Revisiting the Quantum Indifferentiability of Merkle-Damgård: Proof Limitations and Optimal Consistency

Xiaoning Guo[1,2,3] and Chun Guo[1,2,3]

[1]School of Cyber Science and Technology, Shandong University, Qingdao, China
[2]State Key Laboratory of Cryptography and Digital Economy Security, Shandong University, Qingdao, 266237, China
[3]Key Laboratory of Cryptologic Technology and Information Security of Ministry of Education, Shandong University, Qingdao, Shandong, 266237, China

Abstract

The quantum *indifferentiability* of the Merkle-Damgård (MD) domain extender is a foundational problem in post-quantum cryptography. Recently, Gorjan et al. identified critical flaws in existing consistency proofs for this construction, invalidating prior security guarantees.

Quantum indifferentiability proofs typically reduce to two oracle-indistinguishability games: the indistinguishability game and the consistency game. To analyze the systems arising in these games, we formalize and generalize Zhandry's compressed oracle proof technique into a modular quantum game-playing framework for systems constructed from random functions.

On the positive side, we develop an error-propagation technique within our framework to track coherence penalties and bad sampling branches for the consistency game. This yields a tight consistency bound of $O(q_s^{3/2}/2^{n/2})$, matching the optimal complexity of generic quantum collision attacks.

Benefiting from the same formalization, we also identify a fundamental obstruction in the indistinguishability game—a gap also present in Zhandry's original proof. Specifically, under sequential adaptive queries, the simulated quantum state inevitably leaks into the "bad database" subspace and yields a lower bound on the corresponding oracle deviation. This clarifies why completing the full quantum indifferentiability proof requires additional ideas.

Contents

1 Introduction **3**
- 1.1 Our Contributions . . . 4
- 1.2 Related Works . . . 5
- 1.3 Acknowledgments . . . 6

2 Preliminaries **7**
- 2.1 Quantum and Indistinguishable . . . 7
- 2.2 Quantum Indifferentiability . . . 10
- 2.3 Prefix-Free Merkle-Damgård Construction . . . 11

3 The Compressed Oracle Framework **13**
- 3.1 The Compressed Oracle . . . 13
- 3.2 The Quantum Game-Playing Framework . . . 18
- 3.3 Subspace Leakage Analysis . . . 25

4 Quantum Indifferentiability Analysis of Prefix-Free Merkle–Damgård **31**
- 4.1 System Modeling and Simulator . . . 31
- 4.2 Indistinguishability Barrier . . . 33
- 4.3 Optimal Consistency Bound . . . 40

A Deferred Proofs **51**
- A.1 Proof of Lemma 2.4 . . . 51
- A.2 Proof of Lemma 2.14 . . . 51
- A.3 Proof of Lemma 3.1 . . . 52
- A.4 Proof of Lemma 4.2 . . . 52

1 Introduction

The Random Oracle Model (ROM), introduced by Bellare and Rogaway [BR93], is a central methodology in modern cryptography. It idealizes a cryptographic hash function as a truly random function that can only be accessed through black-box queries. This abstraction greatly simplifies security analyses and has led to influential proofs for practical primitives and protocols, such as RSA-based encryption and signatures, the Full-Domain Hash paradigm, and Fiat–Shamir-type signatures [BR95; BR96; PS96]. Nevertheless, a proof in the ROM does not by itself justify replacing the random oracle with an arbitrary concrete hash function.

Real hash functions are structured algorithms built from fixed-size components. For example, the Merkle–Damgård (MD) construction [Mer90; Dam90] iterates a compression function, while the Sponge construction [Ber+07] is based on a permutation with an internal state. This structural gap is not merely a theoretical mismatch. The plain MD iteration is vulnerable to length-extension attacks: from the digest of a message, an attacker may compute the digest of a properly padded extension without knowing the original message. When such a hash function is used naively inside a larger scheme, for example as a secret-prefix MAC of the form $H(K\|M)$, this structural weakness can be lifted to an attack on the entire scheme. Thus, understanding when a structured hash construction can securely replace an ideal random oracle is a fundamental question.

The notion of indifferentiability was introduced by Maurer et al. [MRH04]. Coron et al. then used this notion as a framework for analyzing the random-oracle soundness of hash-function domain extenders [Cor+05]. In an indifferentiability proof, one compares two oracle systems. In the real world, the adversary interacts with a construction $C[f]$ together with its ideal internal primitive f. In the ideal world, the construction is replaced by a true random oracle H, while a simulator $S[H]$ emulates the internal primitive consistently with H. This framework has become a standard way to formalize when a structured hash construction behaves like an ideal random oracle. Representative classical results include the indifferentiability of strengthened or prefix-free variants of Merkle–Damgård [Cor+05; Cha+06], the indifferentiability of the Sponge construction [Ber+08], and the analysis of permutation-based compression functions and tree-based modes [DPP12].

Quantum setting. The emergence of quantum computation changes this picture. Shor's algorithm [Sho94] breaks widely used number-theoretic public-key assumptions, while Grover's algorithm [Gro96] gives a quadratic speedup for generic search. Security proofs that rely on hash functions modeled as random oracles must therefore be revisited in the Quantum Random Oracle Model (QROM) [Bon+11], where adversaries may query the random oracle in quantum superposition. The quantum setting also makes indifferentiability substantially more delicate. In the classical setting, a simulator can observe the adversary's oracle queries, lazily sample answers, and record a classical query transcript. In the quantum setting, however, the simulator must answer superposition queries coherently, without measuring the adversary's state. Consequently, the classical lazy-sampling intuition no longer applies directly.

Zhandry's compressed oracle proof technique provides a powerful way to record quantum queries coherently and has become one of the main tools for analyzing random-function oracle systems in the QROM [Zha19]. Using this technique, prior works proposed quantum indifferentiability proofs for both the prefix-free MD domain extender and the Sponge construction [Zha19; Ala+25]. However, Alagic et al. pointed out a technical flaw in the consistency part of Zhandry's proof for prefix-free MD [Ala+25]. They further suggested that this gap might be repaired using their techniques, but only at the cost of a looser security bound.

Consequently, unlike the Sponge construction, the quantum indifferentiability of prefix-free MD remains unresolved.

Practical relevance. The MD paradigm remains practically important through standardized SHA-2-based hash functions [FIPS 180-4]. In post-quantum cryptography, NIST's stateless hash-based signature standard SLH-DSA explicitly includes SHA2-based parameter sets [FIPS 205], while its stateful signature recommendations for LMS and XMSS heavily rely on SHA-256 instantiations [SP 800-208]. Moreover, SHA-2 is deeply embedded in hardware accelerators, certified cryptographic modules, embedded systems, and long-lived protocol deployments. Replacing such infrastructure is costly and slow, so SHA-2-based constructions are likely to remain in use throughout the post-quantum transition. Since the SHA-2 family follows the MD paradigm, these standardization and deployment facts show that MD-type hash functions continue to serve as critical operational backbones in the post-quantum era. This makes the quantum indifferentiability of prefix-free MD not only theoretically natural, but also practically significant:

Can the prefix-free Merkle–Damgård construction be proven quantum indifferentiable from a random oracle in the QROM?

1.1 Our Contributions

In this work, we revisit the quantum indifferentiability of the prefix-free MD construction by carefully analyzing the two oracle-indistinguishability games underlying the proof: the indistinguishability game and the consistency game. To achieve this, we unify existing compressed oracle proof techniques into a modular quantum game-playing framework. This unified approach greatly simplifies complex security proofs by reducing them to bounding a few core analytical quantities. Using this framework, we successfully establish a tight positive bound for the consistency game, but subsequently identify a fundamentally new barrier in the indistinguishability game. Our specific contributions are detailed as follows.

Modular quantum game-playing framework. First, we formalize the compressed oracle proof technique developed by Zhandry and refined in subsequent works [Zha19; HI19; Chu+21; Ala+25], generalizing it to quantum oracle-indistinguishability problems for systems constructed from random functions. Our framework closely mirrors classical game-playing proof techniques [Sho04; BR06], where lazy sampling is used to record the database of a random function and to define a set of "good" databases shared by the two worlds. In the classical paradigm, one exploits the randomness of lazy sampling and the sparsity of bad database relations to bound the probability that a database becomes bad in a single query step. By propagating this bad-event probability over all queries, one obtains an upper bound on the probability of ending in a bad database, while the remaining task is to bound the operational difference between the two worlds conditioned on the database staying good.

To adapt this intuition to the quantum setting, we use compressed oracles as a coherent analogue of classical lazy sampling and replace the set of good databases with the subspace spanned by good database basis states. Consequently, an oracle-indistinguishability game is reduced to bounding two classes of operator norms: leakage norms, which measure how the oracle operators map states from the good database subspace into the bad database subspace, and deviation norms, which measure the difference between the two

world operators restricted to the good database subspace. Crucially, we show that these operator-norm bounds can often be reduced to classical sparsity properties of the underlying good-database relation.

Optimal consistency bound. Applying this framework to the consistency game, we develop an error-propagation technique to bound the overall simulation deviation. Specifically, we track the evolution of the main state together with the propagation of bad sampling branches and coherence penalty vectors across sequential queries. We show that the main states remain identical in the two worlds, while the accumulated bad branches and penalty vectors remain negligible. This yields a tight bound on the oracle deviation restricted to the good database subspace. Suppressing the dependence on the maximum message length parameter ℓ, our analysis gives a tight consistency bound of $O(q_s^{3/2}/2^{n/2})$, matching the optimal complexity of generic quantum collision attacks [BHT98; Zha15].

Barrier in the indistinguishability game. For the indistinguishability game, our framework allows us to isolate a technical obstruction in the existing proof strategy. The relevant step in Zhandry's original proof, corresponding to Lemma 30 in [Zha19], implicitly relies on a commutativity condition between the Decode/Encode operations and the collision measurement on the simulated database D_f. However, for the MD domain extender, the same compression function is iterated across all chaining steps. As a result, Decode and Encode may actively add or delete entries in D_f, which breaks the commutativity needed by the argument.

To demonstrate this obstruction, we construct a three-stage query sequence that drives the simulated database state toward basis states on which the Decode operation aborts. We emphasize that this query sequence is not an indifferentiability attack against the MD construction itself; rather, it shows that, for the current simulator and within the present framework, the quantum state leaks into the bad database subspace and yields a lower bound on the corresponding oracle deviation, clarifying why completing its full quantum indifferentiability proof requires additional ideas.

1.2 Related Works

Ristenpart et al. [RSS11] showed that the classical indifferentiability composition theorem only applies to single-stage security games. They also showed that domain-extension schemes require stateful simulators in indifferentiability proofs. Carstens et al. [Car+18] studied perfect quantum indifferentiability. Under a quantum-information-theoretic assumption, they related perfect quantum indifferentiability of classical constructions to the existence of classical indifferentiability proofs with stateless simulators. Zhandry [Zha21] studied reset indifferentiability in the quantum setting and analyzed reset quantum indifferentiability for a one-round Sponge construction. Carolan et al. [CPZ24] studied indifferentiability with precomputation and proved that the one-round Sponge is indifferentiable with precomputation from a random oracle. Czajkowski et al. [Cza+19] developed a quantum game-playing approach for quantum indifferentiability based on the one-way-to-hiding technique. Hosoyamada and Iwata [HI19] proved quantum security of the 4-round Luby–Rackoff construction. Bhaumik et al. [Bha+24] pointed out an error in the compressed-oracle-based proof for the 4-round Luby–Rackoff construction. The impossibility result on post-quantum public indifferentiability for Merkle–Damgård [Hos26] showed that plain MD fails to satisfy weak quantum indifferentiability notions, even though the corresponding classical notions can be satisfied. Zhandry [Zha25] studied how to model

unitary oracles and analyzed indistinguishability and indifferentiability issues arising from different quantum oracle interfaces.

1.3 Acknowledgments

This paper was supported by the National Key Research and Development Program of China (grant 2023YFA1011200) and the National Natural Science Foundation of China (grant 62372274). Grammatical revisions, figure illustrations, and reference formatting were assisted by Gemini 3.1 Pro. All technical content, including theorems and cryptographic proofs, was completed independently by the human authors.

2 Preliminaries

Before diving into the specific concepts, we establish standard notations used throughout this paper. Let $\{0,1\}^b$ denote the set of all b-bit binary strings. Let $\{0,1\}^*$ denote the set of all finite non-empty binary strings. For a bit string $M \in \{0,1\}^*$, let $|M|$ denote its bit length. For two bit strings $M, M' \in \{0,1\}^*$, let $M\|M'$ denote their concatenation. For a natural number $i \in \mathbb{N}$, let $\langle i \rangle$ denote its standard binary representation. The set of all finite-length strings composed of b-bit blocks is denoted by $(\{0,1\}^b)^+$, while strings restricted to a maximum of ℓ blocks are denoted by $(\{0,1\}^b)^{\leq \ell}$. For any function f, $\mathsf{Range}(f)$ represents the range, or image, of the function.

2.1 Quantum and Indistinguishable

Let X be a finite set. We represent any finite-dimensional quantum system or register associated with X by the Hilbert space $\mathcal{H}_X \cong \mathbb{C}^{|X|}$. The elements $x \in X$ correspond to an orthonormal basis $\{|x\rangle\}_{x\in X}$ for $\mathcal{H}_X$, termed the computational basis. The norm of any vector $|\psi\rangle$ in this space is induced by the inner product, defined as $\||\psi\rangle\|_2 := \sqrt{\langle\psi|\psi\rangle}$. A pure quantum state is represented by a unit vector, satisfying $\||\psi\rangle\|_2 = 1$. More generally, a mixed quantum system is described by a density operator ρ on $\mathcal{H}_X$, which is a positive semi-definite operator with unit trace ($\mathrm{Tr}(\rho) = 1$). The space of all density operators on $\mathcal{H}_X$ is denoted as $\mathcal{D}(\mathcal{H}_X)$.

Definition 2.1 (Trace Distance)**.** *The trace distance between two density operators $\rho, \sigma \in \mathcal{D}(\mathcal{H}_X)$ is defined as:*

$$\mathsf{d}_t(\rho,\sigma) := \frac{1}{2}\|\rho-\sigma\|_1 = \frac{1}{2}\mathrm{Tr}\left(\sqrt{(\rho-\sigma)^\dagger(\rho-\sigma)}\right).$$

When restricted to pure states $\rho = |\psi\rangle\langle\psi|$ and $\sigma = |\phi\rangle\langle\phi|$, the trace distance strictly reduces to:

$$\mathsf{d}_t(|\psi\rangle\langle\psi|, |\phi\rangle\langle\phi|) = \sqrt{1-|\langle\psi|\phi\rangle|^2}.$$

Lemma 2.2 (Indistinguishability)**.** *Let $\rho, \sigma \in \mathcal{D}(\mathcal{H}_X)$ be two quantum states. The maximum distinguishing advantage (indistinguishability) of any quantum measurement between ρ and σ corresponds exactly to the trace distance $\mathsf{d}_t(\rho,\sigma)$.*

Definition 2.3 (Euclidean Distance)**.** *For any two pure quantum states $|\psi\rangle, |\phi\rangle \in \mathcal{H}_X$, their Euclidean distance $\mathsf{d}_e(|\psi\rangle, |\phi\rangle)$ is defined as the distance induced by the ℓ_2-vector norm:*

$$\mathsf{d}_e(|\psi\rangle, |\phi\rangle) := \||\psi\rangle - |\phi\rangle\|_2.$$

Lemma 2.4. *For any two pure quantum states $|\psi\rangle, |\phi\rangle \in \mathcal{H}_X$, the trace distance $\mathsf{d}_t(|\psi\rangle\langle\psi|, |\phi\rangle\langle\phi|)$ and the Euclidean distance $\mathsf{d}_e(|\psi\rangle, |\phi\rangle)$ satisfy the following relationship:*

$$\mathsf{d}_t(|\psi\rangle\langle\psi|, |\phi\rangle\langle\phi|) \leq \mathsf{d}_e(|\psi\rangle, |\phi\rangle)\sqrt{1-\frac{\mathsf{d}_e^2(|\psi\rangle, |\phi\rangle)}{4}} \leq \mathsf{d}_e(|\psi\rangle, |\phi\rangle).$$

Proof. The proof is deferred to Appendix A.1. □

A linear map $O : \mathcal{H}_X \to \mathcal{H}_Y$ is referred to as an operator from the domain space $\mathcal{H}_X$ to the target space $\mathcal{H}_Y$. When the domain and range are identical ($\mathcal{H}_X = \mathcal{H}_Y$), we simply call it an operator on $\mathcal{H}_X$.

Definition 2.5 (Operator Norm and Distance)**.** *The operator norm (spectral norm) of $O : \mathcal{H}_X \to \mathcal{H}_Y$ is defined by taking the supremum over all non-zero vectors:*

$$\|O\|_\infty := \sup_{|\psi\rangle \neq 0} \frac{\|O|\psi\rangle\|_2}{\||\psi\rangle\|_2}.$$

Consequently, the operator distance between two operators $O, O' : \mathcal{H}_X \to \mathcal{H}_Y$ is defined as:

$$\mathsf{d}_o(O, O') := \|O - O'\|_\infty.$$

To formalize operations on specific sub-structures, for any operator A and a subspace $\mathcal{V} \subseteq \mathcal{H}_X$, the image space under A is defined as the subspace $A(\mathcal{V}) := \text{span}\{A|\psi\rangle : |\psi\rangle \in \mathcal{V}\}$. Furthermore, two subspaces $\mathcal{V}$ and $\mathcal{W}$ are orthogonal, denoted as $\mathcal{V} \perp \mathcal{W}$, if $\langle v|w\rangle = 0$ for all $|v\rangle \in \mathcal{V}$ and $|w\rangle \in \mathcal{W}$.

To systematically manage state space expansions, we rigorously distinguish between two fundamental paradigms: tensor product expansion and direct sum expansion.

Tensor Product Expansion. This method models composite interactive systems. Given disjoint state spaces $\mathcal{H}_X$ and $\mathcal{H}_Y$ (referred to as subsystems), their expanded composite system is the tensor product $\mathcal{H}_{XY} := \mathcal{H}_X \otimes \mathcal{H}_Y$. A pure state in this space is denoted as $|\xi\rangle_{XY}$, and a separable state is factored as $|\phi\rangle_X|\psi\rangle_Y$. Under this expansion, an operator A natively defined on $\mathcal{H}_X$ is extended to the composite space as $A_X \otimes \mathbb{I}_Y$. For brevity, this extended global operator is denoted simply as A_X. Conversely, the inverse reduction effectively isolates the local operator A acting on $\mathcal{H}_X$.

Direct Sum Expansion. In contrast, this method structurally enlarges the dimension of a state space. State spaces are expanded by adjoining orthogonal basis states to form $\mathcal{H}_{X\cup Y} := \mathcal{H}_X \oplus \mathcal{H}_Y$, where $\mathcal{H}_X$ and $\mathcal{H}_Y$ are mutually orthogonal subspaces ($\mathcal{H}_X \perp \mathcal{H}_Y$). An operator O strictly defined on $\mathcal{H}_X$ is extended to $\mathcal{H}_{X\cup Y}$ by mapping the orthogonal complement $\mathcal{H}_Y$ strictly to the zero vector. Conversely, any global operator O on $\mathcal{H}_{X\cup Y}$ can be localized to $\mathcal{H}_X$, denoted as $O|_{\mathcal{H}_X}$, preserving its mapping on $\mathcal{H}_X$ while mapping everything outside to zero.

Measurements and Controlled Operations. Based on the tensor formalism, an orthogonal projection Π on $\mathcal{H}_X$ is a self-adjoint and idempotent operator. A partial measurement restricted strictly to the subsystem $\mathcal{H}_X$ is formalized by extending a complete set of local projections $\{\Pi_i\}$ to the composite system, yielding $\{\Pi_{i,X} \otimes \mathbb{I}_Y\}$. Similarly, a controlled operator $C_{X\to Y}$ applies a conditional transformation on $\mathcal{H}_Y$ dictated by the state of $\mathcal{H}_X$. Defined over the computational basis $\{|x\rangle\}$ of $\mathcal{H}_X$, it is written as $C_{X\to Y} := \sum_x |x\rangle\langle x|_X \otimes U_{x,Y}$.

Lemma 2.6 (Commutativity)**.** *The following commutativity properties natively hold for quantum systems:*

1. *Two controlled operators $C_{X\to Y} = \sum_x |x\rangle\langle x|_X \otimes U_{x,Y}$ and $C'_{X\to Y} = \sum_x |x\rangle\langle x|_X \otimes U'_{x,Y}$ commute if, for every $x \in X$, the corresponding target operations commute (i.e., $[U_{x,Y}, U'_{x,Y}] = 0 \implies [C_{X\to Y}, C'_{X\to Y}] = 0$).*
2. *Tensor product operators $A_X \otimes B_Y$ and $A'_X \otimes B'_Y$ commute if their respective localized subsystem operators commute individually (i.e., $[A_X, A'_X] = 0$ and $[B_Y, B'_Y] = 0 \implies [A_X \otimes B_Y, A'_X \otimes B'_Y] = 0$).*

Isometries and Partial Isometries. Regarding norm-preserving maps, an operator $S : \mathcal{H}_X \to \mathcal{H}_Y$ is an isometry if $S^\dagger S = \mathbb{I}_X$. We distinguish an isometric embedding ($\dim(\mathcal{H}_Y) > \dim(\mathcal{H}_X)$) from an isometric isomorphism ($\dim(\mathcal{H}_Y) = \dim(\mathcal{H}_X)$). A unitary operator U is exclusively a special case of the latter where $\mathcal{H}_X = \mathcal{H}_Y$. When an isometric embedding injects $\mathcal{H}_X$ into a direct sum space $\mathcal{H}_{X\cup Y}$, extending its domain via the aforementioned zero-padding yields a partial isometry W. Formally, W acts as an isometry on its initial space, $W^\dagger(\mathcal{H}_{X\cup Y})$, and maps the orthogonal complement to zero. Its image in the target space is the final space, $W(\mathcal{H}_{X\cup Y})$. Crucially, the products $W^\dagger W$ and $WW^\dagger$ elegantly form the orthogonal projections onto the initial and final spaces, respectively.

To facilitate indistinguishability bounds, we formalize the computation of the operator norm.

Lemma 2.7 (Properties of the Operator Norm)**.** *For any operators A, B, vector $|\psi\rangle$, and appropriately dimensioned Hilbert spaces, the operator norm satisfies:*

1. *Submultiplicativity: $\|AB\|_\infty \leq \|A\|_\infty \|B\|_\infty$.*
2. *Bounded Operator Action: $\|A|\psi\rangle\|_2 \leq \|A\|_\infty \||\psi\rangle\|_2$.*
3. *Contractivity: $\|O\|_\infty \leq 1$ for any contractive operator O.*

It is explicitly noted that orthogonal projections, isometries, and partial isometries all fundamentally belong to the class of contractive operators, and thus their operator norms are inherently bounded by 1.

Lemma 2.8 (Orthogonal Direct Sum Reduction)**.** *Let $\mathcal{H}_X = \bigoplus_i \mathcal{H}_{X_i}$ be an orthogonal direct sum decomposition of the domain space. Let $O : \mathcal{H}_X \to \mathcal{H}_Y$ be a linear operator. Utilizing the local restriction notation, let $O_i := O|_{\mathcal{H}_{X_i}}$. If the image subspaces are mutually orthogonal (i.e., $O(\mathcal{H}_{X_i}) \perp O(\mathcal{H}_{X_j})$ for all $i \neq j$), then the operator norm simplifies exactly to:*

$$\|O\|_\infty = \max_i \|O_i\|_\infty.$$

Quantum Query Model Before defining the quantum query model, we first specify how a classical function is embedded as an operator in a quantum system. Let $f : X \to Y$ be a classical function. To evaluate f reversibly, we define its quantum oracle as a unitary operator O_f acting on the composite system $\mathcal{H}_X \otimes \mathcal{H}_Y$. For all $x \in X$ and $y \in Y$, its action is defined as:

$$O_f|x\rangle|y\rangle = |x\rangle|y + f(x)\rangle,$$

where $+$ denotes the group operation of Y, which is typically the bitwise XOR operation ($\oplus$) over the vector space $\mathbb{F}_2^m$.

When modeling a quantum algorithm that adaptively makes q_i queries to each of n quantum oracles O_i ($i \in \{0, \ldots, n-1\}$), intermediate measurements and adaptive choices can be eliminated via the deferred measurement principle [NC10], encoding techniques, and virtual queries [DFH22]. Accounting for the multi-oracle encoding mechanism, the algorithm is non-adaptively modeled as an ordered sequence of $t = n\sum_{i=0}^{n-1} q_i$ oracle unitaries interleaved with t internal unitaries, starting from an arbitrary initial state $|\Psi_0\rangle$.

The full system is partitioned into three registers: the adversary's input-output register $\mathcal{H}_{XY}$, workspace register $\mathcal{H}_W$, and a simulator ancilla register $\mathcal{H}_C$. Oracle queries act only on $\mathcal{H}_{XY}$, while internal unitaries act globally across $\mathcal{H}_{XY} \otimes \mathcal{H}_W \otimes \mathcal{H}_C$. Given the encoded oracle sequence $(O'_1, \ldots, O'_t)$ where $O'_k \in \{O_0, \ldots, O_{n-1}\}$, the final state prior to measurement is:

$$|\Psi_t\rangle = U_t O'_t U_{t-1} O'_{t-1} \ldots U_1 O'_1 |\Psi_0\rangle.$$

The final measurement is performed only on $\mathcal{H}_W$ via a complete set of computational-basis projectors $\{P_w = |w\rangle\langle w|_W\}$ satisfying $\sum_w P_w = \mathbb{I}_W$. By Born's rule, the algorithm outputs $z = w$ with probability:

$$\Pr[z = w] = \left\| (\mathbb{I}_{XY} \otimes P_w \otimes \mathbb{I}_C) \, |\Psi_t\rangle \right\|_2^2 .$$

The total query complexity is exactly $t = n \sum_{i=0}^{n-1} q_i$. Since n is typically a constant, we generally omit the parameter n in the subsequent discussions.

