf{d}\rangle_D|\mathcal{L}\rangle_L\big) = \mathsf{CStO}\big(|x\rangle_X|y\rangle_Y|\mathbf{d}\rangle_D\big) \otimes |\mathcal{L}\rangle_L.$$

The control depends only on the query address and the fixed list $\mathcal{L}$, so the operation extends linearly to superpositions of query inputs and database states. It preserves both the register L and the restriction of the residual database to unlisted addresses. By the equivalence of the standard and compressed oracles, this construction simulates the random function conditioned on $\mathcal{L}$.

In particular, consider the first query at an unlisted address x. Decompression maps the empty residual database to

$$\mathsf{StdDecomp}_x|\emptyset\rangle_D = \frac{1}{\sqrt{2^n}} \sum_{z\in\{0,1\}^n} |\{(x,z)\}\rangle_D.$$

After the response update and recompression, the joint state is

$$\begin{aligned} &U_{\mathsf{init}}\big(|x\rangle_X|y\rangle_Y|\emptyset\rangle_D|\mathcal{L}\rangle_L\big) \\ &\quad = \frac{1}{\sqrt{2^n}} \sum_{z\in\{0,1\}^n} |x\rangle_X|y\oplus z\rangle_Y \otimes \big(\mathsf{StdDecomp}_x^\dagger|\{(x,z)\}\rangle_D\big) \otimes |\mathcal{L}\rangle_L. \end{aligned}$$

Here $\mathsf{StdDecomp}_x^\dagger$ acts only at address x; all other residual entries remain absent.