

Tight Characterizations for Preprocessing against Cryptographic Salting

Fangqi Dong*

Qipeng Liu[†]

Kewen Wu[‡]

Abstract

Cryptography often considers the strongest yet plausible attacks in the real world. Preprocessing (a.k.a. non-uniform attack) plays an important role in both theory and practice: an efficient online attacker can take advantage of advice prepared by a time-consuming preprocessing stage.

Salting is a heuristic strategy to counter preprocessing attacks by feeding a small amount of randomness to the cryptographic primitive. We present general and tight characterizations of preprocessing against cryptographic salting, with upper bounds matching the advantages of the most intuitive attack. Our result quantitatively strengthens the previous work by Coretti, Dodis, Guo, and Steinberger (EUROCRYPT'18). Our proof exploits a novel connection between the non-uniform security of salted games and direct product theorems for memoryless algorithms.

For quantum adversaries, we give similar characterizations for property finding games, resolving an open problem of the quantum non-uniform security of salted collision resistant hash by Chung, Guo, Liu, and Qian (FOCS'20). Our proof extends the compressed oracle framework of Zhandry (CRYPTO'19) to prove quantum strong direct product theorems for property finding games in the average-case hardness.

*IIS, Tsinghua University. Email: dongfangqi77@gmail.com.

[†]University of California, San Diego. Email: qipengliu0@gmail.com.

[‡]University of California, Berkeley. Email: shlw_kevin@hotmail.com. Supported by a Sloan Research Fellowship and NSF CAREER Award CCF-2145474.

1 Introduction

In modern cryptography, the notion of *security* serves as the cornerstone for evaluating the robustness of cryptographic objects against various adversarial attacks. In most constructions, the security boils down to the specific hardness of cryptographic primitives. For instance, the hardness of hash functions refers to the complexity of finding hash collisions; the security of one-way functions hinges on the intractability of finding pre-images of designated outputs.

Traditionally under the lens of uniform security, the adversary’s capability is assumed to be fixed and oblivious to the actual primitive. However, as cryptographic systems evolve to tackle increasingly sophisticated adversaries, the concept of *non-uniform security* emerges as a critical consideration [Yao90, Unr07, CDGS18], which departs from uniform security by allowing the adversary to adaptively choose its attack strategies based on pre-computed knowledge or advice on the specific cryptographic primitive in question. This departure reflects a more realistic model of real-world attacks, where the adversary performs extensive interactions with the given primitive with significant amount of time and computational power in an offline phase before actually trying to break the protocol in the online phase. Given this offline preprocessing stage, adversaries can sometimes perform much better than their uniform counterparts. For example, finding collisions in a vanilla hash function becomes trivial if the adversary reads through its truth table and store one pair of collisions offline.

Cryptographic *salting* is a fundamental and universal technique to address the above issue. By including a random value, known as a *salt*, from the salt space $[K]$ to the primitive, the output becomes unique even if the input data is the same. Take collision resistant hash as an example. In the salted version, a random oracle $H: [K] \times [M] \rightarrow [N]$ is provided. The goal of the adversary is, given a uniformly random salt k , to find a pair of distinct inputs $x, x' \in [M]$ such that $H(k, x) = H(k, x')$. Intuitively, the randomness of the salt k prevents adversaries from effectively using pre-computed tables unless they store look-up tables for each possible salt value, vastly increasing the computational resources required to mount a successful attack.

The “salting” approach, dating back to a work by Morris and Thompson [MT79], guarantees heuristic non-uniform security. The provable consequences of salting were first investigated by De, Trevisan, and Tulsiani [DTT10]. Later, Chung, Lin, Mahmoody, and Pass [CLMP13] studied salting for collision resistant hash, one of the most important cryptographic applications. Then Mahmoody and Mohammed [MM16] used salting to obtain non-uniform black-box separation results and Dodis, Guo, and Katz [DGK17] established tight bounds for some specific salted cryptographic applications.

More recently, Coretti, Dodis, Guo, and Steinberger [CDGS18] proved that salting generically defeats preprocessing. Let G be a cryptographic game in the random oracle model (ROM) and $\varepsilon_G(T)$ be the optimal winning probability that a T -query algorithm can obtain. Let G_K be the salted game with the salt space $[K]$ and $\varepsilon_{G_K}(S, T)$ be the maximum success probability that a non-uniform algorithm with S -bit advice and T queries can achieve. Then they showed that $\varepsilon_{G_K}(S, T)$ can be bounded by (a multiple of) $\varepsilon_G(T)$ with an additional non-uniform advantage term decaying as the salt space $[K]$ enlarges.

Theorem 1.1 ([CDGS18, Corollary 18 and 19]). *For any cryptographic games G in the ROM, we have*

- **Multiplicative version.** $\varepsilon_{G_K}(S, T) \leq \tilde{O}(\varepsilon_G(T) + ST/K)$.
- **Additive version.** $\varepsilon_{G_K}(S, T) \leq \varepsilon_G(T) + \tilde{O}\left(\sqrt{ST/K}\right)$.

Here we use \tilde{O} to hide low order terms for simplicity.

When K is sufficiently large, the above theorem suggests that non-uniform security of any game is about the same as its uniform security. However quantitatively, their lower bounds does not match the most intuitive algorithm. For instance, when G is the game for collision finding, we have $\varepsilon_G(T) = T^2/N$ by birthday paradox.¹ The optimal non-uniform algorithm should be the following: store collisions for S distinct salts; then in the online stage, it either outputs a stored collision for the challenge salt, or executes a birthday-paradox attack. The algorithm achieves advantage $T^2/N + S/K$, but [CDGS18] only provides an upper bound of $T^2/N + ST/K$.

Furthermore, when a security game with challenges (unlike collision finding) is considered, [Theorem 1.1](#) is even looser. Dodis, Guo, and Katz [DGK17] showed that for the function / permutation inversion problem, the optimal bound is $T/N + (ST)/(KN)$, whereas [Theorem 1.1](#) only offers $T/N + S/K$.

Beyond classical adversaries, the growing power of quantum computing poses significant threats to modern cryptography. Chung, Guo, Liu, and Qian [CGLQ20] and Liu [Liu23] studied how salting could prevent quantum non-uniform attacks and provides bounds similar to [Theorem 1.1](#). However, even for the simplest yet most important cryptographic application — collision finding in the quantum random oracle model (QROM) — they were only able to give an upper bound $T^3/N + ST/K$, whereas the best known attack only achieves $T^3/N + S/K$, where T^3/N comes from the quantum collision finding algorithm [BHT97] and S/K is the same advantage by storing collisions for S salts.

Therefore, in this article, we ask:

*Can we provide tighter characterizations for preprocessing against salting,
in both the classical and quantum world?*

1.1 Our Results

Classical Results. Let G, G_K, ε_G , and ε_{G_K} be defined above. We present our first result.

Theorem 1.2 (Consequence of [Theorem 4.1](#)). *For any cryptographic game G in any idealized model, we have*

- **Multiplicative version.** $\varepsilon_{G_K}(S, T) \leq 2 \cdot \varepsilon_G(T) + 2S/K$.
- **Additive version.** $\varepsilon_{G_K}(S, T) \leq \varepsilon_G(T) + 4\sqrt{S/K}$.

For applications, the multiplicative version of [Theorem 1.2](#) is suited for games with small advantages (i.e., when $\varepsilon_G(T) \ll 1$) and the additive version of [Theorem 1.2](#) should be used for games with noticeable advantages (e.g., $\varepsilon_G(T) \geq 1/2$ when G is a decision game). We also remark that while the addend in the additive version of [Theorem 1.2](#) incurs a square-root-type loss compared with the multiplicative version, similar term appears and is shown to be tight in related works on non-uniform security of cryptographic primitives (see e.g., [GGKL21]).

Our [Theorem 1.2](#) improves [CDGS18] (see [Theorem 1.1](#)) in all aspects:

¹We omit low-order terms in the introduction for conciseness.

- First, our result shaves the extra T factor in front of each S/K in [Theorem 1.1](#). This improvement leads to a sharp characterization for games without challenges (like collision finding or zero-preimage-finding): with S bits of advice, one can store answers for $\approx S$ salts and use T queries to solve challenges for other salts.
- Second, our result works in any idealized model (ROM, the random permutation model, the ideal cipher model, the generic group model, or any distribution of oracles) with any query type (see [Remark 4.2](#) for detail), while [Theorem 1.1](#) only works for ROM.
- Finally, our bound is much cleaner and simpler, providing explicit control of the low order terms. Our bound is even capable of handling small constant regimes of S, T, K and can be useful in practice.

When games with challenges are considered (e.g., salted function inversion), [Theorem 1.2](#) may no longer be tight. Intuitively, different challenges have completely different answers and a piece of S -bit advice cannot guarantee correctness on S salts. We develop a general framework ([Theorem 4.15](#)) for understanding salting for these kinds of games. To demonstrate its versatility and strength, we provide two non-exhaustive examples: we reprove the bound for salted function inversion as in [\[DGK17\]](#) and provide a tight characterization of salting when the advice is large.

Corollary 1.3 ([Corollary 4.20](#) Restated, Salted Function Inversion). *Let Inv be the game, where a challenge is a uniformly sampled $y \in [N]$ and the goal is to find $x \in [M]$ in a random function such that x is a pre-image of y . Let Inv_K be the salted game. Then we have*

$$\varepsilon_{\text{Inv}_K}(S, T) \leq O\left(\frac{T}{N} + \frac{ST}{KN}\right).$$

Corollary 1.4 (Consequence of [Corollary 4.17](#), Large Advice). *When $S = \Omega(K)$, for any cryptographic game G in any idealized model, it holds that*

$$\varepsilon_{G_K}(S, T) \leq O\left(\max_{\substack{n_1, \dots, n_K \in \mathbb{Z}^+ \\ n_1 + \dots + n_K = S}} \sum_{i=1}^K \frac{\delta_G(n_i, T)}{K}\right),$$

where $\delta_G(n, T)$ is the n -th root of the uniform security of the vanilla game G equipped with n independent challenges using nT queries.

It turns out that $\delta_G(n, T)$ is an upper (and usually tight) bound on the non-uniform security of the game G with n -bit advice and T queries [\[Unr07, IK10, CDGS18, CGLQ20, GLLZ21\]](#). In these cases, our [Corollary 1.4](#) is also tight and offers an intuitive strategy to achieve the optimality for large $S = n_1 + \dots + n_K$: reserve n_k bits for each salt k as the best n_k -bit advice. Then the algorithm uses the corresponding advice in the online challenge phase, achieving success probability $(\delta(n_1, T) + \dots + \delta(n_K, T))/K$.

Quantum Results. Next, we turn our attention to the quantum case. [\[CGLQ20, Liu23\]](#) already studied general bounds for non-uniform quantum attacks against salting. They similarly have an extra factor of T in front of S/K , making the bound non-tight even for collision resistant hash.

We start by aiming for a general quantum result as [Theorem 1.2](#), which successfully shaves the additional factor T . However, in the proof of [Theorem 1.2](#), we require an average-case direct

product theorem (DPT) for general games. While a worst-case quantum DPT is known [She11, LR13], an average-case quantum DPT is needed to analyze the idealized model, which turns out to be a long-standing open problem in quantum query complexity.

Therefore, we resort to the compressed oracle technique by Zhandry [Zha19], a powerful framework for proving average-case complexity in the quantum random oracle model (QROM). We show optimal quantum bounds for property finding games, which includes function inversion, collision finding, k -SUM, and more.

Theorem 1.5 (Consequence of [Theorem 5.13](#)). *Let CRHF be the game, where the goal of a quantum-query algorithm is to find $x \neq x' \in [M]$ such that they have the same image under the random oracle. Its quantum non-uniform security $\varepsilon_{\text{CRHF}_K}^*$ (with quantum advice) is at most*

$$\varepsilon_{\text{CRHF}_K}^*(S, T) \leq \tilde{O}\left(\frac{T^3}{N} + \frac{S}{K}\right).$$

At the center of our quantum results is an quantum strong DPT (in fact, we prove a threshold direct product theorem that is even stronger than direct product theorems) in Zhandry’s compressed oracle framework (see [Theorem 5.19](#)). To the best of our knowledge, this is the first average-case DPT for quantum query algorithms, which may be of independent interests.

Organization. We give an overview of our techniques in [Section 2](#). Formal definitions are listed in [Section 3](#). Then [Section 4](#) contains the analysis of cryptographic salting in the classical case and [Section 5](#) is for the quantum random oracle case. Missing proofs can be found in Appendix.

2 Technique Overview

We give a high-level overview of our proof techniques.

2.1 The Classical Setting

Let \mathcal{A} be a classical algorithm with T queries and S bits of advice. Upon receiving a uniformly random salt $k \sim [K]$, \mathcal{A} aims to solve the k -th independent copy of the game G in the salted game G_K . Let ε be the winning probability of \mathcal{A} and our goal is to give an upper bound on ε .

Removing Advice. To eliminate the effect of the non-uniform advice, we consider executing the above procedure multiple times (say, L times). The key idea here is that, the non-uniform advice is given in advance of the random salt, therefore its power is intuitively diluted across the salt space $[K]$ [IK10, Imp11].

Let k_1, \dots, k_L be the L salts chosen. Let f_1, \dots, f_K be the oracles of each salt. For any fixed f_1, \dots, f_K , the advice is also fixed. Thus by the independence of k_1, \dots, k_L , the probability of \mathcal{A} winning on all these L salts conditioned on f_1, \dots, f_K is simply

$$\left(\mathbb{E}_k[\varepsilon_k]\right)^L,$$

where ε_k is the winning probability of \mathcal{A} on salt k . Considering the randomness of f_1, \dots, f_K , we also have

$$\varepsilon = \mathbb{E}_{f_1, \dots, f_K} \left[\mathbb{E}_k [\varepsilon_k] \right].$$

Hence by convexity, it suffices to bound the probability δ of \mathcal{A} winning L salts simultaneously and then $\varepsilon \leq \delta^{1/L}$.

Intuitively, δ should be of order $o(1)^L$ since the definition of δ involves solving L independent trials, where $o(1)$ is the advantage of solving one. When $L \gg S$, it does not hurt if we replace the non-uniform S -bit advice by a random S -bit string, since the latter has a probability of $2^{-S} \gg 2^{-L}$ to be equal to the former.

Formally, let η be the probability of \mathcal{A} winning L salts with a uniformly random S -bit string as advice (which can in turn be hardwired oblivious to f_1, \dots, f_K). Then we have $\delta \leq 2^S \cdot \eta$ and, more importantly, since the oracle-dependent advice is removed, we only need to analyze uniform algorithms for η . This corresponds to (2) in the proof of [Theorem 4.1](#).

Reducing to Direct Product Theorems. Let \mathcal{B} be the T -query algorithm corresponding to \mathcal{A} with hardwired advice. Then η is equivalent to the success probability of the following procedure:

1. Sample salts $k_1, \dots, k_L \sim [K]$.
2. Check whether \mathcal{B} solves the k_i -th independent copy of G for each $i \in [L]$.

Let U be the set of distinct salts in k_1, \dots, k_L . Then conditioned on [Item 1](#), U is fixed and the success of [Item 2](#) implies on the success of \mathcal{B} solving independent copies of G indexed by U . Though \mathcal{B} has access to all f_1, \dots, f_K even when just solving salt k , intuitively only f_k should be relevant. This corresponds to [Claim 4.4](#).

The above desired phenomenon is exactly characterized by *direct product theorems*: assume our best advantage is γ for a single instance (i.e., the game G here) with T budget (i.e., total amount of queries here). Then solving u independent copies (i.e., $u = |U|$ here) of this instance with uT budget should be γ^u , which is simply applying the optimal strategy for each copy independently.

Since γ^u is attainable, such a direct product theorem is the best possible. In many cases, such a strong result does not hold and the exact decay bound remains elusive for many important applications (see e.g., [\[Raz10\]](#)). The study on direct product theorem is rich on its own with numerous works in circuit complexity (e.g., [\[Yao82, Lev85, IJKW08\]](#)), query complexity (e.g., [\[Sha03, Dru12, She11, LR13\]](#)), communication complexity (e.g., [\[JK21, JK22\]](#)), property testing and coding theory (e.g., [\[DDG⁺15, DS14\]](#)), and more.

Our setting (the classical query model) unfortunately prohibits a general strong direct product theorem [\[Sha03\]](#), and the best known bound incurs an inevitable polynomial loss [\[Dru12\]](#) dependent on γ . Nevertheless, we are able to exploit the memoryless structure of the above procedure to obtain a strong direct product theorem through a similar analysis as presented in [\[NRS98\]](#). We explain it after we complete the analysis of the non-uniform security ε of G_K against T -query algorithms with S -bit advice:

- By considering L independent salts, we have $\varepsilon \leq \delta^{1/L}$.
- By guessing and removing the advice, we have $\delta \leq 2^S \cdot \eta$.
- By strong direct product theorems, we have

$$\eta \leq \mathbb{E}_U [\gamma^{|U|}] \leq \left(\gamma + \frac{L}{K} \right)^L,$$

where γ is the uniform security of G against T -query algorithms and $|U|$ is the number of distinct values in a draw-with-replacement experiment with L draws from the universe $[K]$. The second inequality follows from standard martingale analysis (see [Lemma 3.8](#)).

Rearranging and optimizing the choice of L gives the desired bound (see [Theorem 4.1](#) for detail).

Strong Direct Product Theorems for Memoryless Algorithms. Before presenting our argument, we remark that our analysis coincides with a much earlier work by Nisan, Rudich, and Saks [[NRS98](#)] under a different background. While we only became aware of their work after the completion of our paper, we do *not* consider our analysis for this part as our technical contribution, and we keep our argument within our setting just for completeness.

For simplicity, consider the following scenario and the general treatment can be found in [Section 4.1](#): let x^1, \dots, x^n be n i.i.d. random strings given oracle access to, and we execute a T -query algorithm \mathcal{C} sequentially for $1, 2, \dots, n$ to predict $h(x^1), \dots, h(x^n)$. To compare with the analysis above, $n = |U|$ is the number of distinct salts and x^1, \dots, x^n are the corresponding oracles, \mathcal{C} is essentially \mathcal{B} , h is the predicate for the game (e.g., $h(x)$ is the pre-image of zero in x), and γ is the optimal success probability of a T -query algorithm in predicting an *individual* $h(x^i)$. Our goal is to prove that the above joint prediction succeeds with probability at most γ^n .

The literature of direct product theorems usually assumes the joint prediction is provided by a much more powerful query algorithm, one single algorithm of nT queries that makes all the predictions in the end [[Dru12](#), [She11](#), [LR13](#)]. While intuitively it makes no sense for the algorithm to query outside x^i to predict $h(x^i)$, it can actually distribute the total nT queries in a more clever and adaptive way that outperforms the naive strategy to predict $h(x^i)$ individually with T queries on x^i . As a consequence, a strong bound of γ^n turns out to be false [[Sha03](#)] and the best known bound has an inevitable growing loss dependent on γ itself [[Dru12](#)], making it insufficient for our purposes.

The key belief that drives us towards a bound of γ^n in our setting comes from the fact that \mathcal{C} is executed separately to predict each $h(x^i)$ and it does not share memory between separate prediction tasks. In other words, the above scenario has a memoryless structure that the workspace is cleaned repeatedly after each T queries, which intuitively forbids any contrived attacks requiring coordination between predictions of \mathcal{C} . Note that while this intuition checks out, in reality \mathcal{C} can still query, say x^j , outside x^i to predict $h(x^i)$, which makes the predictions $h(x^i), h(x^j)$ correlated and perhaps magically they can be correct simultaneously with higher probability. In fact, examples in this spirit date back to Feige's famous counterexample for the sharp parallel repetition theorem [[Fei91](#)].

To rigorously verify our belief, we start with the work by Shaltiel [[Sha03](#)], which proves the tight γ^n bound for *fair* algorithms. A fair algorithm is allowed to make all predictions in the end, potentially dependent on all the nT queries, however, it is guaranteed that each x^i can only be queried at most T times. Since the queries can still be intertwined among x^i 's, the predictions are, again, correlated. Nevertheless, since there is an upper bound on the queries for each x^i , a simple induction argument can be carried out. In a bit more detail, one can do induction on the n -tuple (T_1, \dots, T_n) where each T_i is the upper bound of the total number of queries on x^i . Then one can show that the optimal strategy for a fair algorithm is spending T_i queries on x^i to predict $h(x^i)$ independently.

Given the sharp bound for fair algorithms ([Theorem 4.9](#)), our goal is to reduce our memoryless algorithms to fair algorithms ([Lemma 4.10](#)). Note that while the memory is cleaned after each

T queries, the whole execution can make $\gg T$ queries on x^i . On the other hand, since a fair algorithm does not have to be memoryless, these two types of algorithms are incomparable to each other. To get a fair algorithm from a memoryless algorithm, we further exploit the memoryless structure: if some $h(x^i)$ is already predicted, then we can privately simulate the future queries to x^i without actually querying x^i , as \mathcal{C} is separately executed for each prediction and these executions commute.

By this observation, we perform the following out-of-order simulation: let $\mathcal{C}_1, \dots, \mathcal{C}_n$ be \mathcal{C} executed to predict $h(x^1), \dots, h(x^n)$ respectively. We initialize $\mathcal{C}_1, \dots, \mathcal{C}_n$ and assume without loss of generality they are deterministic. While there is some prediction not ready, we use a cycle / path elimination procedure to progress: for each i , if \mathcal{C}_i has not terminated and wants to query x^j , then we add a directed edge from i to j . In this directed graph,

- Either we have a directed cycle $i_1 \rightarrow i_2 \rightarrow \dots \rightarrow i_\ell \rightarrow i_1$.
Then we advance $\mathcal{C}_{i_1}, \dots, \mathcal{C}_{i_\ell}$ by one query to $x^{i_2}, \dots, x^{i_\ell}, x^{i_1}$ respectively.
- Or we have a directed path $i_1 \rightarrow i_2 \rightarrow \dots \rightarrow i_\ell$.
Then we advance $\mathcal{C}_{i_1}, \dots, \mathcal{C}_{i_{\ell-2}}$ by one query to $x^{i_2}, \dots, x^{i_{\ell-1}}$ respectively. In addition, we privately simulate the query of $\mathcal{C}_{i_{\ell-1}}$ on x^{i_ℓ} without actually querying x^{i_ℓ} . This is valid since i_ℓ is the endpoint of the path, thus \mathcal{C}_{i_ℓ} is already terminated and the prediction $h(x^{i_\ell})$ is already generated.

By the above procedure, each x^i is queried at most T times, since when \mathcal{C}_i queries, either x^i gets a query from the predecessor in the cycle / path, or the query is privately simulated. Therefore this out-of-order execution is a fair algorithm and we can apply the known strong direct product theorem for fair algorithms.

A Different Analysis. Finally we sketch a different analysis (see detail in [Section 4.2](#)) that improves the previous analysis for games with numerous challenges (e.g., [Corollary 1.3](#)) or large advice ([Corollary 1.4](#)). Recall that in the reduction to direct product theorems, we focus on U , the set of distinct salts in L independent uniform salts k_1, \dots, k_L . This is tight if the challenge distribution of G has low entropy. For example, if the challenge is null (e.g., collision finding requires no challenge description), then winning salts k_1, \dots, k_L is equivalent to winning salts in U . However if the challenge distribution has high entropy (e.g., the function inversion asks to invert a random challenge point), then winning U can be quite off from winning all the salts, since with high probability we get different challenges even if the salts are the same.

To avoid such a potential loss, we instead look directly into k_1, \dots, k_L . For each salt $k \in [K]$, let n_k be the number of its appearances in k_1, \dots, k_L . Then n_1, \dots, n_K are determined by k_1, \dots, k_L . Now for fixed k_1, \dots, k_L , the algorithm \mathcal{B} needs to solve n_k independent challenges of salt k for each $k \in [K]$. Let G^{n_k} be the multi-challenge version of the game G where n_k i.i.d. challenges are asked simultaneously. Then the algorithm aims to win all of them in a similar memoryless fashion: it can use $n_k T$ queries to solve G^{n_k} separately for each $k \in [K]$. Though the games G^{n_1}, \dots, G^{n_K} are not identical, strong direct product theorems for memoryless algorithms can be proved with an almost identical argument. Therefore, we can relate the non-uniform security of salting to the advantages of the multi-challenge version of G (see [Corollary 1.4](#)).

We remark that the multi-challenge version of G is closely related to the non-uniform security of G itself (without salting), and (tight) upper bounds have been established for various games (see e.g., [[Unr07](#), [CDGS18](#), [CGLQ20](#), [GLLZ21](#)]).

2.2 The Quantum Random Oracle Setting

We now move to the quantum case. The aforementioned reduction in the classical setting works in the quantum setting as long as the non-uniform quantum algorithm is given *classical advice*. In the described reduction, we construct a non-uniform algorithm solving multiple instances of a game, by iteratively executing an algorithm with fixed advice many times. Such a reduction falls short when the advice is quantum, as outputting any answer would destroy the advice in an irreversible way.

Quantum Bit-Fixing. To address the quantum non-cloning issue, we resort to the quantum bit-fixing framework, which was first introduced by Guo, Li, Liu, and Zhang [GLLZ21] and later refined by Liu [Liu23]. In their work, they defined the so-called quantum bit-fixing model (QBF) and proved a novel connection between security in the QBF model and non-uniform security in the quantum random oracle model (QROM). Before introducing the QBF model, let us first grasp the intuition behind it (see Section 5.6 for more details).