2.2 Quantum Indifferentiability

When replacing an ideal primitive (such as a Variable-Input-Length Random Oracle, VIL-RO) with a concrete construction (such as the Merkle-Damgård domain extender) in post-quantum security proofs, standard indistinguishability is insufficient because adversaries may have direct oracle access to the underlying primitive (e.g., the compression function). We therefore employ the notion of *Quantum Indifferentiability*, adapted from Maurer et al. [MRH04] and formalized in the quantum setting by Carstens et al. [Car+18].

Definition 2.9. *A stateless quantum circuit C with quantum access to an ideal primitive f is said to be $(q_c, q_p, \ell_s, \epsilon)$-quantum indifferentiable from an ideal random oracle H, if there exists an efficient quantum simulator S, making at most ℓ_s quantum queries to H per simulated query, such that for any quantum distinguisher $\mathcal{D}$ making at most q_c queries to the construction interface and at most q_p queries to the primitive interface, it holds that:*

$$\left| \Pr[\mathcal{D}^{C^f, f}() = 1] - \Pr[\mathcal{D}^{H, S^H}() = 1] \right| \leq \epsilon.$$

The definitions for the multi-oracle variant and the computational security setting are analogous. To simplify the proof, we follow Zhandry's reduction framework, which decomposes the full quantum indifferentiability goal into two requirements: *indistinguishability* and *consistency* [Zha19].

Definition 2.10 (Indistinguishability)**.** *A simulator S is (q_1, ϵ_{ind})-indistinguishable if, for any quantum distinguisher $\mathcal{D}$ that makes at most q_1 queries exclusively to the primitive interface, it holds that:*

$$\left| \Pr[\mathcal{D}^{f}() = 1] - \Pr[\mathcal{D}^{S^H}() = 1] \right| \leq \epsilon_{ind}.$$

This property requires that the simulated primitive S^H, when evaluated on its own, looks just like the real ideal primitive f.

Definition 2.11 (Consistency)**.** *A simulator S is $(q_2, q_3, \epsilon_{con})$-consistent if, for any quantum distinguisher $\mathcal{D}$ making at most q_2 construction queries and q_3 primitive queries, it holds that:*

$$\left| \Pr[\mathcal{D}^{H, S^H}() = 1] - \Pr[\mathcal{D}^{C^{S^H}, S^H}() = 1] \right| \leq \epsilon_{con}.$$

This property ensures that the two simulated interfaces work together correctly. It means the distinguisher cannot tell whether the construction is calculated directly by the ideal oracle H, or by running the concrete circuit C using the simulated primitive S^H.

Lemma 2.12 (Lemma 6 of [Zha19])**.** *Let ℓ_c denote the maximum number of calls the circuit C makes to the primitive f. If there exists an efficient simulator S that is $(\ell_c q_c + q_p, \epsilon_{ind})$-indistinguishable and $(q_c, q_p, \epsilon_{con})$-consistent, then C is $(q_c, q_p, \ell_s, \epsilon)$-quantum indifferentiable from H, where $\epsilon = \epsilon_{ind} + \epsilon_{con}$, and ℓ_s represents the maximum number of queries S evaluates against H per simulated query.*

Proof. The proof is essentially identical to the proof of Lemma 6 in Zhandry [Zha19]. □

A fundamental and powerful property of the indifferentiability framework is its transitivity. This property ensures that if a lower-level construction securely instantiates an intermediate ideal primitive, and a higher-level construction built upon that intermediate primitive securely instantiates a final ideal primitive, then the composed construction securely instantiates the final primitive directly. This modularity is heavily utilized to simplify the security reductions of complex hash functions and domain extenders.

Lemma 2.13. *Let f_1, f_2, and f_3 be ideal primitives. Let C_1 be a stateless quantum circuit with oracle access to f_1, and C_2 be a stateless quantum circuit with oracle access to f_2. Let ℓ_{c2} denote the maximum number of quantum queries C_2 makes to its underlying primitive f_2 per single evaluation. Suppose the following conditions hold:*

- $C_2^{f_2}$ *is $(q_c, \ell_{s1} \cdot q_p, \ell_{s2}, \epsilon_2)$-quantum indifferentiable from f_3 (via an efficient simulator S_2);*
- $C_1^{f_1}$ *is $(\ell_{c2} \cdot q_c, q_p, \ell_{s1}, \epsilon_1)$-quantum indifferentiable from f_2 (via an efficient simulator S_1).*

Then the composed construction $C_2^{C_1^{f_1}}$ is $(q_c, q_p, \ell_{s1} \cdot \ell_{s2}, \epsilon_1 + \epsilon_2)$-quantum indifferentiable from f_3. The composed simulator is explicitly given by $S = S_1^{S_2}$.

Proof. The proof is essentially identical to the composition argument in Alagic et al. [Ala+25]. According to the definition of quantum indifferentiability, the result is obtained through a two-step hybrid experiment and a reduction proof. □

2.3 Prefix-Free Merkle-Damgård Construction

The standard Merkle-Damgård construction is inherently vulnerable to length-extension attacks, which trivially violate indifferentiability from a random oracle. To preclude such structural weaknesses, Coron et al. [Cor+05] established that the message preprocessing step must employ a prefix-free encoding. A padding function $\mathsf{pf} : \{0,1\}^* \to (\{0,1\}^b)^+$ is strictly prefix-free if, for any two distinct messages $M, M' \in \{0,1\}^*$, the encoded sequence $\mathsf{pf}(M)$ is never a prefix of $\mathsf{pf}(M')$. Furthermore, the function pf must be efficiently computable in both directions to allow seamless encoding and decoding. This algebraic property ensures that no valid evaluation chain can ever be a sub-chain of another valid evaluation.

Let $f : \{0,1\}^n \times \{0,1\}^b \to \{0,1\}^n$ be a fixed-input-length compression function, which serves as our ideal primitive. Given a variable-length message $M \in \{0,1\}^*$, we compute its prefix-free encoding and parse it into l blocks such that $\mathsf{pf}(M) = (m_1, \ldots, m_l)$, where each block $m_i \in \{0,1\}^b$. The prefix-free Merkle-Damgård construction $\mathsf{MD}_f(M)$ is evaluated using a fixed initial value $IV \in \{0,1\}^n$ as follows:

$$\begin{aligned} z_0 &= IV, \\ z_i &= f(z_{i-1}, m_i) \quad \text{for } i = 1, \ldots, l, \\ \mathsf{MD}_f(M) &= z_l. \end{aligned}$$

For a concrete instantiation, we define pf by prepending the message length and appending a standard terminal padding. Specifically, the first b-bit block of the padded sequence is exclusively allocated to record the bit-length of the original message M. Following the message body, a single '1' bit and a minimal sequence of '0' bits are appended to ensure the total string length aligns perfectly with the block boundaries. Formally, the prefix-free encoding is mathematically expressed as:

$$\mathsf{pf}(M) = \langle |M| \rangle_b \parallel M \parallel 1 \parallel 0^k,$$

where $\langle |M| \rangle_b$ denotes the b-bit binary representation of the message length, and $k = (-|M| - 1) \bmod b$ is the minimum non-negative integer of appended zeros required to make the total length a multiple of the block size b. Because the length indicator strictly occupies one block and the suffix padding requires at most one additional block, the fully encoded message expands the original input M by a maximum of two blocks.

Lemma 2.14. *Let f be an ideal random function defined exclusively over the range of* pf. *Define the construction C as $C^f(M) := f(\mathsf{pf}(M))$. Then, the construction C is $(q_c, q_p, 1, 0)$-quantum indifferentiable from an ideal random oracle H, where H is a random function defined over all bit strings.*

Proof. The proof is deferred to Appendix A.2. □

Based on the transitivity lemma and the indifferentiability of the prefix-free encoding established above, our primary security goal is significantly simplified. Subsequently, we only need to prove that the core construction $\mathsf{MD}^f(M)$, strictly restricted to valid encoded inputs $M \in \mathsf{Range}(\mathsf{pf})$ of length at most $\ell + 2$ blocks, is $(q_c, q_p, \ell_s, \epsilon)$-quantum indifferentiable from a random function. Provided this holds and the corresponding simulator makes at most ℓ_s queries, the transitivity property seamlessly bridges the two components. Consequently, we can directly conclude that the complete prefix-free Merkle-Damgård construction, evaluated on arbitrary bit strings of length up to ℓ blocks, is inherently $(q_c, q_p, \ell_s, \epsilon)$-quantum indifferentiable from a global random oracle.

3 The Compressed Oracle Framework

This section provides an independent restatement of the compressed oracle framework, drawing on the compressed-oracle techniques developed and used in prior works [Zha19; HI19; Chu+21; Ala+25]. It also establishes a quantum analogue of the fundamental lemma of game-playing, a principle commonly used in classical lazy-sampling proofs [Sho04; BR06]. The synthesis of these two components yields a modular proof framework for quantum oracle indistinguishability, in which the analysis is reduced to controlling bad-database leakage and operator deviation on the good database subspace.

3.1 The Compressed Oracle

Zhandry's compressed oracle introduces a seminal simulation paradigm for the Quantum Random Oracle Model (QROM). This framework is characterized by two foundational properties that are critical for post-quantum security reductions:

1. **Efficiency:** The simulation processes each quantum query in polynomial time, demanding a quantum register size that scales merely linearly with the adversary's query complexity.

2. **Online Extractability:** The framework provides a direct algebraic methodology to extract the precise QROM input-output mappings acquired by the adversary from the internal state of the simulated database, incurring only negligible statistical deviation.

At its core, the simulated evolution proceeds via three conceptual phases. The first phase employs Stinespring dilation (purification) to transition the classical random function sampling into a deterministic, unitary quantum process over an expanded Hilbert space. The second phase further expands the simulator's internal state space to natively encode unqueried evaluations ($\perp$). The third phase executes a basis transformation on the expanded quantum registers. Essentially, the transition to the controlled operator in the simulation is fundamentally driven by this basis transformation. Crucially, the entirety of this linear transformation remains perfectly indistinguishable from the standard QROM from the perspective of any computationally unbounded quantum adversary.

Let $f : X \to Y$ be a uniformly random function. The standard quantum oracle O_f acts as:

$$O_f|x\rangle_X|y\rangle_Y = |x\rangle_X|y + f(x)\rangle_Y.$$

To purify O_f, we define a sequence indexer permutation $s : X \to \{0, 1, \ldots, |X| - 1\}$ and formalize the full function register $\mathcal{H}_F = \bigotimes_{i=0}^{|X|-1} \mathcal{H}_{Y_i}$, where each local register $\mathcal{H}_{Y_i} \cong \mathcal{H}_Y$. The function f is quantumly encoded via its truth table state $|f\rangle_F = \bigotimes_{i=0}^{|X|-1} |f(s^{-1}(i))\rangle_{Y_i}$. Let $\mathsf{ADD}_{Y,Y_{s(x)}}$ denote the reversible addition gate acting on the answer register Y and the $s(x)$-th function register $\mathcal{H}_{Y_{s(x)}}$, whose effect is to add the value of the $s(x)$-th register into the Y register. The purified oracle $\mathsf{p}O$ acting on $\mathcal{H}_X \otimes \mathcal{H}_Y \otimes \mathcal{H}_F$ is defined as an operator controlled by the x register, which applies a tensor-product identity transformation over Y and F multiplied by this specific addition gate:

$$\mathsf{p}O = \sum_{x \in X} |x\rangle\langle x|_X \otimes \Big(\mathsf{ADD}_{Y,Y_{s(x)}} (\mathbb{I}_Y \otimes \mathbb{I}_F)\Big).$$

Furthermore, letting $|*\rangle = \frac{1}{\sqrt{|Y|}} \sum_{y \in Y} |y\rangle$ denote the uniform superposition state over Y, we formalize the extended output space as $\overline{Y} := Y \cup \{\perp\}$. The partial function quantum

system is then formalized as the tensor product $\mathcal{H}_P = \bigotimes_{i=0}^{|X|-1} \mathcal{H}_{\overline{Y}_i}$, where each local register $\mathcal{H}_{\overline{Y}_i} \cong \mathcal{H}_{\overline{Y}}$. We define the fundamental unitary operator $C_{\overline{Y}}$ on $\mathcal{H}_{\overline{Y}}$ and the controlled operator LC_{XP} on $\mathcal{H}_X \otimes \mathcal{H}_P$ respectively as:

$$C_{\overline{Y}} = \mathbb{I}_Y - |*\rangle\langle *| + |\perp\rangle\langle *| + |*\rangle\langle \perp|,$$
$$LC_{XP} = \sum_{x \in X} |x\rangle\langle x|_X \otimes \mathbb{I}_{\overline{Y}}^{\otimes s(x)} \otimes C_{\overline{Y}} \otimes \mathbb{I}_{\overline{Y}}^{\otimes(|X|-s(x)-1)}.$$

Let D denote the classical database space, initialized with the empty database $\emptyset$ and formally defined as:

$$D := \{\emptyset\} \cup \left\{ \{(x_i, y_i)\}_{i=1}^{k} \;\middle|\; k \leq |X|, x_i \in X, y_i \in Y, \text{ and } x_i \neq x_j \text{ for } i \neq j \right\}.$$

Since any partial function $p : X \to \overline{Y}$ uniquely corresponds to a database $d \in D$ via $d = \{(x, y) \mid p(x) = y \neq \perp\}$, we establish an isometric isomorphism $Enc : \mathcal{H}_P \to \mathcal{H}_D$ such that $Enc|p\rangle = |d\rangle$. Its inverse is given by $Dec := Enc^\dagger$, satisfying $Dec|d\rangle = |p\rangle$.

Consider a quantum adversary making q queries to the random oracle, modeled as a sequence of interleaved arbitrary unitaries $(U_1, \ldots, U_q)$. Let $|\Psi_0\rangle$ denote the initial state of the adversary's joint register $\mathcal{H}_X \otimes \mathcal{H}_Y \otimes \mathcal{H}_W$ (auxiliary registers are omitted for brevity). To rigorously formalize the indistinguishability of the simulation, we define six sequential state evolutions:

- **Process 1 (Standard QROM):** A uniformly random function $f : X \to Y$ is sampled. The adversary interacts purely with the standard oracle O_f, yielding the final joint state:
$$|\Psi_q^{(1)}\rangle = U_q O_f U_{q-1} \cdots O_f |\Psi_0\rangle.$$

- **Process 2 (Purified Oracle):** The function register $\mathcal{H}_F$ is initialized to the uniform superposition $|*\rangle_F^{\otimes|X|}$. The adversary interacts with $\mathsf{p}O$ over $\mathcal{H}_X \otimes \mathcal{H}_Y \otimes \mathcal{H}_W \otimes \mathcal{H}_F$, yielding the final state:
$$|\Psi_q^{(2)}\rangle = U_q \mathsf{p}O U_{q-1} \cdots \mathsf{p}O \left(|\Psi_0\rangle \otimes |*\rangle_F^{\otimes|X|} \right).$$

- **Process 3 (Expanded Purified Oracle):** The function register $\mathcal{H}_F$ is expanded to the partial function register $\mathcal{H}_P$, embedding the initial state identically as $|*\rangle_P^{\otimes|X|}$. The purified oracle is formally extended to operate over $\mathcal{H}_X \otimes \mathcal{H}_Y \otimes \mathcal{H}_P$, still denoted as $\mathsf{p}O$. Let $\mathsf{ADD}'_{Y,\overline{Y}_{s(x)}}$ denote the extended addition operator acting on the answer register Y and the $s(x)$-th partial function register $\mathcal{H}_{\overline{Y}_{s(x)}}$. Unlike the standard addition gate, $\mathsf{ADD}'_{Y,\overline{Y}_{s(x)}}$ maps the global state strictly to the zero vector if the $s(x)$-th register holds the undefined symbol $\perp$, while performing the standard reversible addition otherwise. The extended purified oracle is expressed as $\mathsf{p}O = \sum_{x \in X} |x\rangle\langle x|_X \otimes \left(\mathsf{ADD}'_{Y,\overline{Y}_{s(x)}} (\mathbb{I}_Y \otimes \mathbb{I}_P)\right)$, yielding the final state:
$$|\Psi_q^{(3)}\rangle = U_q \mathsf{p}O U_{q-1} \cdots \mathsf{p}O \left(|\Psi_0\rangle \otimes |*\rangle_P^{\otimes|X|} \right).$$

- **Process 4 (Compressed Oracle):** The partial function register $\mathcal{H}_P$ is initialized to $|\perp\rangle_P^{\otimes|X|}$. The adversary queries the conjugated oracle $C_{\overline{Y}}^{\otimes|X|} \mathsf{p}O (C_{\overline{Y}}^\dagger)^{\otimes|X|}$, resulting in the final state:
$$|\Psi_q^{(4)}\rangle = U_q \left(C_{\overline{Y}}^{\otimes|X|} \mathsf{p}O (C_{\overline{Y}}^\dagger)^{\otimes|X|} \right) U_{q-1} \cdots \left(C_{\overline{Y}}^{\otimes|X|} \mathsf{p}O (C_{\overline{Y}}^\dagger)^{\otimes|X|} \right) \left(|\Psi_0\rangle \otimes |\perp\rangle_P^{\otimes|X|} \right).$$

- **Process 5 (Controlled Compressed Oracle):** The register $\mathcal{H}_P$ is initialized to $|\perp\rangle_P^{\otimes|X|}$. The oracle from Process 4 is simplified to the equivalent operator $LC_{XP}\mathsf{p}OLC_{XP}^\dagger$, yielding the final state:

$$|\Psi_q^{(5)}\rangle = U_q \left(LC_{XP}\mathsf{p}OLC_{XP}^\dagger\right) U_{q-1} \cdots \left(LC_{XP}\mathsf{p}OLC_{XP}^\dagger\right) \left(|\Psi_0\rangle \otimes |\perp\rangle_P^{\otimes|X|}\right).$$

- **Process 6 (Efficient Compressed Oracle):** The classical database register $\mathcal{H}_D$ is initialized to $|\emptyset\rangle_D$. Oracle queries are fully simulated via the compressed operator

$$\mathsf{c}O = (EncLC_{XP}^\dagger Dec)(Enc\mathsf{p}ODec)(EncLC_{XP}Dec),$$

producing the final simulated state:

$$|\Psi_q^{(6)}\rangle = U_q \mathsf{c}OU_{q-1} \cdots \mathsf{c}O\left(|\Psi_0\rangle \otimes |\emptyset\rangle_D\right).$$

Lemma 3.1. *Let $\mathcal{M} = \{|w\rangle\langle w|\}_{w\in W}$ denote the standard projective measurement performed on the adversary's working register $\mathcal{H}_W$. The probability distributions induced by $\mathcal{M}$ at the conclusion of all six quantum processes are perfectly identical.*

Proof. The proof is deferred to Appendix A.3. □

Process 6 represents the exact algorithmic instantiation of Zhandry's compressed oracle. While we will subsequently prove that the required quantum memory for this specific simulation is bounded linearly by q, we abstract away the specific efficiency overhead of the intermediate simulation operators.

For any database $d \in D$, its size $|d|$ denotes the total number of recorded input-output pairs. Its evaluation $d(x)$ yields $y \in Y$ if the mapping $(x, y) \in d$, and returns $\perp$ otherwise. Let $D_x \subset D$ define the specific subspace of databases where the entry at x remains undefined ($d(x) = \perp$). We extend the recording semantics to permit the insertion of undefined pairs, defining the union operation such that $d_x \cup \{(x, \perp)\} = d_x$ for any $d_x \in D_x$. Under this extension, fixing an input $x \in X$ establishes a natural bijection between the tuple $(d_x, y) \in D_x \times Y$ and the global database space D.

To streamline the subsequent analysis, we establish several notational conveniences. First, we define the uniform database superposition over the x-th entry as:

$$|d_x \cup \{(x, *)\}\rangle := \frac{1}{\sqrt{|Y|}} \sum_{y\in Y} |d_x \cup \{(x, y)\}\rangle.$$

Correspondingly, we let $|d_x \cup \{(x, *^\perp)\}\rangle$ denote any state residing strictly within the defined-evaluation subspace $\text{span}\{|d_x \cup \{(x, y)\}\rangle\}_{y\in Y}$ that is orthogonal to $|d_x \cup \{(x, *)\}\rangle$. This explicit subspace restriction mathematically guarantees that $|d_x \cup \{(x, *^\perp)\}\rangle$ has identically zero amplitude on the undefined basis state $|d_x \cup \{(x, \perp)\}\rangle$. Second, by utilizing the isometric isomorphism $Enc : \mathcal{H}_P \to \mathcal{H}_D$ and its inverse $Dec := Enc^\dagger$, we map the extended operators directly onto the classical database space. We define $L := EncLC_{XP}Dec$ and the evaluation operator $P := Enc\mathsf{p}ODec$, with the adjoint naturally being $L^\dagger = EncLC_{XP}^\dagger Dec$. Consequently, the efficient compressed oracle from Process 6 is concisely and consistently expressed as $\mathsf{c}O = L^\dagger PL$.

Lemma 3.2. *Since the basis transformation $C_{\overline{Y}}$ is explicitly Hermitian ($C_{\overline{Y}} = C_{\overline{Y}}^\dagger$), the extended simulated unitaries L and $L^\dagger$ are identical self-adjoint operators ($L = L^\dagger$). Their operational action on the complete computational basis of $\mathcal{H}_X \otimes \mathcal{H}_D$ is characterized as follows. For any fixed input $x \in X$ and any database $d_x \in D_x$ where x is undefined:*

- *They symmetrically map the undefined entry to the uniform superposition over all defined entries, and vice versa:*

$$L\big(|x\rangle|d_x \cup \{(x,\perp)\}\rangle\big) = |x\rangle|d_x \cup \{(x,*)\}\rangle,$$
$$L\big(|x\rangle|d_x \cup \{(x,*)\}\rangle\big) = |x\rangle|d_x \cup \{(x,\perp)\}\rangle.$$

- *They act as the identity on the orthogonal complement state $|d_x \cup \{(x,*^{\perp})\}\rangle$:*

$$L|x\rangle|d_x \cup \{(x,*^{\perp})\}\rangle = |x\rangle|d_x \cup \{(x,*^{\perp})\}\rangle.$$

Lemma 3.3. *The operational action of the extended evaluation operator P on the complete computational basis of $\mathcal{H}_X \otimes \mathcal{H}_Y \otimes \mathcal{H}_D$ is characterized as follows:*

- *For any valid database state d where the evaluation at x is defined ($d(x) \neq \perp$), it computes the database evaluation directly into the adversary's answer register:*

$$P|x\rangle|y\rangle|d\rangle = |x\rangle|y + d(x)\rangle|d\rangle.$$

- *For any database state d where x remains undefined ($d(x) = \perp$), it maps the state strictly to the zero vector:*

$$P|x\rangle|y\rangle|d\rangle = 0.$$

Lemma 3.4. *The operational action of the fully simulated compressed oracle cO on the complete computational basis of $\mathcal{H}_X \otimes \mathcal{H}_Y \otimes \mathcal{H}_D$ is characterized as follows. For any fixed input $x \in X$:*

1. *If the database evaluation at x is undefined ($d(x) = \perp$), the resulting quantum state inherently splits into a classical lazy sampling term and a coherence vector $|\Phi^{(1)}_{coh}(x,y,d_x)\rangle$:*

$$\mathsf{cO}|x\rangle|y\rangle|d_x \cup \{(x,\perp)\}\rangle = \underbrace{\frac{1}{\sqrt{|Y|}}\sum_{y' \in Y}|x\rangle|y+y'\rangle|d_x \cup \{(x,y')\}\rangle}_{\textit{Classical Lazy Sampling}} + |\Phi^{(1)}_{coh}(x,y,d_x)\rangle,$$

   where the coherence vector is rigorously formulated as:

$$|\Phi^{(1)}_{coh}(x,y,d_x)\rangle := \frac{1}{\sqrt{|Y|}}|x\rangle|*\rangle_Y\Big(|d_x \cup \{(x,\perp)\}\rangle - |d_x \cup \{(x,*)\}\rangle\Big).$$

   Furthermore, the Euclidean norm of this coherence term evaluates exactly to $\sqrt{2/|Y|}$.

2. *If the database evaluation at x is defined as $z \in Y$ ($d(x) = z \neq \perp$), the operation similarly splits into a classical direct response and an aggregated coherence vector $|\Phi^{(2)}_{coh}(x,y,d_x,z)\rangle$:*

$$\mathsf{cO}|x\rangle|y\rangle|d_x \cup \{(x,z)\}\rangle = \underbrace{|x\rangle|y+z\rangle|d_x \cup \{(x,z)\}\rangle}_{\textit{Direct Response}} + |\Phi^{(2)}_{coh}(x,y,d_x,z)\rangle,$$

   where this aggregated coherence vector is characterized by:

$$|\Phi^{(2)}_{coh}(x,y,d_x,z)\rangle := \frac{1}{\sqrt{|Y|}}|x\rangle|y+z\rangle\Big(|d_x \cup \{(x,\perp)\}\rangle - |d_x \cup \{(x,*)\}\rangle\Big)$$
$$- \frac{1}{\sqrt{|Y|}}\mathsf{cO}\Big(|x\rangle|y\rangle|d_x \cup \{(x,\perp)\}\rangle\Big).$$

   Furthermore, the Euclidean norm of this aggregated coherence term is strictly bounded by $2\sqrt{2/|Y|}$.

Proof. By definition, $\mathrm{c}O = L^\dagger P L$. We isolate the analysis to the fixed input x and invariant entries d_x. First, we rigorously deduce the full algebraic actions of L and its adjoint $L^\dagger$ on any defined database state. Applying Gram-Schmidt orthogonalization with respect to the uniform superposition state $|d_x \cup \{(x,*)\}\rangle$, we decompose any defined state into its orthogonal components as:

$$\begin{aligned} |d_x \cup \{(x,z)\}\rangle &= \langle d_x \cup \{(x,*)\}|d_x \cup \{(x,z)\}\rangle |d_x \cup \{(x,*)\}\rangle + |d_x \cup \{(x,*^\perp)\}\rangle \\ &= \frac{1}{\sqrt{|Y|}}|d_x \cup \{(x,*)\}\rangle + |d_x \cup \{(x,*^\perp)\}\rangle. \end{aligned} \tag{1}$$

Since $L = L^\dagger$, applying these operators linearly to this decomposition and substituting the mappings from Lemma 3.2 yields their identical operational action:

$$\begin{aligned} L|x\rangle|d_x \cup \{(x,z)\}\rangle &= |x\rangle L\left(\frac{1}{\sqrt{|Y|}}|d_x \cup \{(x,*)\}\rangle + |d_x \cup \{(x,*^\perp)\}\rangle\right) && \text{(Eq. 1)} \\ &= |x\rangle\left(\frac{1}{\sqrt{|Y|}}|d_x \cup \{(x,\perp)\}\rangle + |d_x \cup \{(x,*^\perp)\}\rangle\right) && \text{(Lem. 3.2)} \\ &= |x\rangle\left(\frac{1}{\sqrt{|Y|}}|d_x \cup \{(x,\perp)\}\rangle + |d_x \cup \{(x,z)\}\rangle - \frac{1}{\sqrt{|Y|}}|d_x \cup \{(x,*)\}\rangle\right) && \text{(Eq. 1)} \\ &= |x\rangle|d_x \cup \{(x,z)\}\rangle + \frac{1}{\sqrt{|Y|}}|x\rangle\left(|d_x \cup \{(x,\perp)\}\rangle - |d_x \cup \{(x,*)\}\rangle\right) && (2) \end{aligned}$$

Case 1: $d(x) = \perp$**.** We first consider the scenario where the output value corresponding to the input register x remains undefined in the database.

$$\begin{aligned} \mathrm{c}O\big(|x\rangle|y\rangle|d_x \cup \{(x,\perp)\}\rangle\big) &= L^\dagger P L\big(|x\rangle|y\rangle|d_x \cup \{(x,\perp)\}\rangle\big) \\ &= L^\dagger P\big(|x\rangle|y\rangle|d_x \cup \{(x,*)\}\rangle\big) && \text{(Lem. 3.2)} \\ &= L^\dagger\Big(\frac{1}{\sqrt{|Y|}}\sum_{y'\in Y}|x\rangle|y+y'\rangle|d_x \cup \{(x,y')\}\rangle\Big) && \text{(Lem. 3.3)} \\ &= \frac{1}{\sqrt{|Y|}}\sum_{y'\in Y}|x\rangle|y+y'\rangle|d_x \cup \{(x,y')\}\rangle \\ &\quad + \left(\frac{1}{|Y|}\sum_{y'\in Y}|x\rangle|y+y'\rangle\right)\Big(|d_x \cup \{(x,\perp)\}\rangle - |d_x \cup \{(x,*)\}\rangle\Big) && \text{(Eq. 2)} \\ &= \frac{1}{\sqrt{|Y|}}\sum_{y'\in Y}|x\rangle|y+y'\rangle|d_x \cup \{(x,y')\}\rangle + |\Phi^{(1)}_{coh}(x,y,d_x)\rangle. \end{aligned}$$

Case 2: $d(x) = z \neq \perp$**.** We then consider the scenario where the output value corresponding to the input register x is defined as $z \in Y$ in the database. Noting that the extended evaluation operator P maps any database state where the x-th entry is undefined strictly to the zero vector according to Lemma 3.3, the component involving $|d_x \cup \{(x,\perp)\}\rangle$ completely

vanishes upon the application of P, yielding:

$$\begin{aligned}
\mathsf{cO}|x\rangle|y\rangle|d_x \cup \{(x,z)\}\rangle &= L^\dagger P L\big(|x\rangle|y\rangle|d_x \cup \{(x,z)\}\rangle\big) \\
&= L^\dagger P\big(|x\rangle|y\rangle|d_x \cup \{(x,z)\}\rangle\big) \\
&\quad + L^\dagger P\left(\frac{1}{\sqrt{|Y|}}|x\rangle|y\rangle\Big(|d_x \cup \{(x,\perp)\}\rangle - |d_x \cup \{(x,*)\}\rangle\Big)\right) && \text{(Eq. 2)} \\
&= L^\dagger\big(|x\rangle|y+z\rangle|d_x \cup \{(x,z)\}\rangle\big) - \frac{1}{\sqrt{|Y|}} L^\dagger P\big(|x\rangle|y\rangle|d_x \cup \{(x,*)\}\rangle\big) && \text{(Lem. 3.3)} \\
&= |x\rangle|y+z\rangle|d_x \cup \{(x,z)\}\rangle \\
&\quad + \frac{1}{\sqrt{|Y|}}|x\rangle|y+z\rangle\Big(|d_x \cup \{(x,\perp)\}\rangle - |d_x \cup \{(x,*)\}\rangle\Big) \\
&\quad - \frac{1}{\sqrt{|Y|}}\mathsf{cO}\big(|x\rangle|y\rangle|d_x \cup \{(x,\perp)\}\rangle\big) && \text{(Eq. 2)} \\
&= |x\rangle|y+z\rangle|d_x \cup \{(x,z)\}\rangle + |\Phi^{(2)}_{coh}(x,y,d_x,z)\rangle.
\end{aligned}$$

Norm Bounds. To conclude, we evaluate the bounds on the coherence vectors. These bounds are directly obtained by applying the Pythagorean theorem, the triangle inequality, and the contractivity of cO. □

Remark 3.5. *According to the formal definition established in Process 3, the extended purified oracle* pO *mathematically manifests as a partial isometry rather than a strict unitary operator. We do not discuss whether a partial isometry constitutes a valid quantum computing process. Specifically, since the initial state of the simulation process is strictly fixed, the dynamic quantum state immediately preceding any application of* pO *during the actual evolution invariably resides within its initial space. Restricted to this initial space, the partial isometry* pO *locally degenerates into a rigorous unitary operator, meaning that the actual simulation process solely necessitates the implementation of this unitary behavior. Furthermore, it is particularly noteworthy that since partial isometries are contractive, the composite compressed oracle* cO *operates globally as a contractive operator. This global contractivity is a critical algebraic property that enables the rigorous bounded approximations utilized during distance calculation reductions. We reiterate that while* cO *is absolutely contractive over the expanded global space, the strict fixation of the initial simulation state rigorously guarantees that the state evolution in Process 6 perfectly preserves the exact norm. Therefore, the global contractivity of the operator and the strict norm preservation of the simulated algorithmic execution coexist without any mathematical contradiction.*

Remark 3.6. *The first-order compensation term is independent of the value* y*. The parameter* y *is included here only for the purpose of tracking the generation process. Moreover, the uniform state appearing in the two coherent compensation terms is often referred to as an illegal state in prior compressed-oracle analyses [HI19; Chu+21; Ala+25]. In our formulation, this component appears because expressing the main term in the computational basis inevitably introduces a uniform-state component. Therefore, the coherent compensation terms introduce components with the opposite phase, so that this unwanted uniform contribution is canceled coherently. The remaining residual terms are included so that the operator can forget recorded entries while preserving unitarity.*

3.2 The Quantum Game-Playing Framework

We focus on the indistinguishability of quantum algorithms assisted by quantum oracles, which serves as one of the most fundamental theoretical tools in modern cryptography. To rigorously capture this concept, we define the indistinguishability metric for an adversary navigating between real and ideal oracle distributions.

Definition 3.7 (Quantum Oracle Indistinguishability)**.** *Let $\mathcal{A}$ be a quantum adversary interacting with a tuple of n function distribution oracles. Suppose $\mathcal{A}$ makes at most q_i quantum superposition queries to the i-th oracle for $i \in \{1, \ldots, n\}$. Let $\mathbf{O} = (O_1, \ldots, O_n)$ denote the tuple of oracles instantiated in the real world, and $\mathbf{O}' = (O'_1, \ldots, O'_n)$ denote the corresponding tuple of oracles in the ideal world. We say that the oracle tuples $\mathbf{O}$ and $\mathbf{O}'$ are $(q_1, \ldots, q_n, \epsilon)$-quantum indistinguishable if the difference in the binary output probabilities of $\mathcal{A}$ is bounded by:*

$$\left| \Pr[\mathcal{A}^{\mathbf{O}}() = 1] - \Pr[\mathcal{A}^{\mathbf{O}'}() = 1] \right| \leq \epsilon.$$

Remark 3.8. *When the adversary $\mathcal{A}$ is restricted to a polynomial-time computational complexity and ϵ is a negligible function (i.e., asymptotically bounded by the inverse of any polynomial), this formalizes computational quantum indistinguishability. Furthermore, if the adversary is strictly required to perform a full computational-basis measurement on its query input register prior to evaluating certain specific oracles, the fully quantum oracle indistinguishability model naturally collapses into a classical-quantum hybrid oracle indistinguishability model.*

We now formally model the quantum adversary $\mathcal{A}$ interacting with these multiple independent oracles. Specifically, the total query complexity of the algorithm is defined as $t = \sum_{i=1}^{n} q_i$. The global quantum system is formalized as the tensor product space $\mathcal{H}_X \otimes \mathcal{H}_Y \otimes \mathcal{H}_W \otimes \mathcal{H}_C \otimes \mathcal{H}_{C'}$. These constituent registers correspond respectively to the input register, the output register, the adversary's internal working register, the adversary's auxiliary register, and the oracle's auxiliary register.

As a global requirement for the oracle evaluations, we explicitly establish the underlying constructions for both the real and ideal worlds. For the real world, a finite set of m independent uniformly random functions $f_1, \ldots, f_m$ is classically sampled to instantiate the base function oracles $\{O_{f_1}, \ldots, O_{f_m}\}$. Each of the n distinct real quantum oracles O_i is explicitly instantiated as a constructed oracle defined by a fixed quantum circuit structure C_i that inherently makes internal queries to these base random oracles, formally denoted as $O_i := C_i[O_{f_1}, \ldots, O_{f_m}]$. Conversely, for the ideal world, a separate finite set of m' independent uniformly random functions $f'_1, \ldots, f'_{m'}$ is sampled to instantiate the ideal base function oracles $\{O_{f'_1}, \ldots, O_{f'_{m'}}\}$. Each corresponding ideal quantum oracle O'_i is instantiated via a modified quantum circuit structure C'_i, formally denoted as $O'_i := C'_i[O_{f'_1}, \ldots, O_{f'_{m'}}]$. Without loss of generality, we assume that the dimensions of the input and output registers $\mathcal{H}_X \otimes \mathcal{H}_Y$ are sufficiently large to natively encode the evaluation inputs and outputs for all constructed oracles across both worlds.

Remark 3.9. *The circuit-based formulation $C[O_{f_1}, \ldots, O_{f_m}]$ naturally encapsulates compressed oracles for random functions with non-uniform yet independent output distributions. Specifically, the quantum circuit C utilizes control qubits to determine which function's compressed oracle to invoke. By evaluating this routing logic in superposition, the entire composite circuit executes exactly a single query across all compressed oracles simultaneously. Consequently, this single parallel invocation directly yields a unified compressed oracle for random functions with non-uniform yet independent output distributions.*

The algorithm's execution is rigorously modeled as an ordered sequence alternating between internal unitary operators and oracle queries. Let $\{U_1, \ldots, U_t\}$ denote the sequence of arbitrary internal unitary operators, which act globally across the adversary's accessible space $\mathcal{H}_X \otimes \mathcal{H}_Y \otimes \mathcal{H}_W \otimes \mathcal{H}_C$, while executing as a strict identity transformation $\mathbb{I}_{C'}$ on the oracle's auxiliary register $\mathcal{H}_{C'}$. Conversely, the oracle operators exclusively interact with the query registers and their auxiliary register, acting strictly on the subspace $\mathcal{H}_X \otimes \mathcal{H}_Y \otimes \mathcal{H}_{C'}$. Let the ordered sequence of queried oracles be parameterized by the index tuple $(i_1, \ldots, i_t)$,

where each index $i_k \in \{1, \ldots, n\}$ and each specific oracle appears exactly q_i times throughout the entire sequence. Let $|\Psi_0\rangle$ denote the initial quantum state of the global system, where the input register $\mathcal{H}_X$, the output register $\mathcal{H}_Y$, the adversary's auxiliary register $\mathcal{H}_C$, and the oracle's auxiliary register $\mathcal{H}_{C'}$ are all strictly initialized to their respective clean all-zero states $|0\rangle$, while the adversary's internal working register $\mathcal{H}_W$ carries the initial state of the algorithm.

Based on this mathematical framework, we formally define the two global quantum evolution processes. These two processes precisely characterize the interaction tracks of the quantum adversary within the quantum indistinguishability framework across different worlds.

- **Process 0 (Real World Evolution):** The adversary interacts purely with the real constructed oracles $\{O_1, \ldots, O_n\}$ in the prescribed non-adaptive sequence, alternating seamlessly with its internal unitaries. The final global state of the algorithm after all t queries is denoted as $|\Psi^{(0)}\rangle$, mathematically expressed as the sequential operator product:
$$|\Psi^{(0)}\rangle = U_t O_{i_t} U_{t-1} \cdots U_2 O_{i_2} U_1 O_{i_1} |\Psi_0\rangle.$$
- **Process 1 (Ideal World Evolution):** The adversary identically interacts with the ideal constructed oracles $\{O'_1, \ldots, O'_n\}$ governed by the exact same query index sequence $(i_1, \ldots, i_t)$ and internal unitaries. The final global state of the algorithm after all t queries is denoted as $|\Psi^{(1)}\rangle$, mathematically expressed as the sequential operator product:
$$|\Psi^{(1)}\rangle = U_t O'_{i_t} U_{t-1} \cdots U_2 O'_{i_2} U_1 O'_{i_1} |\Psi_0\rangle.$$

To facilitate the subsequent bounding of distinguishing advantages, we seamlessly replace the standard random oracles with their compressed oracle equivalents, effectively embedding the classical evaluation history into the quantum state. Corresponding to the real and ideal worlds respectively, the global quantum system is expanded by adjoining the independent classical database registers. Specifically, we introduce the composite database spaces $\mathcal{H}_{\mathbf{D}} = \bigotimes_{j=1}^{m} \mathcal{H}_{D_j}$ for the real world and $\mathcal{H}_{\mathbf{D}'} = \bigotimes_{j=1}^{m'} \mathcal{H}_{D'_j}$ for the ideal world, where each local register natively encodes the internal database state of its respective compressed oracle. We formally define the simulated evolution processes as follows:

- **Process 0-a (Real Simulated Evolution):** The standard base random oracles are identically replaced by the efficient compressed oracles $\mathsf{c}O_{f_1}, \ldots, \mathsf{c}O_{f_m}$. Each constructed oracle is accordingly updated to operate over the expanded composite space, formally denoted as $O_i^{\mathsf{c}} := C_i[\mathsf{c}O_{f_1}, \ldots, \mathsf{c}O_{f_m}]$. The database registers are strictly initialized to the empty state $|\emptyset\rangle_{\mathbf{D}}$. The final global state of the algorithm is mathematically expressed as:
$$|\Psi^{(0a)}\rangle = U_t O_{i_t}^{\mathsf{c}} U_{t-1} \cdots U_2 O_{i_2}^{\mathsf{c}} U_1 O_{i_1}^{\mathsf{c}} \left(|\Psi_0\rangle \otimes |\emptyset\rangle_{\mathbf{D}}\right).$$
- **Process 1-a (Ideal Simulated Evolution):** Similarly, the ideal base random oracles are replaced by their corresponding efficient compressed oracles $\mathsf{c}O_{f'_1}, \ldots, \mathsf{c}O_{f'_{m'}}$. Each ideal constructed oracle is symmetrically updated to $O_i'^{\mathsf{c}} := C'_i[\mathsf{c}O_{f'_1}, \ldots, \mathsf{c}O_{f'_{m'}}]$. The respective database registers are strictly initialized to the empty state $|\emptyset\rangle_{\mathbf{D}'}$. The final global state of the algorithm is mathematically expressed as:
$$|\Psi^{(1a)}\rangle = U_t O_{i_t}'^{\mathsf{c}} U_{t-1} \cdots U_2 O_{i_2}'^{\mathsf{c}} U_1 O_{i_1}'^{\mathsf{c}} \left(|\Psi_0\rangle \otimes |\emptyset\rangle_{\mathbf{D}'}\right).$$

Lemma 3.10 (Perfect Simulation Indistinguishability). *Let $\mathcal{M} = \{|w\rangle\langle w|\}_{w \in W}$ denote the standard projective measurement performed exclusively on the adversary's internal working register $\mathcal{H}_W$. The probability distributions induced by $\mathcal{M}$ on the standard final state $|\Psi^{(0)}\rangle$ of Process 0 and the simulated final state $|\Psi^{(0a)}\rangle$ of Process 0-a are perfectly identical. Symmetrically, the measurement distributions on the final states $|\Psi^{(1)}\rangle$ and $|\Psi^{(1a)}\rangle$ are identically indistinguishable.*

Proof. This equivalence follows directly from Lemma 3.1. □

Before presenting the core game-playing lemma, we formalize the orthogonal projectors acting on the classical database registers. Recall that the computational basis states of the database register $\mathcal{H}_{\mathbf{D}}$ are natively represented by elements in the composite set $\mathbf{D} = D_1 \times \cdots \times D_m$. Let $\mathbf{S} \subseteq \mathbf{D}$ be an arbitrary subset. We define the corresponding orthogonal projector onto the subspace spanned by this subset as:

$$\Pi_{\mathbf{D}}^{\mathbf{S}} = \sum_{\mathbf{d} \in \mathbf{S}} |\mathbf{d}\rangle\langle\mathbf{d}|_{\mathbf{D}}.$$

Recall that the projector onto its orthogonal complement space is correspondingly denoted as $\neg\Pi_{\mathbf{D}}^{\mathbf{S}} = \mathbb{I} - \Pi_{\mathbf{D}}^{\mathbf{S}}$.

Lemma 3.11 (Fundamental Quantum Game-Playing Lemma). *Let $\Pi_{\mathbf{D}}^{rg_0}, \ldots, \Pi_{\mathbf{D}}^{rg_t}$ be a sequence of orthogonal projectors on the real database register $\mathcal{H}_{\mathbf{D}}$, corresponding respectively to a sequence of subsets $rg_0, \ldots, rg_t \subseteq \mathbf{D}$. Suppose there exists an isometry S mapping distinct computational basis states in $\mathbf{D}$ to distinct computational basis states in the ideal database set $\mathbf{D}'$, which specifically preserves the empty database state such that $S|\emptyset\rangle_{\mathbf{D}} = |\emptyset\rangle_{\mathbf{D}'}$. Furthermore, we require the initial orthogonal projector to strictly contain the empty database state, algebraically satisfying $\Pi_{\mathbf{D}}^{rg_0}|\emptyset\rangle_{\mathbf{D}} = |\emptyset\rangle_{\mathbf{D}}$. We define the corresponding projected subsets in the ideal world as $ig_0, \ldots, ig_t \subseteq \mathbf{D}'$, with their associated orthogonal projectors strictly constructed via the isometry as $\Pi_{\mathbf{D}'}^{ig_k} = S\Pi_{\mathbf{D}}^{rg_k}S^\dagger$ on $\mathcal{H}_{\mathbf{D}'}$. If for every sequential query step $k \in \{1, \ldots, t\}$, the simulated real oracle $O_{i_k}^{\mathsf{c}}$ and the ideal oracle $O'^{\mathsf{c}}_{i_k}$ satisfy the following three conditions, specifically the real bounded leakage condition, the ideal bounded leakage condition, and the bounded deviation condition:*

$$\left\| \neg\Pi_{\mathbf{D}}^{rg_k} O_{i_k}^{\mathsf{c}} \Pi_{\mathbf{D}}^{rg_{k-1}} \right\|_\infty \leq \epsilon_k,$$
$$\left\| \neg\Pi_{\mathbf{D}'}^{ig_k} O'^{\mathsf{c}}_{i_k} \Pi_{\mathbf{D}'}^{ig_{k-1}} \right\|_\infty \leq \epsilon'_k,$$
$$\left\| \Pi_{\mathbf{D}'}^{ig_k} \left(S O_{i_k}^{\mathsf{c}} S^\dagger - O'^{\mathsf{c}}_{i_k} \right) \Pi_{\mathbf{D}'}^{ig_{k-1}} \right\|_\infty \leq \delta_k,$$

then the real oracle tuple $\mathbf{O}$ and the ideal oracle tuple $\mathbf{O}'$ are $(q_1, \ldots, q_n, \epsilon_{\text{total}})$-quantum indistinguishable, where the total distinguishing advantage is bounded by $\epsilon_{\text{total}} = \sum_{k=1}^{t}(\epsilon_k + \epsilon'_k + \delta_k)$.

Proof. We establish this lemma through a rigorous step-by-step hybrid argument. The ultimate objective of our proof is to establish the statistical indistinguishability of the partial measurements performed exclusively on the adversary's internal working register $\mathcal{H}_W$ between the real world state $|\Psi^{(0)}\rangle$ and the ideal world state $|\Psi^{(1)}\rangle$. To systematically achieve this, we construct a progressive sequence of hybrid quantum evolution processes bridging the real and ideal simulated environments. The indistinguishability between adjacent evolution processes in this sequence is rigorously established either by demonstrating the statistical indistinguishability of the $\mathcal{H}_W$ measurements or by directly bounding the Euclidean distance between their entire global quantum states.

First, we initiate the hybrid sequence by leveraging the simulated evolution processes $|\Psi^{(0a)}\rangle$ and $|\Psi^{(1a)}\rangle$. By directly invoking Lemma 3.10, we definitively establish that Process

0 is statistically indistinguishable from Process 0-a, and symmetrically, Process 1 is statistically indistinguishable from Process 1-a.

Starting from Process 0-a, we construct a sequence of intermediate real-world evolution processes, denoted as Process 0-a_k for $k \in \{0, \ldots, t\}$.

- **Process 0-a_0:**

$$|\Psi^{(0\text{-}a_0)}\rangle = U_t O^{\mathsf{c}}_{i_t} U_{t-1} \cdots U_2 O^{\mathsf{c}}_{i_2} U_1 O^{\mathsf{c}}_{i_1} \Pi^{rg_0}_{\mathbf{D}} \left(|\Psi_0\rangle \otimes |\emptyset\rangle_{\mathbf{D}}\right).$$

- **Process 0-a_k** $(1 \leq k < t)$**:**

$$|\Psi^{(0\text{-}a_k)}\rangle = U_t O^{\mathsf{c}}_{i_t} \cdots O^{\mathsf{c}}_{i_{k+1}} U_k \Pi^{rg_k}_{\mathbf{D}} O^{\mathsf{c}}_{i_k} \cdots \Pi^{rg_1}_{\mathbf{D}} O^{\mathsf{c}}_{i_1} \Pi^{rg_0}_{\mathbf{D}} \left(|\Psi_0\rangle \otimes |\emptyset\rangle_{\mathbf{D}}\right).$$

- **Process 0-a_t:**

$$|\Psi^{(0\text{-}a_t)}\rangle = U_t \Pi^{rg_t}_{\mathbf{D}} O^{\mathsf{c}}_{i_t} U_{t-1} \cdots U_1 \Pi^{rg_1}_{\mathbf{D}} O^{\mathsf{c}}_{i_1} \Pi^{rg_0}_{\mathbf{D}} \left(|\Psi_0\rangle \otimes |\emptyset\rangle_{\mathbf{D}}\right).$$

We initially observe that the Euclidean distance between the final global states of Process 0-a and Process 0-a_0 is strictly bounded by the projection error on the initial state, formally bounded as $\| |\Psi^{(0\text{-}a)}\rangle - |\Psi^{(0\text{-}a_0)}\rangle \|_2 \leq \| \neg\Pi^{rg_0}_{\mathbf{D}} (|\Psi_0\rangle \otimes |\emptyset\rangle_{\mathbf{D}}) \|_2$.