Liu [Liu23] considered a quantum-friendly adaptation of direct product games, where the quantum advice can be effectively reused to win “multiple instances” of the original games (formally, the alternating measurement game). This is achieved by a similar idea of the witness preserving amplification of QMA [MW05], to make sure that the quantum advice is preserved throughout the entire game. The S -fold alternating measurement game has S rounds, with each round allowing the algorithm to make T queries. An algorithm wins if and only if it succeeds all rounds in a sequential order, much like the memoryless algorithms mentioned earlier. Since it is not the primary focus of this work, we refer interested readers to [Liu23] for more details of alternating measurement games. Eventually, if any uniform quantum-query algorithm (conducting T queries per round) can win the S -fold alternating measurement game with probability at most η^S , the quantum non-uniform security $\varepsilon_G^*(S, T)$ of the original game is at most η .

Now the goal is to bound the success probability of the S -fold game. For each $i \in [S]$, let \mathbf{W}_i be the event that the algorithm wins the i -th round. We aim to show that for some η , the following holds:

$$\Pr[\mathbf{W}_1 \wedge \mathbf{W}_2 \wedge \cdots \wedge \mathbf{W}_S] < \eta^S.$$

Define $\mathbf{W}_{<i}$ be the event that the algorithm wins the first $(i - 1)$ rounds. Then we have

$$\Pr[\mathbf{W}_1 \wedge \mathbf{W}_2 \wedge \cdots \wedge \mathbf{W}_S] = \prod_{i=1}^S \Pr[\mathbf{W}_i \mid \mathbf{W}_{<i}].$$

Since [Liu23] also showed that $\Pr[\mathbf{W}_i \mid \mathbf{W}_{<i}]$ is monotonically non-decreasing in i , if we can prove $\Pr[\mathbf{W}_S \mid \mathbf{W}_{<S}] < \eta$, it will imply $\Pr[\mathbf{W}_1 \wedge \cdots \wedge \mathbf{W}_S] < \eta^S$. Even further, we can assume $\Pr[\mathbf{W}_{<S}]$ is typical, say, $\Pr[\mathbf{W}_{<S}] \geq (1/N)^S$.² Since otherwise $\Pr[\mathbf{W}_1 \wedge \cdots \wedge \mathbf{W}_S] \leq \Pr[\mathbf{W}_{<S}] < (1/N)^S \leq \eta^S$ is automatically satisfied as the target η is usually $\gg 1/N$.

The security in the QBF model is defined as the maximum of $\Pr[\mathbf{W}_S \mid \mathbf{W}]$ over any T -query strategy for the S -th game and any event \mathbf{W} that can be achieved by an (ST) -quantum-query algorithm with a “non-negligible” probability (e.g., $(1/N)^S$ as above). In the salted case, \mathbf{W}_S itself

²This lower bound $(1/N)^S$ is arbitrarily chosen in this section for clarity in presentation.

precisely represents the event that a uniform T -quantum-query algorithm wins the game G on a random salt.

Two immediate questions arise. First, even without conditioning on \mathbf{W} , bounding the probability of \mathbf{W}_S necessitates tools from average-case quantum query complexity. Moreover, it becomes even more challenging when conditioning is considered.

In this study, we focus solely on the game for which Zhandry’s compressed oracle technique [Zha19] serves as a potent tool for addressing average-case query complexity. We will first provide an exposition of this technique and then introduce the concepts required to handle conditioning.

An Exposition of the Compressed Oracle Technique. In the rest of the overview, we will focus on the collision finding problem. Let us begin with the classical case. A T -classical-query algorithm can acquire information on at most T input-output pairs of a random function. By the principle of deferred decision, one can assume that outputs are only sampled for those queried inputs, and every other output remains uniformly random. This is commonly referred to as “lazy sampling”. Thus, we can say a random oracle is initialized as $D = \emptyset$ (an empty database), meaning every output is uniformly random at the beginning. For every query $x \in [M]$, if there is a pair $(x, y) \in D$, it outputs y ; otherwise it updates $D \leftarrow D \cup \{(x, y)\}$ for a uniformly random y and returns y . The database contains all the information a classical algorithm learns so far. Most importantly, the probability of the algorithm finds a collision is roughly the probability that D contains a collision. Hence, the objective shifts to bounding the probability that, following T queries, D contains a collision.

The seminal work by Zhandry [Zha19] invented a quantum analogue of the classical lazy sampling, called “compressed oracle”. Since quantum queries can be made in superposition, one single classical database is no longer feasible to track all the information. However, the notion of a “database” is still meaningful if the “database” itself is also stored in superposition. A quantum-query algorithm interacting with a quantum-accessible random oracle can be equivalent simulated as follows:

- The database register is initialized as $|\emptyset\rangle$: the algorithm has not yet queried anything.
- When the algorithm makes a quantum query, the database register gets updated in superposition: for a query $|x\rangle$ and a database $|D\rangle$, the simulator looks up x in D . If x is not in D , it initializes an equal superposition $\sum |y\rangle$ (up to normalization), and updates $|D\rangle$ to $\sum_y |D \cup (x, y)\rangle$. It then returns y also in superposition.

The above description provides a high-level idea, albeit slightly inaccurate. Since a quantum algorithm can forget previously acquired information, a formal “compressed oracle” necessitates an additional step to perform these forgetting operations.

The compressed oracle exhibits several favorable properties akin to those enjoyed by classical lazy sampling technique:

1. **Local change.** Every quantum query will only change $|D\rangle$ locally: the addition of a new entry (x, y) , no change, or the removal of an existing entry (corresponding to the ability of forgetting).
2. **Knowledge upper bound.** The probability of a quantum algorithm finding a collision is approximately the same as the probability of measuring the database register and obtaining D with collision.

Hence, for the remainder of the overview, our focus will solely be on detecting collisions within the database register, rather than the output produced by the quantum algorithm.

Probability as Path Integral. Before we approach our goal: the security of salted collision finding in the QBF model (i.e., to upper bound $\Pr[\mathbf{W}_S|\mathbf{W}]$), we first reprove the bound for the vanilla collision finding [Zha19] against uniform quantum query algorithms. We reinterpret the analysis as a form of “path integral”. This alternative view gives a new way to understand [Zha19] and will help us establish a strong (threshold) direct product theorem to upper bound $\Pr[\mathbf{W}_S|\mathbf{W}]$.

For any T -query quantum algorithm, its interaction with a compressed oracle can be written as $\text{cO } U_T \cdots \text{cO } U_2 \text{cO } U_1 |0\rangle \otimes |\emptyset\rangle$, where $|0\rangle$ is the algorithm’s initial state, $|\emptyset\rangle$ is the initial database, each U_i is a local unitary only on the algorithm’s register, and cO is the compressed oracle query. We are interested in the probability of observing a collision in D , i.e., the squared norm of the following (sub-normalized) state:

$$|\phi\rangle := \Lambda \text{cO } U_T \cdots \text{cO } U_2 \text{cO } U_1 |0\rangle \otimes |\emptyset\rangle,$$

where Λ is a projection onto databases with at least one collision.

Feynman’s interpretation of quantum mechanics (path integral), on a high level, postulates that the probability of an event (observing a collision) is given by the squared modulus of a complex number (as $\| |\phi\rangle \|$ above), and this amplitude is given by adding all paths contributing to the complex number. Since our ultimate goal is to observe a collision within the database, the collision must occur as a result of one of the oracle queries cO . Therefore, we contemplate all potential pathways leading to a collision. For any $t \in \{1, 2, \dots, T\}$, we define the following (sub-normalized) state:

$$|\psi_t\rangle := \underbrace{(\Lambda \text{cO}) U_T}_{T\text{-th query}} \underbrace{(\Lambda \text{cO}) U_{T-1}}_{(T-1)\text{-th query}} \cdots \underbrace{(\Lambda \text{cO}) U_t}_{t\text{-th query}} \underbrace{(\mathbb{I} - \Lambda) \text{cO } U_{t-1} \cdots \text{cO } U_2 \text{cO } U_1}_{(t-1)\text{-th query}} |0\rangle \otimes |\emptyset\rangle.$$

The definition of $|\psi_t\rangle$ guarantees that, starting from the t -th query, the database register will always be supported on those databases containing collision. Observe that $|\phi\rangle = |\psi_1\rangle + |\psi_2\rangle + \cdots + |\psi_T\rangle$. Thus $\| |\phi\rangle \| \leq \| |\psi_1\rangle \| + \| |\psi_2\rangle \| + \cdots + \| |\psi_T\rangle \|$. Then by showing $\| |\psi_t\rangle \| \lesssim \sqrt{t/N}$, one deduces $\| |\phi\rangle \|^2 \lesssim T^3/N$.

We can also treat the classical lazy sampling as a path integral, where we similarly define $|\phi\rangle$ (resp., $|\psi_t\rangle$) as the computation at the end (resp., after t queries) projected onto databases with collisions. The only difference is that, all probabilities are summed directly in the classical case: $\| |\phi\rangle \|^2 = \| |\psi_1\rangle \|^2 + \| |\psi_2\rangle \|^2 + \cdots + \| |\psi_T\rangle \|^2 \lesssim T^2/N$. This is because in this path integral for the classical algorithm, the path decomposition $|\psi_1\rangle, \dots, |\psi_T\rangle$ are mutually orthogonal, for which we can use Pythagorean equation to replace the triangle inequality.

Establishing the Security in the QBF Model. Recall that G_K is the salted collision finding problem. The security of G_K in the QBF model requires us to upper bound $\Pr[\mathcal{A} \text{ wins } G_K | \mathbf{W}]$, where \mathcal{A} is a T -quantum-query algorithm and the event \mathbf{W} can be achieved by an ST -quantum-query algorithm with a “non-negligible” probability. We aim to show that the probability above is upper bounded by $O(T^3/N + S/K)$.

What does this conditioning enforce in the context of the compressed oracle? It implies that the database register does not start with $|\emptyset\rangle$, but rather as a superposition of databases. Moreover,

since \mathbf{W} can be achieved by an ST -quantum-query algorithm, according to the “local change” property, the database register will have non-zero support only over databases with at most ST entries. However, this condition alone is not sufficient to derive the desired bound. Consider the following scenario: though every $|D\rangle$ with non-zero support has at most ST entries, it exactly stores collisions for $ST/2$ distinct salts. In such a case, with a random challenge salt, the database with probability ST/K consists of an answer for the salt even if \mathcal{A} makes no further query. Given this example, we further separate the probability:

$$\Pr[\mathcal{A} \text{ wins } G_K \mid \mathbf{W}] \leq \underbrace{\Pr[\mathcal{A} \text{ wins } G_K \mid \mathbf{W} \wedge \mathbf{E}]}_{(i)} + \underbrace{\Pr[\overline{\mathbf{E}} \mid \mathbf{W}]}_{(ii)},$$

where \mathbf{E} denotes the event that the database register contains $|D\rangle$ that stores collisions for at most $\tilde{S} = 2S \log N$ distinct salts. This idea of posing \mathbf{E} was first proposed by Akshima, Guo, and Liu [AGL22] in the context of classical Merkle-Damgård hash. We adapt it to the quantum setting, combining with the compressed oracle technique.

The term (i) follows closely to the ideas in the lower bound of the vanilla collision finding, with the distinction that the database register commences with databases satisfying the requirements outlined by both \mathbf{W} and \mathbf{E} : each database has at most ST entries and collisions for at most \tilde{S} salts. We show that (i) is $\lesssim T^3/N + \tilde{S}/K$.

To bound the term (ii), we show that $\Pr[\overline{\mathbf{E}}] < (T^3/N)^{\tilde{S}} \lesssim (1/N)^{2S}$. Since we assumed $\Pr[\mathbf{W}] \geq (1/N)^S$, we now have $\Pr[\overline{\mathbf{E}} \mid \mathbf{W}] < \Pr[\overline{\mathbf{E}}] / \Pr[\mathbf{W}] < 1/N$. Upon accomplishing this, we establish security in the QBF model, along with the non-uniform security for salted collision finding. Please refer to [Section 5.6](#) for more detail.

Strong (Threshold) Direct Product Theorem in the Compressed Oracle. Now we give more detail for the term (ii). This final step involves bounding $\Pr[\overline{\mathbf{E}}]$. The probability corresponds to the following direct-product-type experiment: there is a random oracle $f : [K] \times [M] \rightarrow [N]$ (simulated as a compressed oracle). An algorithm makes at most ST quantum queries and the objective is to make the database register contain collisions for at least \tilde{S} salts. We let $\Lambda^{\geq r}$ be a projection onto databases with collisions on *at least* r salts and $\Lambda^{=r}$ projects onto database with collisions on *exact* r salts. Then $\Pr[\overline{\mathbf{E}}]$ is the squared norm of the following (un-normalized) state:

$$\begin{aligned} |\phi\rangle &:= \Lambda^{\geq \tilde{S}} \text{cO } U_{ST} \cdots \text{cO } U_2 \text{cO } U_1 |0\rangle \otimes |\emptyset\rangle \\ &= \sum_{r \geq \tilde{S}} \Lambda^{=r} \text{cO } U_{ST} \cdots \text{cO } U_2 \text{cO } U_1 |0\rangle \otimes |\emptyset\rangle. \end{aligned}$$

We list all possible computation paths that lead to databases in $\Lambda^{=r}$ for some $r \geq \tilde{S}$. For every $|D\rangle \in \Lambda^{=r}$, following its evolution, there exists $\mathbf{z} = (z_1, z_2, \dots, z_r)$ such that before the z_i -th query, there are $(i-1)$ salts having a collision that the database will not forget in the future; right after the z_i -th query, there are i salts having a collision that the database will not forget. We denote the computation path by the state $|\psi_{\mathbf{z}}\rangle$. Similar to the vanilla collision finding, for the case of databases in $\Lambda^{=r}$, it holds that

$$\Lambda^{=r} |\phi\rangle = \sum_{\mathbf{z}} |\psi_{\mathbf{z}}\rangle.$$

Finally, we show that for every possible path \mathbf{z} , the state $|\psi_{\mathbf{z}}\rangle$ has a norm $\lesssim (T/N)^{r/2}$. This step necessitates a much more intricate analysis compared to bounding $\|\psi_i\rangle\|$ in the vanilla case. When some salt obtains a collision in the database that will never be forgotten in the future, we must keep track of which salt i it is, as well as how many queries / entries q_i already exist for this particular salt. If we further split $|\psi_{\mathbf{z}}\rangle$ into paths with all possible $(k_1, \dots, k_r), (q_1, \dots, q_r)$ and still use the triangle inequality, we will not be able to obtain the desired bound. One of our novel and key ideas to finalize the proof in the direct product game is to show that all these paths have some form of orthogonality. We can leverage the orthogonality of all these cases³, and show that we can use Pythagorean equation instead of triangle inequality to complete the proof, much like the path integral for classical algorithms.

Since the total number of \mathbf{z} 's is $\binom{ST}{r}$, $\|\Lambda^{=r}|\phi\rangle\|^2 \lesssim \binom{ST}{r}^2 (T/N)^r \lesssim (T^3/N)^r$. By summing over every $r \geq \tilde{S}$, we can conclude that $\|\phi\rangle\|^2 = \sum_{r \geq \tilde{S}} \|\Lambda^{=r}|\phi\rangle\|^2 \lesssim (T^3/N)^r$, from the summation of geometric series for the case of $T^3/N < 1/2$. More details on our strong (threshold) direct product theorem can be found in [Section 5.5](#).

3 Preliminaries

For each positive integer n , we use $[n]$ to denote set $\{1, 2, \dots, n\}$. We use \mathbb{Z} to denote the set of integers, use \mathbb{N} to denote the set of non-negative integers, and use \mathbb{Z}_+ to denote the set of positive integers.

For a distribution μ , we use $x \sim \mu$ to denote a sample x from μ . For a finite set S , we use $x \sim S$ to denote a uniformly random element x from S .

For a complex number $x \in \mathbb{C}$, we denote its norm $|x| = (x\bar{x})^{1/2}$, where \bar{x} is the conjugate of the complex number x . For a complex vector $\mathbf{x} \in \mathbb{C}^n$ or a quantum state $|\phi\rangle = \sum_{i \in [n]} \alpha_i |i\rangle$, we denote the ℓ_2 -norm $\|\mathbf{x}\| = |\mathbf{x}|_2 = (\sum_{i \in [n]} x_i \bar{x}_i)^{1/2}$ and $\|\phi\rangle\| = (\sum_{i \in [n]} \alpha_i \bar{\alpha}_i)^{1/2}$.

Asymptotics. We use the standard $O(\cdot), \Omega(\cdot), \Theta(\cdot)$ notation, and emphasize that in this paper they only hide universal positive constants that do not depend on any parameter.

Games. We start by defining the game in the general oracle model, then instantiate it in various settings.

Definition 3.1 (Game). A game G is described by an oracle distribution μ over functions $f: [M] \rightarrow [N]$,⁴ a challenge distribution π_f specified by each possible oracle f from μ , and an accepting outcome set (aka predicate) $\mathcal{R}_{f,\text{ch}}$ specified by each possible oracle f from μ and each possible challenge from π_f .

We say a query algorithm \mathcal{A} wins the game given oracle access to f and challenge ch if the output of \mathcal{A} is in $\mathcal{R}_{f,\text{ch}}$, i.e., $\mathcal{A}^f(\text{ch}) \in \mathcal{R}_{f,\text{ch}}$. The winning probability of \mathcal{A} is then computed by

$$\Pr \left[\mathcal{A}^f(\text{ch}) \in \mathcal{R}_{f,\text{ch}} \right],$$

where the randomness is over $f \sim \mu, \text{ch} \sim \pi_f$, and potentially the internal randomness of \mathcal{A} .

³We only show that these paths are ‘‘approximately’’ orthogonal instead of perfectly orthogonal, and this suffices to conclude the result.

⁴Throughout the paper, we assume that oracles are functions with domain $[M]$ and range $[N]$ where the actual M and N may change under different contexts. This is without loss of generality due to the flexibility of μ .

Note that in the above definition, R_f is allowed to be empty, in which case the winning probability is simply zero.

During the analysis, we will need the oracle distribution conditioned on challenge. This naturally leads to the notion of plain games: games without any challenge.

Definition 3.2 (Plain Game). *A plain game G is described by an oracle distribution μ and an accepting outcome set \mathcal{R}_f . The winning probability of a query algorithm is defined similarly as in [Definition 3.1](#).*

Since games with or without challenge can be distinguished by whether there is a challenge distribution, we will address both of them as *games* for convenience when context is clear.

Definition 3.3 (Challenge-Conditioned Game). *Given a game $G = (\mu, \{\pi_f\}, \{\mathcal{R}_{f,\text{ch}}\})$ and a possible challenge ch , we use $G_{\text{ch}} = (\mu_{\text{ch}}, \{\mathcal{R}_{f,\text{ch}}\})$ to denote the plain game of G conditioned on ch , where μ_{ch} is the oracle distribution conditioned on the challenge being ch .*

We also have the following easy fact.

Fact 3.4. $\mu = \mathbb{E}_{\text{ch}}[\mu_{\text{ch}}]$, where ch is sampled from π_f for $f \sim \mu$.

Optimal Winning Probabilities. We start by defining the optimal winning probabilities for uniform and non-uniform query algorithms. Let G be a game and $S, T \in \mathbb{N}$.

- **Classical Uniform Game.** In this case, \mathcal{A} is a classical query algorithm with at most T queries given challenge ch (if not plain) and classical oracle access to f . A classical oracle query to $f: [M] \rightarrow [N]$ is specified by a coordinate $x \in [M]$ and gets in return the value $f(x) \in [N]$. We use $\varepsilon_G(T)$ to denote the maximal winning probability of such algorithms.
- **Classical Non-Uniform Game.** In this case, \mathcal{A} is still limited to at most T classical queries to f . With classical non-uniform power, \mathcal{A} now additionally receives an S -bit advice string σ_f that can depend arbitrarily on f (but not on the challenge ch). We use $\mathcal{A}^f(\sigma_f, \text{ch})$ (or $\mathcal{A}^f(\sigma_f)$ if plain) to explicitly denote that \mathcal{A} takes the non-uniform advice as a separate input. We use $\varepsilon_G(S, T)$ to denote the maximal winning probability of such algorithms.
- **Quantum Uniform Game.** Here, \mathcal{A} is a T -quantum-query algorithm. In this paper, we are only interested in the case when f is a uniformly random oracle. More details are provided in [Section 5.2](#). We use $\varepsilon_G^*(T)$ to denote the maximal winning probability of such quantum algorithms.
- **Quantum Non-Uniform Game.** Here, \mathcal{A} is a T -quantum-query algorithm, with an additional piece of advice σ_f . In the quantum case, we consider the most general case: the advice σ_f has S qubits and can arbitrarily depend on f . We use $\varepsilon_G^*(S, T)$ to denote the maximal winning probability of such non-uniform quantum algorithms.

Since the challenges are explicitly provided to the algorithm, the optimal winning probabilities for uniform algorithms are equivalent for a game and the corresponding challenge-conditioned games.

Fact 3.5. *Let $G = (\mu, \{\pi_f\}, \{\mathcal{R}_{f,\text{ch}}\})$ be a game. Then $\varepsilon_G(T) = \mathbb{E}_{\text{ch}}[\varepsilon_{G_{\text{ch}}}(T)]$.*

Salted Games. Salting is a generic way to restrict non-uniform power. The method appends a piece of random data (aka salt) to inputs fed into a hash function (or more generally, an oracle); and the game will be now played under the random data.

Definition 3.6 (Salted Game). *Let G be a game (see [Definition 3.1](#)) specified by $(\mu, \{\pi_f\}_f, \{\mathcal{R}_{f,\text{ch}}\}_{f,\text{ch}})$. For each $K \in \mathbb{Z}_+$, we define the salted game G_K as follows:*

- The oracle distribution μ_K is over functions $g: [KM] \rightarrow [N]$. We identify $[KM]$ as $[K] \times [M]$ and $g = (f_1, \dots, f_K)$ where each $f_i: [M] \rightarrow [N]$. The distribution μ_K is generated by independently sampling $g(i, \cdot) = f_i \sim \pi$ for each $i \in [K]$.
- For each possible oracle $g = (f_1, \dots, f_K)$ from μ_K , the challenge $\text{ch} = (k, \text{ch}_k) \sim \pi_g$ is produced by first picking a uniformly random $k \sim [K]$, then sampling $\text{ch}_k \sim \pi_{f_k}$.
- For each possible oracle $g = (f_1, \dots, f_K)$ from μ_K and challenge $\text{ch} = (k, \text{ch}_k)$ from π_g , the accepting outcome set $\mathcal{R}_{g,\text{ch}}$ is defined as $\mathcal{R}_{f_k,\text{ch}_k}$.

Remark 3.7. *Intuitively, the game G_K samples K independent copies f_1, \dots, f_K from μ as a joint oracle, then samples a uniformly random one f_k for the challenge, followed by a concrete challenge ch_k from π_k and the corresponding accepting outcome set $\mathcal{R}_{f_k,\text{ch}_k}$. In order to win this game, the algorithm \mathcal{A} needs to output a valid answer for $\mathcal{R}_{f_k,\text{ch}_k}$.*

Note that the definition of optimal winning probabilities carries over. In particular, our goal for the classical setting is to relate $\varepsilon_G(T)$ and $\varepsilon_{G_K}(S, T)$. We similarly want to relate $\varepsilon_G^*(T)$ and $\varepsilon_{G_K}^*(S, T)$ in the quantum random oracle model.

Useful Inequalities. [Lemma 3.8](#) is a moment bound for the number of distinct elements in a draw-with-replacement experiment. Its proof can be found in [Appendix A](#).

Lemma 3.8. *Let $K, L \in \mathbb{Z}_+$. Sample $k_1, \dots, k_L \sim [K]$ independently and let ℓ be the number of distinct elements in k_1, \dots, k_L . Then for any $c \geq 0$, we have*

$$\mathbb{E}[c^\ell] \leq \left(c + \frac{L}{K}\right)^L.$$

[Fact 3.9](#) is a simple counting argument and the estimate follows directly from Stirling's formula, i.e. $a! > (a/e)^a$ for $a \in \mathbb{Z}_+$.

Fact 3.9. *For $K, L \in \mathbb{Z}_+$, we have*

$$|\{n_1, \dots, n_K \in \mathbb{N}: n_1 + \dots + n_K = L\}| = \binom{K+L-1}{K-1} \leq \begin{cases} (2e \cdot K/L)^L & L \leq K, \\ (2e \cdot L/K)^K & L \geq K. \end{cases}$$

4 Non-Uniform Security of Salting: the Classical Setting

In this section, we prove the non-uniform security of salting in the classical case: we upper bound the advantage of an adversarial classical algorithm against salted oracles with the help of advice strings.

Theorem 4.1. Let G be a game and let G_K be the salted game for $K \in \mathbb{Z}_+$. Then for any $S, T \in \mathbb{N}$, we have

$$\varepsilon_{G_K}(S, T) \leq \min_{L \in \mathbb{Z}_+} 2^{S/L} \cdot \left(\varepsilon_G(T) + \frac{L}{K} \right).$$

Remark 4.2. Since we do not assume any structure of G in [Theorem 4.1](#) and its bound does not depend on specific range parameters in G , this result is general enough to handle any kind of classical query model.

Let $f: [N] \rightarrow [M]$ be the oracle for G . Let \mathcal{Q} be the set of allowed queries, e.g., bit query (each time queries some i -th bit of some $f(x)$), word query (each time queries some $f(x)$, which is the query model we use here), majority query (each time queries the majority of $\{f(x) : x \in S\}$ for some $S \subseteq [N]$), and so on. Let \mathcal{S} be the range of the possible answers of queries in \mathcal{Q} .