To formally establish the initial divergence, we define the composite operational sequence $W = U_t O^{\mathsf{c}}_{i_t} \cdots U_1 O^{\mathsf{c}}_{i_1}$. Since all internal unitaries preserve the Euclidean norm and the compressed oracles are globally contractive, their sequential composition ensures the operator norm is bounded as $\|W\|_\infty \leq 1$. By factoring out this common composite operator W and applying the definition of the orthogonal complement $\neg\Pi^{rg_0}_{\mathbf{D}} = \mathbb{I} - \Pi^{rg_0}_{\mathbf{D}}$, the Euclidean distance between the final global states of Process 0-a and Process 0-a_0 is strictly computed and bounded as follows:

$$\begin{aligned}
\| |\Psi^{(0\text{-}a)}\rangle - |\Psi^{(0\text{-}a_0)}\rangle \|_2 &= \left\| W \left(|\Psi_0\rangle \otimes |\emptyset\rangle_{\mathbf{D}}\right) - W\Pi^{rg_0}_{\mathbf{D}} \left(|\Psi_0\rangle \otimes |\emptyset\rangle_{\mathbf{D}}\right) \right\|_2 \\
&= \left\| W \left(\mathbb{I} - \Pi^{rg_0}_{\mathbf{D}}\right) \left(|\Psi_0\rangle \otimes |\emptyset\rangle_{\mathbf{D}}\right) \right\|_2 \\
&= \left\| W \neg\Pi^{rg_0}_{\mathbf{D}} \left(|\Psi_0\rangle \otimes |\emptyset\rangle_{\mathbf{D}}\right) \right\|_2 \\
&\leq \|W\|_\infty \left\| \neg\Pi^{rg_0}_{\mathbf{D}} \left(|\Psi_0\rangle \otimes |\emptyset\rangle_{\mathbf{D}}\right) \right\|_2 \\
&\leq \left\| \neg\Pi^{rg_0}_{\mathbf{D}} \left(|\Psi_0\rangle \otimes |\emptyset\rangle_{\mathbf{D}}\right) \right\|_2.
\end{aligned}$$

Claim 3.12. *For any sequence step $k \in \{1, \ldots, t\}$, the Euclidean distance between the final global quantum states of Process 0-a_{k-1} and Process 0-a_k satisfies:*

$$\| |\Psi^{(0\text{-}a_{k-1})}\rangle - |\Psi^{(0\text{-}a_k)}\rangle \|_2 \leq \left\| \neg\Pi^{rg_k}_{\mathbf{D}} O^{\mathsf{c}}_{i_k} \Pi^{rg_{k-1}}_{\mathbf{D}} \right\|_\infty.$$

Proof. To formally establish this step-wise divergence, we decompose the sequence precisely at the k-th oracle evaluation. We define the composite trailing operational sequence

$$W_k = U_t O^{\mathsf{c}}_{i_t} \cdots U_{k+1} O^{\mathsf{c}}_{i_{k+1}} U_k.$$

Since all internal unitaries preserve the Euclidean norm and the compressed oracles are globally contractive, their sequential composition ensures the operator norm is bounded as $\|W_k\|_\infty \leq 1$. Simultaneously, we define the common intermediate quantum state generated immediately after the $(k-1)$-th projection as $|\Phi_{k-1}\rangle = \Pi^{rg_{k-1}}_{\mathbf{D}} O^{\mathsf{c}}_{i_{k-1}} \cdots U_1 O^{\mathsf{c}}_{i_1} \Pi^{rg_0}_{\mathbf{D}} \left(|\Psi_0\rangle \otimes |\emptyset\rangle_{\mathbf{D}}\right)$. Because $|\Phi_{k-1}\rangle$ is explicitly finalized by the projection $\Pi^{rg_{k-1}}_{\mathbf{D}}$, it inherently satisfies $\Pi^{rg_{k-1}}_{\mathbf{D}} |\Phi_{k-1}\rangle =$

$|\Phi_{k-1}\rangle$, and its Euclidean norm is bounded as $\||\Phi_{k-1}\rangle\|_2 \leq 1$. By factoring out the common composite operator W_k and applying the definition of the orthogonal complement $\neg\Pi_{\mathbf{D}}^{rg_k} = \mathbb{I} - \Pi_{\mathbf{D}}^{rg_k}$, the Euclidean distance is strictly computed and bounded as follows:

$$\begin{aligned}
\||\Psi^{(0\text{-}a_{k-1})}\rangle - |\Psi^{(0\text{-}a_k)}\rangle\|_2 &= \left\| W_k O_{i_k}^{\mathsf{c}} |\Phi_{k-1}\rangle - W_k \Pi_{\mathbf{D}}^{rg_k} O_{i_k}^{\mathsf{c}} |\Phi_{k-1}\rangle \right\|_2 \\
&= \left\| W_k \left(\mathbb{I} - \Pi_{\mathbf{D}}^{rg_k}\right) O_{i_k}^{\mathsf{c}} |\Phi_{k-1}\rangle \right\|_2 \\
&= \left\| W_k \neg\Pi_{\mathbf{D}}^{rg_k} O_{i_k}^{\mathsf{c}} \Pi_{\mathbf{D}}^{rg_{k-1}} |\Phi_{k-1}\rangle \right\|_2 \\
&\leq \|W_k\|_\infty \left\| \neg\Pi_{\mathbf{D}}^{rg_k} O_{i_k}^{\mathsf{c}} \Pi_{\mathbf{D}}^{rg_{k-1}} \right\|_\infty \||\Phi_{k-1}\rangle\|_2 \\
&\leq \left\| \neg\Pi_{\mathbf{D}}^{rg_k} O_{i_k}^{\mathsf{c}} \Pi_{\mathbf{D}}^{rg_{k-1}} \right\|_\infty .
\end{aligned}$$

□

Next, we explicitly construct the global basis transformation on the database register to define Process 0-b_t. This transformation explicitly applies the isometry S to the initial state, replaces all subsequent operators with their S-conjugated counterparts, and finally applies the adjoint operator $S^\dagger$ to restore the overall basis:

- **Process 0-$\mathbf{b_t}$:**

$$|\Psi^{(0\text{-}b_t)}\rangle = S^\dagger \left(S U_t S^\dagger\right) \cdots \left(S O_{i_1}^{\mathsf{c}} S^\dagger\right) \left(S \Pi_{\mathbf{D}}^{rg_0} S^\dagger\right) S \left(|\Psi_0\rangle \otimes |\emptyset\rangle_{\mathbf{D}}\right).$$

Since this explicitly conjugated sequence algebraically collapses back to the original operational sequence, Process 0-b_t is mathematically identical to $|\Psi^{(0\text{-}a_t)}\rangle$. Subsequently, we losslessly simplify the evolution process of Process 0-b_t to obtain the following Process 0-c_t. According to the composite system commutativity lemma and the explicitly required initial condition $S|\emptyset\rangle_{\mathbf{D}} = |\emptyset\rangle_{\mathbf{D}'}$, we algebraically eliminate the S and $S^\dagger$ surrounding the internal unitaries U_k, as well as the final adjoint operator $S^\dagger$. By structurally replacing the conjugated projectors $S\Pi_{\mathbf{D}}^{rg_k} S^\dagger$ with their defined ideal equivalents $\Pi_{\mathbf{D}'}^{ig_k}$, we completely transition the sequence into the ideal subspace:

- **Process 0-$\mathbf{c_t}$:**

$$|\Psi^{(0\text{-}c_t)}\rangle = U_t \Pi_{\mathbf{D}'}^{ig_t} \left(S O_{i_t}^{\mathsf{c}} S^\dagger\right) U_{t-1} \cdots U_1 \Pi_{\mathbf{D}'}^{ig_1} \left(S O_{i_1}^{\mathsf{c}} S^\dagger\right) \Pi_{\mathbf{D}'}^{ig_0} \left(|\Psi_0\rangle \otimes |\emptyset\rangle_{\mathbf{D}'}\right).$$

We now construct a sequence of intermediate hybrid processes, denoted as Process 0-c_k for k from $t-1$ down to 0, to progressively replace the conjugated real oracles and their corresponding ideal projectors with the pure ideal oracles from the last step down to the first.

- **Process 0-$\mathbf{c_k}$ ($1 \leq k \leq t-1$):** The operations from the $(k+1)$-th step up to the t-th step are identically replaced by the unprojected ideal oracles O'^{c}, while the earlier steps (1 to k) retain the fully projected conjugated real oracles:

$$|\Psi^{(0\text{-}c_k)}\rangle = U_t O'^{\mathsf{c}}_{i_t} \cdots U_{k+1} O'^{\mathsf{c}}_{i_{k+1}} U_k \Pi_{\mathbf{D}'}^{ig_k} \left(S O_{i_k}^{\mathsf{c}} S^\dagger\right) \cdots \Pi_{\mathbf{D}'}^{ig_0} \left(|\Psi_0\rangle \otimes |\emptyset\rangle_{\mathbf{D}'}\right).$$

- **Process 0-$\mathbf{c_0}$:** This process completely strips all intermediate projectors (except the initial one) and replaces all conjugated oracles with ideal ones, making it mathematically identical to Process 1-a:

$$|\Psi^{(0\text{-}c_0)}\rangle = U_t O'^{\mathsf{c}}_{i_t} \cdots U_1 O'^{\mathsf{c}}_{i_1} \Pi_{\mathbf{D}'}^{ig_0} \left(|\Psi_0\rangle \otimes |\emptyset\rangle_{\mathbf{D}'}\right) = |\Psi^{(1\text{-}a)}\rangle.$$

Claim 3.13. *For any sequence step $k \in \{1, \dots, t\}$, the Euclidean distance between the global quantum states of Process 0-c_k and Process 0-c_{k-1} satisfies:*

$$\| |\Psi^{(0\text{-}c_k)}\rangle - |\Psi^{(0\text{-}c_{k-1})}\rangle \|_2 \leq \epsilon'_k + \delta_k.$$

Proof. We mathematically decompose the sequence precisely at the k-th oracle evaluation. We define the composite ideal trailing operational sequence $W'_k = U_t O'^{\mathsf{c}}_{i_t} \cdots U_{k+1} O'^{\mathsf{c}}_{i_{k+1}} U_k$. Since all subsequent ideal oracles and unitaries are globally contractive, this sequential composition ensures the operator norm is bounded as $\|W'_k\|_\infty \leq 1$. We define the common intermediate quantum state generated immediately before the k-th oracle evaluation as:

$$|\Phi'_{k-1}\rangle = U_{k-1} \Pi^{ig_{k-1}}_{\mathbf{D}'} \left(S O^{\mathsf{c}}_{i_{k-1}} S^\dagger \right) \cdots U_1 \Pi^{ig_1}_{\mathbf{D}'} \left(S O^{\mathsf{c}}_{i_1} S^\dagger \right) \Pi^{ig_0}_{\mathbf{D}'} \left(|\Psi_0\rangle \otimes |\varnothing\rangle_{\mathbf{D}'} \right).$$

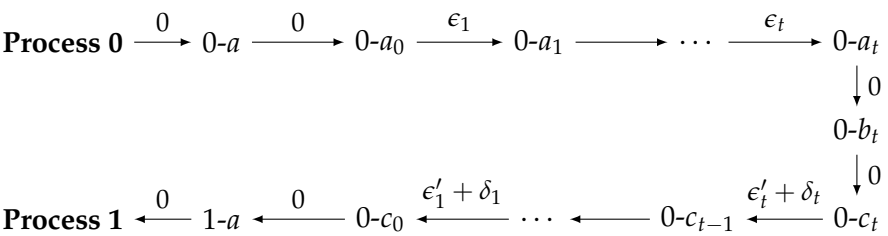
Since the internal unitary U_{k-1} acts strictly as the identity on the database register $\mathcal{H}_{\mathbf{D}'}$, it perfectly commutes with the orthogonal projector $\Pi^{ig_{k-1}}_{\mathbf{D}'}$. Thus, this intermediate state inherently resides within the ideal projected subspace, ensuring $\Pi^{ig_{k-1}}_{\mathbf{D}'} |\Phi'_{k-1}\rangle = |\Phi'_{k-1}\rangle$, and its Euclidean norm is bounded as $\| |\Phi'_{k-1}\rangle \|_2 \leq 1$.

By factoring out the composite operator W'_k and algebraically inserting the identity $\mathbb{I} = \Pi^{ig_k}_{\mathbf{D}'} + \neg \Pi^{ig_k}_{\mathbf{D}'}$ directly onto the ideal oracle $O'^{\mathsf{c}}_{i_k}$, we dynamically extract the sequence bounds using the triangle inequality alongside the established bounded deviation and ideal bounded leakage conditions:

$$\begin{aligned}
\| |\Psi^{(0\text{-}c_k)}\rangle - |\Psi^{(0\text{-}c_{k-1})}\rangle \|_2 &= \left\| W'_k \Pi^{ig_k}_{\mathbf{D}'} \left(S O^{\mathsf{c}}_{i_k} S^\dagger \right) |\Phi'_{k-1}\rangle - W'_k O'^{\mathsf{c}}_{i_k} |\Phi'_{k-1}\rangle \right\|_2 \\
&\leq \left\| \left(\Pi^{ig_k}_{\mathbf{D}'} S O^{\mathsf{c}}_{i_k} S^\dagger - O'^{\mathsf{c}}_{i_k} \right) |\Phi'_{k-1}\rangle \right\|_2 \\
&= \left\| \left(\Pi^{ig_k}_{\mathbf{D}'} S O^{\mathsf{c}}_{i_k} S^\dagger - \left(\Pi^{ig_k}_{\mathbf{D}'} + \neg \Pi^{ig_k}_{\mathbf{D}'} \right) O'^{\mathsf{c}}_{i_k} \right) \Pi^{ig_{k-1}}_{\mathbf{D}'} |\Phi'_{k-1}\rangle \right\|_2 \\
&= \left\| \left(\Pi^{ig_k}_{\mathbf{D}'} \left(S O^{\mathsf{c}}_{i_k} S^\dagger - O'^{\mathsf{c}}_{i_k} \right) - \neg \Pi^{ig_k}_{\mathbf{D}'} O'^{\mathsf{c}}_{i_k} \right) \Pi^{ig_{k-1}}_{\mathbf{D}'} |\Phi'_{k-1}\rangle \right\|_2 \\
&\leq \left\| \Pi^{ig_k}_{\mathbf{D}'} \left(S O^{\mathsf{c}}_{i_k} S^\dagger - O'^{\mathsf{c}}_{i_k} \right) \Pi^{ig_{k-1}}_{\mathbf{D}'} \right\|_\infty \| |\Phi'_{k-1}\rangle \|_2 + \left\| \neg \Pi^{ig_k}_{\mathbf{D}'} O'^{\mathsf{c}}_{i_k} \Pi^{ig_{k-1}}_{\mathbf{D}'} \right\|_\infty \| |\Phi'_{k-1}\rangle \|_2 \\
&\leq \delta_k + \left\| \neg \Pi^{ig_k}_{\mathbf{D}'} O'^{\mathsf{c}}_{i_k} \Pi^{ig_{k-1}}_{\mathbf{D}'} \right\|_\infty \\
&\leq \delta_k + \epsilon'_k.
\end{aligned}$$

□

Since Process 0-c_0 completely replaces all conjugated oracles with ideal oracles without any intermediate projections, its final quantum state is perfectly indistinguishable from the ideal state, yielding $|\Psi^{(0\text{-}c_0)}\rangle = |\Psi^{(1\text{-}a)}\rangle$. The hybrid experimental flow of the entire proof is illustrated as follows:



This completes the proof.

□

3.3 Subspace Leakage Analysis

The fundamental quantum game-playing lemma establishes that the indistinguishability between real and ideal quantum oracles can be strictly bounded by analyzing the real bounded leakage ϵ_k, the ideal bounded leakage ϵ'_k, and the bounded deviation δ_k. This section introduces standard analytical tools to evaluate these leakage bounds.

We develop these tools in two structural phases. First, we provide a sequential decomposition tool that reduces the global leakage generated by multiple queries to compressed oracles within a quantum circuit into an aggregated sum of single-query leakage bounds. Second, we introduce a transition tool that transforms the leakage bound of a single compressed oracle query directly into combinatorial properties regarding database predicates.

Recall the constructed quantum oracle $O^{\mathsf{c}} = C[\mathsf{c}O_{f_1}, \ldots, \mathsf{c}O_{f_m}]$, which acts as a fixed quantum circuit invoking m underlying independent compressed oracles. Suppose the oracle algorithm makes exactly n internal queries to these underlying compressed oracles. The execution of this quantum circuit is rigorously modeled as a fixed alternating sequence of $n+1$ internal unitary operators $V_1, \ldots, V_{n+1}$ and n compressed oracle invocations parameterized by an index sequence $(i_1, \ldots, i_n)$, where $i_k \in \{1, \ldots, m\}$. The internal unitaries V_k act globally across the adversary's query registers, the circuit's own auxiliary registers, and the composite state registers utilized for simulating the compressed oracles. Specifically, the quantum circuit's database state register is formulated as the tensor product space $\mathcal{H}_{\mathbf{D}} = \bigotimes_{j=1}^{m} \mathcal{H}_{D_j}$. Consequently, the complete operational sequence of the constructed oracle is mathematically expressed as:

$$O^{\mathsf{c}} = V_{n+1}\mathsf{c}O_{i_n}V_n \cdots V_2\mathsf{c}O_{i_1}V_1.$$

The global leakage bound of the entire constructed oracle O^{c} is evaluated by the operator norm $\left\|\neg\Pi_{\mathbf{D}}^{g_n} O^{\mathsf{c}} \Pi_{\mathbf{D}}^{g_0}\right\|_\infty$, where $\Pi_{\mathbf{D}}^{g_0}$ and $\Pi_{\mathbf{D}}^{g_n}$ represent the initial and final orthogonal projectors onto the defined safe database subspaces $g_0 \subseteq \mathbf{D}$ and $g_n \subseteq \mathbf{D}$. The following lemma demonstrates that this global multi-step leakage is bounded by the algebraic sum of individual single-step transitions across intermediate target subspaces.

Lemma 3.14 (Sequential Leakage Decomposition)**.** *Let $g_0, g_1, \ldots, g_n \subseteq \mathbf{D}$ be a chosen sequence of intermediate database subsets, and let $\Pi_{\mathbf{D}}^{g_0}, \ldots, \Pi_{\mathbf{D}}^{g_n}$ denote their corresponding orthogonal subspace projectors on $\mathcal{H}_{\mathbf{D}}$. If for every sequential query step $k \in \{1, \ldots, n\}$, the single-step transition leakage of the corresponding compressed oracle is rigorously bounded by:*

$$\left\|\neg\Pi_{\mathbf{D}}^{g_k}\mathsf{c}O_{i_k}\Pi_{\mathbf{D}}^{g_{k-1}}\right\|_\infty \leq x_k,$$

then the global transition leakage of the fully constructed oracle O^{c} is strictly upper-bounded by the sum of the individual step bounds:

$$\left\|\neg\Pi_{\mathbf{D}}^{g_n} O^{\mathsf{c}} \Pi_{\mathbf{D}}^{g_0}\right\|_\infty \leq \sum_{k=1}^{n} x_k.$$

Proof. By substituting $O^{\mathsf{c}} = V_{n+1}\mathsf{c}O_{i_n} \cdots V_1$, noting that V_{n+1} commutes with $\neg\Pi_{\mathbf{D}}^{g_n}$, and inserting the identity resolution $\mathbb{I}_{\mathbf{D}} = \Pi_{\mathbf{D}}^{g_{n-1}} + \neg\Pi_{\mathbf{D}}^{g_{n-1}}$ prior to the final oracle evaluation, we apply the triangle inequality and the global contractivity of all internal operators to

recursively bound the operator norm:

$$\begin{aligned}
\left\|\neg\Pi_{\mathbf{D}}^{g_n} O^{\mathsf{c}} \Pi_{\mathbf{D}}^{g_0}\right\|_\infty &= \left\|V_{n+1}\neg\Pi_{\mathbf{D}}^{g_n}\mathsf{c}O_{i_n}\left(\Pi_{\mathbf{D}}^{g_{n-1}} + \neg\Pi_{\mathbf{D}}^{g_{n-1}}\right) V_n \cdots \mathsf{c}O_{i_1} V_1 \Pi_{\mathbf{D}}^{g_0}\right\|_\infty \\
&\leq \left\|V_{n+1}\neg\Pi_{\mathbf{D}}^{g_n}\mathsf{c}O_{i_n}\Pi_{\mathbf{D}}^{g_{n-1}} V_n \cdots \Pi_{\mathbf{D}}^{g_0}\right\|_\infty + \left\|V_{n+1}\neg\Pi_{\mathbf{D}}^{g_n}\mathsf{c}O_{i_n}\neg\Pi_{\mathbf{D}}^{g_{n-1}} V_n \cdots \Pi_{\mathbf{D}}^{g_0}\right\|_\infty \\
&\leq \|V_{n+1}\|_\infty \left\|\neg\Pi_{\mathbf{D}}^{g_n}\mathsf{c}O_{i_n}\Pi_{\mathbf{D}}^{g_{n-1}}\right\|_\infty \left\|V_n \cdots \Pi_{\mathbf{D}}^{g_0}\right\|_\infty \\
&\quad + \left\|V_{n+1}\neg\Pi_{\mathbf{D}}^{g_n}\mathsf{c}O_{i_n}\right\|_\infty \left\|\neg\Pi_{\mathbf{D}}^{g_{n-1}} V_n \cdots \mathsf{c}O_{i_1} V_1 \Pi_{\mathbf{D}}^{g_0}\right\|_\infty \\
&\leq x_n + \left\|\neg\Pi_{\mathbf{D}}^{g_{n-1}}\left(V_n \mathsf{c}O_{i_{n-1}} \cdots \mathsf{c}O_{i_1} V_1\right)\Pi_{\mathbf{D}}^{g_0}\right\|_\infty \\
&\leq \cdots \leq \sum_{k=1}^{n} x_k.
\end{aligned}$$

This completes the proof. □

To systematically evaluate the single-step leakage bound, we map the quantum state transitions directly onto the combinatorial properties of the classical database.

For typographical brevity and structural clarity, we formally associate any classical database predicate $\mathcal{P}_{\mathbf{S}}$ with the explicit subset of database configurations $\mathbf{S} \subseteq \mathbf{D}$ that satisfy it. Throughout the reductions, we will interchangeably use the logical evaluation $\mathcal{P}_{\mathbf{S}}(\mathbf{d}) = 1$ and the set membership $\mathbf{d} \in \mathbf{S}$. The orthogonal subspace projector corresponding to $\mathcal{P}_{\mathbf{S}}$ is denoted by $\Pi^{\mathbf{S}}$, and its orthogonal complement by $\neg\Pi^{\mathbf{S}}$.

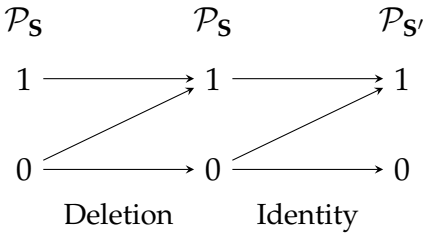
Let $\mathbf{d}_{i,x}$ denote a composite database configuration where the x-th entry of the i-th specific database explicitly remains undefined. Let $\mathbf{d}_{i,x} \cup \{(i, x, y)\}$ denote the resulting database configuration after securely appending the defined value $y \in Y_i$. We define the union operation with the undefined assignment such that it strictly evaluates to the exact identity: $\mathbf{d}_{i,x} \cup \{(i, x, \perp)\} = \mathbf{d}_{i,x}$.

Definition 3.15 (Predicate Leakage Monotonicity)**.** *A transition from predicate $\mathcal{P}_{\mathbf{S}}$ to $\mathcal{P}_{\mathbf{S}'}$ satisfies leakage monotonicity if it unconditionally fulfills the following two criteria:*

1. *For any non-empty database $\mathbf{d}$, if $\mathcal{P}_{\mathbf{S}}(\mathbf{d}) = 1$, then $\mathcal{P}_{\mathbf{S}}(\mathbf{d} \setminus \{(*, *, *)\}) = 1$.*
2. *For any database $\mathbf{d}$, if $\mathcal{P}_{\mathbf{S}}(\mathbf{d}) = 1$, then $\mathcal{P}_{\mathbf{S}'}(\mathbf{d}) = 1$.*

Lemma 3.16. *The transition from $\mathcal{P}_{\mathbf{S}}$ to $\mathcal{P}_{\mathbf{S}'}$ satisfies leakage monotonicity if and only if the complementary subspace strictly maintains upward closure under record insertion (i.e., $\mathbf{d} \notin \mathbf{S} \implies \mathbf{d} \cup \{(i, x, y)\} \notin \mathbf{S}$), and the base database residing in $\mathbf{S}$ strictly satisfies $\mathbf{S}'$, algebraically enforcing that the intersection between the prerequisite subspace and the target escape subspace is strictly empty (i.e., $\mathbf{S} \cap \neg\mathbf{S}' = \emptyset$).*

Proof. The combinatorial state transitions governed by leakage monotonicity can be visually represented by the following logical transition diagram. The deletion operation illustrates the downward-closure of the prerequisite set, while the identity operation illustrates strict subset inclusion.



This completes the proof. □

Definition 3.17 (Maximum Escape Volume and Rate)**.** *Let $\Delta(i, \mathcal{P}_{\mathbf{S}} \to \mathcal{P}_{\mathbf{S}'})$ denote the maximum escape volume induced by random sampling. It is mathematically defined as the maximum number of evaluations y that force the database to transition out of the safe target subset $\mathbf{S}'$, maximized over all valid inputs x and databases residing within the prerequisite subset $\mathbf{S}$ left undefined on x:*

$$\Delta(i, \mathcal{P}_{\mathbf{S}} \to \mathcal{P}_{\mathbf{S}'}) = \max_{\substack{x \in X_i \\ \mathbf{d}_{i,x} \in \mathbf{S}}} \left| \{ y \in Y_i \mid \mathbf{d}_{i,x} \cup \{(i, x, y)\} \notin \mathbf{S}' \} \right|.$$

Consequently, the maximum escape rate is formally defined as the ratio of this maximum escape volume to the cardinality of the output space $|Y_i|$:

$$\textit{Maximum Escape Rate} = \frac{\Delta(i, \mathcal{P}_{\mathbf{S}} \to \mathcal{P}_{\mathbf{S}'})}{|Y_i|}.$$

Lemma 3.18 (Single-Query Leakage Bound)**.** *Let $\mathcal{P}_{\mathbf{S}}$ and $\mathcal{P}_{\mathbf{S}'}$ be database predicates satisfying leakage monotonicity. For any single compressed oracle invocation $\mathsf{c}O_{i_k}$, the transition leakage operator norm is strictly bounded by:*

$$\left\| \neg \Pi^{\mathbf{S}'} \mathsf{c}O_{i_k} \Pi^{\mathbf{S}} \right\|_\infty \leq \sqrt{13 \frac{\Delta(i_k, \mathcal{P}_{\mathbf{S}} \to \mathcal{P}_{\mathbf{S}'})}{|Y_{i_k}|}}.$$

Proof. The operator $\mathsf{c}O_{i_k}$ and the database projectors inherently decompose into a direct sum over the orthogonal invariant subspaces spanned by the fixed query inputs $x \in X_{i_k}$ and the base undefined database configurations $\mathbf{d}_{i_k,x} \in \mathbf{D}$. Consequently, the global operator norm is determined entirely by the maximum local operator norm across these localized subspaces.