Then we can view each oracle f as a function $f_{\mathcal{Q}}: \mathcal{Q} \rightarrow \mathcal{S}$. In particular, $f_{\mathcal{Q}}(q)$ returns the query answer on f given a query $q \in \mathcal{Q}$. Then the allowed query \mathcal{Q} on f becomes the word query on $f_{\mathcal{Q}}$. Therefore, we can simply treat $f_{\mathcal{Q}}$ as the oracle in the game G , change the accepting outcome set accordingly, and apply [Theorem 4.1](#).⁵

Remark 4.3. Though [Theorem 4.1](#) is generally tight (see [Theorem 1.2](#)), it can be further strengthened in some setting. Intuitively, if the game has numerous possible challenges, one cannot store sufficient amount of answers to cover all the challenges, especially given the presence of salts. In other words, answers in a look-up table need to not only specify the salt, but also specify the exact challenge corresponding to the answer.

Given the above intuition, in [Section 4.2](#) we modify the proof of [Theorem 4.1](#) to get a different estimate, which improves [Theorem 4.1](#) when there are many possible challenges. We also explicitly compute some examples to demonstrate the tightness.

The strategy to prove [Theorem 4.1](#) is as follows: let g be the salted oracle sampled from $g \sim \mu_K$.

- We first relate the success probability of winning a single challenge $\text{ch} \sim \pi_g$ to the success probability of winning independent challenges $\text{ch}_1, \dots, \text{ch}_L \sim \pi_g$.
- Then we show that the S -bit advice string does not boost the winning probability much when we have multiple independent challenges in the sense that a random string of S -bit can be taken as the advice.
- Now that the advice string is substituted by a random string independent of g , we only need to bound the winning probability of a uniform query algorithm against multiple independent challenges. This will be handled by a direct product theorem.

One can view the direct product theorem as a type of concentration results, where one analyze the high order advantages of algorithms. This reduction from non-uniform security to uniform security with concentration is analogous to earlier works of Impagliazzo and Kabanets [[IK10](#), [Imp11](#)], and has been revived in recent works [[GGKL21](#), [ACDW20](#), [GK22](#), [AGL22](#), [Liu23](#)].

We emphasize that the most general direct product theorems [[Dru12](#)] incur a necessary loss in parameters and cannot produce sharp bounds that [Theorem 4.1](#) needs. Instead, we exploit the *memoryless* property of our algorithms and show a strong direct product theorem for this class of algorithms.

⁵One subtlety here is that f and $f_{\mathcal{Q}}$ may not be a one-to-one correspondence. Though f defines $f_{\mathcal{Q}}$, it is possible that the answers to the queries does not uniquely identify the oracle. In this case, the distribution of $f_{\mathcal{Q}}$ is properly induced the distribution of f and the predicate (aka accepting outcome set) will be probabilistic given $f_{\mathcal{Q}}$ and algorithm's output. The argument still works out with minor changes and we leave this as an exercise for interested readers.

Proof of Theorem 4.1. Let G be specified by $(\mu, \{\pi_f\}_f, \{\mathcal{R}_{f,\text{ch}}\}_{f,\text{ch}})$. Recall [Definition 3.6](#) and [Remark 3.7](#). The game G_K samples K independent copies of G and focuses on a uniformly random one of them. As such, let $f_1, \dots, f_K \sim \mu$ and $k \sim [K]$ be independent, then $g = (f_1, \dots, f_K)$ is the joint oracle sampled from μ_K for G_K . Let $\text{ch}_k \sim \pi_{f_k}$, then $\text{ch} = (k, \text{ch}_k)$ is the challenge sampled from π_g for G_K .

Fix an arbitrary $L \in \mathbb{Z}_+$. Let \mathcal{A} be an arbitrary non-uniform query algorithm for G_K that takes an S -bit advice string σ_g and makes at most T queries to g . Then the winning probability of \mathcal{A} is $\Pr[\mathcal{A}^g(\sigma_g, \text{ch}) \in \mathcal{R}_{g,\text{ch}}]$, where $\mathcal{R}_{g,\text{ch}} = \mathcal{R}_{f_k,\text{ch}_k}$. Since we only make classical queries, oracle access to g is equivalent to oracle access to f_1, \dots, f_K separately. Therefore it suffices to prove

$$\Pr_{f_1, \dots, f_K} \left[\mathcal{A}^{f_1, \dots, f_K}(\sigma_g, k, \text{ch}_k) \in \mathcal{R}_{f_k, \text{ch}_k} \right] \leq 2^{S/L} \cdot \left(\varepsilon_G(T) + \frac{L}{K} \right), \quad (1)$$

where we assume that \mathcal{A} is a deterministic query algorithm given σ_g, k, ch_k by fixing the optimal randomness.

To remove the non-uniform advice from the LHS of (1), we use a tensor trick. Let challenges $(k_1, \text{ch}_{k_1}), \dots, (k_L, \text{ch}_{k_L}) \sim \pi_g$ be independent copies of (k, ch_k) . Then for any fixed g (i.e., f_1, \dots, f_K), the advice σ_g is also fixed and we have

$$\Pr_{(k_i, \text{ch}_{k_i}), i \in [L]} \left[\mathcal{A}^{f_1, \dots, f_K}(\sigma_g, k_i, \text{ch}_{k_i}) \in \mathcal{R}_{f_{k_i}, \text{ch}_{k_i}}, \forall i \in [L] \right] = \Pr_{k, \text{ch}_k} \left[\mathcal{A}^{f_1, \dots, f_K}(\sigma_g, k, \text{ch}_k) \in \mathcal{R}_{f_k, \text{ch}_k} \right]^L.$$

Hence

$$\begin{aligned} \text{LHS of (1)} &= \mathbb{E}_{f_1, \dots, f_K} \left[\Pr_{k, \text{ch}_k} \left[\mathcal{A}^{f_1, \dots, f_K}(\sigma_g, k, \text{ch}_k) \in \mathcal{R}_{f_k, \text{ch}_k} \right] \right] \\ &\leq \mathbb{E}_{f_1, \dots, f_K} \left[\Pr_{k, \text{ch}_k} \left[\mathcal{A}^{f_1, \dots, f_K}(\sigma_g, k, \text{ch}_k) \in \mathcal{R}_{f_k, \text{ch}_k} \right]^L \right]^{1/L} && \text{(by convexity)} \\ &= \Pr_{\substack{f_1, \dots, f_K \\ (k_i, \text{ch}_{k_i}), i \in [L]}} \left[\mathcal{A}^{f_1, \dots, f_K}(\sigma_g, k_i, \text{ch}_{k_i}) \in \mathcal{R}_{f_{k_i}, \text{ch}_{k_i}}, \forall i \in [L] \right]^{1/L}. \end{aligned}$$

Now observe that if σ_g is replaced by a uniformly random string $\sigma \sim \{0, 1\}^S$, then σ still has a probability of 2^{-S} to be identical with σ_g , from which we recover the true non-uniform algorithm. Putting this into the above calculation, we have

$$\begin{aligned} \text{LHS of (1)} &\leq \left(2^S \cdot \Pr_{\substack{f_1, \dots, f_K \\ (k_i, \text{ch}_{k_i}), i \in [L] \\ \sigma \sim \{0, 1\}^S}} \left[\mathcal{A}^{f_1, \dots, f_K}(\sigma, k_i, \text{ch}_{k_i}) \in \mathcal{R}_{f_{k_i}, \text{ch}_{k_i}}, \forall i \in [L] \right] \right)^{1/L} \\ &\leq 2^{S/L} \cdot \Pr_{\substack{f_1, \dots, f_K \\ (k_i, \text{ch}_{k_i}), i \in [L]}} \left[\mathcal{B}^{f_1, \dots, f_K}(k_i, \text{ch}_{k_i}) \in \mathcal{R}_{f_{k_i}, \text{ch}_{k_i}}, \forall i \in [L] \right]^{1/L}, \quad (2) \end{aligned}$$

where \mathcal{B} is a uniform T -query algorithm obtained by fixing the optimal choice of σ for \mathcal{A} .

For each possible (k_1, \dots, k_L) , we fix an arbitrary⁶ $U \subseteq [L]$ such that U contains indices of distinct salts in k_1, \dots, k_L . Note that U is deterministic given (k_1, \dots, k_L) . Now we further analyze the RHS of (2):

$$\begin{aligned}
& \Pr_{\substack{f_1, \dots, f_K \\ (k_i, \text{ch}_{k_i}), i \in [L]}} \left[\mathcal{B}^{f_1, \dots, f_K}(k_i, \text{ch}_{k_i}) \in \mathcal{R}_{f_{k_i}, \text{ch}_{k_i}}, \forall i \in [L] \right] \\
& \leq \Pr_{\substack{f_1, \dots, f_K \\ (k_i, \text{ch}_{k_i}), i \in [L]}} \left[\mathcal{B}^{f_1, \dots, f_K}(k_i, \text{ch}_{k_i}) \in \mathcal{R}_{f_{k_i}, \text{ch}_{k_i}}, \forall i \in U \right] \\
& = \mathbb{E}_{\substack{U \\ (k_i, \text{ch}_{k_i}), i \in U}} \left[\Pr_{f_1, \dots, f_K} \left[\mathcal{B}^{f_1, \dots, f_K}(k_i, \text{ch}_{k_i}) \in \mathcal{R}_{f_{k_i}, \text{ch}_{k_i}}, \forall i \in U \mid U, \{(k_i, \text{ch}_{k_i})\}_{i \in U} \right] \right]. \quad (3)
\end{aligned}$$

Observe that $(U, \{k_i\}_{i \in U})$ is independent from (f_1, \dots, f_K) and each ch_{k_i} only distorts the distribution of f_{k_i} from μ to $\mu_{\text{ch}_{k_i}}$ (recall [Definition 3.3](#)). Hence, conditioned on $(U, \{(k_i, \text{ch}_{k_i})\}_{i \in U})$, the oracles f_1, \dots, f_K are still independent; and in particular, f_{k_i} has distribution $\mu_{\text{ch}_{k_i}}$ for $i \in U$ and other oracles have marginal distribution μ . Therefore, the RHS of (3) can be simplified as

$$\mathbb{E}_{\substack{U \\ (k_i, \text{ch}_{k_i}), i \in U \\ f_j \sim \mu, j \notin \{k_i : i \in U\}}} \left[\Pr_{f_{k_i} \sim \mu_{\text{ch}_{k_i}}, i \in U} \left[\mathcal{B}^{f_1, \dots, f_K}(k_i, \text{ch}_{k_i}) \in \mathcal{R}_{f_{k_i}, \text{ch}_{k_i}}, \forall i \in U \right] \right]. \quad (4)$$

We will prove the following claim by a strong direct product theorem in [Section 4.1](#).

Claim 4.4. *For any fixed $(U, \{(k_i, \text{ch}_{k_i})\}_i, \{f_j\}_j)$, we have*

$$\Pr_{f_{k_i} \sim \mu_{\text{ch}_{k_i}}, i \in U} \left[\mathcal{B}^{f_1, \dots, f_K}(k_i, \text{ch}_{k_i}) \in \mathcal{R}_{f_{k_i}, \text{ch}_{k_i}}, \forall i \in U \right] \leq \prod_{i \in U} \varepsilon_{G_{\text{ch}_{k_i}}}(T).$$

Assuming [Claim 4.4](#), we can upper bound (4):

$$\begin{aligned}
& \mathbb{E}_{\substack{U \\ (k_i, \text{ch}_{k_i}), i \in U \\ f_j \sim \mu, j \notin \{k_i : i \in U\}}} \left[\prod_{i \in U} \varepsilon_{G_{\text{ch}_{k_i}}}(T) \right] = \mathbb{E}_{k_i, i \in U} \left[\mathbb{E}_{\text{ch}_{k_i}, i \in U} \left[\prod_{i \in U} \varepsilon_{G_{\text{ch}_{k_i}}}(T) \right] \right] \\
& = \mathbb{E}_{k_i, i \in U} \left[\prod_{i \in U} \mathbb{E}_{\text{ch}_{k_i}} \left[\varepsilon_{G_{\text{ch}_{k_i}}}(T) \right] \right] \quad (\text{by independence}) \\
& = \mathbb{E}_{k_i, i \in U} \left[\prod_{i \in U} \varepsilon_G(T) \right] = \mathbb{E}_U \left[\varepsilon_G(T)^{|U|} \right] \quad (\text{by Fact 3.5}) \\
& \leq \left(\varepsilon_G(T) + \frac{L}{K} \right)^L. \quad (\text{by Lemma 3.8})
\end{aligned}$$

Combining with (2) and (3), this establishes (1) and thus concludes the proof of [Theorem 4.1](#). \square

⁶For example, if $(k_1, k_2) = (1, 2)$, then $U = \{1, 2\}$; if $(k_1, k_2) = (1, 1)$, then U can be either $\{1\}$ or $\{2\}$.

4.1 Strong Direct Product Theorems for Memoryless Algorithms

In this section, we will present the strong direct product theorem we use in [Claim 4.4](#). This theorem was proved within the context of classical complexity theory in a much earlier work by Nisan, Rudich, and Saks [[NRS98](#)]. While we only became aware of their work after the completion of our paper, we do *not* consider our analysis in [Section 4.1](#) as our technical contribution, and we prove the theorem here within our model for completeness.

We start by defining the direct product of games.

Definition 4.5 (Direct Product of Games). *Let G_1, \dots, G_k be k plain games specified by μ_1, \dots, μ_k and collections of predicates $\mathcal{R}_{1,f_1}, \dots, \mathcal{R}_{k,f_k}$. The direct product of these games is denoted by $G^\times = G_1 \times G_2 \times \dots \times G_k$ with $\mu^\times = \mu_1 \times \mu_2 \times \dots \times \mu_k$ and predicate $\mathcal{R}_{f_1, \dots, f_k}^\times = \mathcal{R}_{1,f_1} \times \dots \times \mathcal{R}_{k,f_k}$.*

The probability that a query algorithm \mathcal{A} wins G^\times is $\Pr[\mathcal{A}^{f_1, \dots, f_k} \in \mathcal{R}_{1,f_1} \times \dots \times \mathcal{R}_{k,f_k}]$, where \mathcal{A} can make queries to any f_1, \dots, f_k . The direct product theorem intuitively says that the winning probability of kT -query algorithms for G^\times should be roughly the product of the winning probabilities of T -query algorithms for each G_i .

To capture the algorithms in [Claim 4.4](#), we focus on memoryless algorithms for direct product of games. In these algorithms, the workspace is repeatedly refreshed, which forbids long-range computation.

Definition 4.6 (Memoryless Algorithm). *Let $G^\times = G_1 \times \dots \times G_k$ be a direct product of games. A query algorithm \mathcal{A} for G^\times is memoryless if it is deterministic⁷ and there exist query algorithms $\mathcal{A}_1, \dots, \mathcal{A}_k$ such that $\mathcal{A}^{f_1, \dots, f_k} \equiv \mathcal{A}_1^{f_1, \dots, f_k}, \dots, \mathcal{A}_k^{f_1, \dots, f_k}$. That is, \mathcal{A} runs $\mathcal{A}_1, \dots, \mathcal{A}_k$ sequentially for $\mathcal{R}_{1,f_1}, \dots, \mathcal{R}_{k,f_k}$ and does not share memory between two executions. Moreover, we say \mathcal{A} is (T_1, \dots, T_k) -memoryless if each \mathcal{A}_i makes at most T_i queries.*

Note that, while (T_1, \dots, T_k) -memoryless algorithms are guaranteed to make at most T_i queries for each predicate \mathcal{R}_{i,f_i} , the queries can be interlaced and highly dependent on the query history. Therefore, it is generally not the case that a (T_1, \dots, T_k) -memoryless algorithm would simply execute a T_1 -query algorithm on f_1 , followed by a T_2 -query algorithm on f_2 , and so on. However, as shown in [Theorem 4.7](#), this special strategy turns out to be optimal.

Theorem 4.7 (Strong Direct Product Theorem; [[NRS98](#), Theorem 3.1]). *Let $G^\times = G_1 \times \dots \times G_k$ be a direct product of games. If \mathcal{A} is a (T_1, \dots, T_k) -memoryless algorithm for G^\times , then its winning probability is at most $\prod_{i=1}^k \varepsilon_{G_i}(T_i)$.*

Given [Theorem 4.7](#), we can easily deduce [Claim 4.4](#).

Proof of [Claim 4.4](#). With $(U, \{(k_i, \text{ch}_{k_i})\}_i, \{f_j\}_j)$ fixed, each $\mathcal{B}^{f_1, \dots, f_k}(k_i, \text{ch}_{k_i})$ is a deterministic T -query algorithm on $\{f_{k_i}\}_i$. Overall, they form a (T, \dots, T) -query algorithm to solve challenge-conditioned games $\{G_{\text{ch}_{k_i}}\}_i$ in a memoryless way. Hence the desired bound follows from [Theorem 4.7](#). \square

To prove [Theorem 4.7](#), we will convert memoryless algorithms to another special type of algorithms for which strong direct product theorems are easier to establish.

⁷The assumption about being deterministic is not necessary as one can fix the randomness of the algorithm. However assuming so is easier for our analysis.

Definition 4.8 (Fair Algorithms). Let $G^\times = G_1 \times \dots \times G_k$ be a direct product of games. We say a query algorithm \mathcal{A} for G^\times is (T_1, \dots, T_k) -fair if it is deterministic⁸ and makes at most T_i queries to each f_i in any execution.

We emphasize that a (T_1, \dots, T_k) -fair algorithm does not need to be memoryless. Moreover, the queries to different oracles can be mixed up as well. Therefore, fair algorithms and memoryless algorithms are incomparable. By an induction on the total number of queries, one can prove a sharp direct product theorem for fair algorithms [Sha03].

Theorem 4.9 ([Sha03, Theorem 5.2]). Let $G^\times = G_1 \times \dots \times G_k$ be a direct product of games. If \mathcal{A} is a (T_1, \dots, T_k) -fair algorithm for G^\times , then its winning probability is at most $\prod_{i=1}^k \varepsilon_{G_i}(T_i)$.

We remark that [Sha03, Theorem 5.2] proves for the case where $T_1 = \dots = T_k$ and each G_i is a decision game (i.e., each \mathcal{R}_{i,f_i} is a singleton set of a binary bit). However the proof can be extended to our setting with minimal changes. We also note that Drucker [Dru12] gives an alternative proof using martingale analysis, which generalizes beyond fair algorithms with a necessary loss on the bounds.

In light of Theorem 4.9 and to prove Theorem 4.7, it suffices to present a reduction from memoryless algorithms to fair algorithms which we present below. Note that Theorem 4.7 follows immediately by the combination of Lemma 4.10 and Theorem 4.9.

Lemma 4.10. Let $G^\times = G_1 \times \dots \times G_k$ be a direct product of games. Let \mathcal{A} be any (T_1, \dots, T_k) -memoryless algorithm for G^\times . Then there exists a (T_1, \dots, T_k) -fair algorithm \mathcal{B} with winning probability at least the winning probability of \mathcal{A} .

To prove Lemma 4.10, we will frequently use the notion of partial assignments.

Definition 4.11 (Partial assignment). Any string $u \in ([M] \cup \{*\})^N$ is a partial assignment for functions of range $[M]$ and domain $[M]$. For a function $f: [M] \rightarrow [N]$, we say it is consistent with u (or $f \triangleleft u$) if $f(x) = u_x$ holds for any x that $u_x \neq *$.

Now we present the reduction and prove Lemma 4.10.

Proof of Lemma 4.10. Assume without loss of generality that the oracle of each game G_i is a function from $[M]$ to $[N]$. Recall Definition 4.6. There exists k algorithms $\mathcal{A}_1, \dots, \mathcal{A}_k$, each making at most T_1, \dots, T_k queries respectively and \mathcal{A} executes as follows:

Algorithm: The memoryless algorithm \mathcal{A}

- \mathcal{A} executes the algorithm one by one:
 1. It runs \mathcal{A}_1 and obtains an output e_1 ;
 2. It runs \mathcal{A}_2 and obtains an output e_2 ;
 - ...
 - k. It runs \mathcal{A}_k and obtains an output e_k .
- Eventually, \mathcal{A} outputs e_1, \dots, e_k .

⁸The assumption about being deterministic is not necessary as one can fix the randomness of the algorithm. However assuming so is easier for our analysis.

Since all the algorithms $\mathcal{A}_1, \dots, \mathcal{A}_k$ commute with each other, \mathcal{A} can run an “out-of-order execution” of them, so long as the query order within the same algorithm \mathcal{A}_i does not change. More precisely, the following algorithm $\tilde{\mathcal{A}}$ has the same output distribution as \mathcal{A} :

Algorithm: The “out-of-order execution” algorithm $\tilde{\mathcal{A}}$

- $\tilde{\mathcal{A}}$ initializes $\mathcal{A}_1, \dots, \mathcal{A}_k$, each \mathcal{A}_i makes no query yet.
- As long as there exists $j \in [k]$ such that \mathcal{A}_j has not terminated yet.
 1. $\tilde{\mathcal{A}}$ picks a subset $\mathcal{J} \subseteq [k]$, where \mathcal{A}_j has not terminated for every $j \in \mathcal{J}$.
 2. For each $j \in \mathcal{J}$, it runs \mathcal{A}_j by one query further.
- $\tilde{\mathcal{A}}$ outputs the results e_1, \dots, e_k of $\mathcal{A}_1, \dots, \mathcal{A}_k$.

The subset \mathcal{J} constructed by $\tilde{\mathcal{A}}$ at Step 1 can be arbitrary and adaptive. Before presenting its construction, we note the following important observation.

Lemma 4.12. *Assume \mathcal{A}_j terminates at some stage of $\tilde{\mathcal{A}}$. Let u be the partial assignment of f_j that $\tilde{\mathcal{A}}$ learned so far. Consider the following two cases:*

1. **The real execution.** $\tilde{\mathcal{A}}$ finishes the rest of its execution with the real f_j . Let the winning probability in the case be p .
2. **The simulated execution.** Given u , $\tilde{\mathcal{A}}$ picks some function f'_j consistent with u and finishes the rest of its execution with the fake f'_j . Let the winning probability in the case be p' .

Then there always exists a choice of f'_j such that $p' \geq p$.

Proof of Lemma 4.12. Without loss of generality assume $j = 1$. Since \mathcal{A}_1 already terminates, its outcome e_1 is fixed. We have

$$p = \sum_{f_1 \triangleleft u} [e_1 \in \mathcal{R}_{1,f_1}] \cdot \Pr_{f \sim \mu_1^{(u)}} [f = f_1] \cdot \Pr_{f_2, \dots, f_k} [\mathcal{A}_2, \dots, \mathcal{A}_k \text{ win} | f_1],$$

where $[e_1 \in \mathcal{R}_{1,f_1}]$ is the binary indicator. On the other hand, we have

$$p' = \sum_{f_1 \triangleleft u} [e_1 \in \mathcal{R}_{1,f_1}] \cdot \Pr_{f \sim \mu_1^{(u)}} [f = f_1] \cdot \Pr_{f_2, \dots, f_k} [\mathcal{A}_2, \dots, \mathcal{A}_k \text{ win} | f'_1].$$

Therefore, we can pick $f'_1 := \operatorname{argmax}_{f_1 \triangleleft u} \Pr_{f_2, \dots, f_k} [\mathcal{A}_2, \dots, \mathcal{A}_k \text{ win} | f_1]$ and then $p' \geq p$. \square

With the above observation, we now consider the following algorithm \mathcal{B} that is almost identical to $\tilde{\mathcal{A}}$ except that whenever some \mathcal{A}_j terminates, it picks a function f'_j consistent with the partial assignment learned so far and runs the rest of the protocol using f'_j . By Lemma 4.12, \mathcal{B} wins with probability at least that of $\tilde{\mathcal{A}}$.

Algorithm: The algorithm \mathcal{B}

- \mathcal{B} initializes $\mathcal{A}_1, \dots, \mathcal{A}_k$, each \mathcal{A}_i makes no query yet.
- As long as there exists $j \in [k]$ such that \mathcal{A}_j has not terminated yet.
 1. \mathcal{B} picks a subset $\mathcal{J} \subseteq [k]$, where \mathcal{A}_j has not terminated for every $j \in \mathcal{J}$.
 2. For every $j \in \mathcal{J}$, it runs \mathcal{A}_j by one query further.

If \mathcal{A}_j terminates, \mathcal{B} picks f'_j (as in [Lemma 4.12](#)). In the rest of the execution, whenever f_j is queried, \mathcal{B} uses f'_j .

- \mathcal{B} outputs the results of $\mathcal{A}_1, \dots, \mathcal{A}_k$.

Finally, we show that there exists a way of picking \mathcal{J} such that at every stage and for each $i \in [k]$, the number of queries to f_i is at most the number of queries made by \mathcal{A}_i . This is in general not true, for instance \mathcal{A}_1 and \mathcal{A}_2 can make all their $T_1 + T_2$ queries to f_1 . However by the reduction above and the memoryless property, we can now rearrange queries by carefully designing \mathcal{J} to avoid these cases.

Lemma 4.13. *There exists a strategy of picking \mathcal{J} such that at any stage and for each $i \in [k]$, we have $t_i \leq q_i$, where q_i is the number of queries made by \mathcal{A}_i and t_i be the number of queries made to (the real) f_i .*

A direct consequence of the above lemma is that, when we use the strategy to pick \mathcal{J} for \mathcal{B} , it becomes a fair algorithm since $t_i \leq q_i \leq T_i$.

Proof of Lemma 4.13. We prove the statement by induction. It is obviously true at the initialization stage where $t_i = q_i = 0$ for every $i \in [k]$.

For inductive hypothesis, assume that each \mathcal{A}_i already makes q_i queries and each f_i gets queried t_i times with $t_i \leq q_i$. For the next stage, we define a directed graph (V, E) as follows:

- Let $V \subseteq [k]$ be the set of \mathcal{A}_i 's that has not terminated yet.
 - For every $r \in V$,
 - if A_r makes the next query to f_v and $v \in V$, then a direct edge $\langle r, v \rangle \in E$;
 - otherwise A_r makes the next query to f_v and $v \notin V$, then no edge will be added.
- Since $v \notin V$ means \mathcal{A}_v already terminates and thus every query to f_v will be directed to the simulated oracle f'_v constructed by [Lemma 4.12](#).