Fix an arbitrary input $x \in X_{i_k}$ and a base database $\mathbf{d}_{i_k,x}$. Let $\overline{Y}_{i_k} = Y_{i_k} \cup \{\perp\}$. We define an arbitrary normalized target state $|\psi\rangle$ within this specific invariant subspace as:

$$|\psi\rangle = \sum_{y \in Y_{i_k}} \sum_{y' \in \overline{Y}_{i_k}} \alpha_{y,y'} |x\rangle|y\rangle \big| \mathbf{d}_{i_k,x} \cup \{(i_k, x, y')\} \big\rangle, \quad \text{where} \quad \sum_{y,y'} |\alpha_{y,y'}|^2 = 1.$$

We assume the base configuration satisfies $\mathbf{d}_{i_k,x} \in \mathbf{S}$; otherwise, by leakage monotonicity, all augmented configurations remain outside $\mathbf{S}$, forcing all amplitudes $\alpha_{y,y'} = 0$ and trivially reducing the local norm to 0. Under this valid assumption, the prerequisite condition $\Pi^{\mathbf{S}}|\psi\rangle = |\psi\rangle$ algebraically enforces that the amplitude $\alpha_{y,y'} = 0$ for any specific y' where the resulting configuration $\mathbf{d}_{i_k,x} \cup \{(i_k, x, y')\} \notin \mathbf{S}$.

We first orthogonally decompose the initial state $|\psi\rangle$ into its undefined components ($y' = \perp$) and defined components ($y' = z \in Y_{i_k}$). Evaluating the distinct functional actions of $\mathsf{c}O_{i_k}$ and applying the triangle inequality, we isolate the escape leakage norm into four structurally distinct terms:

$$
\begin{aligned}
\left\|\neg\Pi^{\mathbf{S}'}\mathsf{c}O_{i_k}|\psi\rangle\right\|_2 \leq{}& \left\|\neg\Pi^{\mathbf{S}'}\mathsf{c}O_{i_k}\sum_{y\in Y_{i_k}}\alpha_{y,\perp}|x\rangle|y\rangle|\mathbf{d}_{i_k,x}\rangle\right\|_2 \\
&+\left\|\neg\Pi^{\mathbf{S}'}\mathsf{c}O_{i_k}\sum_{y\in Y_{i_k}}\sum_{z\in Y_{i_k}}\alpha_{y,z}|x\rangle|y\rangle\left|\mathbf{d}_{i_k,x}\cup\{(i_k,x,z)\}\right\rangle\right\|_2 \\
&\leq\left\|\neg\Pi^{\mathbf{S}'}\sum_{y\in Y_{i_k}}\frac{\alpha_{y,\perp}}{\sqrt{|Y_{i_k}|}}\sum_{y''\in Y_{i_k}}|x\rangle|y+y''\rangle\left|\mathbf{d}_{i_k,x}\cup\{(i_k,x,y'')\}\right\rangle\right\|_2 && \text{(Laz. Samp.), (3)} \\
&+\left\|\neg\Pi^{\mathbf{S}'}\sum_{y\in Y_{i_k}}\sum_{z\in Y_{i_k}}\alpha_{y,z}|x\rangle|y+z\rangle\left|\mathbf{d}_{i_k,x}\cup\{(i_k,x,z)\}\right\rangle\right\|_2 && \text{(Dir. Resp.), (4)} \\
&+\left\|\neg\Pi^{\mathbf{S}'}\sum_{y\in Y_{i_k}}\alpha_{y,\perp}|\Phi^{(1)}_{coh}(x,y,\mathbf{d}_{i_k,x})\rangle\right\|_2 && \text{(Coh. 1), (5)} \\
&+\left\|\neg\Pi^{\mathbf{S}'}\sum_{y\in Y_{i_k}}\sum_{z\in Y_{i_k}}\alpha_{y,z}|\Phi^{(2)}_{coh}(x,y,\mathbf{d}_{i_k,x},z)\rangle\right\|_2. && \text{(Coh. 2), (6)}
\end{aligned}
$$

By the structural implications of leakage monotonicity, any defined initial configuration $\mathbf{d}_{i_k,x}\cup\{(i_k,x,z)\}$ that satisfies $\mathbf{S}$ unconditionally satisfies $\mathbf{S}'$, forcing $\mathbf{S}\cap\neg\mathbf{S}'=\emptyset$. Thus, the state in the direct response term is perfectly annihilated by the escape projector: $\||4\rangle\|_2=0$.

For the lazy sampling term, evaluating the escape transitions yields:

$$
\||3\rangle\|_2\leq\sqrt{\sum_{y\in Y_{i_k}}|\alpha_{y,\perp}|^2\sum_{\substack{y''\in Y_{i_k}\\ \mathbf{d}_{i_k,x}\cup\{(i_k,x,y'')\}\notin\mathbf{S}'}}\frac{1}{|Y_{i_k}|}}\leq\sqrt{\frac{\Delta(i_k,\mathcal{P}_{\mathbf{S}}\to\mathcal{P}_{\mathbf{S}'})}{|Y_{i_k}|}}\sqrt{\sum_{y\in Y_{i_k}}|\alpha_{y,\perp}|^2}.
$$

For the first coherence term $\||5\rangle\|_2$, the output register strictly resolves to $|*\rangle_Y$. Because the base unaugmented state $\mathbf{d}_{i_k,x}$ resides in $\mathbf{S}$, it strictly resides in $\mathbf{S}'$, meaning $\neg\Pi^{\mathbf{S}'}$ perfectly annihilates the undefined superposition branch. The projector subsequently isolates the escape outcomes from the defined branches:

$$
\begin{aligned}
\||5\rangle\|_2 &= \left\|\neg\Pi^{\mathbf{S}'}\left(\sum_{y\in Y_{i_k}}\alpha_{y,\perp}\right)\frac{1}{\sqrt{|Y_{i_k}|}}|x\rangle|*\rangle_Y\Big(\left|\mathbf{d}_{i_k,x}\cup\{(i_k,x,\perp)\}\right\rangle-\left|\mathbf{d}_{i_k,x}\cup\{(i_k,x,*)\}\right\rangle\Big)\right\|_2 \\
&= \left|\sum_{y\in Y_{i_k}}\alpha_{y,\perp}\right|\frac{1}{\sqrt{|Y_{i_k}|}}\left\|\neg\Pi^{\mathbf{S}'}\left|\mathbf{d}_{i_k,x}\cup\{(i_k,x,*)\}\right\rangle\right\|_2 \\
&\leq \left(\sqrt{|Y_{i_k}|}\sqrt{\sum_{y\in Y_{i_k}}|\alpha_{y,\perp}|^2}\right)\frac{1}{\sqrt{|Y_{i_k}|}}\sqrt{\frac{\Delta(i_k,\mathcal{P}_{\mathbf{S}}\to\mathcal{P}_{\mathbf{S}'})}{|Y_{i_k}|}} \\
&= \sqrt{\frac{\Delta(i_k,\mathcal{P}_{\mathbf{S}}\to\mathcal{P}_{\mathbf{S}'})}{|Y_{i_k}|}}\sqrt{\sum_{y\in Y_{i_k}}|\alpha_{y,\perp}|^2}.
\end{aligned}
$$

For the second coherence term $\||6\rangle\|_2$, we apply the triangle inequality to bifurcate the structural definition of $|\Phi_{coh}^{(2)}\rangle$:

$$\||6\rangle\|_2 \leq \left\| \neg\Pi^{\mathbf{S}'} \sum_{y,z\in Y_{i_k}} \frac{\alpha_{y,z}}{\sqrt{|Y_{i_k}|}} |x\rangle|y+z\rangle \Big(\big|\mathbf{d}_{i_k,x} \cup \{(i_k,x,\perp)\}\big\rangle - \big|\mathbf{d}_{i_k,x} \cup \{(i_k,x,*)\}\big\rangle \Big) \right\|_2 \tag{7}$$

$$+ \left\| \neg\Pi^{\mathbf{S}'} \frac{1}{\sqrt{|Y_{i_k}|}} \sum_{y,z\in Y_{i_k}} \alpha_{y,z} \mathrm{cO}_{i_k} |x\rangle|y\rangle \big|\mathbf{d}_{i_k,x} \cup \{(i_k,x,\perp)\}\big\rangle \right\|_2. \tag{8}$$

To bound $\||7\rangle\|_2$, we map $u = y + z$ to merge overlapping basis elements:

$$\begin{aligned}
\||7\rangle\|_2^2 &= \sum_{u\in Y_{i_k}} \left\| \frac{1}{|Y_{i_k}|} \left(\sum_{z\in Y_{i_k}} \alpha_{u-z,z} \right) \sum_{y''} \neg\Pi^{\mathbf{S}'} \big|\mathbf{d}_{i_k,x} \cup \{(i_k,x,y'')\}\big\rangle \right\|_2^2 \\
&\leq \sum_{u\in Y_{i_k}} \frac{1}{|Y_{i_k}|^2} \left(|Y_{i_k}| \sum_{z\in Y_{i_k}} |\alpha_{u-z,z}|^2 \right) \Delta(i_k, \mathcal{P}_{\mathbf{S}} \to \mathcal{P}_{\mathbf{S}'}) \\
&= \frac{\Delta(i_k, \mathcal{P}_{\mathbf{S}} \to \mathcal{P}_{\mathbf{S}'})}{|Y_{i_k}|} \sum_{y,z\in Y_{i_k}} |\alpha_{y,z}|^2 \implies \||7\rangle\|_2 \leq \sqrt{\frac{\Delta(i_k, \mathcal{P}_{\mathbf{S}} \to \mathcal{P}_{\mathbf{S}'})}{|Y_{i_k}|}} \sqrt{\sum_{y,z\in Y_{i_k}} |\alpha_{y,z}|^2}.
\end{aligned}$$

To bound $\||8\rangle\|_2$, we define $\beta_y = \sum_{z\in Y_{i_k}} \alpha_{y,z}$. The oracle acts strictly on the undefined configurations, echoing the established limits for $\||3\rangle\|_2$ and $\||5\rangle\|_2$:

$$\begin{aligned}
\||8\rangle\|_2 &= \frac{1}{\sqrt{|Y_{i_k}|}} \left\| \neg\Pi^{\mathbf{S}'} \mathrm{cO}_{i_k} \sum_{y\in Y_{i_k}} \beta_y |x\rangle|y\rangle \big|\mathbf{d}_{i_k,x} \cup \{(i_k,x,\perp)\}\big\rangle \right\|_2 \\
&\leq \frac{1}{\sqrt{|Y_{i_k}|}} \left(2\sqrt{\frac{\Delta(i_k, \mathcal{P}_{\mathbf{S}} \to \mathcal{P}_{\mathbf{S}'})}{|Y_{i_k}|}} \sqrt{\sum_{y\in Y_{i_k}} |\beta_y|^2} \right) \\
&\leq \frac{2}{\sqrt{|Y_{i_k}|}} \sqrt{\frac{\Delta(i_k, \mathcal{P}_{\mathbf{S}} \to \mathcal{P}_{\mathbf{S}'})}{|Y_{i_k}|}} \left(\sqrt{|Y_{i_k}| \sum_{y,z\in Y_{i_k}} |\alpha_{y,z}|^2} \right) = 2\sqrt{\frac{\Delta(i_k, \mathcal{P}_{\mathbf{S}} \to \mathcal{P}_{\mathbf{S}'})}{|Y_{i_k}|}} \sqrt{\sum_{y,z\in Y_{i_k}} |\alpha_{y,z}|^2}.
\end{aligned}$$

Summing $\||7\rangle\|_2$ and $\||8\rangle\|_2$ gives $\||6\rangle\|_2 \leq 3\sqrt{\frac{\Delta(i_k, \mathcal{P}_{\mathbf{S}} \to \mathcal{P}_{\mathbf{S}'})}{|Y_{i_k}|}} \sqrt{\sum_{y,z\in Y_{i_k}} |\alpha_{y,z}|^2}$.

Synthesizing all derived bounds via the Cauchy-Schwarz inequality across the orthogonal state amplitudes, we resolve the absolute escape transition limit:

$$\begin{aligned}
\left\| \neg\Pi^{\mathbf{S}'} \mathrm{cO}_{i_k} |\psi\rangle \right\|_2 &\leq \||3\rangle\|_2 + \||4\rangle\|_2 + \||5\rangle\|_2 + \||6\rangle\|_2 \\
&\leq 2\sqrt{\frac{\Delta(i_k, \mathcal{P}_{\mathbf{S}} \to \mathcal{P}_{\mathbf{S}'})}{|Y_{i_k}|}} \sqrt{\sum_{y\in Y_{i_k}} |\alpha_{y,\perp}|^2} + 3\sqrt{\frac{\Delta(i_k, \mathcal{P}_{\mathbf{S}} \to \mathcal{P}_{\mathbf{S}'})}{|Y_{i_k}|}} \sqrt{\sum_{y,z\in Y_{i_k}} |\alpha_{y,z}|^2} \\
&\leq \sqrt{2^2 + 3^2} \sqrt{\sum_{y\in Y_{i_k}} |\alpha_{y,\perp}|^2 + \sum_{y,z\in Y_{i_k}} |\alpha_{y,z}|^2} \sqrt{\frac{\Delta(i_k, \mathcal{P}_{\mathbf{S}} \to \mathcal{P}_{\mathbf{S}'})}{|Y_{i_k}|}} \\
&= \sqrt{13 \frac{\Delta(i_k, \mathcal{P}_{\mathbf{S}} \to \mathcal{P}_{\mathbf{S}'})}{|Y_{i_k}|}}.
\end{aligned}$$

This strictly establishes the maximum norm bound over all localized subspaces, completing the proof. □

Corollary 3.19. *Suppose the quantum circuit $C[\mathsf{c}O_{f_1}, \ldots, \mathsf{c}O_{f_m}]$ is controlled by an oracle label register $\mathcal{H}_I$, where $I = \{1, \ldots, m\}$, to perform a single query across all underlying compressed oracles, mathematically defined as $C[\mathsf{c}O_{f_1}, \ldots, \mathsf{c}O_{f_m}] = \sum_{i \in I} |i\rangle\langle i|_I \otimes \mathsf{c}O_{f_i}$. For predicates $\mathcal{P}_{\mathbf{S}}$ and $\mathcal{P}_{\mathbf{S}'}$ satisfying leakage monotonicity, the transition leakage operator norm is bounded by:*

$$\left\| \neg\Pi^{\mathbf{S}'} C[\mathsf{c}O_{f_1}, \ldots, \mathsf{c}O_{f_m}] \Pi^{\mathbf{S}} \right\|_\infty \leq \max_{i \in \{1,\ldots,m\}} \sqrt{13 \frac{\Delta(i, \mathcal{P}_{\mathbf{S}} \to \mathcal{P}_{\mathbf{S}'})}{|Y_i|}}.$$

Proof. The globally controlled quantum circuit $C[\mathsf{c}O_{f_1}, \ldots, \mathsf{c}O_{f_m}]$ inherently decomposes into a direct sum over the orthogonal invariant subspaces spanned by the control register indices $i \in I$, the specific query inputs $x \in X_i$, and the base undefined database configurations $\mathbf{d}_{i,x} \in \mathbf{D}$. Consequently, the global operator norm is determined entirely by the absolute maximum of the local operator norms across these localized subspaces. Applying the established single-query bound to each orthogonal branch directly yields the supremum:

$$\begin{aligned} \left\| \neg\Pi^{\mathbf{S}'} C[\mathsf{c}O_{f_1}, \ldots, \mathsf{c}O_{f_m}] \Pi^{\mathbf{S}} \right\|_\infty &= \max_{i \in I} \left\| \neg\Pi^{\mathbf{S}'} \mathsf{c}O_{f_i} \Pi^{\mathbf{S}} \right\|_\infty \\ &\leq \max_{i \in \{1,\ldots,m\}} \sqrt{13 \frac{\Delta(i, \mathcal{P}_{\mathbf{S}} \to \mathcal{P}_{\mathbf{S}'})}{|Y_i|}}. \end{aligned}$$

This completes the proof. □

4 Quantum Indifferentiability Analysis of Prefix-Free Merkle–Damgård

In this section, we analyze the prefix-free Merkle–Damgård (MD) domain extender within the quantum indifferentiability framework, without claiming a complete quantum indifferentiability proof. Our goal is to separate the two oracle-indistinguishability components arising from the reduction and to identify which parts can be controlled by the compressed-oracle game-playing framework. The consistency component admits a positive bound under our analysis, whereas the indistinguishability component reveals a structural obstruction for the present simulator and proof strategy. This analysis clarifies why additional ideas are required to complete a full quantum indifferentiability proof for prefix-free MD.

4.1 System Modeling and Simulator

Following the established quantum game-playing framework, the definition of quantum indifferentiability, and the simplification afforded by the prefix-free Merkle-Damgård construction, our security objective strictly reduces to a two-oracle quantum indistinguishability problem.

Real World. The adversary is granted quantum access to two oracles:

1. The construction oracle $\mathsf{MD}_f : (\{0,1\}^b)^{\leq l+2} \to \{0,1\}^n$. The valid message set strictly belongs to $\mathsf{Range}(\mathsf{pf})$, which rigorously ensures that for any two messages M and M', M is never a prefix of M'. This construction is instantiated as a fixed, polynomial-size, stateless quantum circuit.
2. The primitive oracle $f : \{0,1\}^n \times \{0,1\}^b \to \{0,1\}^n$, representing a standard fixed-input-length uniformly random function.

Ideal World. The adversary interacts with:

1. The ideal construction oracle H, which is a true random function defined over the identical valid prefix-free input space $(\{0,1\}^b)^{\leq l+2} \to \{0,1\}^n$.
2. The simulated primitive oracle, governed by a polynomial-time stateful quantum simulator $\mathcal{S}$ that has query access to the ideal construction oracle H.

Quantum Circuit for the MD_f Construction. The global quantum circuit evaluating the MD_f construction is conceptually decomposed into a coherent superposition of length-controlled quantum sub-circuits. For a specific sub-circuit governed by a fixed message length k, denoted as the controlled operator MD_k, the circuit acts on the input message registers $X_1, \ldots, X_k$, the output register Y, and the auxiliary chaining registers $Z_0, Z_1, \ldots, Z_k$. Since all auxiliary registers naturally initialize to the $|0\rangle$ state in our model, the evaluation directly proceeds as follows:

1. **Forward Evaluation:** A dedicated unitary U_{IV} is initially applied to Z_0 to prepare the state $|IV\rangle$. For each block index i from 1 to k, the controlled primitive oracle f is applied to coherently compute the chaining values across the auxiliary registers.[1]

[1]Technically, what is required here is a controlled-O_f. This can be perfectly simulated using a standard unconditional O_f alongside two controlled-SWAP operators that exchange the output register with an unentangled uniform superposition register. **??** Thus, a single unconditional invocation of O_f flawlessly simulates a controlled-O_f without incurring any penalty in query complexity.

2. **Output Extraction:** Apply a CNOT cascade to XOR the terminal chaining value residing in Z_k into the target response register Y.

3. **Coherent Uncomputation:** Apply the inverse unitary evaluations in reverse sequential order (from k down to 1) and finally apply $U_{\mathrm{IV}}^{\dagger}$ to cleanly uncompute all auxiliary chaining registers $Z_0, Z_1, \ldots, Z_k$ entirely back to $|0\rangle$, preventing any residual entanglement.

To formalize the core evaluation logic, we explicitly present the instantiated quantum sub-circuit MD_2 when the message length is exactly 2 blocks (i.e., $k = 2$) in Figure 1.

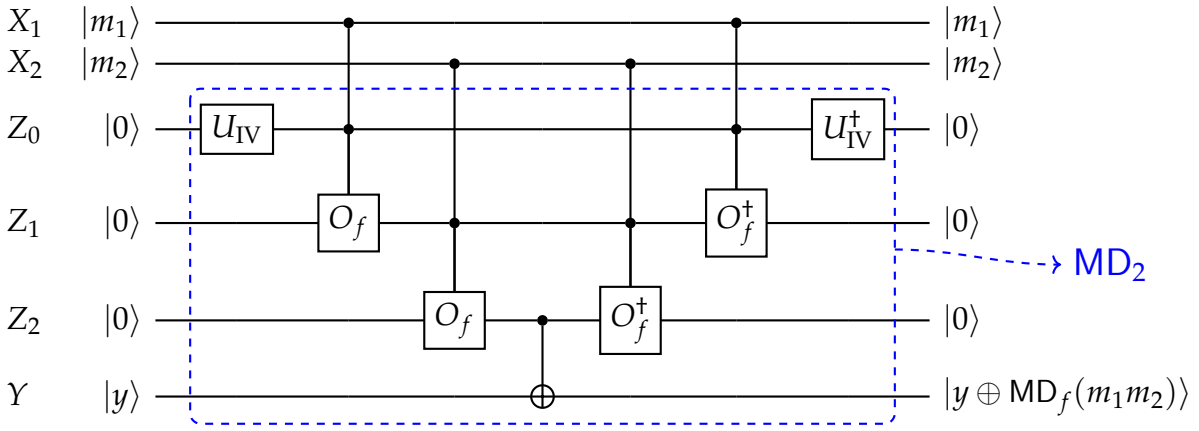


Figure 1: Quantum sub-circuit evaluating the MD_f construction strictly conditioned on a 2-block message length. It utilizes U_{IV} to set the initial vector, cascades the ideal primitive O_f across the chaining registers, extracts the final hash, and subsequently executes a perfect coherent uncomputation to restore all auxiliary registers back to $|0\rangle$.

Real Simulated Evolution (Process 0-a). Mirroring the game-playing framework, the real-world execution is conceptually shifted into the compressed setting. The underlying standard primitive oracle O_f evaluated within the MD_f circuit is identically replaced by its efficient compressed oracle equivalent $\mathsf{c}O_f$. This formally establishes the compressed real oracles:

- $O_1^{\mathsf{c}} := C_1[\mathsf{c}O_f]$: The real construction oracle, instantiated by the MD_f circuit substituting all internal evaluations of O_f with the compressed primitive $\mathsf{c}O_f$.
- $O_2^{\mathsf{c}} := \mathsf{c}O_f$: The real primitive oracle, evaluating the standard fixed-input-length compressed random function.

The central database register $\mathcal{H}_{D_f}$ is strictly initialized to the empty state $|\varnothing\rangle$. Throughout the execution, the adversary's state evolves through an interleaved sequence of internal unitaries and queries to O_1^{c} or O_2^{c}, coherently updating the real database configurations. The final global state of the algorithm is mathematically expressed as:

$$|\Psi^{(0a)}\rangle = U_t O_{i_t}^{\mathsf{c}} U_{t-1} \cdots U_2 O_{i_2}^{\mathsf{c}} U_1 O_{i_1}^{\mathsf{c}} \left(|\psi_0\rangle_A \otimes |\varnothing\rangle_{D_f} \right).$$

Quantum Simulator Algorithm. In the ideal world, the simulator $\mathcal{S}$ must mimic the primitive f while remaining perfectly consistent with the ideal construction oracle H. Recall that $\mathcal{S}$ has query access to H; additionally, here $\mathcal{S}$ utilizes the compressed oracle $\mathsf{c}O_f$ and database $\mathcal{H}_{D_f}$ to simulate a random function $f : \{0,1\}^n \times \{0,1\}^b \to \{0,1\}^n$ for its own internal invocation. Upon receiving a primitive query state $|z\rangle|m\rangle|y\rangle$, $\mathcal{S}$ processes the evaluation dynamically:

1. **Path Extraction:** $\mathcal{S}$ applies a coherent extraction unitary Ext to trace the chaining variable z backwards through its local database D_f. If z traces successfully back to the initial vector IV and the appended block m perfectly completes a valid prefix-free message sequence $M = (m_1, \ldots, m_k, m)$, Ext extracts M into a dedicated message register and flips an internal success flag $|s\rangle$ to $|1\rangle$. Otherwise, the flag remains $|0\rangle$.
2. **Conditional Routing:** Controlled strictly by the superposition flag $|s\rangle$:
   - If $|s\rangle = |1\rangle$ (Completion): $\mathcal{S}$ invokes the external ideal construction oracle H on the extracted full message M and XORs the result directly into $|y\rangle$.
   - If $|s\rangle = |0\rangle$ (Internal): $\mathcal{S}$ invokes its internal compressed oracle update $\mathsf{c}O_f$ on the localized input (z, m) to update D_f and fetch the response.
3. **Inverse Uncomputation:** $\mathcal{S}$ applies the inverse extraction unitary $\mathsf{Ext}^\dagger$ to cleanly uncompute the extracted message and the flag $|s\rangle$ back to zero, preventing any decoherence.

Remark 4.1. *It is important to note that, in this setting, the database D_f may be modified during the simulation. Consequently, the decoding operation is not guaranteed to succeed on every basis state.*

Ideal Simulated Evolution (Process 1-a). Symmetrically to the real-world execution, the ideal-world execution is conceptually shifted into the compressed setting. The underlying ideal construction oracle H is identically replaced by its efficient compressed oracle equivalent $\mathsf{c}O_H$. This formally establishes the compressed ideal oracles:

- $O'^{\mathsf{c}}_1 := \mathsf{c}O_H$: The ideal construction oracle, directly evaluating the compressed variable-input-length random oracle.
- $O'^{\mathsf{c}}_2 := \mathcal{S}^{\mathsf{c}O_H}$: The ideal primitive oracle, logically evaluated by the simulator $\mathcal{S}$ operating over its local database and interacting dynamically with $\mathsf{c}O_H$.

The central database registers, comprising the simulator's internal state $\mathcal{H}_{D_f}$ and the ideal state $\mathcal{H}_{D_H}$, are strictly initialized to the empty state $|\emptyset\rangle$. Throughout the execution, the adversary's state evolves through an interleaved sequence of internal unitaries and queries to O'^{c}_1 or O'^{c}_2, coherently updating the ideal database configurations. The final state of the algorithm is mathematically expressed as:

$$|\Psi^{(1a)}\rangle = U_t O'^{\mathsf{c}}_{i_t} U_{t-1} \cdots U_2 O'^{\mathsf{c}}_{i_2} U_1 O'^{\mathsf{c}}_{i_1} \left(|\psi_0\rangle_A \otimes |\emptyset\rangle_{D_f, D_H} \right).$$

4.2 Indistinguishability Barrier

Recall that the proof of quantum indifferentiability is rigorously reduced to bounding the adversary's advantage in two distinct components: the indistinguishability game and the

consistency game. Within the context of the indistinguishability game, the quantum adversary attempts to distinguish between the real compressed primitive oracle O_2^{c} and the ideal simulated primitive oracle O'^{c}_2, making a total of at most q quantum queries.