Observe that every node in the graph has out-degree at most one and $V \neq \emptyset$. We complete the proof of [Lemma 4.13](#) by the following case analysis.

There exists as a directed cycle. If $r_1 \rightarrow r_2 \rightarrow \dots \rightarrow r_\ell \rightarrow r_1$ forms a directed cycle, then we define $\mathcal{J} = \{r_1, \dots, r_\ell\}$. Therefore each \mathcal{A}_{r_i} will make a query and each function f_{r_i} gets one more query. The induction hypothesis holds.

There is no directed cycle. Then there exists a direct path $r_1 \rightarrow r_2 \rightarrow \dots \rightarrow r_\ell$ where $\ell \geq 1$. Assume A_{r_ℓ} makes the next query to f_{v_ℓ} . Since r_ℓ is the endpoint of the path, we know that \mathcal{A}_{v_ℓ} already terminates and we have replaced f_{v_ℓ} with the simulated oracle f'_{v_ℓ} by [Lemma 4.12](#).

Now let $\mathcal{J} = \{r_1, \dots, r_\ell\}$. Each \mathcal{A}_{r_i} makes a new query, and functions $f_{r_2}, \dots, f_{r_\ell}$ get one more query respectively. Note that f_1 is not queried and the query of \mathcal{A}_{r_ℓ} is simulated by f'_{v_ℓ} without querying the actual oracles. Thus, the induction hypothesis holds. \square

By [Lemma 4.13](#), the algorithm \mathcal{B} is indeed (T_1, \dots, T_k) -fair. Then [Lemma 4.12](#) guarantees that \mathcal{B} has at least the winning probability as that of \mathcal{A} , which completes the proof of [Lemma 4.10](#). \square

4.2 Improved Estimates for Games with Numerous Challenges

In this section, we provide an alternative way of analyzing the RHS of (2), which implies improved bounds for games with numerous challenges.

For a game with challenge, we define the multi-challenge version of it by asking the algorithm to solve multiple independent challenges. Note that the repetition here is on challenges instead of salts.

Definition 4.14 (Multi-Challenge Game). *Let G be a game specified by $(\mu, \{\pi_f\}_f, \{\mathcal{R}_{f,\text{ch}}\}_{f,\text{ch}})$. For each $n \in \mathbb{N}$, we define the multi-challenge game G^n as follows:*

- The oracle distribution is the same μ .
- For each possible oracle f from μ , the challenge $\text{ch} = (\text{ch}^1, \dots, \text{ch}^n)$ is produced by sampling independent $\text{ch}^1, \dots, \text{ch}^n \sim \pi_f$.
- For each possible oracle f and challenge $\text{ch} = (\text{ch}^1, \dots, \text{ch}^n)$, the accepting outcome set $\mathcal{R}_{f,\text{ch}} = \mathcal{R}_{f,\text{ch}^1} \times \dots \times \mathcal{R}_{f,\text{ch}^n}$.
In other words, the algorithm needs to output accepting configuration for each challenge.

We remark that if $n = 0$, then the multi-challenge game G^0 is a trivial game.

Note that the definition of optimal winning probabilities carries over. In particular, $\varepsilon_{G^n}(nT)$ refers to the optimal winning probability of a nT -query classical algorithm for the multi-challenge game G^n . Oftentimes, $\varepsilon_{G^n}(nT)^{1/n}$ is a good upper bound for the non-uniform security of G against algorithms with n -bit advice and T queries (see e.g., [IK10, Imp11, GGKL21]).

Using this notion, we provide the following variant of Theorem 4.1 and then instantiate it on the function / permutation inversion problem (see e.g., the survey by Corrigan-Gibbs and Kogan [CK19]).

Theorem 4.15. *Let G be a game and let G_K be the salted game for $K \in \mathbb{Z}_+$. Then for any $S, T \in \mathbb{N}$, we have*

$$\varepsilon_{G_K}(S, T) \leq \min_{L \in \mathbb{Z}_+} 2^{S/L} \cdot \left(\frac{L!}{K^L} \sum_{\substack{n_1, \dots, n_K \in \mathbb{N} \\ n_1 + \dots + n_K = L}} \prod_{k=1}^K \frac{\varepsilon_{G^{n_k}}(n_k T)}{n_k!} \right)^{1/L}.$$

Proof. The argument deviates from the proof of Theorem 4.1 on the analysis of the RHS of (2). We will prove the following alternative bound which completes the proof of Theorem 4.15:

$$\Pr_{\substack{f_1, \dots, f_K \\ (k_i, \text{ch}_{k_i}), i \in [L]}} \left[\mathcal{B}^{f_1, \dots, f_K}(k_i, \text{ch}_{k_i}) \in \mathcal{R}_{f_{k_i}, \text{ch}_{k_i}}, \forall i \in [L] \right] \leq \frac{L!}{K^L} \sum_{\substack{n_1, \dots, n_K \in \mathbb{N} \\ n_1 + \dots + n_K = L}} \prod_{k=1}^K \frac{\varepsilon_{G^{n_k}}(n_k T)}{n_k!} \quad (5)$$

To establish (5), for each $k \in [K]$ we define $n_k = |\{i \in [L] : k_i = k\}|$ to be the number of salt k 's appearances in k_1, \dots, k_L . Note that n_1, \dots, n_K depend only on k_1, \dots, k_L . Then the following procedure is equivalent to the LHS of (5):

1. First sample n_1, \dots, n_K and k_1, \dots, k_L .
For each $k \in [K]$, let $S_k \subseteq [L]$ of size n_k be the appearances of salt k in k_1, \dots, k_L .
2. For each $k \in [K]$, we sample f_k and execute the following subroutine: for each $i \in S_k$, sample ch_{k_i} and execute $\mathcal{B}^{f_1, \dots, f_K}(k, \text{ch}_{k_i})$ to check if the outcome is in $\mathcal{R}_{f_k, \text{ch}_{k_i}}$.

Recall that \mathcal{B} is a T -query algorithm. Therefore, conditioned on [Item 1](#), each subroutine in [Item 2](#) is an $n_k T$ -query algorithm aiming to solve the multi-challenge game $G^{n_k} = |S_k|$. In particular, their combination is a memoryless algorithm. Therefore we have the following result analogous to [Claim 4.4](#).

Claim 4.16. *Conditioned on [Item 1](#), [Item 2](#) succeeds with probability at most $\prod_{k=1}^K \varepsilon_{G^{n_k}}(n_k T)$.*

Proof of Claim 4.16. The reduction to [Theorem 4.7](#) is similar to the proof of [Claim 4.4](#). We only remark that while [Theorem 4.7](#) is proved for plain games (i.e., games without challenge), the same result naturally extends to games with challenge by conditioning on the challenge as in the proof of [Theorem 4.1](#). \square

Given [Claim 4.16](#), we now verify (5):

$$\begin{aligned}
\text{LHS of (5)} &= \mathbb{E}_{\substack{n_1, \dots, n_K \\ k_1, \dots, k_L}} \left[\prod_{k=1}^K \varepsilon_{G^{n_k}}(n_k T) \right] = \mathbb{E}_{n_1, \dots, n_K} \left[\prod_{k=1}^K \varepsilon_{G^{n_k}}(n_k T) \right] && \text{(by Claim 4.16)} \\
&= \sum_{\substack{n_1, \dots, n_K \in \mathbb{N} \\ n_1 + \dots + n_K = L}} \binom{L}{n_1, \dots, n_K} \cdot \frac{1}{K^L} \cdot \prod_{k=1}^K \varepsilon_{G^{n_k}}(n_k T) \\
&= \sum_{\substack{n_1, \dots, n_K \in \mathbb{N} \\ n_1 + \dots + n_K = L}} \frac{L!}{K^L} \cdot \prod_{k=1}^K \frac{\varepsilon_{G^{n_k}}(n_k T)}{n_k!} = \text{RHS of (5)} && \square
\end{aligned}$$

Now we present a variant of [Theorem 4.15](#) that is potentially easier to use and is in fact tight for the large advice regime (see [Corollary 1.4](#)) and some specific important games (see [Corollary 4.20](#)).

Corollary 4.17. *Let G be a game and let G_K be the salted game for $K \in \mathbb{Z}_+$. Then for any $S, T \in \mathbb{N}$, we have*

$$\varepsilon_{G_K}(S, T) \leq 2e^2 \cdot \min_{L \in \mathbb{Z}_+} 2^{S/L} \max_{\substack{n_1, \dots, n_K \in \mathbb{N} \\ n_1 + \dots + n_K = L}} \sum_{k=1}^K \frac{\varepsilon_{G^{n_k}}(n_k T)^{1/n_k}}{\min\{K, L\}}.$$

Proof. For any fixed n_1, \dots, n_K with $n_1 + \dots + n_K = L$, we have

$$\begin{aligned}
\left(\prod_{k=1}^K \frac{\varepsilon_{G^{n_k}}(n_k T)}{n_k^{n_k}} \right)^{1/L} &= \left(\prod_{k=1}^K \left(\frac{\varepsilon_{G^{n_k}}(n_k T)^{1/n_k}}{n_k} \right)^{n_k} \right)^{1/L} \\
&\leq \sum_{k=1}^K \frac{\varepsilon_{G^{n_k}}(n_k T)^{1/n_k}}{n_k} \cdot \frac{n_k}{L} && \text{(by the weighted AM-GM inequality)} \\
&= \sum_{k=1}^K \frac{\varepsilon_{G^{n_k}}(n_k T)^{1/n_k}}{L}, && (6)
\end{aligned}$$

where we note that $\varepsilon_{G^{n_k}}(n_k T) = 1$ (and thus $\varepsilon_{G^{n_k}}(n_k T)^{1/n_k} = 1$) if $n_k = 0$. Therefore for any fixed $L \in \mathbb{Z}_+$, we have

$$\begin{aligned}
\varepsilon_{G_K}(S, T) &\leq 2^{S/L} \cdot \left(\frac{L!}{K^L} \sum_{\substack{n_1, \dots, n_K \in \mathbb{N} \\ n_1 + \dots + n_K = L}} \prod_{k=1}^K \frac{\varepsilon_{G^{n_k}}(n_k T)}{n_k!} \right)^{1/L} && \text{(by Theorem 4.15)} \\
&\leq 2^{S/L} \cdot \left(\frac{L^L}{K^L} \sum_{\substack{n_1, \dots, n_K \in \mathbb{N} \\ n_1 + \dots + n_K = L}} \prod_{k=1}^K \frac{e^{n_k} \cdot \varepsilon_{G^{n_k}}(n_k T)}{n_k^{n_k}} \right)^{1/L} && \text{(since } a! \geq (a/e)^a) \\
&\leq e \cdot 2^{S/L} \cdot \frac{L}{K} \cdot \binom{L+K-1}{K-1}^{1/L} \cdot \max_{\substack{n_1, \dots, n_K \in \mathbb{N} \\ n_1 + \dots + n_K = L}} \left(\prod_{k=1}^K \frac{\varepsilon_{G^{n_k}}(n_k T)}{n_k^{n_k}} \right)^{1/L} \\
&\leq e \cdot 2^{S/L} \cdot \frac{L}{K} \cdot \binom{L+K-1}{K-1}^{1/L} \cdot \max_{\substack{n_1, \dots, n_K \in \mathbb{N} \\ n_1 + \dots + n_K = L}} \sum_{k=1}^K \frac{\varepsilon_{G^{n_k}}(n_k T)^{1/n_k}}{L}. && \text{(by (6))}
\end{aligned}$$

By [Fact 3.9](#), we know that if $L \leq K$, then $\binom{L+K-1}{K-1}^{1/L} \leq 2e \cdot \frac{K}{L}$ and

$$\varepsilon_{G_K}(S, T) \leq 2e^2 \cdot 2^{S/L} \max_{\substack{n_1, \dots, n_K \in \mathbb{N} \\ n_1 + \dots + n_K = L}} \sum_{k=1}^K \frac{\varepsilon_{G^{n_k}}(n_k T)^{1/n_k}}{L}$$

and otherwise $L \geq K$, then $\binom{L+K-1}{K-1}^{1/L} \leq (2e \cdot \frac{L}{K})^{K/L} \leq 2e$ and

$$\varepsilon_{G_K}(S, T) \leq 2e^2 \cdot 2^{S/L} \max_{\substack{n_1, \dots, n_K \in \mathbb{N} \\ n_1 + \dots + n_K = L}} \sum_{k=1}^K \frac{\varepsilon_{G^{n_k}}(n_k T)^{1/n_k}}{K}$$

as desired. \square

To demonstrate the tightness of [Theorem 4.15](#) and [Corollary 4.17](#), we consider the function / permutation inversion game G , for which an upper bound of $\varepsilon_{G^n}(nT)$ is known.

Definition 4.18 (Function / Permutation Inversion Game). *The function inversion game Inv is the following game:*

- The oracle is a uniformly random function $f: [N] \rightarrow [N]$.
- The challenge ch is a uniformly random element in $[N]$.
- An answer $a \in [N]$ is an accepting configuration if $f(a) = \text{ch}$.

The permutation inversion game can be similarly defined where the oracle is then a uniformly random permutation.⁹

⁹Since [Fact 4.19](#) holds for both function inversion game and permutation inversion game, we do not distinguish them here with extra notation. However we remark that [Fact 4.19](#) is known to be tight for permutation inversion, but conjectured not tight for function inversion. Strengthening the latter will establish the non-uniform security of function inversion, which is a long-standing open problem in cryptography with deep connections to circuit lower bounds and data structure lower bounds [[CK19](#)].

That is, an adversary wins Inv iff it successfully inverts a random point in a random function / permutation. For our purpose, we quote a known upper bound on the multi-challenge version of Inv from [GGKL21].

Fact 4.19 ([GGKL21]). $\varepsilon_{\text{Inv}^n}(nT) \leq O\left(\frac{nT}{N}\right)^n$.

Plugging [Fact 4.19](#) into [Corollary 4.17](#) and choosing $L = S$, we get the following immediate corollary, which gives another proof of the results in [DGK17].

Corollary 4.20. *For any $S, T \in \mathbb{N}$, we have*

$$\varepsilon_{\text{Inv}_K}(S, T) \leq O\left(\frac{ST}{N \cdot \min\{S, K\}}\right) = O\left(\frac{T}{N} + \frac{ST}{KN}\right).$$

We remark that using [Theorem 1.2](#), we can only get a weaker bound of

$$O\left(\frac{T}{N} + \frac{S}{K}\right).$$

Indeed, the message from [Corollary 4.20](#) is much stronger: salting can eliminate the additional advantages of non-uniform advice once the number of salts exceeds the length of the advice string.

5 Non-Uniform Security of Salting in the QROM

In this section, we prove the non-uniform security of salting in the quantum random oracle model (QROM) for property finding problems.

5.1 Preliminaries: Database and Property on Databases

We consider a random oracle $f : [M] \rightarrow [N]$. For convenience, let $\mathbf{X} = [M]$ be the domain of the random oracle. Let $\mathbf{Y} = [N]$ be the range of the random oracle. For $0 \leq \nu \leq M$, let $\binom{\mathbf{X}}{\nu}$ be the set of all subsets of \mathbf{X} with size ν . We use $\mathbf{x} \in \binom{\mathbf{X}}{\nu}$ to denote a vector \mathbf{x} of length ν such that $1 \leq x_1 < x_2 < \dots < x_\nu \leq M$.

Definition 5.1 (Database). *A database $D = \{(x_i, y_i)\}_{i \in [\nu]}$ is defined by $\mathbf{x} \in \binom{\mathbf{X}}{\nu}$, $\mathbf{y} \in \mathbf{Y}^\nu$ of the same length. $|D| = \nu$ is the length/size of the database. An empty database is denoted by \emptyset .*

Loosely speaking, a database D defines knowledge about a random oracle: namely, the random oracle f will output y_i on x_i for $1 \leq i \leq \nu$; its behavior on other inputs is undefined.

Remark 5.2. *Although a database D is similar to a partial assignment, we will use the word “database” only in the quantum setting. D is used in the compressed oracle technique by Zhandry [Zha19]; even if it intuitively describes a partial assignment that an algorithm learns about a function, databases are “stored” in superposition and can also be unlearned by a quantum algorithm. Due to the differences between classical and quantum algorithms, we will adopt the notation “database” in Zhandry’s work for consistency in our analysis.*

For a pair $(x, y) \in \mathbf{X} \times \mathbf{Y}$, $(x, y) \in D$ denotes that D contains this pair; for an input $x \in \mathbf{X}$, $x \in_{\mathbf{X}} D$ denotes that there exists $y \in \mathbf{Y}$ such that $(x, y) \in D$. We similarly define $(x, y) \notin D$, $x \notin_{\mathbf{X}} D$, $y \in_{\mathbf{Y}} D$, and $y \notin_{\mathbf{Y}} D$. For $(x, y) \in D$, we define $D(x) := y$; otherwise $x \notin_{\mathbf{X}} D$, $D(x)$ is set as a special symbol \perp .

We say a database $D \subseteq D'$ if every $(x_i, y_i) \in D$ is also contained in D' . For $x \notin_{\mathbf{X}} D$ and $y \in \mathbf{Y}$, define $D \cup \{(x, y)\}$ as the database with (x, y) added into D .

Define $\mathcal{D}_{\mathbf{X}, \mathbf{Y}}$ as the set of all databases of a random oracle:

$$\mathcal{D}_{\mathbf{X}, \mathbf{Y}} = \left\{ D = \{(x_i, y_i)\}_{i \in [\nu]} : \mathbf{x} \in \binom{\mathbf{X}}{\nu}, \mathbf{y} \in \mathbf{Y}^\nu, \forall \nu \in [M] \right\}$$

We will ignore the subscripts when \mathbf{X} and \mathbf{Y} are clear from the context and simply write \mathcal{D} .

Definition 5.3 (Property of Databases). *A property P of databases \mathcal{D} is a subset of \mathcal{D} . For a property P , let $P^{(\nu)}$ denote all databases with property P of size ν and $P^{(\leq \nu)}$ denote all databases with property P of size at most ν .*

We say a property P is *monotone* if for any $D \in P$ and $D \subseteq D'$, we have $D' \in P$.

5.2 Preliminaries: Quantum Random Oracle and Compressed Oracle

Here we recall the background of quantum random oracle model and the compressed oracle technique introduced by [Zha19]. This section is mostly taken verbatim from Section 2 of [GLLZ21], and we make some changes whenever these changes work better with our strong threshold direct product theorem.

Quantum Random Oracle Model. An oracle-aided quantum algorithm can perform quantum computation as well as quantum oracle queries. A quantum oracle query for an oracle $f : [M] \rightarrow [N]$ is modeled as a unitary $U_f : |x\rangle |u\rangle = |x\rangle |u + f(x)\rangle$, where $+$ here is the addition modulo N .

A random oracle is a random function $f : [M] \rightarrow [N]$. The random function f is chosen at the beginning. A quantum algorithm making T oracle queries to f can be modeled as the following: it has three registers $|x\rangle, |u\rangle, |z\rangle$, where $x \in [M]$, $u \in [N]$, and z is the algorithm's internal working memory. It starts with some input state $|0\rangle_{\mathcal{A}} = |0\rangle |0\rangle |\psi\rangle$, then it applies a sequence of unitary to the state: $U_1, U_f, U_2, U_f, \dots, U_T, U_f, U_{\text{final}}$ and a final measurement over the computational basis. Each U_f is the quantum oracle query unitary: $U_f |x\rangle |u\rangle = |x\rangle |u + f(x)\rangle$. Each U_i is a local quantum computation independent of f , and U_{final} is the local quantum computation (independent of f) after the last query. For a uniform algorithm, we have $|0\rangle_{\mathcal{A}} = |0\rangle |0\rangle |0\rangle$ as the algorithm can initialize its internal register by U_1 . By postponing measurement, we assume that the measurement only happens in the end and is in the computational basis.

Compressed Oracle. Compressed oracle is an analogy of the classical lazy sampling method. To simulate a random oracle, one can sample $f(x)$ for all inputs x and store everything in quantum accessible registers. Such an implementation of a random oracle is inefficient and hard to analyze. Instead of recording all the information of f in the registers, Zhandry provides a solution to argue the amount of information an algorithm knows about the random oracle.

The oracle register records a database that contains the output on each input x ; the output is an element in $[N] \cup \{\perp\}$, where \perp is a special symbol denoting that the value is "uninitialized".

The database is initialized as an empty list D_0 of length M , i.e., the pure state $|\emptyset\rangle_f := |\perp, \perp, \dots, \perp\rangle$. Let $|D|$ denote the number of non- \perp entries in D . Define $D(x)$ to be the x -th entry of D . $D(x)$ can be seen as the output of the oracle on x if $D(x) \neq \perp$; otherwise, the oracle's output on x is still undetermined.

For the compressed oracle cO , we have several oracle variations as defined in [Zha19, Section 3]. They are equivalent, and the only difference is that the oracle registers and/or the query registers are encoded in different ways between queries. Therefore we use cO to denote the compressed oracle query, and when we do the simulation, we can implement cO by any of its variations.

Here we define two oracle variations, the compressed standard oracle CStO and the compressed phase oracle CPhsO .

$$\begin{aligned}\text{CStO} &:= \text{StdDecomp} \cdot \text{CStO}' \cdot \text{StdDecomp}, \\ \text{CPhsO} &:= \text{StdDecomp} \cdot \text{CPhsO}' \cdot \text{StdDecomp}.\end{aligned}$$

They are both unitary operators that operates on the joint system of the algorithm registers \mathcal{A} and oracle's registers f .

- $\text{CStO}' |x, u\rangle |D\rangle = |x, u + D(x)\rangle |D\rangle$ when $D(x) \neq \perp$, which writes the output of x defined in D to the u register. This operator will never be applied on an x, D where $D(x) = \perp$.
- $\text{CPhsO}' |x, u\rangle |D\rangle = (-1)^{u \cdot D(x)} |x, u\rangle |D\rangle$ when $D(x) \neq \perp$, which writes the output of x defined in D to the phase. This operator will never be applied on an x, D where $D(x) = \perp$.
- $\text{StdDecomp}(|x\rangle \otimes |D\rangle) = |x\rangle \otimes \text{StdDecomp}_x |D\rangle$, where $\text{StdDecomp}_x |D\rangle$ works on the x -th register of the database $D(x)$ and swaps a uniform superposition $\frac{1}{\sqrt{N}} \sum_y |y\rangle$ with $|\perp\rangle$ on the x -th register:
 - If $D(x) = \perp$, StdDecomp_x maps $|\perp\rangle$ to $\frac{1}{\sqrt{N}} \sum_y |y\rangle$, or equivalently, $\text{StdDecomp}_x |D\rangle = \frac{1}{\sqrt{N}} \sum_y |D \cup (x, y)\rangle$. Intuitively, if the database does not contain information about x , it samples a fresh y as the output of x .
 - If $D(x) \neq \perp$, StdDecomp_x is an identity on $\frac{1}{\sqrt{N}} \sum_y \omega_N^{uy} |y\rangle$ for all $u \neq 0$ and maps the uniform superposition $\frac{1}{\sqrt{N}} \sum_y |y\rangle$ to $|\perp\rangle$, where ω_N is the N -th root of unity. More formally, for D' with $D'(x) = \perp$, we have

$$\text{StdDecomp}_x \frac{1}{\sqrt{N}} \sum_y \omega_N^{uy} |D' \cup (x, y)\rangle = \frac{1}{\sqrt{N}} \sum_y \omega_N^{uy} |D' \cup (x, y)\rangle \text{ for any } u \neq 0,$$

and

$$\text{StdDecomp}_x \frac{1}{\sqrt{N}} \sum_y |D' \cup (x, y)\rangle = |D'\rangle.$$

Since all $\frac{1}{\sqrt{N}} \sum_y \omega_N^{uy} |y\rangle$ and $|\perp\rangle$ form a basis, these define a unique unitary operation.

A quantum algorithm making T oracle queries to a compressed oracle can be modeled as follows: it has three registers $|x\rangle, |u\rangle, |z\rangle$, where $x \in [M]$, $u \in [N]$, and z is the algorithm's internal working memory. Started with some input state $|0\rangle_{\mathcal{A}} = |0\rangle |0\rangle |\psi\rangle$, the joint state of the algorithm and the compressed oracle is $|0\rangle_{\mathcal{A}} \otimes |\emptyset\rangle_f$. It then applies a sequence of unitary to the state: $U_1,$

$cO, U_2, cO, \dots, U_T, cO, U_{\text{final}}$, and a final measurement over computational basis. For a uniform algorithm, we have $|0\rangle_{\mathcal{A}} = |0\rangle |0\rangle |0\rangle$ as the algorithm can initialize its internal register by U_1 .

Zhandry [Zha19] proved that the quantum random oracle model and the compressed standard oracle model are perfectly indistinguishable by any *unbounded* quantum algorithm. Due to the equivalence between oracle variations, this also holds for the compressed phase oracle model.

In this work, we only consider query complexity, and thus simulation efficiency is irrelevant to us. We simulate a random oracle as a compressed phase oracle, i.e., $cO = \text{CPhsO}$, to help us analyze security games with the help from the following lemmas due to [Zha19, CGLQ20]. They are proved for the compressed standard oracle model, and due to the equivalence, they also hold for the compressed phase oracle.

The first lemma gives a general formulation of the overall state of \mathcal{A} and the compressed oracle after \mathcal{A} makes T queries, even conditioned on arbitrary measurement results. It provides the intuition that one oracle query can append at most one element to the database register in superposition.