Following the quantum game-playing proof framework, to formally bound the divergence between the real and ideal simulated evolutions, we must logically bridge their respective underlying database state spaces. In the following, we proceed to define a formal isometric injection mapping from the real database space $\mathcal{H}_{D_f}$ to the joint ideal database space $\mathcal{H}_{D_f} \otimes \mathcal{H}_{D_H}$, and subsequently formalize a sequence of *good databases* defined over D_f.

Good Databases and Injection. To formally analyze the indistinguishability, we first define the notion of *good databases*. A real database $D_f \in \mathcal{H}_{D_f}$ is considered a good database if it contains no collisions and its outputs never hit the initial vector IV. Mathematically, this is precisely expressed by the following two conditions for any records $(z, m) \in D_f$:

- No collisions: $D_f(z, m) \neq D_f(z', m')$ for any $(z', m') \in D_f$ such that $(z, m) \neq (z', m')$.
- No hitting IV: $D_f(z, m) \neq IV$.

Based on these conditions, we formalize the structural properties of good databases using graph theory. Let $G_{D_f} = (V, E)$ be a directed graph associated with D_f. The vertex set V consists exclusively of the query records $v = (z, m, y) \in D_f$. A directed edge $(v, v') \in E$ exists if and only if the output of the first record exactly equals the chaining input of the second, i.e., $y = z'$.

Lemma 4.2 (Combinatorial Structure of Good Databases)**.** *For any real good database D_f, the subgraph containing a node whose chaining input is IV must be a directed tree, and this node is the unique root of the tree. All nodes in these trees are path-reachable nodes. Moreover, for every path-reachable node, the path from the root to this node is unique. If the message string induced by a path-reachable node lies in the range of the prefix-free encoding, then we call this node the terminal node of a complete path. Such a terminal node cannot be used to form any other complete path.*

Proof. The proof is deferred to Appendix A.4. □

Based on the aforementioned combinatorial properties, we formally define an injection function V mapping from the good databases on D_f to the joint database space $D_f \times D_H$. For any real good database d_f, the injection V systematically removes all terminal records (z, m, y) that form complete prefix-free evaluation paths in d_f. Simultaneously, for each removed terminal record, V inserts a corresponding record (M, y) into the ideal database D_H, where $M = (m_1, \ldots, m_k, m)$ denotes the fully concatenated message.

The reversibility of this mapping is strictly guaranteed by the properties of good databases. For any image $V(d_f)$ derived from a good database d_f, every record $(M, y) \in D_H$ allows us to uniquely and deterministically trace the evaluation path for the prefix $m_1, \ldots, m_k$ within the remaining records of D_f. This trace perfectly recovers the chaining variable z, thereby enabling the exact restoration of the removed record (z, m, y) back into D_f. Consequently, we formally define the image $V(d_f)$ of a real good database d_f as an *ideal good database* over $D_f \times D_H$.

Verbally, the definition of this ideal good database is strictly equivalent to the following conditions:

(a) No output collisions occur anywhere within the system (including internal collisions strictly within D_f or D_H, as well as cross-collisions between D_f and D_H).

(b) No outputs ever hit the initial vector IV (strictly applicable to all outputs within D_f, D_H, and across both).

(c) If a record (M, y) exists in D_H where $M = (m_1, \ldots, m_k, m)$, then an evaluation path corresponding sequentially to the message inputs $(m_1, \ldots, m_k)$ must strictly exist in D_f.

(d) No complete evaluation paths exist anywhere in D_f.

By arbitrarily extending this injection to the set of non-good databases, we obtain a complete total injection over the entire classical database space on D_f. Through continuous linear extension, this classical mapping uniquely induces an isometric embedding operator mapping from $\mathcal{H}_{D_f}$ to $\mathcal{H}_{D_f} \otimes \mathcal{H}_{D_H}$, which we overload with the same notation V.

Finally, we define the intermediate real good database sequence $\mathsf{rg}_0, \ldots, \mathsf{rg}_q$, where each rg_i represents the subspace spanned by real good databases containing at most i query records. Correspondingly, we define the ideal good sequence as $\mathsf{ig}_i = V(\mathsf{rg}_i)$. Because the injection V strictly preserves the total number of records across the bipartite system (removing one from D_f and adding one to D_H), ig_i naturally represents the subspace of ideal good databases of size at most i.

Remark 4.3. *It is worth noting that the formal definition of the ideal good database presented here is technically incomplete; a fully rigorous definition should concurrently enforce that the simulator's auxiliary label register is strictly maintained at $|s\rangle = |0\rangle$. Because the subsequent analysis in this indistinguishability section remains algebraically unaffected by potential uncomputation failures of this label register, we deliberately omit this constraint to streamline the proof. As will be explicitly demonstrated in the subsequent consistency analysis section, accommodating these label uncomputation failures only introduces a bounded, negligible error to the overall global bounds.*

Lemma 4.4. *For any $i \in \{1, \ldots, q\}$,*

$$\|\neg\Pi^{\mathsf{rg}_i} O_2^{\mathsf{c}} \Pi^{\mathsf{rg}_{i-1}}\| \leq \sqrt{\frac{13i}{2^n}}.$$

Proof. The transition from $\mathcal{P}_{\mathsf{rg}_{i-1}}$ to $\mathcal{P}_{\mathsf{rg}_i}$ satisfies strict monotonicity. The maximum escape rate is bounded by $i/2^n$. Therefore, the leakage condition is upper-bounded by $\sqrt{13i/2^n}$. □

Next, we need to analyze the ideal leakage bound. However, recall the third condition defining the ideal good database introduced earlier: if a record exists in D_H, the corresponding records used to form this evaluation path must strictly exist in D_f. This specific condition causes the structural transition from ig_{i-1} to ig_i to no longer satisfy predicate monotonicity. For instance, deleting an intermediate path record such as $(m_1, IV, *)$ will directly cause the database to become bad. Consequently, we can no longer utilize the established maximum escape rate tools for this proof, and must instead perform a refined, direct derivation. Unfortunately, we find that such leakage occurs quite easily.

Theorem 4.5. *There exists a polynomial-time quantum adversary $\mathcal{A}$ making exactly three sequential queries to the ideal simulated primitive oracle O'^{c}_2, such that the final global quantum state $|\Psi'_3\rangle$ inherently escapes the ideal good subspace ig_3, satisfying the strict lower bound:*

$$\left\| \neg\Pi^{\mathsf{ig}_3} |\Psi'_3\rangle \right\|_2 \geq 1 - \sqrt{\frac{32}{|Y|}}.$$

Proof. The distinguishing attack algorithm is defined as follows:

Algorithm 1 Quantum Adversary $\mathcal{A}$ Execution Sequence

1: **Initialize** $|\Psi_0\rangle = |IV, m_1\rangle_{X_1}|0\rangle_{Y_1}|m_2\rangle_{X_2}|0\rangle_{Y_2} \otimes |\emptyset\rangle_{D_f}|\emptyset\rangle_{D_H}$, where $m_1 \neq m_2$, $m_1 \parallel m_2 \in \mathsf{Range}(\mathsf{pf})$, and $m_1 \parallel m_1, m_2 \parallel m_1, m_2 \notin \mathsf{Range}(\mathsf{pf})$[2].
2: **Query 1:** Apply $O_2'^{\mathsf{c}}$ on static input X_1, XORing the response into Y_1.
3: **Query 2:** Apply $O_2'^{\mathsf{c}}$ on dynamically composed (Y_1, X_2), XORing response into Y_2.
4: **Query 3:** Apply $O_2'^{\mathsf{c}}$ on static input X_1 again, XORing the response into Y_1.

For the first query on (IV, m_1), the adversary applies $O_2'^{\mathsf{c}}$ to the static input X_1. Because the database is initially empty, the uncomputation of the path extraction is omitted. The simulator evaluates the local oracle via lazy sampling into D_f. Let $|\Psi_1'\rangle$ denote the quantum state after the first query, partitioned as $|\Psi_1'\rangle = |\Psi_{1,\text{main}}'\rangle + |\Psi_{1,\text{rem}}'\rangle$, where:

$$
\begin{aligned}
|\Psi_{1,\text{main}}'\rangle &= \frac{1}{\sqrt{|Y|}} \sum_{y_1 \in Y} |IV, m_1\rangle_{X_1}|y_1\rangle_{Y_1}|m_2\rangle_{X_2}|0\rangle_{Y_2}\big|\{(IV, m_1, y_1)\}\big\rangle_{D_f}|\emptyset\rangle_{D_H}, \\
|\Psi_{1,\text{rem}}'\rangle &= |m_2\rangle_{X_2}|0\rangle_{Y_2}|\emptyset\rangle_{D_H} \otimes |\Phi_{coh}^{(1)}((IV, m_1), 0, \emptyset)\rangle_{X_1, Y_1, D_f}.
\end{aligned} \tag{9}
$$

For the second query, the adversary applies $O_2'^{\mathsf{c}}$ on the dynamically composed input (y_1, m_2). The simulator traces the database D_f and concatenates the path $m_1 \parallel m_2$. Because $m_1 \parallel m_2 \in \mathsf{Range}(\mathsf{pf})$, the simulator intercepts the query, bypasses D_f, and directly queries the macroscopic hash oracle H. The state evaluates to $|\Psi_2'\rangle = |\Psi_{2,\text{main}}'\rangle + |\Psi_{2,\text{rem}}'\rangle$, where:

$$
\begin{aligned}
|\Psi_{2,\text{main}}'\rangle &= \frac{1}{|Y|} \sum_{y_1, y_2 \in Y} |IV, m_1\rangle_{X_1}|y_1\rangle_{Y_1}|m_2\rangle_{X_2}|y_2\rangle_{Y_2}\big|\{(IV, m_1, y_1)\}\big\rangle_{D_f}\big|\{(m_1 \parallel m_2, y_2)\}\big\rangle_{D_H}, \\
|\Psi_{2,\text{rem}}'\rangle &= \frac{1}{\sqrt{|Y|}} \sum_{y_1 \in Y} |IV, m_1\rangle_{X_1}|y_1\rangle_{Y_1}\big|\{(IV, m_1, y_1)\}\big\rangle_{D_f} \otimes |\Phi_{coh}^{(1)}(m_1 \parallel m_2, 0, \emptyset)\rangle_{D_H} \\
&\quad + \mathsf{c}O_H|\Psi_{1,\text{rem}}'\rangle.
\end{aligned} \tag{10}
$$

For the third query, the adversary targets (IV, m_1) again to dynamically uncompute Y_1. This triggers a direct response from D_f, uncomputing y_1 back to $|0\rangle$ and cleanly emptying D_f. The collapsed state evaluates to $|\Psi_3'\rangle = |\Psi_{3,\text{main}}'\rangle + |\Psi_{3,\text{rem}}'\rangle$, where:

$$
\begin{aligned}
|\Psi_{3,\text{main}}'\rangle &= \frac{1}{\sqrt{|Y|}} \sum_{y_2 \in Y} |IV, m_1\rangle_{X_1}|0\rangle_{Y_1}|m_2\rangle_{X_2}|y_2\rangle_{Y_2}|\emptyset\rangle_{D_f}\big|\{(m_1 \parallel m_2, y_2)\}\big\rangle_{D_H}, \\
|\Psi_{3,\text{rem}}'\rangle &= \frac{1}{|Y|} \sum_{y_1, y_2 \in Y} |m_2\rangle_{X_2}|y_2\rangle_{Y_2}\big|\{(m_1 \parallel m_2, y_2)\}\big\rangle_{D_H} \otimes |\Phi_{coh}^{(2)}((IV, m_1), y_1, \emptyset, y_1)\rangle_{X_1, Y_1, D_f} \\
&\quad + \mathsf{c}O_f|\Psi_{2,\text{rem}}'\rangle.
\end{aligned} \tag{11}
$$

[2]This condition is not stringent. For instance, assuming a prefix-free encoding of the form $\mathsf{pf}(m) = \langle|m|\rangle \parallel m \parallel 10^{b-\langle|m|\rangle-1}$ for $|m| < b$, the concatenated message $m_1 \parallel m_2 = 0^{b-3}100 \parallel 111110^{b-5} \in \mathsf{Range}(\mathsf{pf})$, whereas $m_2 \notin \mathsf{Range}(\mathsf{pf})$, and similarly $m_1 \parallel m_1, m_2 \parallel m_1 \notin \mathsf{Range}(\mathsf{pf})$. Furthermore, this requirement can be significantly relaxed; we introduce this condition strictly to streamline the algebraic evaluation of the quantum states, rather than as a fundamentally necessary restriction.

We bound the cumulative error norm of $|\Psi'_{3,\mathrm{rem}}\rangle$ by recursively applying the triangle inequality and the contractivity of O'^{c}_2:

$$\begin{aligned}\||\Psi'_{3,\mathrm{rem}}\rangle\|_2 &\le \||(11)\rangle\|_2 + \|O'^{\mathsf{c}}_2|\Psi'_{2,\mathrm{rem}}\rangle\|_2 \\ &\le \||(11)\rangle\|_2 + \||\Psi'_{2,\mathrm{rem}}\rangle\|_2 \\ &\le \||(11)\rangle\|_2 + \||(10)\rangle\|_2 + \|O'^{\mathsf{c}}_2|\Psi'_{1,\mathrm{rem}}\rangle\|_2 \\ &\le \||(11)\rangle\|_2 + \||(10)\rangle\|_2 + \||(9)\rangle\|_2.\end{aligned}$$

We evaluate the norms of the individual remainder components:

$$\begin{aligned}\||(9)\rangle\|_2 &= \sqrt{\frac{2}{|Y|}}, \\ \||(10)\rangle\|_2 &= \sqrt{\sum_{y_1\in Y}\left(\frac{1}{\sqrt{|Y|}}\right)^2\frac{2}{|Y|}} = \sqrt{\frac{2}{|Y|}}, \\ \||(11)\rangle\|_2 &\le \sqrt{\sum_{y_1,y_2\in Y}\left(\frac{1}{|Y|}\right)^2\frac{8}{|Y|}} = 2\sqrt{\frac{2}{|Y|}}.\end{aligned}$$

Synthesizing the exact bounds yields:

$$\||\Psi'_{3,\mathrm{rem}}\rangle\|_2 \le 2\sqrt{\frac{2}{|Y|}} + \sqrt{\frac{2}{|Y|}} + \sqrt{\frac{2}{|Y|}} = 4\sqrt{\frac{2}{|Y|}} = \sqrt{\frac{32}{|Y|}}.$$

According to the third structural condition defining the ideal good database ig_3, if any output $(m_1 \parallel m_2, y_2)$ is populated in D_H, the continuous evaluation path corresponding to the prefix must strictly exist within D_f. Because the principal branch $|\Psi'_{3,\mathrm{main}}\rangle$ firmly possesses a terminal hash record in D_H while leaving $D_f = \emptyset$, it resides exclusively within the bad subspace defined by ig_3. Consequently, measuring the escape projector $\neg\Pi^{\mathrm{ig}_3}$ on the normalized principal branch yields:

$$\left\|\neg\Pi^{\mathrm{ig}_3}|\Psi'_{3,\mathrm{main}}\rangle\right\|_2 = \||\Psi'_{3,\mathrm{main}}\rangle\|_2 = 1.$$

Applying the reverse triangle inequality under the orthogonal projection, we mathematically establish the definitive lower bound on the transition leakage:

$$\begin{aligned}\left\|\neg\Pi^{\mathrm{ig}_3}|\Psi'_3\rangle\right\|_2 &= \left\|\neg\Pi^{\mathrm{ig}_3}\left(|\Psi'_{3,\mathrm{main}}\rangle + |\Psi'_{3,\mathrm{rem}}\rangle\right)\right\|_2 \\ &\ge \left\|\neg\Pi^{\mathrm{ig}_3}|\Psi'_{3,\mathrm{main}}\rangle\right\|_2 - \left\|\neg\Pi^{\mathrm{ig}_3}|\Psi'_{3,\mathrm{rem}}\rangle\right\|_2 \\ &\ge 1 - \||\Psi'_{3,\mathrm{rem}}\rangle\|_2 \\ &\ge 1 - \sqrt{\frac{32}{|Y|}}.\end{aligned}$$

This completes the proof. □

Theorem 4.6. *Let $|\Psi'_2\rangle$ be the global quantum state after the second query of the distinguishing attack algorithm. During the third sequential query, the transition deviation between the conjugated real simulated primitive oracle and the ideal simulated primitive oracle severely violates the bounded deviation condition. Specifically, the deviation norm is strictly bounded below by:*

$$\left\|\Pi^{\mathrm{ig}_3}\left(VO^{\mathsf{c}}_2V^\dagger - O'^{\mathsf{c}}_2\right)\Pi^{\mathrm{ig}_2}|\Psi'_2\rangle\right\|_2 \ge 1 - \frac{3}{|Y|} - \frac{6\sqrt{2}+\sqrt{3}}{\sqrt{|Y|}}.$$

Proof. We trace the algebraic divergence of the quantum state $|\Psi_2'\rangle$ under the two disparate oracle operators. From the previous evaluation, the state immediately after the second query is partitioned as $|\Psi_2'\rangle = |\Psi_{2,\text{main}}'\rangle + |\Psi_{2,\text{rem}}'\rangle$, where the principal component is:

$$|\Psi_{2,\text{main}}'\rangle = \frac{1}{|Y|} \sum_{y_1,y_2 \in Y} |IV, m_1\rangle_{X_1} |y_1\rangle_{Y_1} |m_2\rangle_{X_2} |y_2\rangle_{Y_2} \big|\{(IV, m_1, y_1)\}\big\rangle_{D_f} \big|\{(m_1 \parallel m_2, y_2)\}\big\rangle_{D_H}. \tag{12}$$

Crucially, the summation over $(y_1, y_2) \in Y^2$ inherently includes bad configurations where the outputs collide or hit the initial vector IV. These specific configurations structurally violate the definition of the ideal good database ig_2 and must be filtered. We define the disjoint subset of bad configurations $S_{\text{bad}} \subset Y^2$ comprising three mutually exclusive cases:

1. $y_1 = y_2$ (which inherently covers $y_1 = y_2 = IV$). This case contains $|Y|$ elements.
2. $y_1 = IV$ and $y_2 \neq IV$. This case contains $|Y| - 1$ elements.
3. $y_2 = IV$ and $y_1 \neq IV$. This case contains $|Y| - 1$ elements.

The exact cardinality is $|S_{\text{bad}}| = |Y| + 2(|Y| - 1) = 3|Y| - 2$. The remaining safe configurations form the subset $S_{\text{safe}} = Y^2 \setminus S_{\text{bad}}$, consisting of exactly $(|Y| - 1)(|Y| - 2)$ elements where $y_1 \neq y_2$, $y_1 \neq IV$, and $y_2 \neq IV$.

We rigorously partition $|(12)\rangle = |\Psi_{2,\text{main}}''\rangle + |\Psi_{2,\text{bad}}\rangle$, absorbing the bad branches directly into an updated total remainder $|\Psi_{2,\text{rem}}''\rangle = |\Psi_{2,\text{rem}}'\rangle + |\Psi_{2,\text{bad}}\rangle$. The states are explicitly formulated as:

$$|\Psi_{2,\text{main}}''\rangle = \frac{1}{|Y|} \sum_{(y_1,y_2) \in S_{\text{safe}}} |IV, m_1\rangle_{X_1} |y_1\rangle_{Y_1} |m_2\rangle_{X_2} |y_2\rangle_{Y_2} \big|\{(IV, m_1, y_1)\}\big\rangle_{D_f} \big|\{(m_1 \parallel m_2, y_2)\}\big\rangle_{D_H}, \tag{13}$$

$$|\Psi_{2,\text{bad}}\rangle = \frac{1}{|Y|} \sum_{(y_1,y_2) \in S_{\text{bad}}} |IV, m_1\rangle_{X_1} |y_1\rangle_{Y_1} |m_2\rangle_{X_2} |y_2\rangle_{Y_2} \big|\{(IV, m_1, y_1)\}\big\rangle_{D_f} \big|\{(m_1 \parallel m_2, y_2)\}\big\rangle_{D_H}. \tag{14}$$

The exact Euclidean norm of the isolated bad penalty evaluates to:

$$\| |(14)\rangle \|_2 = \sqrt{\sum_{(y_1,y_2) \in S_{\text{bad}}} \left(\frac{1}{|Y|}\right)^2} = \sqrt{\frac{3|Y| - 2}{|Y|^2}} \leq \sqrt{\frac{3}{|Y|}}.$$

By applying the triangle inequality, the global norm of the updated total remainder is strictly bounded by:

$$\| |\Psi_{2,\text{rem}}''\rangle \|_2 \leq \| |\Psi_{2,\text{rem}}'\rangle \|_2 + \| |(14)\rangle \|_2 \leq \sqrt{\frac{2}{|Y|}} + \sqrt{\frac{3}{|Y|}}. \tag{15}$$

Because $|(13)\rangle$ is strictly purged of collisions and IV-hitting outputs, and perfectly satisfies the path-dependency condition for its D_H records, it entirely resides within the ideal good subspace ig_2. Thus, the projection operator acts as the strict identity on this component: $\Pi^{\text{ig}_2} |\Psi_{2,\text{main}}''\rangle = |\Psi_{2,\text{main}}''\rangle$.

We evaluate the conjugated real oracle operator $V O_2^{\varsigma} V^{\dagger}$ acting on $|\Psi_{2,\text{main}}''\rangle$:

1. **Adjoint Mapping ($V^\dagger$):** The isometry $V^\dagger$ seamlessly maps the completed path record in D_H back to the real chaining record in D_f.

2. **Real Oracle Uncomputation (O_2^{c}):** The third query targets (IV, m_1) to uncompute the external register Y_1. Extracting the coherence penalty $|\Phi_{coh}^{\text{real}}\rangle$, the principal real state correctly evaluates to:

$$|\Psi_{3,\text{main}}^{\text{real}}\rangle = \frac{1}{|Y|} \sum_{(y_1,y_2)\in S_{\text{safe}}} |IV, m_1\rangle_{X_1}|0\rangle_{Y_1}|m_2\rangle_{X_2}|y_2\rangle_{Y_2}\big|\{(IV, m_1, y_1), (y_1, m_2, y_2)\}\big\rangle_{D_f}|\emptyset\rangle_{D_H}. \tag{16}$$

3. **Forward Mapping (V):** Because the internal database D_f rigidly preserves the complete continuous path from IV, the isometry V successfully compresses this path. It transitions the final output y_2 into the macroscopic hash database D_H while retaining the prefix record in D_f:

$$V|\Psi_{3,\text{main}}^{\text{real}}\rangle = \frac{1}{|Y|} \sum_{(y_1,y_2)\in S_{\text{safe}}} |IV, m_1\rangle_{X_1}|0\rangle_{Y_1}|m_2\rangle_{X_2}|y_2\rangle_{Y_2}\big|\{(IV, m_1, y_1)\}\big\rangle_{D_f}\big|\{(m_1 \parallel m_2, y_2)\}\big\rangle_{D_H}.$$

4. **Final Projection (Π^{ig_3}):** Crucially, because this compressed state correctly retains the requisite prefix evaluation path in D_f to support the terminal record in D_H, and because S_{safe} strictly guarantees no collisions or IV-hits, it perfectly satisfies all conditions of the ideal good database. Thus, the projection perfectly preserves it: $\Pi^{\text{ig}_3} V|\Psi_{3,\text{main}}^{\text{real}}\rangle = V|\Psi_{3,\text{main}}^{\text{real}}\rangle$.

In parallel, we evaluate the ideal simulated primitive oracle O'^{c}_2 acting on $|\Psi''_{2,\text{main}}\rangle$. Because the inverse Fourier transform inherently requires the complete uniform superposition over Y to perfectly annihilate the prefix record $D_f(IV, m_1)$ back to the vacuum state $|\emptyset\rangle_{D_f}$, directly evaluating the truncated state $|\Psi''_{2,\text{main}}\rangle$ is algebraically unwieldy. Instead, we exploit the linearity of the oracle by rewriting the principal component back to its full summation: $|\Psi''_{2,\text{main}}\rangle = |\Psi'_{2,\text{main}}\rangle - |\Psi_{2,\text{bad}}\rangle$.

Applying O'^{c}_2 to the intact superposition $|\Psi'_{2,\text{main}}\rangle$ allows the dynamic uncomputation to cleanly purge D_f, while strictly preserving the macroscopic record in D_H. Extracting the ideal coherence penalty $|\Phi_{coh}^{\text{ideal}}\rangle$, this yields:

$$|\Psi_{3,\text{main}}^{\text{ideal}}\rangle = \frac{1}{\sqrt{|Y|}} \sum_{y_2\in Y} |IV, m_1\rangle_{X_1}|0\rangle_{Y_1}|m_2\rangle_{X_2}|y_2\rangle_{Y_2}|\emptyset\rangle_{D_f}\big|\{(m_1 \parallel m_2, y_2)\}\big\rangle_{D_H}. \tag{17}$$

Because the macroscopic record exists securely in D_H while its requisite prefix path is completely absent from D_f, |(17)⟩ violates the continuous evaluation path condition (c) of the ideal good database. The projection consequently cleanly annihilates this state: $\Pi^{\text{ig}_3}|\Psi_{3,\text{main}}^{\text{ideal}}\rangle = 0$. The subtracted component $-O'^{\mathsf{c}}_2|\Psi_{2,\text{bad}}\rangle$ propagates forward as a strictly bounded new unitary error term.