Lemma 5.4. *If \mathcal{A} makes at most T queries to a compressed oracle, assuming the overall state of \mathcal{A} and the compressed oracle is $\sum_{x,u,z,D} \alpha_{x,u,z,D} |x, u, z\rangle_{\mathcal{A}} \otimes |D\rangle_f$ where $|x, u, z\rangle$ is \mathcal{A} 's registers and $|D\rangle$ is the oracle's registers, then it only has support on all D such that $|D| \leq T$. In other words, the overall state can be written as*

$$\sum_{x,u,z,D:|D|\leq T} \alpha_{x,u,z,D} |x, u, z\rangle_{\mathcal{A}} \otimes |D\rangle_f.$$

The second lemma provides a quantum analogue of lazy sampling in the classical ROM. It gives us a way to connect the compressed oracle with the original oracle U_f , and therefore we can upper bound the winning probability of a quantum algorithm with access to a random oracle by analyzing the database / oracle register.

Lemma 5.5 ([Zha19, Lemma 5]). *Let $f: [M] \rightarrow [N]$ be a random oracle. Consider a quantum algorithm \mathcal{A} making queries to the standard oracle and outputting tuples $(x_1, \dots, x_c, y_1, \dots, y_c, z)$. Suppose the random function f is measured after \mathcal{A} produces its output. Let R be an arbitrary set of such tuples. Suppose with probability p , \mathcal{A} outputs a tuple such that (1) the tuple is in R and (2) $f(x_i) = y_i$ for all i . Now consider running \mathcal{A} with the compressed standard oracle CStO , and suppose the database D is measured after \mathcal{A} produces its output. Let p' be the probability that (1) the tuple is in R and (2) $D(x_i) = y_i$ (in particular, $D(x_i) \neq \perp$) for all i . Then $\sqrt{p} \leq \sqrt{p'} + \sqrt{c/N}$.*

This is also true for the compressed phase oracle CPhsO .

5.3 Classical and Quantum Property Finding Problem

Classical Property Finding. The optimal winning probability of many important games can be captured by the following property finding problem.

Definition 5.6 (Classical Property Finding for Property P). *Let P be a monotone property.¹⁰ We say a query algorithm wins if $D \in P$, where D is the database corresponding to the partial assignment a classical algorithm learns by querying a random oracle.*

¹⁰We are only interested in the case of P being monotone, i.e., learning more information is always better.

Some Examples. In the pre-image search problem, the goal is to find an input x such that $f(x) = 0$. The property is defined as $P = \{D : 0 \in_Y D\}$. For the collision finding problem, we want to find a pair of distinct inputs that map to the same output. In this case, $P = \{D : D(x) = D(x') \text{ for some distinct } x, x' \in_X D\}$.

Definition 5.7 (Transition Probability). *For a monotone property P , we define \bar{P} to be the complement of P , and the probability p_ν as*

$$p_\nu = \max_{\substack{D \in \bar{P}^{(\leq \nu)} \\ x \notin_X D}} \Pr_{y \leftarrow \mathbf{Y}} [D \cup \{(x, y)\} \in P],$$

where we recall that $\bar{P}^{(\leq \nu)}$ denotes all databases in \bar{P} of size at most ν . By the monotonicity of P , p_ν is non-decreasing on ν .

Lemma 5.8. *Let P be a monotone property with transition probability p_t . For any classical T -query algorithm, the probability of winning the property finding problem of P is at most $p_0 + p_1 + \dots + p_{T-1} \leq T \cdot p_T$.*

Since we are not interested in the above classical lemma for the purpose of this work, we just offer an illustrative example to compare with the quantum analog of it.

Proof Sketch of Lemma 5.8. The probability of an algorithm reaching $D \in P$ is upper bounded by the summation of its probability of reaching D at the i -th step for every $i \in \{1, 2, \dots, T\}$. Thus, the probability is at most $p_0 + \dots + p_{T-1} \leq T \cdot p_T$. \square

Quantum Property Finding Problem. In the quantum case, we work in the compressed oracle framework. At any stage of an algorithm, the joint system of the algorithm and the random oracle is described by a state over the algorithm register and the database register.

let Λ_P be a projection that in superposition checks whether the database is in P :

$$\Lambda_P = \sum_{D \in P} |D\rangle \langle D|.$$

We can similarly define the quantum version of property finding problem in the quantum random oracle model.

Definition 5.9 (Quantum Property Finding). *The quantum property finding problem for property P is defined as follows: a quantum algorithm interacts with a compressed oracle; when the algorithm stops, a binary-valued measurement $\{\Lambda_P, \mathbb{I} - \Lambda_P\}$ is applied to the compressed oracle register and the algorithm wins if the measurement outcome is 0 (corresponding to Λ_P , it “finds” the property).*

Theorem 5.10 ([CFHL21, Theorem 5.7]). *Let P be a monotone property with transition probability p_t . For any quantum T -query algorithm, the probability of its winning the quantum property finding problem for property P is at most $\gamma_T = O(\sqrt{p_0} + \sqrt{p_1} + \dots + \sqrt{p_{T-1}})^2 = O(T^2 \cdot p_T)$.*

5.4 Main Quantum Theorem

Definition 5.11. Let R be a collection of tuples in $X^c \times Y^c$. Define P_R to be the property on \mathcal{D} such that $D \in P_R$ if and only if D consists of a tuple in R :

$$D \in P_R \iff \exists (x_1, \dots, x_c, y_1, \dots, y_c) \in R \text{ such that } D(x_i) = y_i, \forall i \in \{1, 2, \dots, c\}.$$

It is easy to see that P_R is monotone for any R .

The game we consider in the QROM is the following.

Definition 5.12 (Game G_R). Let R be a collection of tuples. G_R asks a quantum-query algorithm to output $(x_1, \dots, x_c, y_1, \dots, y_c) \in R$.

Here we are interested in the case when c is small. Some examples of G_R include: finding pre-images of zero, collision finding, k -SUM. Problems like computing the parity of the function has large c and therefore are not covered in our main theorem. We now present our main theorem in the quantum case.

Theorem 5.13. Let R be a collection of tuples and $R \subseteq \mathbf{X}^c \times \mathbf{Y}^c$. Let G_R be the game defined above. Let p_T be the transition probability, and γ_T be an upper bound on the winning probability of any T -quantum-query algorithm, derived from Zhandry's compressed oracle technique ([Theorem 5.10](#)); $\gamma_T = \Theta(T^2 \cdot p_T)$. If the transition probability p_t is polynomial in t , i.e., $p_t = \Theta(t^{c_1}/N^{c_2})$ for some constants $c_1, c_2 \geq 0$, then for every $T \geq c$,

$$\varepsilon_{G_R, K}^*(S, T) \leq \tilde{O} \left(\gamma_T + \frac{S}{K} + \frac{c}{N} \right),$$

where \tilde{O} hides low order terms in $\log N$.

Since an algorithm needs to output a c -tuple of input-output pairs that are consistent with a random oracle, we can without loss of generality only consider the case when $T \geq c$. For G_R that we are interested in the work, the compressed oracle technique shows that $\varepsilon_{G_R}^*(T) \leq O(\gamma_T + \frac{c}{N})$. When the technique is tight, our theorem is also tight. Zhandry [[Zha19](#)] demonstrated the tightness of [Theorem 5.10](#) for some special cases. Thus, we get the following consequences.

Corollary 5.14 (Consequences of [Theorem 5.13](#)). For CRHF whose goal is to find two distinct inputs whose images are identical, we have

$$\varepsilon_{\text{CRHF}_K}^*(S, T) \leq \tilde{O} \left(\frac{T^3}{N} + \frac{S}{K} \right).$$

For Inv_0 whose goal is to find a pre-image of zero, we have

$$\varepsilon_{\text{Inv}_0, K}^*(S, T) \leq \tilde{O} \left(\frac{T^2}{N} + \frac{S}{K} \right).$$

For k SUM whose goal is to find k distinct inputs whose images sum up to zero, we have

$$\varepsilon_{k\text{SUM}_K}^*(S, T) \leq \tilde{O} \left(\frac{T^{k+1}}{N} + \frac{S}{K} \right).$$

5.5 Quantum Direct Product Theorems for Property Finding Problem

To prove our main theorem [Theorem 5.13](#), one crucial step is to give a (threshold) strong direct product theorem for the property finding problem. In this part, we introduce the notion of quantum direct product of property finding problems, and present the strong direct product theorems (SDPT) in this case.

Quantum Direct Product of Property Finding Problem. Now a random oracle f is treated as $[K] \times \mathbf{X} \rightarrow \mathbf{Y}$. For convenience, $\tilde{\mathbf{X}} := [K] \times \mathbf{X}$. The random oracle can be viewed as a concatenation of K independent random oracles; namely, for each $k \in [K]$, $f_k = f(k, \cdot)$ is a random oracle $\mathbf{X} \rightarrow \mathbf{Y}$.

An input is denoted by $(k, x) \in \tilde{\mathbf{X}}$ for $k \in [K]$ and $x \in \mathbf{X}$. A database D is similarly defined by $\mathbf{x} \in \binom{\tilde{\mathbf{X}}}{\nu}$ and $\mathbf{y} \in \mathbf{Y}^\nu$ of the same length ν . It is also nature to define $((k, x), y) \in D, (k, x) \in \tilde{\mathbf{X}}, y \in \mathbf{Y} \implies D$ and $D \subseteq D'$.

Definition 5.15 (Restriction of Databases). For a database $D \in \mathcal{D}_{\tilde{\mathbf{X}}, \mathbf{Y}}$ and an index $k \in [K]$, we define the restriction of D to k as $D|_k \in \mathcal{D}_{\mathbf{X}, \mathbf{Y}}$:

$$D|_k = \{(x, y) : ((k, x), y) \in D\}.$$

In other words, $D|_k$ is a subset of D with prefix k removed.

Definition 5.16 (Direct Product of Property). Let P be a property on $\mathcal{D}_{\mathbf{X}, \mathbf{Y}}$. Let $\kappa \in \mathbb{Z}_+$ and $\tilde{\mathbf{X}} = [\kappa] \times \mathbf{X}$. Define the direct product property on $\mathcal{D}_{\tilde{\mathbf{X}}, \mathbf{Y}}$ as $P^{\times \kappa}$:

$$D \in P^{\times \kappa} \iff D|_k \in P, \forall k \in [\kappa].$$

In the quantum setting, when the database register is a superposition state of databases, $D|_k \in P$ corresponds to the projected state where every superposition satisfies $D|_k \in P$.

Definition 5.17 (Success on Salts). Let P be a property on $\mathcal{D}_{\mathbf{X}, \mathbf{Y}}$. If $D|_k \in P$, we say that the database D succeeds on salt k . If \mathcal{K} is the set of k such that $D|_k \in P$, we say that the database D succeeds on $|\mathcal{K}|$ salts.

Quantum Direct Product Theorems. We have the following direct product theorem for quantum property finding problem:

Theorem 5.18 (SDPT For Quantum Property Finding). Let P be a monotone property with transition probability p_t . Let γ_T be an upper bound on the winning probability of any T -quantum-query algorithm, derived from Zhandry's compressed oracle technique ([Theorem 5.10](#)).

Let $\kappa \in \mathbb{Z}_+$ and $P^{\times \kappa}$ be the direct product. If p_T is polynomial in T , i.e., $p_T = \Theta(T^{c_1}/N^{c_2})$ for some constants $c_1, c_2 \geq 0$, then with $\gamma_T = \Theta(T^2 \cdot p_T)$, the maximum winning probability of any κT -query quantum algorithm for the quantum property finding problem for $P^{\times \kappa}$ is at most $O(1)^\kappa \cdot \gamma_T^\kappa$.

Actually we can have a threshold version of [Theorem 5.18](#).

Theorem 5.19 (Threshold SDPT for Quantum Property Finding). *Let P be a monotone property with transition probability p_t . Let γ_T be an upper bound on the winning probability of any T -quantum-query algorithm, derived from Zhandry's compressed oracle technique (Theorem 5.10).*

For salt space $[K]$ and threshold $\kappa \leq K$, if p_T is polynomial in T , i.e., $p_T = \Theta(T^{c_1}/N^{c_2})$ for some constants $c_1, c_2 \geq 0$, the maximum winning probability for any B -query quantum algorithm ($B \in \mathbb{Z}_+$) to succeed on at least κ salts in this quantum property finding problem is at most $O(1)^\kappa \cdot (\gamma_{B/\kappa})^\kappa$.

Remark 5.20. *For Theorem 5.19, we can rewrite p_T as the polynomial $p(T) = \Theta(T^{c_1}/N^{c_2})$, and in this way we define $p_{B/\kappa} := p(B/\kappa)$ when B/κ is not an integer. Similarly we define $\gamma_{B/\kappa} = t^2 \cdot p_t = t^2 \cdot p(t)$ for $t = B/\kappa$.*

The proof of Theorem 5.18 and Theorem 5.19 can be found in Appendix C, where we prove the strong direct product theorem for quantum property finding, and the threshold version as a corollary.

5.6 Proof of Theorem 5.13

In this section, we consider the following two-stage quantum algorithms $(\mathcal{B}, \mathcal{A})$.

- Let \mathcal{B} be a 2ST-quantum-query algorithm. First, \mathcal{B} executes and outputs a binary predicate b , together with a quantum state.
- The quantum state is then given to \mathcal{A} . The goal of \mathcal{A} is to win $G_{R,K}$ within T quantum queries and with the state received from \mathcal{B} .

We denote the event of $b = 0$ by \mathbf{W} . To prove Theorem 5.13, we only need to prove the following theorem.

Theorem 5.21. *Let $G_{R,K}$ be a game and $c \in \mathbb{Z}^+$ as defined in Theorem 5.13. Assume there exists $\delta > \frac{1}{N}$ such that for every two-stage quantum algorithms $(\mathcal{B}, \mathcal{A})$ with 2ST and T queries each, we have $\Pr[\mathbf{W}] > 1/N^S$ and*

$$\Pr_{\substack{f \sim \mu \\ k \sim [K]}} \left[\mathcal{A}^f(k) \text{ wins} \mid \mathbf{W} \right] \leq \delta.$$

Then $\varepsilon_{G_{R,K}}^*(S, T) \leq \delta$.

The proof is intuitively explained in the technical overview (Quantum Bit-Fixing section) and a formal proof is deferred to Appendix B.1.

Thus, our goal becomes to give an upper bound for $\Pr[\mathcal{A}^f(k) \text{ wins} \mid \mathbf{W}]$ when $\Pr[\mathbf{W}]$ is not “too small”. We define \mathbf{E} to be the event that the database registers succeed on at most $2S \log N$ salts. With \mathbf{E} , we can give the following bound of $\Pr[\mathcal{A}^f(k) \text{ wins} \mid \mathbf{W}]$.

Claim 5.22. *Define \mathbf{E} to be the event that when the algorithm \mathcal{B} finishes, the database registers succeed on at most $2S \log N$ salts. Define $\overline{\mathbf{E}}$ to be the complement of \mathbf{E} where the database registers can succeed on more than $2S \log N$ salts. Then,*

$$\Pr_{\substack{f \sim \mu \\ k \sim [K]}} [\mathcal{A}^f(k) \text{ wins} \mid \mathbf{W}] \leq \Pr_{\substack{f \sim \mu \\ k \sim [K]}} [\mathcal{A}^f(k) \text{ wins} \mid \mathbf{W} \wedge \mathbf{E}] + \frac{\Pr_{f \sim \mu}[\overline{\mathbf{E}}]}{\Pr_{f \sim \mu}[\mathbf{W}]} \tag{7}$$

From [Claim 5.22](#), we can give an upper bound for the two terms in RHS of (7) respectively. This leads us to the proof of [Theorem 5.13](#). The main idea of the proof is as follows:

- For the first term $\Pr[\mathcal{A}^f(k) \text{ wins } | \mathbf{W} \wedge \mathbf{E}]$,
 1. We apply [Lemma 5.5](#) and reduce it to giving an upper bound for the quantum property finding problem, as in [Claim 5.23](#).
 2. Then we upper bound that quantum property finding problem, as in [Claim 5.24](#).
- For the second term $\Pr[\overline{\mathbf{E}}]/\Pr[\mathbf{W}]$, since $\Pr[\mathbf{W}] \geq 1/N^S$, we only need to give an upper bound for $\Pr[\overline{\mathbf{E}}]$. This corresponds to a threshold strong direct product theorem for the quantum property finding problem, as in [Claim 5.25](#).

Claim 5.23. For \mathcal{A}^f that outputs answers as a tuple $(x_1, \dots, x_c, y_1, \dots, y_c)$, we have the following relation between the original game and its corresponding property finding problem:

$$\Pr_{\substack{f \sim \mu \\ k \sim [K]}} [\mathcal{A}^f(k) \text{ wins } | \mathbf{W} \wedge \mathbf{E}] \leq 2 \cdot \Pr_{\substack{f \sim \mu \\ k \sim [K]}} [(D|_k \in P) | \mathbf{W} \wedge \mathbf{E}] + \frac{2c}{N}. \quad (8)$$

Here D corresponds to the database registers, and $D|_k \in P$ corresponds to the projected state that $D|_k \in P$ for every superposition; it means that \mathcal{A}^f succeeds on salt k for the quantum property finding problem of P .

Claim 5.24. Conditioned on $\mathbf{W} \wedge \mathbf{E}$, we have the following upper bound of the maximum winning probability for the property finding problem on a random salt $k \sim [K]$:

$$\Pr_{\substack{f \sim \mu \\ k \sim [K]}} [(D|_k \in P) | \mathbf{W} \wedge \mathbf{E}] \leq O(\gamma_{2T}) + \tilde{O}\left(\frac{S}{K}\right). \quad (9)$$

In the theorem statement, γ_t is assumed to be a polynomial in t . Thus, $\gamma_{2T} = O(\gamma_T)$.

Claim 5.25. For $\overline{\mathbf{E}}$ as defined in [Claim 5.22](#),

$$\Pr_{f \sim \mu} [\overline{\mathbf{E}}] \leq \Pr_{f \sim \mu} [\mathbf{W}] \cdot \frac{1}{N}. \quad (10)$$

With [Claim 5.22](#), [Claim 5.23](#), [Claim 5.24](#), and [Claim 5.25](#), we can conclude an upper bound for $\Pr[\mathcal{A}^f(k) \text{ wins } | \mathbf{W}]$. Together with [Theorem 5.21](#), this gives us an upper bound for the non-uniform security as in [Theorem 5.13](#).

Proofs of all the claims and [Theorem 5.13](#) can be found in [Appendix B](#).

Acknowledgement

KW wants to thank Guangxu Yang and Penghui Yao for references in direct product theorems.

References

- [ACDW20] Akshima, David Cash, Andrew Drucker, and Hoeteck Wee. “Time-space tradeoffs and short collisions in merkle-damgård hash functions”. In: *Annual International Cryptology Conference*. Springer, 2020, pp. 157–186 (cit. on p. 16).

- [AGL22] Akshima, Siyao Guo, and Qipeng Liu. “Time-space lower bounds for finding collisions in Merkle-Damgård hash functions”. In: *Annual International Cryptology Conference*. Springer. 2022, pp. 192–221 (cit. on pp. 12, 16).
- [BHT97] Gilles Brassard, Peter Hoyer, and Alain Tapp. “Quantum algorithm for the collision problem”. In: *arXiv preprint quant-ph/9705002* (1997) (cit. on p. 3).
- [CDGS18] Sandro Coretti, Yevgeniy Dodis, Siyao Guo, and John P. Steinberger. “Random Oracles and Non-uniformity”. In: *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part I*. Vol. 10820. 2018, pp. 227–258. DOI: [10.1007/978-3-319-78381-9_9](https://doi.org/10.1007/978-3-319-78381-9_9) (cit. on pp. 2–4, 8).
- [CFHL21] 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: *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer. 2021, pp. 598–629 (cit. on pp. 30, 40).
- [CGLQ20] Kai-Min Chung, Siyao Guo, Qipeng Liu, and Luowen Qian. “Tight quantum time-space tradeoffs for function inversion”. In: *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2020, pp. 673–684 (cit. on pp. 3, 4, 8, 29).
- [CK19] Henry Corrigan-Gibbs and Dmitry Kogan. “The function-inversion problem: Barriers and opportunities”. In: *Theory of Cryptography Conference*. Springer. 2019, pp. 393–421 (cit. on pp. 23, 25).
- [CLMP13] Kai-Min Chung, Huijia Lin, Mohammad Mahmoody, and Rafael Pass. “On the power of nonuniformity in proofs of security”. In: *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*. 2013, pp. 389–400 (cit. on p. 2).
- [DDG⁺15] Roei David, Irit Dinur, Elazar Goldenberg, Guy Kindler, and Igor Shinkar. “Direct sum testing”. In: *Proceedings of the 2015 Conference on innovations in theoretical computer science*. 2015, pp. 327–336 (cit. on p. 6).
- [DGK17] Yevgeniy Dodis, Siyao Guo, and Jonathan Katz. “Fixing Cracks in the Concrete: Random Oracles with Auxiliary Input, Revisited”. In: *Advances in Cryptology - EUROCRYPT 2017 - 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Paris, France, April 30 - May 4, 2017, Proceedings, Part II*. Vol. 10211. 2017, pp. 473–495. DOI: [10.1007/978-3-319-56614-6_16](https://doi.org/10.1007/978-3-319-56614-6_16) (cit. on pp. 2–4, 26).
- [Dru12] Andrew Drucker. “Improved direct product theorems for randomized query complexity”. In: *computational complexity* 21.2 (2012), pp. 197–244 (cit. on pp. 6, 7, 16, 20).
- [DS14] Irit Dinur and David Steurer. “Direct product testing”. In: *2014 IEEE 29th Conference on Computational Complexity (CCC)*. IEEE. 2014, pp. 188–196 (cit. on p. 6).
- [DTT10] Anindya De, Luca Trevisan, and Madhur Tulsiani. “Time space tradeoffs for attacks against one-way functions and PRGs”. In: *Annual Cryptology Conference*. Springer. 2010, pp. 649–665 (cit. on p. 2).
- [Fei91] Uriel Feige. “On the success probability of the two provers in one-round proof systems”. In: *1991 Proceedings of the Sixth Annual Structure in Complexity Theory Conference*. IEEE Computer Society. 1991, pp. 116–117 (cit. on p. 7).

- [GGKL21] Nick Gravin, Siyao Guo, Tsz Chiu Kwok, and Pinyan Lu. “Concentration bounds for almost k -wise independence with applications to non-uniform security”. In: *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2021, pp. 2404–2423 (cit. on pp. 3, 16, 23, 26).
- [GK22] Ashrujit Ghoshal and Ilan Komargodski. “On time-space tradeoffs for bounded-length collisions in Merkle-Damgård hashing”. In: *Annual International Cryptology Conference*. Springer. 2022, pp. 161–191 (cit. on p. 16).
- [GLLZ21] Siyao Guo, Qian Li, Qipeng Liu, and Jiapeng Zhang. “Unifying Presampling via Concentration Bounds”. In: *Theory of Cryptography Conference*. Springer. 2021, pp. 177–208 (cit. on pp. 4, 8, 9, 27).
- [IJKW08] Russell Impagliazzo, Ragesh Jaiswal, Valentine Kabanets, and Avi Wigderson. “Uniform direct product theorems: simplified, optimized, and derandomized”. In: *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*. 2008, pp. 579–588 (cit. on p. 6).
- [IK10] Russell Impagliazzo and Valentine Kabanets. “Constructive proofs of concentration bounds”. In: *International Workshop on Randomization and Approximation Techniques in Computer Science*. Springer. 2010, pp. 617–631 (cit. on pp. 4, 5, 16, 23).
- [Imp11] Russell Impagliazzo. “Relativized separations of worst-case and average-case complexities for NP”. In: *2011 IEEE 26th Annual Conference on Computational Complexity*. IEEE. 2011, pp. 104–114 (cit. on pp. 5, 16, 23).
- [JK21] Rahul Jain and Srijita Kundu. “A Direct Product Theorem for One-Way Quantum Communication”. In: *36th Computational Complexity Conference*. 2021 (cit. on p. 6).
- [JK22] Rahul Jain and Srijita Kundu. “A direct product theorem for quantum communication complexity with applications to device-independent QKD”. In: *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2022, pp. 1285–1295 (cit. on p. 6).
- [Lev85] Leonid A Levin. “One-way functions and pseudorandom generators”. In: *Proceedings of the seventeenth annual ACM symposium on Theory of computing*. 1985, pp. 363–365 (cit. on p. 6).
- [Liu23] Qipeng Liu. “Non-uniformity and Quantum Advice in the Quantum Random Oracle Model”. In: *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer. 2023, pp. 117–143 (cit. on pp. 3, 4, 9, 16, 37, 38).
- [LR13] Troy Lee and Jérémie Roland. “A strong direct product theorem for quantum query complexity”. In: *computational complexity* 22 (2013), pp. 429–462 (cit. on pp. 5–7).
- [MM16] Mohammad Mahmoody and Ameer Mohammed. “On the power of hierarchical identity-based encryption”. In: *Advances in Cryptology—EUROCRYPT 2016: 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8–12, 2016, Proceedings, Part II* 35. Springer. 2016, pp. 243–272 (cit. on p. 2).
- [MT79] Robert Morris and Ken Thompson. “Password security: A case history”. In: *Communications of the ACM* 22.11 (1979), pp. 594–597 (cit. on p. 2).