By synthesizing the operators over the complete projected state, the deviation state vector expands definitively as:

$$\begin{aligned}
&\Pi^{\text{ig}_3}\left(VO_2^{\mathsf{c}}V^\dagger - O'^{\mathsf{c}}_2\right)\Pi^{\text{ig}_2}|\Psi'_2\rangle \\
&= \Pi^{\text{ig}_3}VO_2^{\mathsf{c}}V^\dagger\left(|\Psi''_{2,\text{main}}\rangle + \Pi^{\text{ig}_2}|\Psi'_{2,\text{rem}}\rangle\right) - \Pi^{\text{ig}_3}O'^{\mathsf{c}}_2\left(|\Psi'_{2,\text{main}}\rangle - |\Psi_{2,\text{bad}}\rangle + \Pi^{\text{ig}_2}|\Psi'_{2,\text{rem}}\rangle\right) \\
&= V|\Psi_{3,\text{main}}^{\text{real}}\rangle - \Pi^{\text{ig}_3}V|\Phi_{coh}^{\text{real}}\rangle + \Pi^{\text{ig}_3}|\Phi_{coh}^{\text{ideal}}\rangle \\
&\quad + \Pi^{\text{ig}_3}O'^{\mathsf{c}}_2|\Psi_{2,\text{bad}}\rangle + \Pi^{\text{ig}_3}VO_2^{\mathsf{c}}V^\dagger\Pi^{\text{ig}_2}|\Psi'_{2,\text{rem}}\rangle - \Pi^{\text{ig}_3}O'^{\mathsf{c}}_2\Pi^{\text{ig}_2}|\Psi'_{2,\text{rem}}\rangle.
\end{aligned} \tag{18}$$

To establish the absolute lower bound, we apply the reverse triangle inequality directly on $|(18)\rangle$. Because V is an isometry, the Euclidean norm of the preserved real principal state strictly computes to:

$$\|V|\Psi^{\text{real}}_{3,\text{main}}\rangle\|_2 = \sqrt{\frac{(|Y|-1)(|Y|-2)}{|Y|^2}} \geq 1 - \frac{3}{|Y|}.$$

Because orthogonal projectors and unitaries are globally contractive, substituting the established remainder norms evaluates the rigorous final bound:

$$\begin{aligned}\left\|\Pi^{\text{ig}_3}\left(VO_2^{\mathsf{c}}V^\dagger - O'^{\mathsf{c}}_2\right)\Pi^{\text{ig}_2}|\Psi'_2\rangle\right\|_2 &\geq \|V|\Psi^{\text{real}}_{3,\text{main}}\rangle\|_2 - \||\Phi^{\text{real}}_{coh}\rangle\|_2 - \||\Phi^{\text{ideal}}_{coh}\rangle\|_2 \\ &\quad - \||\Psi_{2,\text{bad}}\rangle\|_2 - 2\||\Psi'_{2,\text{rem}}\rangle\|_2 \\ &\geq \left(1 - \frac{3}{|Y|}\right) - 2\sqrt{\frac{2}{|Y|}} - 2\sqrt{\frac{2}{|Y|}} - \sqrt{\frac{3}{|Y|}} - 2\left(\sqrt{\frac{2}{|Y|}}\right) \\ &= 1 - \frac{3}{|Y|} - \frac{6\sqrt{2}+\sqrt{3}}{\sqrt{|Y|}}.\end{aligned}$$

This completes the proof. □

4.3 Optimal Consistency Bound

Recall that in the real world, the adversary interacts with the oracle pair $(O_1^{\mathsf{c}}, O_2^{\mathsf{c}}) = (\mathsf{MD}_f[S^{\mathsf{c}O_H}], S^{\mathsf{c}O_H})$, whereas in the ideal world, the adversary interacts with the pair $(O'^{\mathsf{c}}_1, O'^{\mathsf{c}}_2) = (\mathsf{c}O_H, S^{\mathsf{c}O_H})$. We assume the adversary makes at most q_c queries to the construction oracle (the first oracle) and q_s queries to the primitive oracle (the second oracle). Let ℓ denote the maximum input block length for the construction oracle.

To rigorously bound the deviation between the real and ideal executions, we define a sequence of projections $\Pi^{g_0}, \ldots, \Pi^{g_{q_S}}$ acting on the joint Hilbert space of the local database D_f and the simulator's auxiliary label register $|s\rangle$. Specifically, g_i denotes the subspace spanned by all computational basis states wherein the database D_f contains at most i records, contains no collisions, no output hits the initial vector IV, and the label bit is strictly $|s\rangle = |0\rangle$.

We interleave this projection sequence into the combined query execution as follows: we insert the initial projector Π^{g_0} immediately before the very first oracle query. Subsequently, immediately after the i-th query to the primitive oracle (O_2^{c} or O'^{c}_2), we insert the projector Π^{g_i}. For any query to the construction oracle (O_1^{c} or O'^{c}_1) that occurs after the i-th primitive query but before the $(i+1)$-th primitive query, we append the current projector Π^{g_i} immediately after it. This sequential application continues until the entire execution concludes.

Remark 4.7. *It is worth noting that the label bit $|s\rangle$ of the simulator S does not necessarily return to $|0\rangle$ after the uncomputation phase, and thus cannot be trivially ignored. If S fails to find a valid evaluation path within D_f, it operates dynamically on the control bits of the label computation function (i.e., the configuration of D_f). This inherent operation introduces the fundamental possibility of an uncomputation failure.*

Before establishing the bounding lemma for the primitive oracle, we first show that the ideal macroscopic construction oracle introduces identically zero leakage with respect to our defined good subspace g_i.

Lemma 4.8. *For any $i \in \{0, \ldots, q_s\}$, the ideal macroscopic oracle $O'^{\mathsf{c}}_1 = \mathsf{c}O_H$ strictly preserves the good subspace g_i. That is,*

$$\left\|\neg\Pi^{g_i}O'^{\mathsf{c}}_1\Pi^{g_i}\right\|_\infty = 0.$$

Proof. By Definition, the good subspace g_i imposes constraints strictly on the local database D_f (record count, collisions, IV-hits) and the simulator's label bit $|s\rangle$. The ideal random oracle $\mathsf{c}O_H$ evaluates purely into the macroscopic database D_H and acts as the identity on both D_f and $|s\rangle$. Consequently, $\mathsf{c}O_H$ commutes completely with Π^{g_i}, ensuring no probability mass can escape to $\neg g_i$. □

Next, we rigorously bound the cardinality of the local bad output set that could trigger a violation of g_i when a new record is sampled. Let x be the current queried input and $d_x = D_f \setminus \{x\}$ be the local database excluding the query x.

Lemma 4.9. *Assume d_x contains no collisions, no IV-hits, and no complete evaluation path terminating at x. If adding a new record (x, y') fails to form a complete path, then deleting it clearly maintains this property. We define $L_{d_x} \subset Y$ as the set of modified outputs y'' such that the updated database $d_x \cup \{(x, y'')\}$ either forms a complete path terminating at x, produces a collision, or hits IV. The cardinality of this bad set is strictly bounded by:*

$$|L_{d_x}| \leq 2|d_x| + 2.$$

Proof. The outputs $y'' \in L_{d_x}$ arise from two mutually exclusive categories of bad events:

1. **Collisions and IV-hits:** The new output y'' either hits IV or collides with the output of an existing record in d_x. Because d_x contains $|d_x|$ records, there are at most $|d_x|$ existing outputs. Including the IV, this category contributes at most $|d_x| + 1$ elements to L_{d_x}.

2. **Path Formation:** Assuming the first event does not occur, for y'' to enable $d_x \cup \{(x, y'')\}$ to suddenly form a continuous evaluation path, the new record must bridge an existing disconnected graph. Because it does not connect to existing outputs (no collisions), y'' must collide directly with the chaining input (the state component) of some existing node. The available chaining inputs comprise exactly the inputs of the $|d_x|$ records in d_x, plus the chaining input of x itself. This imposes at most $|d_x| + 1$ constraints.

Summing the maximum bounds of both categories, we obtain $|L_{d_x}| \leq (|d_x| + 1) + (|d_x| + 1) = 2|d_x| + 2$. □

Equipped with this structural bound, we establish the leakage bound for a single query to the primitive oracle.

Lemma 4.10. *For any query index $i \in \{1, \ldots, q_s\}$, the transition leakage of the primitive oracle is bounded by:*

$$\|\neg\Pi^{g_i} O_2^{\mathsf{c}} \Pi^{g_{i-1}}\|_\infty = \left\|\neg\Pi^{g_i} O'^{\mathsf{c}}_2 \Pi^{g_{i-1}}\right\|_\infty \leq \sqrt{\frac{26i}{|Y|}}.$$

Proof. Because both the real primitive oracle O_2^{c} and the ideal primitive oracle O'^{c}_2 utilize the identical underlying simulator $S^{\mathsf{c}O_H}$, and because the definition of g_i is completely independent of the macroscopic hash oracle $\mathsf{c}O_H$, we only need to bound the transition leakage of $\mathsf{c}O_f$.

Consider an arbitrary fixed evaluation branch defined by the query x and the surrounding database configuration d_x. We only need to consider the subspace where the database has support on $d_x \cup \{(x, y')\}$, this joint configuration and the label bit strictly satisfy g_{i-1}, and the database structurally cannot form a complete evaluation path terminating at x.

Let $\overline{Y} = Y \cup \{\perp\}$. The initial normalized state constrained within this valid subspace is explicitly defined as:

$$|\Psi_0\rangle = \sum_{y \in Y} \sum_{\substack{y' \in \overline{Y} \text{ s.t. } d_x \cup \{(x,y')\} \\ \text{satisfies } g_{i-1} \text{ and} \\ \text{lacks complete path at } x}} \alpha_{y,y'} |x\rangle_X |y\rangle_Y \big| d_x \cup \{(x,y')\} \big\rangle_{D_f}.$$

We orthogonally decompose this initial state based on whether the local record at x is currently undefined ($y' = \perp$) or already defined ($y' = z \in Y$):

$$|\Psi_0^{\perp}\rangle = \sum_{y \in Y} \alpha_{y,\perp} |x\rangle_X |y\rangle_Y \big| d_x \cup \{(x,\perp)\} \big\rangle_{D_f},$$
$$|\Psi_0^{\neq \perp}\rangle = \sum_{y \in Y} \sum_{z \in Y} \alpha_{y,z} |x\rangle_X |y\rangle_Y \big| d_x \cup \{(x,z)\} \big\rangle_{D_f}.$$

Applying the compressed oracle yields the updated quantum state $|\Psi_1\rangle = \mathsf{c}O_f |\Psi_0^{\perp}\rangle + \mathsf{c}O_f |\Psi_0^{\neq \perp}\rangle$. According to the fundamental action of $\mathsf{c}O_f$, this linearly expands into four distinct structural terms (lazy sampling, direct response, and two coherent penalty terms):

$$\mathsf{c}O_f |\Psi_0^{\perp}\rangle = \underbrace{\sum_{y \in Y} \frac{\alpha_{y,\perp}}{\sqrt{|Y|}} \sum_{\hat{y} \in Y} |x\rangle |y + \hat{y}\rangle \big| d_x \cup \{(x,\hat{y})\} \big\rangle}_{|\Psi_{\text{samp}}\rangle} + \underbrace{\sum_{y \in Y} \alpha_{y,\perp} |\Phi_{coh}^{(1)}(x,y,d_x)\rangle}_{|\Psi_{\text{coh1}}\rangle},$$
$$\mathsf{c}O_f |\Psi_0^{\neq \perp}\rangle = \underbrace{\sum_{y,z \in Y} \alpha_{y,z} |x\rangle |y + z\rangle \big| d_x \cup \{(x,z)\} \big\rangle}_{|\Psi_{\text{dir}}\rangle} + \underbrace{\sum_{y,z \in Y} \alpha_{y,z} |\Phi_{coh}^{(2)}(x,y,d_x,z)\rangle}_{|\Psi_{\text{coh2}}\rangle}.$$

We subsequently apply the orthogonal projector $\neg\Pi^{g_i}$ to evaluate the transition leakage. By the triangle inequality, the leakage norm is strictly bounded by the sum of the individually projected components:

$$\big\| \neg\Pi^{g_i} \mathsf{c}O_f |\Psi_0\rangle \big\|_2 \leq \big\| \neg\Pi^{g_i} |\Psi_{\text{samp}}\rangle \big\|_2 + \| \neg\Pi^{g_i} |\Psi_{\text{dir}}\rangle \|_2 + \| \neg\Pi^{g_i} |\Psi_{\text{coh1}}\rangle \|_2 + \| \neg\Pi^{g_i} |\Psi_{\text{coh2}}\rangle \|_2.$$

We analyze the projection of each component into the bad subspace defined by the local bad set L_{d_x}:

1. For the direct response $|\Psi_{\text{dir}}\rangle$, because the initial configuration inherently satisfied g_{i-1} and the unmodified output z triggers no new structural violations, it remains perfectly within g_i. Thus, it is cleanly annihilated: $\| \neg\Pi^{g_i} |\Psi_{\text{dir}}\rangle \|_2 = 0$.

2. For the lazy sampling term $|\Psi_{\text{samp}}\rangle$, the branches where $\hat{y} \notin L_{d_x}$ identically preserve the structural integrity required by g_i and are projected away. Only the amplitude corresponding to $\hat{y} \in L_{d_x}$ escapes: $\big\| \neg\Pi^{g_i} |\Psi_{\text{samp}}\rangle \big\|_2 \leq \sqrt{\frac{|L_{d_x}|}{|Y|}} \sqrt{\sum |\alpha_{y,\perp}|^2}$.

3. For the coherence terms, their constituent superpositions strictly follow the established evaluation mechanics of the compressed oracle. The unaugmented base branches inherently satisfy g_i and are annihilated, meaning the residual coherence leakage scales proportionally to the intersection with the bad set L_{d_x}. Specifically, the projected norms are analytically bounded by $\| \neg\Pi^{g_i} |\Psi_{\text{coh1}}\rangle \|_2 \leq \sqrt{\frac{|L_{d_x}|}{|Y|}} \sqrt{\sum |\alpha_{y,\perp}|^2}$ and $\| \neg\Pi^{g_i} |\Psi_{\text{coh2}}\rangle \|_2 \leq 3\sqrt{\frac{|L_{d_x}|}{|Y|}} \sqrt{\sum |\alpha_{y,z}|^2}$.

Synthesizing these exact fractional bounds and applying the Cauchy-Schwarz inequality across the orthogonal state amplitudes, we resolve the absolute escape transition limit:

$$\begin{aligned}\left\|\neg\Pi^{g_i}\mathrm{c}O_f|\Psi_0\rangle\right\|_2 &\le 2\sqrt{\frac{|L_{d_x}|}{|Y|}}\sqrt{\sum_{y\in Y}|\alpha_{y,\perp}|^2}+3\sqrt{\frac{|L_{d_x}|}{|Y|}}\sqrt{\sum_{y,z\in Y}|\alpha_{y,z}|^2}\\ &\le \sqrt{2^2+3^2}\sqrt{\sum_{y}|\alpha_{y,\perp}|^2+\sum_{y,z}|\alpha_{y,z}|^2}\sqrt{\frac{|L_{d_x}|}{|Y|}}\\ &= \sqrt{13\frac{|L_{d_x}|}{|Y|}}.\end{aligned}$$

Substituting the cardinality bound $|L_{d_x}| \le 2|d_x|+2$ from Lemma 4.9, and enforcing the structural constraint $|d_x| \le i-1$ explicitly guaranteed by the initial projection g_{i-1}, we secure the absolute upper bound $|L_{d_x}| \le 2(i-1)+2 = 2i$. Substituting this directly into the synthesized norm bound perfectly evaluates to:

$$\left\|\neg\Pi^{g_i}\mathrm{c}O_f\Pi^{g_{i-1}}\right\|_\infty \le \sqrt{\frac{13\cdot 2i}{|Y|}} = \sqrt{\frac{26i}{|Y|}}.$$

This completes the proof. □

Lemma 4.11. *For any query index $i \in \{0,\dots,q_s\}$, the transition leakage of the real construction oracle and the bounded deviation between the real and ideal construction oracles are strictly upper-bounded by the maximal evaluation depth ℓ:*

$$\begin{aligned}\|\neg\Pi^{g_i}O_1^{\mathsf{c}}\Pi^{g_i}\|_\infty &\le 2\sum_{k=1}^{\ell-1}\sqrt{\frac{26(i+k)}{|Y|}},\\ \left\|\Pi^{g_i}\left(O_1^{\mathsf{c}}-O'^{\mathsf{c}}_1\right)\Pi^{g_i}\right\|_\infty &\le 2\sum_{k=1}^{\ell-1}\sqrt{\frac{26(i+k)}{|Y|}}.\end{aligned}$$

Proof. We meticulously analyze the internal operational mechanics of the real construction oracle $O_1^{\mathsf{c}} = \mathsf{MD}_f[S^{\mathrm{c}O_H}]$. To rigorously trace the exact algebraic evolution and error propagation without loss of generality, we fix an arbitrary prefix-free encoded message consisting of exactly three blocks, $M = (m_1, m_2, m_3)$. The complete circuit evaluation entails exactly three forward evaluations and two backward uncomputations. Therefore, the global operation translates algebraically to five sequential invocations of the underlying simulator, denoted as $S := S^{\mathrm{c}O_H}$, such that $O_1^{\mathsf{c}}|\Psi_0\rangle = S^5|\Psi_0\rangle$.

Let the initial normalized state, residing entirely within the prerequisite good subspace g_i, be explicitly formalized as:

$$|\Psi_0\rangle = \sum_{y\in Y}\sum_{d\in g_i}\alpha_{y,d}|IV,m_1\rangle_{X_1}|0\rangle_{Y_1}|m_2\rangle_{X_2}|0\rangle_{Y_2}|m_3\rangle_{X_3}|y\rangle_{Y_3}|d\rangle_{D_f}|0\rangle_S.$$

First Forward Query (Step 1). The simulated primitive acts on $|\Psi_0\rangle$ targeting $X_1 = (IV, m_1)$. Because $d \in g_i$, the initial vector fundamentally admits no complete evaluation path, causing S to bypass $\mathrm{c}O_H$ and execute the internal primitive $\mathrm{c}O_f$. Orthogonally decomposing $|\Psi_0\rangle = |\Psi_0^{\perp}\rangle + |\Psi_0^{\neq\perp}\rangle$ and applying $\mathrm{c}O_f$, the state expands into five structurally distinct

terms. Let $|\Phi_{\text{rest}}\rangle$ denote the unperturbed spectator registers:

$$|\Psi_{1,\text{dir}}\rangle = \sum_y \sum_{d(x_1)=z} \alpha_{y,d} |x_1\rangle_{X_1} |z\rangle_{Y_1} |\Phi_{\text{rest}}\rangle |d\rangle_{D_f},$$

$$|\Psi_{1,\text{samp_safe}}\rangle = \sum_y \sum_{d(x_1)=\perp} \frac{\alpha_{y,d}}{\sqrt{|Y|}} \sum_{\hat{y}\notin L_d} |x_1\rangle_{X_1} |\hat{y}\rangle_{Y_1} |\Phi_{\text{rest}}\rangle \big|d \cup \{(x_1,\hat{y})\}\big\rangle_{D_f},$$

$$|\Psi_{1,\text{samp_bad}}\rangle = \sum_y \sum_{d(x_1)=\perp} \frac{\alpha_{y,d}}{\sqrt{|Y|}} \sum_{\hat{y}\in L_d} |x_1\rangle_{X_1} |\hat{y}\rangle_{Y_1} |\Phi_{\text{rest}}\rangle \big|d \cup \{(x_1,\hat{y})\}\big\rangle_{D_f},$$

$$|\Psi_{1,\text{coh1}}\rangle = \sum_y \sum_{d(x_1)=\perp} \alpha_{y,d} |\Phi^{(1)}_{coh}(x_1,0,d)\rangle_{X_1,Y_1,D_f} \otimes |\Phi_{\text{rest}}\rangle,$$

$$|\Psi_{1,\text{coh2}}\rangle = \sum_y \sum_{d(x_1)=z} \alpha_{y,d} |\Phi^{(2)}_{coh}(x_1,0,d,z)\rangle_{X_1,Y_1,D_f} \otimes |\Phi_{\text{rest}}\rangle.$$

Crucially, for the structurally sound branches ($|\Psi_{1,\text{dir}}\rangle$ and $|\Psi_{1,\text{samp_safe}}\rangle$), the integrity of the prefix-free path guarantees the label bit safely uncomputes to $|s\rangle = |0\rangle$. We define this preserved principal state as $|\Psi_{1,\text{main}}\rangle = |\Psi_{1,\text{dir}}\rangle + |\Psi_{1,\text{samp_safe}}\rangle$. The three residual components ($|\Psi_{1,\text{samp_bad}}\rangle, |\Psi_{1,\text{coh1}}\rangle, |\Psi_{1,\text{coh2}}\rangle$) are directly aggregated into the first forward error round $|E_{1f}\rangle$. It is worth noting that due to structural violations (e.g., collisions creating unexpected cycles), the label bit $|s\rangle$ within $|E_{1f}\rangle$ might not revert to $|0\rangle$. However, since our subsequent analysis strictly bounds the maximum Euclidean norm of this error vector, we can safely ignore the internal algebraic configuration of the label bit. Thus, we have:

$$S|\Psi_0\rangle = |\Psi_{1,\text{main}}\rangle + |E_{1f}\rangle.$$

Second Forward Query (Step 2). Operating on $|\Psi_{1,\text{main}}\rangle$, the simulator queries $X_2 = m_2$. The trace identifies $(IV, m_1) \to m_2$. Due to prefix-free encoding, $m_1 \parallel m_2$ is not a complete message, maintaining $|s\rangle = |0\rangle$ and enforcing another $\mathsf{c}O_f$ query. Identically decomposing $|\Psi_{1,\text{main}}\rangle$ based on X_2 and expanding via $\mathsf{c}O_f$ yields a safe principal component $|\Psi_{2,\text{main}}\rangle = |\Psi_{2,\text{dir}}\rangle + |\Psi_{2,\text{samp_safe}}\rangle$, alongside a second error aggregation $|E_{2f}\rangle$. Therefore:

$$S|\Psi_{1,\text{main}}\rangle = |\Psi_{2,\text{main}}\rangle + |E_{2f}\rangle.$$

Third Ideal Query (Step 3). Operating on $|\Psi_{2,\text{main}}\rangle$, the simulator traces the complete prefix-free message $m_1 \parallel m_2 \parallel m_3$. The label bit deterministically flips to $|s\rangle = |1\rangle$, routing the evaluation exclusively to the macroscopic ideal oracle $\mathsf{c}O_H$. Because $\mathsf{c}O_H$ commutes with D_f, the structural topology is perfectly preserved. The updated principal state is defined as $|\tilde{\Psi}_{2,\text{main}}\rangle$, yielding:

$$S|\Psi_{2,\text{main}}\rangle = |\tilde{\Psi}_{2,\text{main}}\rangle.$$

First Uncomputation (Step 4). The simulator S now uncomputes X_2, Y_2 on $|\tilde{\Psi}_{2,\text{main}}\rangle$. Tracing the non-terminal prefix $m_1 \parallel m_2$ ensures no label flip, executing a $\mathsf{c}O_f$ uncomputation. We introduce an algebraic correction term by simultaneously adding and subtracting the corresponding bad sampling branches:

$$S|\tilde{\Psi}_{2,\text{samp_safe}}\rangle = \mathsf{c}O_f\Big(|\tilde{\Psi}_{2,\text{samp_safe}}\rangle + |\tilde{\Psi}_{2,\text{samp_bad}}\rangle\Big) - \mathsf{c}O_f|\tilde{\Psi}_{2,\text{samp_bad}}\rangle.$$

By the involutory mechanics of the compressed oracle, the intact sampling and direct components cleanly uncompute to the prior states $|\tilde{\Psi}_1^{\perp}\rangle$ and $|\tilde{\Psi}_1^{\neq\perp}\rangle$, generating coherence penalties $|\tilde{\Psi}'_{2,\text{coh1}}\rangle$ and $|\tilde{\Psi}'_{2,\text{coh2}}\rangle$. Reassembling the principal state $|\tilde{\Psi}_{1,\text{main}}\rangle = |\tilde{\Psi}_1^{\perp}\rangle + |\tilde{\Psi}_1^{\neq\perp}\rangle$ and

assigning the correction and coherence penalties to the first backward error round $|E_{2b}\rangle$, we establish:

$$S|\tilde{\Psi}_{2,\text{main}}\rangle = |\tilde{\Psi}_{1,\text{main}}\rangle - |E_{2b}\rangle.$$

Second Uncomputation (Step 5). Finally, S uncomputes X_1, Y_1 on $|\tilde{\Psi}_{1,\text{main}}\rangle$. Applying identical correction logic with $|\tilde{\Psi}_{1,\text{samp_bad}}\rangle$, the principal state perfectly collapses back to $|\tilde{\Psi}_0\rangle$. Forming the final error round $|E_{1b}\rangle$, we obtain:

$$S|\tilde{\Psi}_{1,\text{main}}\rangle = |\tilde{\Psi}_0\rangle - |E_{1b}\rangle.$$

Global Error Accumulation and Final Bounds. Crucially, $|\tilde{\Psi}_0\rangle$ is identically the initial state $|\Psi_0\rangle$ perfectly updated by the macroscopic hash evaluation $\mathsf{c}O_H$. Thus, $|\tilde{\Psi}_0\rangle = O'^{\mathsf{c}}_1|\Psi_0\rangle$.

By linearly expanding the global circuit evaluation $O^{\mathsf{c}}_1|\Psi_0\rangle = S^5|\Psi_0\rangle$ through iterative substitution of the stepwise state transitions, we rigorously capture the exact error propagation:

$$\begin{aligned}
S|\Psi_0\rangle &= |\Psi_{1,\text{main}}\rangle + |E_{1f}\rangle,\\
S^2|\Psi_0\rangle &= |\Psi_{2,\text{main}}\rangle + |E_{2f}\rangle + S|E_{1f}\rangle,\\
S^3|\Psi_0\rangle &= |\tilde{\Psi}_{2,\text{main}}\rangle + S|E_{2f}\rangle + S^2|E_{1f}\rangle,\\
S^4|\Psi_0\rangle &= |\tilde{\Psi}_{1,\text{main}}\rangle - |E_{2b}\rangle + S^2|E_{2f}\rangle + S^3|E_{1f}\rangle,\\
S^5|\Psi_0\rangle &= |\tilde{\Psi}_0\rangle - |E_{1b}\rangle - S|E_{2b}\rangle + S^3|E_{2f}\rangle + S^4|E_{1f}\rangle.
\end{aligned}$$

Substituting $|\tilde{\Psi}_0\rangle = O'^{\mathsf{c}}_1|\Psi_0\rangle$, the global action is exactly formulated as:

$$O^{\mathsf{c}}_1|\Psi_0\rangle = O'^{\mathsf{c}}_1|\Psi_0\rangle + S^4|E_{1f}\rangle + S^3|E_{2f}\rangle - S|E_{2b}\rangle - |E_{1b}\rangle.$$

Because the underlying simulator is fundamentally a contractive operator ($\|S\|_\infty \leq 1$), the sequential applications S^m strictly do not amplify the Euclidean norms of the propagated error vectors (e.g., $\|S^4|E_{1f}\rangle\|_2 \leq \||E_{1f}\rangle\|_2$). Furthermore, because the ideal oracle perfectly preserves the good subspace (Lemma 4.8), $\neg\Pi^{g_i}O'^{\mathsf{c}}_1|\Psi_0\rangle = 0$.

Applying the reverse triangle inequality directly yields the simultaneous bounds for transition leakage and deviation:

$$\begin{aligned}
\|\neg\Pi^{g_i}O^{\mathsf{c}}_1|\Psi_0\rangle\|_2 &\leq \||E_{1f}\rangle\|_2 + \||E_{2f}\rangle\|_2 + \||E_{2b}\rangle\|_2 + \||E_{1b}\rangle\|_2,\\
\left\|\Pi^{g_i}\left(O^{\mathsf{c}}_1 - O'^{\mathsf{c}}_1\right)|\Psi_0\rangle\right\|_2 &\leq \||E_{1f}\rangle\|_2 + \||E_{2f}\rangle\|_2 + \||E_{2b}\rangle\|_2 + \||E_{1b}\rangle\|_2.
\end{aligned}$$

Generalization to Maximum Block Length ℓ. While the algebraic trace above was explicitly derived for a 3-block message to illustrate the structural evolution, the inductive framework generalizes seamlessly to any valid prefix-free message $M = (m_1, \ldots, m_L)$ of arbitrary block length $L \leq \ell$. For a valid message of length L, the circuit triggers exactly $L-1$ internal forward queries to $\mathsf{c}O_f$ and $L-1$ corresponding backward uncomputations, with the terminal L-th query interacting exclusively with the macroscopic ideal oracle $\mathsf{c}O_H$.

At the k-th internal query (where $1 \leq k \leq L-1$), the local database has been safely augmented by at most $k-1$ sampled records. Because the initial state satisfies g_i, the database size is strictly bounded by $|d| \leq i+k-1$. Consequently, the maximum volume of the local bad set dynamically expands to $|L_d| \leq 2|d|+2 \leq 2(i+k-1)+2 = 2(i+k)$. By the established Cauchy-Schwarz synthesis (Lemma 4.10), the Euclidean norms for the k-th forward

evaluation error $|E_{kf}\rangle$ and its paired backward uncomputation error $|E_{kb}\rangle$ are identically bounded by:

$$\||E_{kf}\rangle\|_2 = \||E_{kb}\rangle\|_2 \leq \sqrt{13\frac{|L_d|}{|Y|}} \leq \sqrt{\frac{26(i+k)}{|Y|}}.$$

Summing the accumulated error rounds yields the absolute deviation bound for a generic message of length L:

$$2\sum_{k=1}^{L-1}\sqrt{\frac{26(i+k)}{|Y|}}.$$

Because this leakage aggregates monotonically with respect to the query sequence depth, the strict global supremum across all possible adversarial queries is naturally achieved at the maximum allowed block length $L = \ell$:

$$2\sum_{k=1}^{\ell-1}\sqrt{\frac{26(i+k)}{|Y|}}.$$

Since $|\Psi_0\rangle$ is an arbitrary normalized state within g_i, this supremum perfectly characterizes the global operator norm. This completes the proof. □

Theorem 4.12 (Consistency Bound)**.** *The prefix-free Merkle-Damgård construction* MD_f *and its corresponding simulator* $S^{\mathsf{c}O_H}$ *are explicitly* $(q_c, q_s, \epsilon_{con})$*-consistent. The distinguishing advantage of any quantum adversary making at most* q_c *construction queries and* q_s *primitive queries is strictly bounded by:*

$$\epsilon_{con} \leq 2\sum_{i=1}^{q_s}\sqrt{\frac{26i}{2^n}} + 4q_c\sum_{k=1}^{\ell-1}\sqrt{\frac{26(q_s+k)}{2^n}},$$

where n *is the output bit-length of the underlying hash function, and* ℓ *is the maximum block length of the prefix-free encoded message.*

Proof. According to the formal definition of the consistency game, we must bound the distinguishing advantage between the real oracle pair $\mathbf{O} = (O_1^{\mathsf{c}}, O_2^{\mathsf{c}}) = (\mathsf{MD}_f[S^{\mathsf{c}O_H}], S^{\mathsf{c}O_H})$ and the corresponding ideal oracle pair $\mathbf{O}' = (O'^{\mathsf{c}}_1, O'^{\mathsf{c}}_2) = (\mathsf{c}O_H, S^{\mathsf{c}O_H})$. A critical structural observation here is that the primitive oracles in both respective worlds are absolutely identical: $O_2^{\mathsf{c}} = O'^{\mathsf{c}}_2 = S^{\mathsf{c}O_H}$.

To systematically bound the aggregate divergence, we invoke the Fundamental Quantum Game-Playing Lemma (Lemma 3.11), which establishes that the total distinguishing advantage ϵ_{con} is strictly bounded by the sum of the transition leakages and deviation norms across all query steps in the sequence. Throughout the execution, the adversary sequentially makes q_s primitive queries and q_c construction queries. We interleave the sequence of orthogonal projectors $\Pi^{g_0}, \ldots, \Pi^{g_{q_s}}$ such that the subspace index i increments exclusively upon the execution of a primitive query.

We first evaluate the bounds for the primitive queries. For any i-th query to the primitive oracle (where $1 \leq i \leq q_s$), the transition leakages and deviation are determined as:

- **Real Bounded Leakage:** $\epsilon_{s,i} = \|\neg\Pi^{g_i} O_2^{\mathsf{c}} \Pi^{g_{i-1}}\|_\infty \leq \sqrt{\frac{26i}{2^n}}$ (Derived from Lemma 4.10, substituting $|Y| = 2^n$).
- **Ideal Bounded Leakage:** $\epsilon'_{s,i} = \left\|\neg\Pi^{g_i} O'^{\mathsf{c}}_2 \Pi^{g_{i-1}}\right\|_\infty \leq \sqrt{\frac{26i}{2^n}}$ (Derived identically from Lemma 4.10).

- **Bounded Deviation:** $\delta_{s,i} = \left\| \Pi^{g_i} \left(O_2^{\mathsf{c}} - O_2'^{\mathsf{c}} \right) \Pi^{g_{i-1}} \right\|_\infty = 0$, strictly because the real and ideal primitive operators are algebraically identical.

Summing these specific errors across all q_s primitive queries evaluates exactly to an aggregate primitive error of $2\sum_{i=1}^{q_s} \sqrt{\frac{26i}{2^n}}$.

We subsequently evaluate the bounds for the construction queries. For any query to the construction oracle that is executed when the current database projector is Π^{g_i} (meaning exactly i primitive queries have been evaluated so far, where $0 \leq i \leq q_s$), the transition leakages and deviation are determined as:

- **Real Bounded Leakage:** $\epsilon_{c,i} = \| \neg\Pi^{g_i} O_1^{\mathsf{c}} \Pi^{g_i} \|_\infty \leq 2\sum_{k=1}^{\ell-1} \sqrt{\frac{26(i+k)}{2^n}}$ (Derived from Lemma 4.11).
- **Ideal Bounded Leakage:** $\epsilon'_{c,i} = \left\| \neg\Pi^{g_i} O_1'^{\mathsf{c}} \Pi^{g_i} \right\|_\infty = 0$ (Derived from Lemma 4.8, establishing the ideal oracle's perfect preservation of g_i).
- **Bounded Deviation:** $\delta_{c,i} = \left\| \Pi^{g_i} \left(O_1^{\mathsf{c}} - O_1'^{\mathsf{c}} \right) \Pi^{g_i} \right\|_\infty \leq 2\sum_{k=1}^{\ell-1} \sqrt{\frac{26(i+k)}{2^n}}$ (Derived from Lemma 4.11).

The consolidated single-query error for an arbitrary construction query is thus strictly bounded by $4\sum_{k=1}^{\ell-1} \sqrt{\frac{26(i+k)}{2^n}}$.

Because this specific bound monotonically increases with respect to the primitive query index i (due to the progressively expanding local bad set L_d), the absolute supremum across all possible adversarial interleavings of queries natively occurs in the worst-case scenario. This worst-case scenario is actualized when the adversary strategically defers all q_c construction queries to the very end of the sequential evaluation, corresponding precisely to the index $i = q_s$. Under this strict maximal configuration, the accumulated error for all q_c construction queries is bounded unconditionally by $4q_c\sum_{k=1}^{\ell-1} \sqrt{\frac{26(q_s+k)}{2^n}}$.

Summing the evaluated aggregate limits for both the primitive and construction oracle sequences comprehensively envelops the entire execution space, definitively yielding the final consistency advantage bound:

$$\epsilon_{\text{con}} \leq 2\sum_{i=1}^{q_s} \sqrt{\frac{26i}{2^n}} + 4q_c \sum_{k=1}^{\ell-1} \sqrt{\frac{26(q_s+k)}{2^n}}.$$

This completes the proof. □

Remark 4.13. *By evaluating the derived consistency bound asymptotically, the total distinguishing advantage of the quantum adversary is bounded by:*

$$\epsilon_{con} = O\left(\frac{q_s^{3/2} + q_c \ell \sqrt{q_s + \ell}}{2^{n/2}} \right).$$

In the standard cryptographic setting where the maximum message block length ℓ is considered a constant ($\ell = O(1)$), and assuming $q_c \leq q_s$, the bound is strictly dominated by the $O(q_s^{3/2}/2^{n/2})$ term. Consequently, for any quantum adversary to achieve a non-negligible distinguishing advantage ($\epsilon_{con} = \Theta(1)$), the required primitive query complexity must reach $q_s = \Omega(2^{n/3})$. This theoretical query lower bound perfectly matches the optimal complexity of generic quantum collision attacks (e.g., the Brassard-Høyer-Tapp algorithm **??***).*

References

[Dam90] Ivan Bjerre Damgård. "A Design Principle for Hash Functions". In: *Advances in Cryptology – CRYPTO '89*. Springer, 1990, pp. 416–427.

[Mer90] Ralph C. Merkle. "One Way Hash Functions and DES". In: *Advances in Cryptology – CRYPTO '89*. Springer, 1990, pp. 428–446.

[BR93] Mihir Bellare and Phillip Rogaway. "Random Oracles are Practical: A Paradigm for Designing Efficient Protocols". In: *Proceedings of the 1st ACM Conference on Computer and Communications Security*. ACM, 1993, pp. 62–73.

[Sho94] Peter W. Shor. "Algorithms for Quantum Computation: Discrete Logarithms and Factoring". In: *Proceedings 35th Annual Symposium on Foundations of Computer Science*. IEEE. 1994, pp. 124–134.

[BR95] Mihir Bellare and Phillip Rogaway. "Optimal Asymmetric Encryption". In: *Advances in Cryptology – EUROCRYPT '94*. Springer, 1995, pp. 92–111.

[BR96] Mihir Bellare and Phillip Rogaway. "The Exact Security of Digital Signatures: How to Sign with RSA and Rabin". In: *Advances in Cryptology – EUROCRYPT '96*. Springer, 1996, pp. 399–416.

[Gro96] Lov K. Grover. "A Fast Quantum Mechanical Algorithm for Database Search". In: *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*. ACM, 1996, pp. 212–219.

[PS96] David Pointcheval and Jacques Stern. "Security Proofs for Signature Schemes". In: *Advances in Cryptology – EUROCRYPT '96*. Springer, 1996, pp. 387–398.

[BHT98] Gilles Brassard, Peter Høyer, and Alain Tapp. "Quantum Cryptanalysis of Hash and Claw-Free Functions". In: *LATIN '98: Theoretical Informatics*. Springer, 1998, pp. 163–169. DOI: 10.1007/BFb0054319.

[MRH04] Ueli Maurer, Renato Renner, and Clemens Holenstein. "Indifferentiability, Impossibility Results on Reductions, and Applications to the Random Oracle Methodology". In: *Theory of Cryptography Conference*. Springer, 2004, pp. 21–39.

[Sho04] Victor Shoup. *Sequences of Games: A Tool for Taming Complexity in Security Proofs*. Cryptology ePrint Archive, Report 2004/332. 2004. URL: https://eprint.iacr.org/2004/332.

[Cor+05] Jean-Sébastien Coron, Yevgeniy Dodis, Cécile Malinaud, and Prashant Puniya. "Merkle–Damgård Revisited: How to Construct a Hash Function". In: *Advances in Cryptology – CRYPTO 2005*. Springer, 2005, pp. 430–448.

[BR06] Mihir Bellare and Phillip Rogaway. "The Security of Triple Encryption and a Framework for Code-Based Game-Playing Proofs". In: *Advances in Cryptology – EUROCRYPT 2006*. Springer, 2006, pp. 409–426. DOI: 10.1007/11761679_25.

[Cha+06] Donghoon Chang, Sangjin Lee, Mridul Nandi, and Moti Yung. "Indifferentiability of Alternatives to Merkle–Damgård and Damgård's Construction". In: *Advances in Cryptology – ASIACRYPT 2006*. Springer, 2006, pp. 1–16.

[Ber+07] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. "Sponge Functions". In: *ECRYPT Hash Workshop*. 2007.

[Ber+08] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. "On the Indifferentiability of the Sponge Construction". In: *Advances in Cryptology – EUROCRYPT 2008*. Springer, 2008, pp. 181–197.

[NC10] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. 10th Anniversary Edition. Cambridge University Press, 2010.

[Bon+11] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. "Random Oracles in a Quantum World". In: *Advances in Cryptology – ASIACRYPT 2011*. Springer, 2011, pp. 41–69.

[RSS11] Thomas Ristenpart, Hovav Shacham, and Thomas Shrimpton. "Careful with Composition: Limitations of the Indifferentiability Framework". In: *Advances in Cryptology – EUROCRYPT 2011*. Springer, 2011, pp. 487–506. DOI: 10.1007/978-3-642-20465-4_27.

[DPP12] Yevgeniy Dodis, Krzysztof Pietrzak, and Prashant Puniya. "Indifferentiability of Permutation-Based Compression Functions and Tree-Based Modes of Operation, with Applications to SHA-3". In: *Advances in Cryptology – EUROCRYPT 2012*. Springer, 2012, pp. 37–54.

[FIPS 180-4] NIST. *Secure Hash Standard (SHS)*. Tech. rep. FIPS PUB 180-4. National Institute of Standards and Technology, 2015. DOI: 10.6028/NIST.FIPS.180-4. URL: https://doi.org/10.6028/NIST.FIPS.180-4.

[Zha15] Mark Zhandry. "A Note on the Quantum Collision and Set Equality Problems". In: *Quantum Information and Computation* 15.7–8 (2015), pp. 557–567. DOI: 10.26421/QIC15.7-8-2.

[Car+18] Tore Vincent Carstens, Ehsan Ebrahimi, Gelo Noel Tabia, and Dominique Unruh. *On Quantum Indifferentiability*. Cryptology ePrint Archive, Paper 2018/257. 2018. URL: https://eprint.iacr.org/2018/257.

[Cza+19] Jan Czajkowski, Christian Majenz, Christian Schaffner, and Sebastian Zur. *Quantum Lazy Sampling and Game-Playing Proofs for Quantum Indifferentiability*. Cryptology ePrint Archive, Paper 2019/428. 2019. URL: https://eprint.iacr.org/2019/428.

[HI19] Akinori Hosoyamada and Tetsu Iwata. "4-Round Luby–Rackoff Construction is a qPRP". In: *Advances in Cryptology – ASIACRYPT 2019*. Springer, 2019, pp. 145–174. DOI: 10.1007/978-3-030-34578-5_6.

[Zha19] Mark Zhandry. "How to Record Quantum Queries, and Applications to Quantum Indifferentiability". In: *Advances in Cryptology – CRYPTO 2019*. Springer, 2019, pp. 239–268.

[SP 800-208] David A. Cooper, Daniel C. Apon, Quynh H. Dang, Michael S. Davidson, Morris Dworkin, and Carl A. Miller. *Recommendation for Stateful Hash-Based Signature Schemes*. Tech. rep. NIST Special Publication 800-208. National Institute of Standards and Technology, 2020. DOI: 10.6028/NIST.SP.800-208. URL: https://doi.org/10.6028/NIST.SP.800-208.

[Chu+21] Kai-Min Chung, Serge Fehr, Yu-Hsuan Huang, and Tai-Ning Liao. "On the Compressed-Oracle Technique, and Post-Quantum Security of Proofs of Sequential Work". In: *Advances in Cryptology – EUROCRYPT 2021*. Springer, 2021, pp. 598–629. DOI: 10.1007/978-3-030-77886-6_21.

[Zha21] Mark Zhandry. "Redeeming Reset Indifferentiability and Applications to Post-Quantum Security". In: *Advances in Cryptology – ASIACRYPT 2021*. Springer, 2021, pp. 518–548. DOI: 10.1007/978-3-030-92062-3_18.

[DFH22] Jelle Don, Serge Fehr, and Yu-Hsuan Huang. "Adaptive Versus Static Multi-oracle Algorithms, and Quantum Security of a Split-Key PRF". In: *Theory of Cryptography – 20th International Conference, TCC 2022, Chicago, IL, USA, November 7–10, 2022, Proceedings, Part I*. Vol. 13747. Lecture Notes in Computer Science. Springer, 2022, pp. 33–51. DOI: 10.1007/978-3-031-22318-1_2.

[Bha+24] Ritam Bhaumik, Benoît Cogliati, Jordan Ethan, and Ashwin Jha. "Mind the Bad Norms: Revisiting Compressed Oracle-Based Quantum Indistinguishability Proofs". In: *Advances in Cryptology – ASIACRYPT 2024*. Springer, 2024, pp. 215–247. DOI: 10.1007/978-981-96-0947-5_8.

[CPZ24] Joseph Carolan, Alexander Poremba, and Mark Zhandry. *(Quantum) Indifferentiability and Pre-Computation*. Cryptology ePrint Archive, Paper 2024/1727. 2024. URL: https://eprint.iacr.org/2024/1727.

[FIPS 205] NIST. *Stateless Hash-Based Digital Signature Standard*. Tech. rep. FIPS PUB 205. National Institute of Standards and Technology, 2024. DOI: 10.6028/NIST.FIPS.205. URL: https://doi.org/10.6028/NIST.FIPS.205.

[Ala+25] Gorjan Alagic, Joseph Carolan, Christian Majenz, and Saliha Tokat. *The Sponge is Quantum Indifferentiable*. 2025. arXiv: 2504.16887. URL: https://arxiv.org/abs/2504.16887.

[Zha25] Mark Zhandry. *How to Model Unitary Oracles*. Cryptology ePrint Archive, Paper 2025/1072. 2025. URL: https://eprint.iacr.org/2025/1072.

[Hos26] Akinori Hosoyamada. *The Impossibility of Post-Quantum Public Indifferentiability for Merkle–Damgård*. Cryptology ePrint Archive, Paper 2026/128. 2026. URL: https://eprint.iacr.org/2026/128.

A Deferred Proofs

A.1 Proof of Lemma 2.4

Proof. Let the inner product of the two pure states be expressed in terms of its real and imaginary parts as $\langle\psi|\phi\rangle = a + ib$, where $a = \mathrm{Re}(\langle\psi|\phi\rangle)$ and $b = \mathrm{Im}(\langle\psi|\phi\rangle)$. By expanding the square of the Euclidean distance, we have:

$$\mathsf{d}_e^2(|\psi\rangle, |\phi\rangle) = \||\psi\rangle - |\phi\rangle\|_2^2 = \langle\psi|\psi\rangle + \langle\phi|\phi\rangle - 2\mathrm{Re}(\langle\psi|\phi\rangle) = 2 - 2a.$$

Thus, $a = 1 - \frac{1}{2}\mathsf{d}_e^2(|\psi\rangle, |\phi\rangle)$. On the other hand, the square of the trace distance for pure states satisfies:

$$\mathsf{d}_t^2(|\psi\rangle\langle\psi|, |\phi\rangle\langle\phi|) = 1 - |\langle\psi|\phi\rangle|^2 = 1 - (a^2 + b^2).$$

Substituting a into the equation gives:

$$\mathsf{d}_t^2(|\psi\rangle\langle\psi|, |\phi\rangle\langle\phi|) = 1 - \left(1 - \frac{1}{2}\mathsf{d}_e^2(|\psi\rangle, |\phi\rangle)\right)^2 - b^2.$$

Expanding the square, we get:

$$\mathsf{d}_t^2(|\psi\rangle\langle\psi|, |\phi\rangle\langle\phi|) = \mathsf{d}_e^2(|\psi\rangle, |\phi\rangle) - \frac{\mathsf{d}_e^4(|\psi\rangle, |\phi\rangle)}{4} - b^2.$$

By dropping the non-negative term $b^2 = (\mathrm{Im}(\langle\psi|\phi\rangle))^2 \geq 0$, we obtain:

$$\mathsf{d}_t^2(|\psi\rangle\langle\psi|, |\phi\rangle\langle\phi|) \leq \mathsf{d}_e^2(|\psi\rangle, |\phi\rangle) - \frac{\mathsf{d}_e^4(|\psi\rangle, |\phi\rangle)}{4}.$$

Taking the square root on both sides gives:

$$\mathsf{d}_t(|\psi\rangle\langle\psi|, |\phi\rangle\langle\phi|) \leq \mathsf{d}_e(|\psi\rangle, |\phi\rangle)\sqrt{1 - \frac{\mathsf{d}_e^2(|\psi\rangle, |\phi\rangle)}{4}}.$$

Since $\sqrt{1 - \frac{\mathsf{d}_e^2(|\psi\rangle, |\phi\rangle)}{4}} \leq 1$, we further obtain:

$$\mathsf{d}_t(|\psi\rangle\langle\psi|, |\phi\rangle\langle\phi|) \leq \mathsf{d}_e(|\psi\rangle, |\phi\rangle).$$

This completes the proof. □

A.2 Proof of Lemma 2.14

Proof. We construct an efficient quantum simulator S equipped with quantum oracle access to the global random oracle H. Because the primitive f is defined exclusively over the range of pf, any valid query sequence X submitted to the simulator inherently possesses a unique valid preimage. Upon receiving a superposition query on a sequence X, the simulator directly applies the efficient inverse algorithm pf^{-1} to coherently decode the input register and recover the original message $M = \mathsf{pf}^{-1}(X)$. The simulator then explicitly routes the query to $H(M)$ and returns the exact result. In the real world, the construction output is evaluated as $C^f(M) = f(\mathsf{pf}(M))$. In the simulated ideal world, the simulated primitive queries perfectly mirror the true oracle evaluations $H(M)$. Because pf is globally injective, there are no structural collisions, ensuring that the joint distributions of both interfaces (C^f, f) and (H, S^H) are perfectly identical. For every primitive query, the simulator evaluates the external oracle H exactly once, establishing the parameter $\ell_s = 1$. Consequently, the distinguishing advantage is unconditionally zero, yielding $\epsilon = 0$. This completes the proof. □

A.3 Proof of Lemma 3.1

Proof. We compare the six quantum processes step by step.

$1 \to 2$: The measurement on the register $\mathcal{F}$ is postponed indefinitely.

$2 \to 3$: The acting space is extended to a larger Hilbert space, and this extension does not affect the quantum evolution on the registers measured by $\mathcal{M}$.

$3 \to 4$: A change of computational basis is applied to the entire evolution. The basis-change operations commute through the adversary's internal operations and cancel pairwise. The final basis restoration is postponed until after the measurement $\mathcal{M}$.

$4 \to 5$: The redundant operations are removed using commutativity and pairwise cancellation.

$5 \to 6$: This transition is the same as the transition from Process 3 to Process 4. A global change of basis is applied, the basis-change operations cancel through commutativity, and the final basis restoration is postponed until after the measurement.

Therefore, each adjacent pair of processes induces the same probability distribution under $\mathcal{M}$. Hence, the probability distributions induced by $\mathcal{M}$ at the conclusion of all six quantum processes are perfectly identical. □

A.4 Proof of Lemma 4.2

Proof. Let $\mathrm{indeg}(v)$ denote the in-degree of a vertex $v = (z, m, y) \in V$. We call a vertex with chaining input $z = IV$ an initial node. Due to the "No hitting IV" condition, no record outputs IV. Hence, every initial node has in-degree zero.

By the "No collisions" condition, for any chaining input z' of a vertex v', there exists at most one record $v \in D_f$ whose output satisfies $y = z'$. Therefore, every non-initial node has in-degree at most one.

It follows that the connected subgraph containing an initial node is a directed tree rooted at that initial node. In particular, this initial node is the unique root of the tree. Consequently, for any path-reachable node v in this tree, the path from the root to v is unique. This unique path determines a unique message block sequence $M = (m_1, m_2, \ldots, m_k)$.

Suppose that the message string induced by a path-reachable node lies in $\mathsf{Range}(\mathsf{pf})$. We then call this node the terminal node of a complete path. We claim that such a terminal node cannot be used to form any other complete path.

Indeed, if this terminal node were used to form another complete path, then the unique path ending at this node would be a strict prefix of the path forming another complete message. Equivalently, the corresponding encoded message string would be a strict prefix of another encoded message string in $\mathsf{Range}(\mathsf{pf})$. This contradicts the prefix-free property of pf. Therefore, a terminal node of a complete path cannot be used to form any other complete path. This completes the proof. □