- [MW05] Chris Marriott and John Watrous. “Quantum arthur–merlin games”. In: *computational complexity* 14.2 (2005), pp. 122–152 (cit. on p. 9).
- [NRS98] Noam Nisan, Steven Rudich, and Michael Saks. “Products and help bits in decision trees”. In: *SIAM Journal on Computing* 28.3 (1998), pp. 1035–1050 (cit. on pp. 6, 7, 19).
- [Raz10] Ran Raz. “Parallel repetition of two prover games (invited survey)”. In: *2010 IEEE 25th Annual Conference on Computational Complexity*. IEEE. 2010, pp. 3–6 (cit. on p. 6).
- [Sha03] R Shaltiel. “Towards proving strong direct product theorems”. In: *Computational Complexity* 12.2-Jan (2003), pp. 1–22 (cit. on pp. 6, 7, 20).
- [She11] Alexander A Sherstov. “Strong direct product theorems for quantum communication and query complexity”. In: *Proceedings of the forty-third annual ACM symposium on Theory of computing*. 2011, pp. 41–50 (cit. on pp. 5–7).
- [Unr07] Dominique Unruh. “Random Oracles and Auxiliary Input”. In: *Advances in Cryptology - CRYPTO 2007, 27th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2007, Proceedings*. Vol. 4622. 2007, pp. 205–223. DOI: [10.1007/978-3-540-74143-5_12](https://doi.org/10.1007/978-3-540-74143-5_12) (cit. on pp. 2, 4, 8).
- [Yao82] Andrew C Yao. “Theory and application of trapdoor functions”. In: *23rd Annual Symposium on Foundations of Computer Science (SFCS 1982)*. IEEE. 1982, pp. 80–91 (cit. on p. 6).
- [Yao90] AC-C Yao. “Coherent functions and program checkers”. In: *Proceedings of the twenty-second annual ACM symposium on Theory of computing*. 1990, pp. 84–94 (cit. on p. 2).
- [Zha19] Mark Zhandry. “How to record quantum queries, and applications to quantum indistinguishability”. In: *Annual International Cryptology Conference*. Springer. 2019, pp. 239–268 (cit. on pp. 5, 10, 11, 26–29, 31, 39).

A Missing Proofs in Section 3

Proof of Lemma 3.8. We use standard martingale analysis. For each $i = 0, 1, \dots, L$, let ℓ_i be the number of distinct elements in k_1, \dots, k_i . Then $\ell_0 = 0$ and $\ell_L = \ell$. Since each ℓ_i is at most L , we have

$$\mathbb{E} \left[c^{\ell_i} \right] = \mathbb{E} \left[\mathbb{E} \left[c^{\ell_i} \mid k_1, \dots, k_{i-1} \right] \right] = \mathbb{E} \left[c^{\ell_{i-1}} \cdot \left(\frac{\ell_{i-1}}{K} + c \cdot \left(1 - \frac{\ell_{i-1}}{K} \right) \right) \right] \leq \left(c + \frac{L}{K} \right) \cdot \mathbb{E} \left[c^{\ell_{i-1}} \right].$$

The desired bound follows by applying the above inequality L times. □

B Missing Proofs in Section 5.6

B.1 Proof of Theorem 5.21

Proof. We first apply the result in [Liu23] to relate the winning probability of a game against non-uniform quantum algorithms in QROM with that in the quantum bit-fixing (QBF) model. Then we analyze the maximum winning probability in the QBF model by connecting it to another game in the QROM.

We first recall the QBF model. The definition is adapted from [Liu23].

Definition B.1 (Games in the ϱ -BF-QROM). Games in the ϱ -BF-QROM are similar to games in the standard QROM, except now the random function f has a different distribution.

- Before a game starts, a ϱ -query quantum algorithm \mathcal{B} (having no input) is picked and fixed by an adversary.
- **Rejection Sampling Stage.** A random oracle f is picked uniformly at random conditioned on \mathcal{B}^f outputs 0. In other words, the distribution of f is defined by a rejection sampling:
 1. $f \sim \mu$, where μ is the uniform distribution over $\{f : [KM] \rightarrow [N]\}$.
 2. Run \mathcal{B}^f and output a binary outcome b together with a quantum state $|\sigma\rangle$.
 3. Restart from step 1 if $b \neq 0$.
- **Online Stage.** The game is then executed with oracle access to f , and an (online) algorithm \mathcal{A} starts with the state $|\sigma\rangle$.

Furthermore, they showed that we only need to care in the case when $\Pr[\mathbf{W}]$ is not very small (at least $1/N^S$), where \mathbf{W} is defined as an event for $b = 0$. Otherwise, then the RHS in the inequality [Liu23, Lemma 6.5] is already bounded by $1/N^S < \delta^S$, for any δ that we are interested in the theorem statement.

We directly apply the result in [Liu23] to go from non-uniform quantum algorithms to algorithms in ϱ -BF-QROM.

Theorem B.2 ([Liu23, Theorem 6.1]). Let G be any game with $T_{\text{Samp}}, T_{\text{Verify}}$ being the number of queries made by Samp and Verify defined in [Liu23, Definition 4.3]. For any S, T , let $\varrho = S \cdot (T + T_{\text{Verify}} + T_{\text{Samp}})$.

If G has security $\mathcal{V}(\varrho, T)$ in the ϱ -BF-QROM, then it has maximum winning probability $\varepsilon_G^*(S, T) \leq 2 \cdot \mathcal{V}(\varrho, T)$ against (S, T) non-uniform quantum algorithms with quantum advice.

In the game $G_{R,K}$, we have $T_{\text{Samp}} = 0$ (since there is no challenge) and $T_{\text{Verify}} = c$ (recall c is the size of interested tuples). Since we assume $T \geq c$ in the theorem statement, $\varrho = S \cdot (T + c) \leq 2ST$. Theorem B.2 gives us a way to upper bound $\varepsilon_G^*(S, T)$ by upper bounding $\mathcal{V}(2ST, T)$.

Thus, we are interested in the probability,

$$\mathcal{V}(2ST, T) = \Pr_{\substack{f \sim \mu \\ k \sim [K]}} \left[\mathcal{A}^f(k) \text{ wins} \mid \mathbf{W} \right], \quad (11)$$

where $k \sim [K]$ is a uniformly sampled salt, for any $\Pr[\mathbf{W}] \geq 1/N^S$. This corresponds to the probability that an algorithm first runs \mathcal{B} by making $2ST$ queries, and conditioned that \mathcal{B}^f outputs 0, it receives a uniformly sampled salt k and runs $\mathcal{A}^f(k)$. We can use the compressed oracle technique to simulate the random oracle due to their perfect indistinguishability, where we will analyze the joint state of the algorithm's registers and the database/oracle registers. \square

B.2 Proof of Claim 5.22

Proof of Claim 5.22. This is by the following inequality: for every event \mathbf{A}, \mathbf{B} and \mathbf{C} that occur with probability > 0 , $\Pr[\mathbf{B} \wedge \mathbf{C}] > 0$,

$$\begin{aligned} \Pr[\mathbf{A}|\mathbf{B}] &= \frac{\Pr[\mathbf{A} \wedge \mathbf{B}]}{\Pr[\mathbf{B}]} = \frac{\Pr[\mathbf{A} \wedge \mathbf{B} \wedge \mathbf{C}]}{\Pr[\mathbf{B}]} + \frac{\Pr[\mathbf{A} \wedge \mathbf{B} \wedge \overline{\mathbf{C}}]}{\Pr[\mathbf{B}]} \\ &\leq \frac{\Pr[\mathbf{A} \wedge \mathbf{B} \wedge \mathbf{C}]}{\Pr[\mathbf{B} \wedge \mathbf{C}]} + \frac{\Pr[\mathbf{B} \wedge \overline{\mathbf{C}}]}{\Pr[\mathbf{B}]} \end{aligned}$$

$$\leq \Pr[\mathbf{A}|\mathbf{B} \wedge \mathbf{C}] + \frac{\Pr[\overline{\mathbf{C}}]}{\Pr[\mathbf{B}]}.$$

Letting \mathbf{A} be the event that $\mathcal{A}^f(k)$ wins, $\mathbf{B} = \mathbf{W}$, and $\mathbf{C} = \mathbf{E}$, we can obtain [Claim 5.22](#). \square

B.3 Proof of Claim 5.23

Proof of Claim 5.23. Note that the conditioning on $\mathbf{W} \wedge \mathbf{E}$ only projects the algorithm's initial state $|\sigma\rangle$ to the space where $\mathbf{W} \wedge \mathbf{E}$ holds, and this does not affect the correctness of the analysis in [\[Zha19, Lemma 5\]](#). That is, for $c = 1$, we can still write the final state for $\mathcal{A}^f(k)$ conditioned on $\mathbf{W} \wedge \mathbf{E}$ in the form of

$$\sum_{x,u,z,D} \alpha_{x,u,z,D,0} |x, u, z, D\rangle + \sum_{r \neq 0^n} \alpha_{x,u,z,D,r} \frac{1}{\sqrt{N}} \sum_y \omega_N^{y \cdot r} |x, u, z, D \cup (x, y)\rangle,$$

with the compressed oracle. The generalization for $c \geq 1$ in [\[Zha19, Lemma 5\]](#) also holds. Therefore,

$$\begin{aligned} \sqrt{\Pr_{\substack{f \sim \mu \\ k \sim [K]}} [\mathcal{A}^f(k) \text{ wins} | \mathbf{W} \wedge \mathbf{E}]} &\leq \sqrt{\Pr_{\substack{f \sim \mu \\ k \sim [K]}} [(D|_k \in P) | \mathbf{W} \wedge \mathbf{E}]} + \sqrt{\frac{c}{N}} \\ &\leq \sqrt{2 \cdot \left(\Pr_{\substack{f \sim \mu \\ k \sim [K]}} [(D|_k \in P) | \mathbf{W} \wedge \mathbf{E}] + \frac{c}{N} \right)} \\ &\quad \text{(since } \sqrt{a} + \sqrt{b} \leq \sqrt{2a + 2b}\text{)} \end{aligned}$$

as desired. \square

B.4 Proof of Claim 5.24

Proof of Claim 5.24. For an algorithm that first succeeds on \mathbf{W} and \mathbf{E} , and then succeeds on salt k after T queries, we can denote its final state as:

$$|\phi\rangle = \Gamma_k^P \cdot cO \cdot U_T \cdots cO \cdot U_2 \cdot cO \cdot U_1 |\sigma\rangle.$$

Here Γ_k^P is the projection on the database register that $D|_k \in P$. $|\sigma\rangle$ is the state after projecting to the space where \mathcal{B}^f outputs 0 and \mathbf{E} happens, and thus every superposition of $|\sigma\rangle$ succeeds on at most $2S \log N$ salts. Here the algorithm $\mathcal{A}^f(k)$ starts from $|\sigma\rangle$ and makes T oracle queries.

With $|\phi\rangle$, we can have

$$\Pr_{\substack{f \sim \mu \\ k \sim [K]}} [(D|_k \in P) | \mathbf{W} \wedge \mathbf{E}] = \mathbb{E}_{k \sim [K]} [\| |\phi\rangle \|^2],$$

and now our goal is to upper bound the expected norm of $|\phi\rangle$.

We define $\Gamma_k^{>T}$ as the projection on the database register satisfying that the size of $D|_k$ is larger than T and that $D|_k \in \bar{P}$. Define $\bar{\Gamma} = \mathbb{I} - \Gamma_k^P - \Gamma_k^{>T}$, which is the projection onto the remaining cases. Now we can write $|\phi\rangle$ as the following:

$$|\phi\rangle = \Gamma_k^P \cdot cO \cdot U_T \cdots cO \cdot U_1 |\sigma\rangle$$

$$\begin{aligned}
&= \Gamma_k^P \cdot \text{cO} \cdot U_T \cdots \text{cO} \cdot U_1 \cdot (\Gamma_k^P + \Gamma_k^{>T} + \bar{\Gamma}) |\sigma\rangle \\
&= |\phi_1\rangle + |\phi_2\rangle + |\phi_3\rangle.
\end{aligned}$$

Here $|\phi_1\rangle$ is the part of $|\phi\rangle$ that starts from $\Gamma_k^P |\sigma\rangle$, $|\phi_2\rangle$ is the one that starts from $\Gamma_k^{>T} |\sigma\rangle$, and $|\phi_3\rangle$ is the one that starts from $\bar{\Gamma} |\sigma\rangle$.

Then we can upper bound expected $\| |\phi\rangle \|^2$ by upper bounding expected $\| |\phi_i\rangle \|^2$, $i = 1, 2, 3$:

$$\mathbb{E}_{k \sim [K]} [\| |\phi\rangle \|^2] \leq \mathbb{E}_{k \sim [K]} [3 \cdot (\| |\phi_1\rangle \|^2 + \| |\phi_2\rangle \|^2 + \| |\phi_3\rangle \|^2)].$$

For $|\phi_1\rangle$. Note that every superposition of $|\sigma\rangle$ succeeds on at most $2S \log N$ salts,

$$\begin{aligned}
\mathbb{E}_{k \sim [K]} [\| |\phi_1\rangle \|^2] &= \mathbb{E}_{k \sim [K]} [\| \Gamma_k^P \cdot \text{cO} \cdot U_T \cdots \text{cO} \cdot U_1 \cdot \Gamma_k^P |\sigma\rangle \|^2] \\
&\leq \mathbb{E}_{k \sim [K]} [\| \Gamma_k^P |\sigma\rangle \|^2] = \frac{1}{K} \cdot \sum_{k \in [K]} \| \Gamma_k^P |\sigma\rangle \|^2 \\
&\leq \frac{2S \log N}{K} = \tilde{O}\left(\frac{S}{K}\right). \tag{12}
\end{aligned}$$

For $|\phi_2\rangle$. Note that $|\sigma\rangle$ is obtained by \mathcal{B}^f making $2ST$ queries. Thus for every superposition in $|\sigma\rangle$, the database D can have at most $2ST/(T+1) = O(S)$ salts i such that the size of $D|_i$ is larger than T . Therefore,

$$\begin{aligned}
\mathbb{E}_{k \sim [K]} [\| |\phi_2\rangle \|^2] &= \mathbb{E}_{k \sim [K]} [\| \Gamma_k^P \cdot \text{cO} \cdot U_T \cdots \text{cO} \cdot U_1 \cdot \Gamma_k^{>T} |\sigma\rangle \|^2] \\
&\leq \mathbb{E}_{k \sim [K]} [\| \Gamma_k^{>T} |\sigma\rangle \|^2] = \frac{1}{K} \cdot \sum_{k \in [K]} \| \Gamma_k^{>T} |\sigma\rangle \|^2 \\
&\leq \frac{2S}{K} = O\left(\frac{S}{K}\right). \tag{13}
\end{aligned}$$

For $|\phi_3\rangle$. $|\phi_3\rangle$ starts from $\bar{\Gamma} |\sigma\rangle$, where $\bar{\Gamma}$ means that at that time it does not succeed on salt k , and it does not contain more than T elements of salt k in the database. For every fixed k , we define a sequence of properties P_t ($j = T, T+1, \dots, 2T$) such that the complement of each property $\neg P_t$ is defined as $P \cup \text{SZ}_{\leq t}$ (similar to [CFHL21, Remark 5.8]), where $\text{SZ}_{\leq t}$ corresponds to the property that the size of the restricted database $D|_k$ is no larger than t . Here the upper bound $2T$ of t comes from the fact that $\bar{\Gamma} |\sigma\rangle$ starts with at most T elements and the algorithm makes at most T extra queries.

Now we apply [CFHL21, Lemma 5.6] and the maximum winning probability is bounded by

$$\llbracket \neg P_T \xrightarrow{T} P_{2T} \rrbracket \leq \sum_{t=T+1}^{2T} \llbracket \neg P_{t-1} \rightarrow P_t \rrbracket,$$

where $\llbracket P \rightarrow P' \rrbracket$ corresponds to the maximum transition capacity (quantum analogue of transition probability) to go from property P to P' after one query, and $\llbracket P \xrightarrow{T} P' \rrbracket$ is the case for T queries.

We use the following fact about an upper bound for the transition capacity $\llbracket \neg P_t \rightarrow P_{t+1} \rrbracket$. The fact is restated and proved in [Appendix D.3](#).

Fact B.3. Let P be a monotone property with transition probability p_t . For any joint state $|\psi\rangle$ of the algorithm's registers and the oracle registers, define $\Gamma_k^{\leq t}$ as the projection that projects every superposition $|i, x, u, z\rangle |D\rangle$ to the states such that $i = k$ and the size of the database $D|_k$ is at most t . Define Γ_k^P as the projection that projects to $D|_k \in P$. Then we can have the following result for the decrease of norm during the transition from $D|_k \in \bar{P}$ to $D|_k \in P$:

$$\left\| \Gamma_k^P \cdot \text{CPhsO} \cdot (\mathbb{I} - \Gamma_k^P) \Gamma_k^{\leq t} |\psi\rangle \right\| \leq \sqrt{8 \cdot p_t} \left\| (\mathbb{I} - \Gamma_k^P) \Gamma_k^{\leq t} |\psi\rangle \right\|.$$

This implies $\llbracket -P_t \rightarrow P_{t+1} \rrbracket \leq \sqrt{8p_t}$, and gives us the following upper bound of the maximum winning probability for every $k \in [K]$:

$$O(\sqrt{p_T} + \sqrt{p_{T+1}} + \cdots + \sqrt{p_{2T-1}})^2 \leq O(\gamma_{2T}), \quad (14)$$

where we used the definition of γ_T from [Theorem 5.10](#).

Therefore, from [\(12\)](#), [\(13\)](#), and [\(14\)](#),

$$\mathbb{E}_{k \sim [K]} [\|\phi\|^2] \leq 3 \cdot \mathbb{E}_{k \sim [K]} [\|\phi_1\|^2 + \|\phi_2\|^2 + \|\phi_3\|^2] \leq O(\gamma_{2T}) + \tilde{O}\left(\frac{S}{K}\right). \quad \square$$

B.5 Proof of Claim 5.25

Proof of Claim 5.25. Without loss of generality, we assume $\gamma_T = o(1)$. By letting $B = 2ST$ and $k = 2S \log N$ in [Theorem 5.19](#), we can upper bound $\Pr_{f \sim \mu}[\bar{\mathbf{E}}] \leq (O(1) \cdot \gamma_{T/\log N})^{2S \log N} \leq (1/2)^{2S \log N} = (1/N)^{2S}$. As $\Pr_{f \sim \mu}[\mathbf{W}] \geq 1/N^S$, we now obtain

$$\Pr_{f \sim \mu}[\bar{\mathbf{E}}] \leq \frac{\Pr_{f \sim \mu}[\mathbf{W}]}{N^S} \leq \frac{\Pr_{f \sim \mu}[\mathbf{W}]}{N}. \quad \square$$

B.6 Proof of Theorem 5.13

Proof of Theorem 5.13. We put together the claims in [Section 5.6](#) and conclude [Theorem 5.13](#):

$$\begin{aligned} \Pr_{\substack{f \sim \mu \\ k \sim [K]}} [\mathcal{A}^f(k) \text{ wins } |\mathbf{W}] &\leq \Pr_{\substack{f \sim \mu \\ k \sim [K]}} [\mathcal{A}^f(k) \text{ wins } |\mathbf{W} \wedge \mathbf{E}] + \frac{\Pr_{f \sim \mu}[\bar{\mathbf{E}}]}{\Pr_{f \sim \mu}[\mathbf{W}]} && \text{(by Claim 5.22)} \\ &\leq 2 \cdot \Pr_{\substack{f \sim \mu \\ k \sim [K]}} [(D|_k \in P) | \mathbf{W} \wedge \mathbf{E}] + \frac{2c}{N} + \frac{\Pr_{f \sim \mu}[\bar{\mathbf{E}}]}{\Pr_{f \sim \mu}[\mathbf{W}]} && \text{(by Claim 5.23)} \\ &\leq O(\gamma_{2T}) + \tilde{O}\left(\frac{S}{K}\right) + \frac{2c}{N} + \frac{\Pr_{f \sim \mu}[\bar{\mathbf{E}}]}{\Pr_{f \sim \mu}[\mathbf{W}]} && \text{(by Claim 5.24)} \\ &\leq O(\gamma_{2T}) + \tilde{O}\left(\frac{S}{K}\right) + O\left(\frac{c}{N}\right). && \text{(by Claim 5.25)} \end{aligned}$$

According to [\(11\)](#), this gives an upper bound for $\mathcal{V}(2ST, T)$. Therefore, since γ_T is polynomial in T , by [Theorem B.2](#) we can obtain

$$\varepsilon_{G_{R,K}}^*(S, T) \leq \tilde{O}\left(\gamma_T + \frac{S}{K} + \frac{c}{N}\right). \quad \square$$

C Proof of SDPT for Quantum Property Finding

In this section, we present the proof of our strong direct product theorem [Theorem 5.18](#) for quantum property finding problem. This proof can also capture a threshold version ([Theorem 5.19](#)) as a corollary.

We first characterize the winning probability we need to compute for quantum direct product algorithms. Then we present how to use the idea of “path integral” to give an upper bound of this probability: split the final state into paths, and then give an upper bound of the maximum norm for all paths.

The Winning Probability of a Quantum Direct Product Algorithm. For quantum direct product of property finding problems, we model the result state of a quantum algorithm that makes kT queries to the oracle using the compressed oracle technique. We extend the algorithm’s first register from $|x\rangle$ to $|k, x\rangle$ with salt register $|k\rangle$, and the salt space here is $[\kappa]$. After κT queries, we will apply the binary-value measurement $\{\Lambda_{P^{\times\kappa}}, \mathbb{I} - \Lambda_{P^{\times\kappa}}\}$ on the oracle register to check in superposition whether the database is in $P^{\times\kappa}$:

$$\Lambda_{P^{\times\kappa}} = \sum_{D \in P^{\times\kappa}} |D\rangle \langle D|.$$

The algorithm wins if the measurement outcome of $\{\Lambda_{P^{\times\kappa}}, \mathbb{I} - \Lambda_{P^{\times\kappa}}\}$ is 0. Since the final local computation U_{final} after the κT -th query is independent of the oracle register, we can ignore that operator, and the winning probability for this algorithm can be written as

$$\Pr[\mathcal{A} \text{ wins } P^{\times\kappa}] = \|(\mathbb{I}_{\mathcal{A}} \otimes \Lambda_{P^{\times\kappa}}) \cdot \text{cO} \cdot U_{\kappa T} \cdot \text{cO} \cdot U_{\kappa T-1} \cdots \text{cO} \cdot U_1 |\psi_0\rangle\|^2.$$

We say that a state (or a projector) succeeds on salt $k \in [\kappa]$ if the state (or every state in the image of the projector) is in the image of $\sum_{D|k \in P} |D\rangle \langle D|$.

Here we fix this property P . We further define $\Lambda^{=r}$ to be the projection onto the states that can succeed on exactly r salts, and define $\Lambda^{\geq r}$ to be the projection that succeeds on at least r salts, i.e., $\Lambda^{\geq r'} = \sum_{r=r'}^{\kappa} \Lambda^{=r}$ for salt space $[\kappa]$. We also define Λ_k to be the projection onto the states that can succeed on salt k . Therefore, for salt space $[\kappa]$, $\Lambda^{\geq\kappa} = \Lambda^{=\kappa} = \prod_{k \in [\kappa]} \Lambda_k = \mathbb{I}_{\mathcal{A}} \otimes \Lambda_{P^{\times\kappa}}$.

We define the final winning state after the measurement as $|\psi_{\text{win}}\rangle$,

$$|\psi_{\text{win}}\rangle = \Lambda^{=\kappa} \cdot \text{cO} \cdot U_{\kappa T} \cdot \text{cO} \cdot U_{\kappa T-1} \cdots \text{cO} \cdot U_1 |\psi_0\rangle, \quad (15)$$

and the winning probability of a κT -query algorithm can be written as

$$\Pr[\mathcal{A} \text{ wins } P^{\times\kappa}] = \| |\psi_{\text{win}}\rangle \|^2. \quad (16)$$

To prove SDPT for quantum property finding, our goal is to give an upper bound for (16).

C.1 Splitting the Final Winning State

To give an upper bound for (16), we will first split the final state $|\psi_{\text{win}}\rangle$ into a sum of several “path” terms according to the “measurement results” after each oracle query.

We define \mathcal{P}_{κ} to be the set of all permutations of salts $[\kappa]$. From $|\psi_{\text{win}}\rangle$ as in (15), we define a set of $\{|\psi_{\mathbf{z}, \pi}\rangle\}$ based on the result after each oracle query. Here $\mathbf{z} = (z_1, z_2, \dots, z_{\kappa})$ satisfies

$1 \leq z_1 < z_2 < \dots < z_\kappa \leq \kappa T$ and $\pi = (\pi_1, \pi_2, \dots, \pi_\kappa) \in \mathcal{P}_\kappa$ is a permutation all salts $[\kappa]$. $|\psi_{\mathbf{z},\pi}\rangle$ is defined as follows:

$$|\psi_{\mathbf{z},\pi}\rangle = \left(\mathcal{E}_{\mathbf{z},\pi}^{(\kappa T)} \cdot \text{cO} \cdot U_{\kappa T} \right) \cdot \left(\mathcal{E}_{\mathbf{z},\pi}^{(\kappa T-1)} \cdot \text{cO} \cdot U_{\kappa T-1} \right) \cdots \left(\mathcal{E}_{\mathbf{z},\pi}^{(1)} \cdot \text{cO} \cdot U_1 \right) |\psi_0\rangle, \quad (17)$$

where $\mathcal{E}_{\mathbf{z},\pi}^{(t)}$ is a projection on the database register and it will be applied after the t -th oracle query. More precisely, $\mathcal{E}_{\mathbf{z},\pi}^{(t)}$ is a product of projections on each $D|_k$:

$$\mathcal{E}_{\mathbf{z},\pi}^{(t)} = \mathcal{E}_{\mathbf{z},\pi,1}^{(t)} \cdot \mathcal{E}_{\mathbf{z},\pi,2}^{(t)} \cdots \mathcal{E}_{\mathbf{z},\pi,\kappa}^{(t)},$$

where $\mathcal{E}_{\mathbf{z},\pi,k}^{(t)} = \begin{cases} \Lambda_k, & \text{if } t \geq z_r, \text{ for } r \in [\kappa] \text{ being the location of } k \text{ in } \pi, \text{ i.e., } \pi_r = k, \\ \mathbb{I} - \Lambda_k, & \text{if } t = z_r - 1, \text{ for } r \text{ as defined above,} \\ \mathbb{I}, & \text{if } t < z_r - 1, \text{ for } r \text{ as defined above.} \end{cases}$

Here we give an example for $\kappa = 2, T = 3$, with $\mathbf{z} = (2, 5)$ and $\pi = (\pi_1, \pi_2) = (1, 2)$:

$$\begin{aligned} |\psi_{\mathbf{z},\pi}\rangle &= (\Lambda_1 \Lambda_2 \cdot \text{cO} \cdot U_6) \cdot (\Lambda_1 \Lambda_2 \cdot \text{cO} \cdot U_5) \cdot (\Lambda_1 (\mathbb{I} - \Lambda_2) \cdot \text{cO} \cdot U_4) \\ &\quad \cdot (\Lambda_1 \cdot \text{cO} \cdot U_3) \cdot (\Lambda_1 \cdot \text{cO} \cdot U_2) \cdot ((\mathbb{I} - \Lambda_1) \cdot \text{cO} \cdot U_1) |\psi_0\rangle \end{aligned}$$

Intuitively, by fixing \mathbf{z} and π , we obtain $|\psi_{\mathbf{z},\pi}\rangle$ as a part of $|\psi_{\text{win}}\rangle$, where we fix the z_r -th query as the latest time when the state moves from non-success to success on salt $k = \pi_r$, and always succeeds on salt π_r afterwards, for every $r \in [\kappa]$.

This idea comes from a way of partitioning the winning event in the classical case. For classical direct product of property finding problem, the algorithm has a database D corresponding to the partial assignments it learns. It starts from an empty database, and appends at most one element (x, y) to D after one oracle query. It wins if its final database succeeds on every salt $k \in [\kappa]$ for property P .

- Following the evolution of the algorithm's database when we actually run an algorithm, we can record the number of queries z_r when it succeeds on r salts, and at that time which new salt $\pi_r \in [\kappa]$ it succeeds on. Note that the algorithm's database will never forget the elements it learns, $\pi = (\pi_1, \pi_2, \dots, \pi_\kappa)$ will form a permutation of $[\kappa]$, and $z_1 < z_2 < \dots < z_\kappa$ since the algorithm can only learn one element after one oracle query. These give us definitions of $\mathbf{z} = (z_1, z_2, \dots, z_\kappa) \in Z$ and $\pi = (\pi_1, \pi_2, \dots, \pi_\kappa) \in \mathcal{P}_\kappa$.
- The probability of winning the property finding game can be written as a sum of the winning probability with respect to every (\mathbf{z}, π) . In light of the symmetry of different salts, we can sum over all π and only consider the contribution with respect to \mathbf{z} . This corresponds to the probability of "first time succeeding on r salts after z_r -th query for each $r \in [\kappa]$ ".

When it comes to the quantum case, querying an element that is already in the database might lead to "forgetting", which means that each superposition of the database might not always append one element after one oracle query. This might result in the case that there are multiple times when the database moves from non-success to success on some salt. Note that the final winning state is in the image of $\Lambda^{\text{win}} = \prod_{k \in [\kappa]} \Lambda_k$. Therefore we can define a term with respect to (\mathbf{z}, π) , where we think of the z_r -th query as the latest time when it moves from non-success to success on salt π_r after this oracle query, and always succeeds on salt π_r afterwards. This corresponds to the definition in (17).

Similarly to the classical case, we do not distinguish between permutations and define a term $|\psi_{\mathbf{z}}\rangle$ with respect to every \mathbf{z} as "a path" by considering all possible π 's.

Definition C.1 (Paths in $|\psi_{\text{win}}\rangle$). Let Z to be the set of all $\mathbf{z} = (z_1, z_2, \dots, z_\kappa)$ such that $1 \leq z_1 < z_2 < \dots < z_\kappa \leq \kappa T$. Recall (17). We define $|\psi_{\mathbf{z}}\rangle$ to be the path corresponding to $\mathbf{z} = (z_1, z_2, \dots, z_\kappa)$:

$$|\psi_{\mathbf{z}}\rangle = \sum_{\pi \in \mathcal{P}_\kappa} |\psi_{\mathbf{z}, \pi}\rangle.$$

Here we claim that all the paths form a perfect partition of $|\psi_{\text{win}}\rangle$ as formalized in [Claim C.2](#). The proof of [Claim C.2](#) can be found in [Appendix D.1](#).

Claim C.2. With the definition of \mathbf{z} , Z , and $|\psi_{\mathbf{z}}\rangle$ in [Definition C.1](#), we have

$$|\psi_{\text{win}}\rangle = \sum_{\mathbf{z} \in Z} |\psi_{\mathbf{z}}\rangle.$$

C.2 Maximum Norm of a Path

With the path splitting, [Claim C.2](#) gives us a way to upper bound the norm of $|\psi_{\text{win}}\rangle$ by considering the maximum norm of each path. This follows from the triangle inequality:

$$\| |\psi_{\text{win}}\rangle \| \leq \sum_{\mathbf{z} \in Z} \| |\psi_{\mathbf{z}}\rangle \| \leq |Z| \cdot \max_{\mathbf{z} \in Z} \| |\psi_{\mathbf{z}}\rangle \|.$$

For each path, we can give an upper bound of its maximum norm.

Claim C.3. Recall that p_t is the transition probability defined in [Definition 5.7](#). We have

$$\max_{\mathbf{z} \in Z} \| |\psi_{\mathbf{z}}\rangle \|^2 \leq 8^\kappa \cdot \max_{\substack{t_1 + \dots + t_\kappa = \kappa T \\ t_1, \dots, t_\kappa \in \mathbb{Z}_+}} \left\{ \prod_{i=1}^{\kappa} p_{t_i-1} \right\}.$$

To prove [Claim C.3](#), we analyze the evolution of $|\psi_{\mathbf{z}}\rangle$. A natural idea is to upper bound the norm of each $|\psi_{\mathbf{z}, \pi}\rangle$ for all $\pi \in \mathcal{P}_\kappa$ and make use of some orthogonality relation between different π 's. However, the issue is that we cannot trivially obtain even pairwise orthogonality in $\{ |\psi_{\mathbf{z}, \pi}\rangle \}_{\pi \in \mathcal{P}_\kappa}$ by looking at the support on computational basis. An example is that for 3-query final states $|\psi_{\mathbf{z}, (1,2,3)}\rangle$ and $|\psi_{\mathbf{z}, (2,1,3)}\rangle$, they can share the same database register, which means they are not orthogonal. Therefore, we need to make detailed analysis following the evolution of $|\psi_{\mathbf{z}, \pi}\rangle$ to explore more subtle orthogonality relation during the process of reaching the final state.

Now we fix some $\mathbf{z} = (z_1, z_2, \dots, z_\kappa)$ and analyze this path.

We define $\Lambda_{\mathcal{S}}$ to be the projection that succeeds on salts in \mathcal{S} , $\Lambda_{\mathcal{S}} = \prod_{k \in \mathcal{S}} \Lambda_k$. We also define τ_k as a projection on the query register that projects to queries on salt k . That is, $\tau_k |k, x, u, z, D\rangle = |k, x, u, z, D\rangle$ and $\tau_k |k', x, u, z, D\rangle = 0$ for $k' \neq k$.

For every $\pi \in \mathcal{P}_\kappa$, we first rewrite $|\psi_{\mathbf{z}, \pi}\rangle$. Note that $\mathcal{E}_{\mathbf{z}, \pi}^{(t)}$ only applies on the database register and U_t only applies on the query register (and the algorithm's local memory). Therefore $\mathcal{E}_{\mathbf{z}, \pi}^{(t)}$ and $U_{t'}$ commutes, and from (17) we can rewrite $|\psi_{\mathbf{z}, \pi}\rangle$ as

$$\begin{aligned} |\psi_{\mathbf{z}, \pi}\rangle &= \left(\mathcal{E}_{\mathbf{z}, \pi}^{(\kappa T)} \cdot \text{cO} \cdot U_{\kappa T} \right) \cdot \left(\mathcal{E}_{\mathbf{z}, \pi}^{(\kappa T-1)} \cdot \text{cO} \cdot U_{\kappa T-1} \right) \cdots \left(\mathcal{E}_{\mathbf{z}, \pi}^{(1)} \cdot \text{cO} \cdot U_1 \right) |\psi_0\rangle \quad (\text{defined in (17)}) \\ &= \zeta_{\mathbf{z}, \pi}^{(\kappa T)} \cdot \zeta_{\mathbf{z}, \pi}^{(\kappa T-1)} \cdots \zeta_{\mathbf{z}, \pi}^{(2)} \cdot \zeta_{\mathbf{z}, \pi}^{(1)} |\psi_0\rangle, \end{aligned} \quad (18)$$

$$\text{where } \zeta_{\mathbf{z},\pi}^{(t)} = \begin{cases} \text{cO} \cdot U_t, & \text{if } t < z_1, \\ \Lambda_{\{\pi_1, \dots, \pi_r\}} \cdot \text{cO} \cdot U_t, & \text{if } z_r < t < z_{r+1} \text{ for some } r \in [\kappa - 1], \\ \Lambda_{\{\pi_1, \dots, \pi_r\}} \cdot \text{cO} \cdot (\mathbb{I} - \Lambda_{\pi_r}) \cdot \tau_{\pi_r} \cdot U_t, & \text{if } t = z_r \text{ for some } r \in [\kappa]. \end{cases}$$

This is by applying the following for every $r \in [\kappa]$ of the $(z_r - 1)$ -th query, moving all $(\mathbb{I} - \Lambda_{\pi_r})$ factor from $\mathcal{E}_{\mathbf{z},\pi}^{(z_r-1)}$ to a projector of the z_r -th query,

$$\begin{aligned} \left(\mathcal{E}_{\mathbf{z},\pi}^{(z_r)} \cdot \text{cO} \cdot U_{z_r} \right) \cdot \left(\mathcal{E}_{\mathbf{z},\pi}^{(z_r-1)} \cdot \text{cO} \cdot U_{z_r-1} \right) &= \left(\mathcal{E}_{\mathbf{z},\pi}^{(z_r)} \cdot \text{cO} \cdot U_{z_r} \right) \cdot \left((\mathbb{I} - \Lambda_{\pi_r}) \Lambda_{\{\pi_1, \dots, \pi_{r-1}\}} \cdot \text{cO} \cdot U_{z_r-1} \right) \\ &= \left(\mathcal{E}_{\mathbf{z},\pi}^{(z_r)} \cdot \text{cO} \cdot (\mathbb{I} - \Lambda_{\pi_r}) \cdot U_{z_r} \right) \cdot \left(\Lambda_{\{\pi_1, \dots, \pi_{r-1}\}} \cdot \text{cO} \cdot U_{z_r-1} \right), \end{aligned}$$

and this gives us the definition of $\zeta_{\mathbf{z},\pi}^{(t)}$. As for every superposition of the queried salt, one oracle query can change the database only restricted to that salt, we know $\Lambda_k \cdot \text{cO} \cdot (\mathbb{I} - \Lambda_k) \tau_{k'} = 0$ for $k' \neq k$. For $t = z_r$, we can write $\Lambda_{\{\pi_1, \dots, \pi_r\}} \cdot \text{cO} \cdot (\mathbb{I} - \Lambda_{\pi_r}) \cdot U_t = \Lambda_{\{\pi_1, \dots, \pi_r\}} \cdot \text{cO} \cdot (\mathbb{I} - \Lambda_{\pi_r}) \cdot \tau_{\pi_r} \cdot U_t$.

Inside the definition of $\zeta_{\mathbf{z},\pi}^{(z_r)}$ is a transition pattern. Before the z_r -th oracle query, the database register of the state is projected to $(\mathbb{I} - \Lambda_{\pi_r})$, i.e., $D|_{\pi_r} \in \bar{P}$ and \bar{P} is the complement of the property P (from Definition 5.7). After the oracle query cO , it is projected on Λ_{π_r} , where $D|_{\pi_r} \in P$. If we know the size u of $D|_{\pi_r}$, it will fit with the notion of the transition probability p_u (recall Definition 5.7), and there might be a decrease on the norm with respect to this transition probability, from the intuition of “lazy sampling” of the compressed oracle.

We follow from this intuition and analyze the evolution of the process. We define $|\psi_{\mathbf{z},\pi}^t\rangle$ from (18) as applying only the first t operators $\zeta_{\mathbf{z},\pi}^{(\theta)}$ for $\theta \in [t]$. Note that for $z_r \leq t < z_{r+1}$, these operators will only depend on the first r terms (π_1, \dots, π_r) of π , and thus for π, π' that share the same (π_1, \dots, π_r) , they share the same $|\psi_{\mathbf{z},\pi}^t\rangle$. Let $\mathcal{P}_r(\kappa)$ be the set of all permutations of sets of size r chosen from salt space $[\kappa]$. In this way, we can redefine $|\psi_{\mathbf{z},\iota}^t\rangle$ ($z_r \leq t < z_{r+1}$) for $\iota \in \mathcal{P}_r(\kappa)$:

$$|\psi_{\mathbf{z},\iota}^t\rangle = \zeta_{\mathbf{z},\iota}^{(t)} \cdot \zeta_{\mathbf{z},\iota}^{(t-1)} \cdots \zeta_{\mathbf{z},\iota}^{(2)} \cdot \zeta_{\mathbf{z},\iota}^{(1)} |\psi_0\rangle. \quad (19)$$

Here $\zeta_{\mathbf{z},\iota}^{(\theta)}$ is still well-defined, as $\zeta_{\mathbf{z},\pi}^{(\theta)}$ only makes use of the first r elements of π , and now ι is of length r as a truncated permutation. By adding all possible $\iota \in \mathcal{P}_r(\kappa)$, we can similarly define $|\psi_{\mathbf{z}}^t\rangle$:

$$|\psi_{\mathbf{z}}^t\rangle = \sum_{\iota \in \mathcal{P}_r(\kappa)} |\psi_{\mathbf{z},\iota}^t\rangle.$$

In addition, we define $|\psi_{\mathbf{z},\mathcal{S}}^t\rangle$ for $\mathcal{S} \subset [\kappa]$ and $|\mathcal{S}| = r$ as the sum of $|\psi_{\mathbf{z},\iota}^t\rangle$ over all permutations ι of \mathcal{S} .

We will need the following notion of used elements as a proxy to analyze the progress the algorithm makes towards finding the property. It intuitively quantifies the number of queries for each winning salt, where an additional $t - |D|$ factor comes from the quantum power of forgetting.

Definition C.4 (Number of Used Elements). *For state $|\psi_{\mathbf{z}}^t\rangle$ ($z_r \leq t < z_{r+1}$ for some $r \in [\kappa - 1]$), we define the number of used elements for the database register of each superposition of $|\psi_{\mathbf{z},\iota}^t\rangle$ as*

$$t - |D| + \sum_{k \in \{\iota_1, \dots, \iota_r\}} |D_k|,$$

where $D_k = D|_k$ is the restriction of D to salt k .

That is, for each state $|\psi_{\mathbf{z},S}^t\rangle$ (which equals $\sum |\psi_{\mathbf{z},\iota}^t\rangle$ that sums over every permutation ι of S), all the used elements in a database include “empty” entries and elements with salt index in $S = \{\iota_1, \dots, \iota_r\}$. Here the number of “empty” entries correspond to the number of forgotten entries; that is, a database is “supposed” to have t entries after t queries, but it might have $|D| < t$ entries due to the ability of forgetting in QROM, and the number of the “empty” entries are $t - |D|$ in this case.

Definition C.5. We define Q_B as a projection on the database register onto the database states that contain B used elements, and $Q_{\leq B} = \sum_{j=0}^B Q_j$.

For $r \in [\kappa]$, $j, B \in \mathbb{N}$, $S \subseteq [\kappa]$, and $|S| = r$, we define

$$g_{r,S}^j = \|Q_j |\psi_{\mathbf{z},S}^{z_r}\rangle\|^2$$

as the square of the norm of $|\psi_{\mathbf{z},S}^{z_r}\rangle$ with exactly j used elements in the database.

Remark C.6. Here the definition of $g_{r,S}^j$ depends on \mathbf{z} . As we are fixing \mathbf{z} throughout this section, \mathbf{z} does not show up in the notation of $g_{r,S}^j$ for simplicity.

By defining $z_0 = 0$, we extend the definition to $g_{0,S}^B$ (for $S = \emptyset$). Thus $\sum_{j=0}^B g_{0,S}^j = g_{0,S}^0 = 1$.

Definition C.7. For $\ell \in [\kappa]$, $B \in \mathbb{N}$, and the transition probability p_t from [Definition 5.7](#), we define

$$P_\ell^B = \max_{\substack{t_1 + \dots + t_\ell = B \\ t_1, \dots, t_\ell \in \mathbb{Z}_+}} \left\{ \prod_{i=1}^{\ell} p_{t_i-1} \right\}.$$

With the above definitions, we have the following claim of the evolution of the norm.

Claim C.8. With the definition of $g_{r,S}^j$ and P_ℓ^B , for $r \in [\kappa]$ and any $\ell \in [r]$, for summation over all $S, S' \subseteq [\kappa]$ such that $|S| = r$ and $|S'| = r - \ell$,

$$\sum_S \sum_{j \leq B} g_{r,S}^j \leq 8^\ell \cdot \sum_{S'} \sum_{j=0}^B P_\ell^{B-j} \cdot g_{r-\ell,S'}^j.$$

As a corollary, by letting $\ell = r$, we have $\sum_S \sum_{j \leq B} g_{r,S}^j \leq 8^r \cdot P_r^B$.

The proof of [Claim C.8](#) can be found in [Appendix D.2](#). Now we can prove [Claim C.3](#) given [Claim C.8](#).

Proof of Claim C.3. Let $r = \kappa$, $\ell = \kappa$, and $B = \kappa T$. Note that for the direct product case, salt space is $[\kappa]$, and $|S| = \kappa$ means that $S = [\kappa]$. Thus for all $\mathbf{z} \in Z$, we can have

$$\begin{aligned} \|\psi_{\mathbf{z}}\|^2 &\leq \|\psi_{\mathbf{z}}^{z_\kappa}\|^2 = \sum_{j=0}^{\kappa T} \|Q_j |\psi_{\mathbf{z},S}^{z_\kappa}\rangle\|^2 = \sum_{j=0}^{\kappa T} g_{r,S}^j && \text{(by orthogonality on } j) \\ &\leq 8^\kappa \cdot P_\kappa^{\kappa T} && \text{(by Claim C.8)} \\ &= 8^\kappa \cdot \max_{\substack{t_1 + \dots + t_\kappa = \kappa T \\ t_1, \dots, t_\kappa \in \mathbb{Z}_+}} \left\{ \prod_{i=1}^{\kappa} p_{t_i-1} \right\}. && \square \end{aligned}$$

Corollary (Theorem 5.18 Restated). *With every path bounded by Claim C.3, if the transition probability p_t is polynomial in t ,¹¹ then with $\gamma_T = \Theta(T^2 \cdot p_T)$, we have a strong direct product for quantum property finding:*

$$\Pr[\mathcal{A} \text{ wins } P^{\times \kappa}] \leq O(1)^\kappa \cdot \gamma_T^\kappa.$$

Proof. The size of set Z is $|Z| = \binom{\kappa T}{\kappa} \leq (e \cdot T)^\kappa$, and thus

$$\begin{aligned} \Pr[\mathcal{A} \text{ wins } P^{\times \kappa}] &= \|\psi_{\text{win}}\|^2 \leq |Z|^2 \cdot \max_{\mathbf{z} \in Z} \|\psi_{\mathbf{z}}\|^2 \\ &\leq e^{2\kappa} \cdot T^{2\kappa} \cdot \max_{\substack{t_1 + \dots + t_\kappa = \kappa T \\ t_1, \dots, t_\kappa \in \mathbb{Z}_+}} \left\{ \prod_{i=1}^{\kappa} p_{t_i-1} \right\}. \end{aligned}$$

Recall that $p_t = \Theta(t^{c_1}/N^{c_2})$ is polynomial in t and thus $\gamma_t = \Theta(t^{2+c_1}/N^{c_2})$. Therefore

$$\begin{aligned} \Pr[\mathcal{A} \text{ wins } P^{\times \kappa}] &\leq O(1)^\kappa \cdot T^{2\kappa} \cdot \max_{\substack{t_1 + \dots + t_\kappa = \kappa T \\ t_1, \dots, t_\kappa \in \mathbb{Z}_+}} \left\{ \prod_{i=1}^{\kappa} p_{t_i-1} \right\} \\ &= O(1)^\kappa \cdot T^{2\kappa} \cdot \max_{\substack{t_1 + \dots + t_\kappa = \kappa T \\ t_1, \dots, t_\kappa \in \mathbb{Z}_+}} \left\{ \prod_{i=1}^{\kappa} \frac{t_i^{c_1}}{N^{c_2}} \right\} \\ &\leq O(1)^\kappa \cdot T^{2\kappa} \cdot (T^{c_1}/N^{c_2})^\kappa \\ &= O(1)^\kappa \cdot \gamma_T^\kappa \end{aligned}$$

as desired. □

C.3 Threshold SDPT

For the threshold version, we consider the case of finding properties for at least κ salts among all $K \geq \kappa$ salts, with $B \in \mathbb{Z}_+$ queries to the oracle. To this end, we extend (15)

$$|\psi_{\text{win}}\rangle = \Lambda^{\geq \kappa} \cdot \text{cO} \cdot U_B \cdot \text{cO} \cdot U_{B-1} \cdots \text{cO} \cdot U_1 |\psi_0\rangle.$$

By orthogonality, we project $|\psi_{\text{win}}\rangle$ onto different spaces where the database register succeeds on exactly $\kappa \leq r \leq K$ salts as $\Lambda^{\geq \kappa} = \sum_{r=\kappa}^K \Lambda^{=r}$, and

$$\|\psi_{\text{win}}\|^2 = \sum_{r=\kappa}^K \|\Lambda^{=r} \psi_{\text{win}}\|^2.$$

For every r , we apply Claim C.2 on $\Lambda^{=r} |\psi_{\text{win}}\rangle$ and obtain the triangle inequality:

$$\|\Lambda^{=r} \psi_{\text{win}}\| \leq \sum_{\mathbf{z} \in Z} \|\Lambda^{=r} \psi_{\mathbf{z}}\| \leq |Z| \cdot \max_{\mathbf{z} \in Z} \|\Lambda^{=r} \psi_{\mathbf{z}}\|.$$

Remark C.9. *Here Z is defined to be the set of all $\mathbf{z} = (z_1, \dots, z_r)$ such that $1 \leq z_1 < z_2 < \dots < z_r \leq B$. The definition of Z depends on both the number of winning salts r and the number of all queries B .*

¹¹Recall that this means $p_t = \Theta(t^{c_1}/N^{c_2})$ for some constants $c_1, c_2 \geq 0$.

We also have a different version of [Claim C.3](#):

Claim C.10. *With \mathbf{z} , Z , and $\{\Lambda^{=r} |\psi_{\mathbf{z}}\rangle\}_{\mathbf{z} \in Z}$ defined for this case, we have*

$$\max_{\mathbf{z} \in Z} \|\Lambda^{=r} |\psi_{\mathbf{z}}\rangle\|^2 \leq 8^r \cdot \max_{\substack{t_1 + \dots + t_r = B \\ t_1, \dots, t_r \in \mathbb{Z}_+}} \left\{ \prod_{i=1}^r p_{t_i-1} \right\},$$

where p_t is the transition probability defined in [Definition 5.7](#).

Proof. For every $\mathbf{z} \in Z$, we split the final state $\Lambda^{=r} |\psi_{\mathbf{z}}\rangle$ according to the set of salts the database register succeeds on. That is, for all $\mathcal{S} \subseteq [K]$ such that $|\mathcal{S}| = r$,

$$\begin{aligned} \|\Lambda^{=r} |\psi_{\mathbf{z}}\rangle\|^2 &= \sum_{\mathcal{S}} \|\Lambda^{=r} |\psi_{\mathbf{z}, \mathcal{S}}\rangle\|^2 && \text{(by orthogonality)} \\ &\leq \sum_{\mathcal{S}} \|\psi_{\mathbf{z}, \mathcal{S}}^{z_r}\rangle\|^2 = \sum_{\mathcal{S}} \sum_{j=0}^B \|Q_j |\psi_{\mathbf{z}, \mathcal{S}}^{z_r}\rangle\|^2 = \sum_{\mathcal{S}} \sum_{j=0}^B g_{r, \mathcal{S}}^j \\ &\leq 8^r \cdot \max_{\substack{t_1 + \dots + t_r = B \\ t_1, \dots, t_r \in \mathbb{Z}_+}} \left\{ \prod_{i=1}^r p_{t_i-1} \right\}, \end{aligned}$$

where the last inequality follows from [Claim C.8](#). □

We now show how to obtain the threshold SDPT [Theorem 5.19](#).

Corollary ([Theorem 5.19](#) Restated). *[Theorem 5.18](#) also holds for the threshold version. That is, for salt space $[K]$ and an algorithm \mathcal{A} that makes $B \in \mathbb{Z}_+$ queries, if $p_t = \Theta(t^{c_1}/N^{c_2})$ for some constants $c_1, c_2 \geq 0$, then with $\gamma_t = \Theta(t^2 \cdot p_t)$, the probability of succeeding on more than $\kappa \leq K$ salts will have the following upper bound:*

$$\Pr[\mathcal{A} \text{ wins } P^{(\geq \kappa)}] \leq O(1)^\kappa \cdot (\gamma_{B/\kappa})^\kappa,$$

where we recall $\gamma_{B/r}$ from [Remark 5.20](#) when B/r is fractional.

Proof. We can assume without loss of generality that

$$B/\kappa \leq O\left(N^{c_2/(c_1+2)}\right). \tag{20}$$

This is because if otherwise

$$\gamma_{B/\kappa} = \Theta\left(\frac{(B/\kappa)^{c_1+2}}{N^{c_2}}\right) \geq 1$$

and hence the statement holds trivially.

Now we proceed to proving [Theorem 5.19](#) given (20). Similar to [Appendix C.2](#),¹² we first have

$$\|\Lambda^{=r} |\psi_{\text{win}}\rangle\|^2 \leq |Z|^2 \cdot \max_{\mathbf{z} \in Z} \|\Lambda^{=r} |\psi_{\mathbf{z}}\rangle\|^2 \leq O(1)^r \cdot (\gamma_{B/r})^r \tag{21}$$

¹²Actually the analysis of [Appendix C.2](#) only holds for $r \leq B$. In the case of $r > B$, we have $\|\Lambda^r |\psi_{\text{win}}\rangle\| = 0$, which is trivially bounded by $O(1)^r \cdot (\gamma_{B/r})^r$. This is because now the database register cannot contain r elements, and thus no adversary can win the property finding game on r salts (recall [Definition 5.9](#) for quantum property finding).

$$\begin{aligned}
&\leq \Theta\left(\frac{(B/r)^{c_1+2}}{N^{c_2}}\right)^r \\
&\leq 2^{-r}. \tag{by (20)}
\end{aligned}$$

Then we sum over different r for $\|\psi_{\text{win}}\|^2$:

$$\begin{aligned}
\|\psi_{\text{win}}\|^2 &= \sum_{r=\kappa}^K \|\Lambda^r \psi_{\text{win}}\|^2 \\
&\leq 2 \cdot \|\Lambda^{\kappa} \psi_{\text{win}}\|^2 && \text{(since } \|\Lambda^r \psi_{\text{win}}\|^2 \leq 2^{-r} \text{ by above)} \\
&\leq O(1)^\kappa \cdot (\gamma_{B/\kappa})^\kappa && \text{(by (21))}
\end{aligned}$$

as desired. \square

D Missing Proofs in Appendix C

D.1 Proof of Claim C.2

Proof of Claim C.2. For $|\psi_{\text{win}}\rangle$ defined in (15) and $\mathbf{z}, Z, |\psi_{\mathbf{z}}\rangle$ defined in Definition C.1, we generalize Z to be the set of all $\mathbf{z} = (z_1, \dots, z_\kappa)$ such that $1 \leq z_1 < \dots < z_\kappa \leq B$, and we write $|\psi_{\text{win}}\rangle$ for κ salts and B queries as $|\psi_{\text{win}}^{\kappa, B}\rangle$. We also generalize the salt space to $[K]$, while $|\psi_{\text{win}}^{\kappa, B}\rangle$ is redefined as the final state where every superposition of database succeeds on *exactly* κ salts. We define $|\psi_{\mathbf{z}}^{\kappa, B}\rangle$ as in Definition C.1 correspondingly.

Now we prove that Claim C.2 holds for all $\kappa, B \in \mathbb{Z}_+$ with $\kappa \leq \min\{B, K\}$. In fact, we prove the following: for every $\kappa, B \in \mathbb{Z}_+$ with $\kappa \leq \min\{B, K\}$, and Z as the set of all (z_1, \dots, z_κ) with $1 \leq z_1 < \dots < z_\kappa \leq B$, we have

$$|\psi_{\text{win}}^{\kappa, B}\rangle = \sum_{\mathbf{z} \in Z} |\psi_{\mathbf{z}}^{\kappa, B}\rangle.$$

From this we can immediately obtain Claim C.2 by letting $\kappa = K$ and $B = \kappa T$.

We additionally define $|\psi_{\text{win}}^{\mathcal{S}, B}\rangle$ for $\mathcal{S} \subset [K]$, $|\mathcal{S}| = \kappa$ to be the part of $|\psi_{\text{win}}^{\kappa, B}\rangle$ that succeeds on salts *exactly* from \mathcal{S} . Thus $|\psi_{\text{win}}^{\kappa, B}\rangle = \sum_{\mathcal{S}} |\psi_{\text{win}}^{\mathcal{S}, B}\rangle$ for summation over $\mathcal{S} \subset [K]$ such that $|\mathcal{S}| = \kappa$. We only need to prove that for every $\mathcal{S} \subset [K]$ and $|\mathcal{S}| = \kappa$, $|\psi_{\text{win}}^{\mathcal{S}, B}\rangle$ also has this kind of splitting $\{|\psi_{\mathbf{z}}^{\mathcal{S}, B}\rangle\}_{\mathbf{z} \in Z}$. We prove it by induction on (κ, B) .

Base case of $(1, 1)$. We start from $(\kappa, B) = (1, 1)$. Without loss of generality, we can let $\mathcal{S} = \{1\}$. By definition, $|\psi_{\text{win}}^{\mathcal{S}, 1}\rangle = \Lambda_1 \cdot \text{cO} \cdot U_1 |\psi_0\rangle$. As $1 \leq z_1 \leq B$, there is only one possible path with $\mathbf{z} = (z_1)$, $z_1 = 1$, and $Z = \{\mathbf{z}\}$. Thus $|\psi_{\text{win}}^{\mathcal{S}, 1}\rangle = \sum_{\mathbf{z} \in Z} |\psi_{\mathbf{z}}^{\mathcal{S}, 1}\rangle$ for $(\kappa, B) = (1, 1)$.

From $(1, B)$ to $(1, B+1)$. Now we assume that it holds for $(1, B)$ and prove the case of $(1, B+1)$. Without loss of generality, we can assume $\mathcal{S} = [1]$. Now $Z = \{(t)\}_{t=1}^{B+1}$. Here we denote the state just after the t -th query as $|\psi^t\rangle$, and here $|\psi_{\text{win}}^{[1], B+1}\rangle = \Lambda_1 |\psi^{B+1}\rangle$. When it comes to the path, we have the following:

$$|\psi_{\text{win}}^{[1], B+1}\rangle = \Lambda_1 \cdot \text{cO} \cdot U_{B+1} (\Lambda_1 + \mathbb{I} - \Lambda_1) |\psi^B\rangle$$

$$\begin{aligned}
&= \Lambda_1 \cdot \text{cO} \cdot U_{B+1} |\psi_{\text{win}}^{[1],B}\rangle + (\Lambda_1 \cdot \text{cO} \cdot U_{B+1}) \cdot (\mathbb{I} - \Lambda_1) |\psi^B\rangle \\
&= \left(\sum_{\mathbf{z}'=(t), t \leq B} \Lambda_1 \cdot \text{cO} \cdot U_{B+1} |\psi_{\mathbf{z}'}^{[1],B}\rangle \right) + |\psi_{(B+1)}^{[1],B+1}\rangle \quad (\text{by induction hypothesis}) \\
&= \sum_{\mathbf{z} \in Z} |\psi_{\mathbf{z}}^{[1],B+1}\rangle.
\end{aligned}$$

Therefore we conclude that the case of $(1, B)$ holds for all $B \in \mathbb{Z}_+$.

From $(\kappa - 1, \kappa - 1)$ to (κ, κ) . Now we assume that the formula holds for $(\kappa - 1, \kappa - 1)$ and prove the case of (κ, κ) . For (κ, κ) and $\mathcal{S} \subset [K]$ of size κ , without loss of generality, we can let $\mathcal{S} = [\kappa]$. Now the set of path has only one element defined by $\mathbf{z} = (1, 2, \dots, \kappa)$, $Z = \{\mathbf{z}\}$, as we require $z_1 < z_2 < \dots < z_\kappa$. When it comes to the path, we have the following:

$$\begin{aligned}
|\psi_{\text{win}}^{[\kappa],\kappa}\rangle &= \Lambda^{=\kappa} \cdot \text{cO} \cdot U_\kappa (\Lambda^{=\kappa} + \mathbb{I} - \Lambda^{=\kappa}) |\psi^{\kappa-1}\rangle \\
&= \Lambda^{=\kappa} \cdot \text{cO} \cdot U_\kappa |\psi_{\text{win}}^{[\kappa],\kappa-1}\rangle + \sum_{k \in [\kappa]} \Lambda^{=\kappa} \cdot \text{cO} \cdot U_\kappa |\psi_{\text{win}}^{[\kappa] \setminus \{k\}, \kappa-1}\rangle \\
&= \sum_{k \in [\kappa]} \Lambda^{=\kappa} \cdot \text{cO} \cdot U_\kappa |\psi_{\text{win}}^{[\kappa] \setminus \{k\}, \kappa-1}\rangle \\
&= \sum_{k \in [\kappa]} \Lambda^{=\kappa} \cdot \text{cO} \cdot U_\kappa |\psi_{\mathbf{z}'}^{[\kappa] \setminus \{k\}, \kappa-1}\rangle = |\psi_{\mathbf{z}}^{[\kappa],\kappa}\rangle \quad (\mathbf{z}' = (1, \dots, \kappa - 1) \text{ by induction hypothesis}) \\
&= \sum_{\mathbf{z} \in Z} |\psi_{\mathbf{z}}^{[\kappa],\kappa}\rangle.
\end{aligned}$$

For the second line uses $\Lambda^{=\kappa} \cdot \text{cO} \cdot U_\kappa (\mathbb{I} - \Lambda^{=\kappa}) |\psi^{\kappa-1}\rangle = \sum_{k \in [\kappa]} \Lambda^{=\kappa} \cdot \text{cO} \cdot U_\kappa |\psi_{\text{win}}^{[\kappa] \setminus \{k\}, \kappa-1}\rangle$ as we must win exactly $\kappa - 1$ salts before the final query to win all of them. The third line is by the fact that we will need at least κ queries to satisfy the property for each salt in $[\kappa]$, and thus the first term in the second line is 0.

Therefore we can conclude that the case of (κ, κ) holds for all $\kappa \leq K$.

From $(\kappa - 1, B), (\kappa, B)$ to $(\kappa, B + 1)$. Now we assume that it holds for any (κ, b) that $2 \leq \kappa \leq K$ and all $b \leq B$, for some $K, B \in \mathbb{Z}_+$. We prove the case of $(\kappa, B + 1)$ for all $\kappa \leq K$. Without loss of generality, we let $\mathcal{S} = [\kappa]$.

In this case, $1 \leq z_1 < \dots < z_\kappa \leq B + 1$, we start with the value of z_κ . If $z_\kappa \leq B$, we can work with $Z_0 = \{\mathbf{z} : 1 \leq z_1 < \dots < z_\kappa \leq B\}$, and it is handled by the case of (κ, B) . If $z_\kappa = B + 1$, we can work with $Z_1 = \{\mathbf{z} : 1 \leq z_1 < \dots < z_{\kappa-1} \leq B\}$, and it is handled by the case of $(\kappa - 1, B)$. Note that $Z = \{\mathbf{z} : \mathbf{z} = \mathbf{z}_0 \text{ for } \mathbf{z}_0 \in Z_0 \text{ or } \mathbf{z} = \mathbf{z}_1 \cup \{B + 1\} \text{ for } \mathbf{z}_1 \in Z_1\}$. Thus Z is the disjoint union of Z_0, Z_1' , where $Z_1' = \{\mathbf{z} \cup \{B + 1\} : \mathbf{z} \in Z_1\}$.

When it comes to the paths, we can decompose $|\psi_{\text{win}}^{[\kappa],B+1}\rangle$ as follows:

$$\begin{aligned}
|\psi_{\text{win}}^{[\kappa],B+1}\rangle &= \Lambda^{=\kappa} \cdot \text{cO} \cdot U_{B+1} (\Lambda^{=\kappa} + \mathbb{I} - \Lambda^{=\kappa}) |\psi^B\rangle \\
&= \Lambda^{=\kappa} \cdot \text{cO} \cdot U_{B+1} |\psi_{\text{win}}^{[\kappa],B}\rangle + \sum_{k \in [\kappa]} \Lambda^{=\kappa} \cdot \text{cO} \cdot U_{B+1} |\psi_{\text{win}}^{[\kappa] \setminus \{k\}, B}\rangle
\end{aligned}$$

$$\begin{aligned}
&= \sum_{\mathbf{z} \in Z_0} \Lambda^{\kappa} \cdot c\mathcal{O} \cdot U_{B+1} |\psi_{\mathbf{z}}^{[\kappa], B}\rangle + \sum_{k \in [\kappa]} \sum_{\mathbf{z}' \in Z_1} \Lambda^{\kappa} \cdot c\mathcal{O} \cdot U_{B+1} |\psi_{\mathbf{z}'}^{[\kappa] \setminus \{k\}, B}\rangle \\
&= \sum_{\mathbf{z} \in Z_0 \cup Z_1'} |\psi_{\mathbf{z}}^{[\kappa], B+1}\rangle.
\end{aligned}$$

The second line is by the fact that we can only succeed on one more salt after one query, and thus $|\psi^B\rangle$ should succeed on at least $\kappa - 1$ salts. The third line is by the assumptions of (κ, B) and $(\kappa - 1, B)$.

Conclusion. By induction, we can conclude that the splitting of paths holds for all (κ, B) , $\kappa \leq \min\{B, K\}$ as desired. \square

D.2 Proof of Claim C.8

We first present the following lemma.

Lemma D.1. Assume a_1, \dots, a_n and b_1, \dots, b_n satisfy $\sum_{i=1}^k a_i \leq \sum_{i=1}^k b_i$ for all $k \in [n]$. Then for every non-increasing $c_1 \geq c_2 \geq \dots \geq c_n \geq 0$, we have

$$\sum_{i=1}^k c_i \cdot a_i \leq \sum_{i=1}^k c_i \cdot b_i.$$

Proof. We denote $A_k = \sum_{i=1}^k a_i$ and $B_k = \sum_{i=1}^k b_i$. Thus we have $A_k \leq B_k$ for all $k \in [n]$ and

$$\sum_{i=1}^k c_i \cdot a_i = c_k \cdot A_k + \sum_{i=1}^{k-1} (c_i - c_{i+1}) \cdot A_i \leq c_k \cdot B_k + \sum_{i=1}^{k-1} (c_i - c_{i+1}) \cdot B_i = \sum_{i=1}^k c_i \cdot b_i. \quad \square$$

Now we present the proof for Claim C.8.

Proof of Claim C.8. We extend Definition C.7 to $\ell = 0$ and define $P_0^B = 1$ for all $B \in \mathbb{N}$. The case of $r = 0$ is true by definition, as both sides are identical to $\sum_{j \leq B} g_{0, \emptyset}^j = g_{0, \emptyset}^0 = 1$.

For any $r \in \mathbb{Z}_+$, we prove the statement by induction on $\ell \in [r] \cup \{0\}$.

Base case $\ell = 0$. In this case, since $P_0^{B-j} = 1$, we have

$$\sum_S \sum_{j=0}^B g_{r, S}^j = \sum_S \sum_{j=0}^B P_0^{B-j} \cdot g_{r, S}^j.$$

From ℓ to $\ell + 1$. Now we assume that it holds for any value $\leq \ell$ and prove the case of $\ell + 1$. To go from ℓ to $\ell + 1$, we will at least need to analyze how much the norm is reduced when the database just succeeds on one more salt. Here we need to build some relations between $|\psi_{\mathbf{z}, S}^{z_r - \ell}\rangle$ and $|\psi_{\mathbf{z}, S'}^{z_r - \ell - 1}\rangle$. We consider the case for $|\psi_{\mathbf{z}, S}^{z_r}\rangle$ and $|\psi_{\mathbf{z}, S'}^{z_r - 1}\rangle$. We remark that here the size of S automatically corresponds to the index r of z_r in $|\psi_{\mathbf{z}, S}^{z_r}\rangle$. That is, $|S| = r$ and $|S'| = r - 1$.

We start from $|\psi_{\mathbf{z}, S}^{z_r}\rangle$. By definition, $g_{r, S}^j = \|Q_j |\psi_{\mathbf{z}, S}^{z_r}\rangle\|^2$. To analyze the transition pattern, we can further decompose $g_{r, S}^B$ into a sum of $g_{r, S, k}^{j, \nu}$ to indicate the number of elements ν on that

queried salt k in the database. Informally here j corresponds to the number of used elements that are not of the z_r -th queried salt k . Here we define Q^ν as the projection that projects every superposition $|k, x, u, z\rangle |D\rangle$ to the states such that the size of the database $D|_k$ (restricted on that salt k) is ν . Recall that τ_k is the projection on the query register that projects to queries on salt k and $Q_{j+\nu}$ projects to states that the database contain $j + \nu$ used elements (see [Definition C.5](#)). Then we define $g_{r,\mathcal{S},k}^{j,\nu}$ ($k \in \mathcal{S}$) as the following:

$$\begin{aligned} g_{r,\mathcal{S},k}^{j,\nu} &= \|\tau_k \cdot Q_{j+\nu} Q^\nu |\psi_{\mathbf{z},\mathcal{S}}^{z_r}\rangle\|^2 \\ &= \|\tau_k \cdot Q_{j+\nu} Q^\nu \cdot \Lambda_{\mathcal{S}} \cdot \text{cO}(\mathbb{I} - \Lambda_k) U_{z_r} |\psi_{\mathbf{z},\mathcal{S}\setminus\{k\}}^{z_r-1}\rangle\|^2 && \text{(by the definition of } |\psi_{\mathbf{z},\mathcal{S}}^{z_r}\rangle) \\ &= \|Q_{j+\nu} Q^\nu \Lambda_{\mathcal{S}} \cdot \text{cO} \cdot (\mathbb{I} - \Lambda_k) \tau_k \cdot U_{z_r} |\psi_{\mathbf{z},\mathcal{S}\setminus\{k\}}^{z_r-1}\rangle\|^2. \end{aligned} \quad (22)$$

For the last equality above, we note that $Q_{j+\nu} Q^\nu \Lambda_{\mathcal{S}}$ is still a projection on the database register, it commutes with τ_k , and cO commutes with τ_k . For fixed \mathcal{S} and by orthogonality, we have

$$g_{r,\mathcal{S}}^B = \sum_{k \in \mathcal{S}} \sum_{j=0}^B g_{r,\mathcal{S},k}^{j,B-j}.$$

Now we look at $(\mathbb{I} - \Lambda_k) \tau_k \cdot U_{z_r} |\psi_{\mathbf{z},\mathcal{S}\setminus\{k\}}^{z_r-1}\rangle$, which is the state just before the query that succeeds on a new salt k . We want to analyze the norm of that state.

Let $\tilde{Q}_{j+\nu}^\nu$ be the projection onto the the state that every superposition can contribute to falling in $Q_{j+\nu} Q^\nu$ after one oracle query cO . That is, $Q_{j+\nu} Q^\nu \cdot \text{cO} \cdot (\mathbb{I} - \tilde{Q}_{j+\nu}^\nu) |\psi\rangle = 0$ for any state $|\psi\rangle$. Then, for any $\mathcal{S}' \subset [\kappa]$ such that $|\mathcal{S}'| = r - 1$, we define

$$h_{r,\mathcal{S}',k}^{j,\nu} = \|\tilde{Q}_{j+\nu}^\nu \cdot \Lambda_{\mathcal{S}'} (\mathbb{I} - \Lambda_k) \cdot \tau_k \cdot U_{z_r} |\psi_{\mathbf{z},\mathcal{S}'}^{z_r-1}\rangle\|^2. \quad (23)$$

It corresponds to the norm of the state just before the z_r -th query, such that each of its superposition queries on salt k and can contribute to falling in $g_{r,\mathcal{S}' \cup \{k\},k}^{j,\nu}$.

Remark D.2. *In fact, for every superposition $|k, x, u, z\rangle |D\rangle$ that lies in the image of $\tilde{Q}_{j+\nu}^\nu$ in (23), it must belong to either one of these two cases (we can also define $\tilde{Q}_{j+\nu}^\nu$ in this way):*

- (1) $(k, x) \notin D$, the size of $D|_k$ is $\nu - 1$, and the number of used elements (with respect to \mathcal{S}' , as in (23)) in D is j .

This corresponds to the case when the z_r -th query is on an entry (k, x) that is not in the database $(k, x) \notin D$. The z_r -th query will append one element to the database, adding the size of $D|_k$ by 1 and the number of used elements by ν .

- (2) $(k, x) \in D$, the size of $D|_k$ is ν , and the number of used elements (with respect to \mathcal{S}') is $j - 1$.

This corresponds to the case when the z_r -th query is on an entry (k, x) that is already in the database $(k, x) \in D$. After the z_r -th query, as the state moves from $(\mathbb{I} - \Lambda_k)$ to Λ_k , $|D|$ will remain the same, the size of $D|_k$ remaining ν and the number of used elements increased by $\nu + 1$.

We can also obtain orthogonality for $\tilde{Q}_{j+\nu}^\nu$ for different j, ν (with the same \mathcal{S}') from the above property, by looking at their support. An immediate result is $\|\sum_{k \notin \mathcal{S}} \sum_{\nu \in \mathbb{N}} \sum_{j \in \mathbb{N}} \tau_k \cdot \tilde{Q}_{j+\nu}^\nu |\psi_{\mathbf{z},\mathcal{S}}\rangle\| \leq \|\psi_{\mathbf{z},\mathcal{S}}\|$ for any state $|\psi_{\mathbf{z},\mathcal{S}}\rangle$ with used elements with respect to \mathcal{S} .

We have the following two facts about $g_{r,\mathcal{S},k}^{j,\nu}$, $h_{r,\mathcal{S}',k}^{j,\nu}$, $g_{r-1,\mathcal{S}'}^j$, which gives us a way to go from the z_r -th query back to the z_{r-1} -th.

Fact D.3. *With the above definition, by the compressed oracle technique,*

$$g_{r,S,k}^{j,\nu} \leq 8 \cdot p_{\nu-1} \cdot h_{r,S \setminus \{k\},k}^{j,\nu}.$$

Fact D.3 formalizes the way to go from z_r to $z_r - 1$. Its proof follows from the analysis of the compressed oracle and is left in [Appendix D.3](#).

Fact D.4. *For $r \in [\kappa]$, $B \in \mathbb{N}$, and $S' \subset [\kappa]$ such that $|S'| = r - 1$,*

$$\sum_{j=0}^B \sum_{\nu \in \mathbb{N}} \sum_{k \notin S'} h_{r,S',k}^{j,\nu} \leq \sum_{j=0}^B g_{r-1,S'}^j.$$

Fact D.4 gives us a way to go from $z_r - 1$ back to z_{r-1} . The intuition is the following: for every superposition $|k, x, u, z\rangle |D\rangle$ in $|\psi_{r-1,S'}^{z_{r-1}}\rangle$, the number of the used elements in (each superposition of) D , which is with respect to the set S' , will not decrease during the evolution from the z_{r-1} -th query to the $(z_r - 1)$ -th query. This is because during the evolution for this path, we apply U_t , cO , and $\Lambda_{S'}$ on the state just after the z_{r-1} -th query. U_t performs local computation and it does not change the database register; $\Lambda_{S'}$ will be a projection on the computational basis and it does not reduce the number of used elements in a database. cO performs an oracle query, and it can only add, change, or delete an element. While adding and changing do not reduce the number of used elements, deleting does not reduce either, as empty entry is also considered as a used element. Therefore, for any superposition that is counted in $g_{r-1,S'}^j$, the number of used elements in the database register will be at least j after it applies the $(z_r - 1)$ -th query, and it will “fall into” $\{h_{r,S',k}^{j',\nu} \mid j' \geq j, \nu \in \mathbb{N}, k \notin S'\}$.

Proof of Fact D.4. Following from the non-decreasing property of the number of used elements with respect to S' ($S' \subset [\kappa]$, $|S'| = r - 1$), we can obtain $\|Q_{\leq B} |\psi_{z,S'}^{z_r-1}\rangle\| \leq \|Q_{\leq B} |\psi_{z,S'}^{z_{r-1}}\rangle\|$, and thus

$$\begin{aligned} \sum_{j=0}^B \sum_{\nu \in \mathbb{N}} \sum_{k \notin S'} h_{r,S',k}^{j,\nu} &= \sum_{j=0}^B \sum_{\nu \in \mathbb{N}} \sum_{k \notin S'} \|\tilde{Q}_{j+\nu}^\nu \cdot (\mathbb{I} - \Lambda_k) \tau_k \cdot U_{z_r} |\psi_{z,S'}^{z_r-1}\rangle\|^2 \\ &\leq \sum_{j=0}^B \sum_{\nu \in \mathbb{N}} \sum_{k \notin S'} \|\tilde{Q}_{j+\nu}^\nu \cdot \tau_k \cdot U_{z_r} |\psi_{z,S'}^{z_r-1}\rangle\|^2 = \sum_{j=0}^B \left\| \sum_{\nu \in \mathbb{N}} \sum_{k \notin S'} \tilde{Q}_{j+\nu}^\nu \cdot \tau_k \cdot U_{z_r} |\psi_{z,S'}^{z_r-1}\rangle \right\|^2 \\ &\leq \sum_{j=0}^B \|Q_j \cdot U_{z_r} |\psi_{z,S'}^{z_r-1}\rangle\|^2 = \|Q_{\leq B} |\psi_{z,S'}^{z_r-1}\rangle\|^2 \\ &\leq \|Q_{\leq B} |\psi_{z,S'}^{z_{r-1}}\rangle\|^2 = \sum_{j=0}^B g_{r-1,S'}^j. \quad \square \end{aligned}$$

With [Fact D.3](#) and [Fact D.4](#), we have the following:

$$\begin{aligned} \sum_S \sum_{j=0}^B P_\ell^{B-j} \cdot g_{r-\ell,S}^j &= \sum_S \sum_{j=0}^B P_\ell^{B-j} \cdot \left(\sum_{\nu=0}^j \sum_{k \in S} g_{r-\ell,S,k}^{\nu,j-\nu} \right) \\ &\leq \sum_S \sum_{j=0}^B P_\ell^{B-j} \cdot \left(\sum_{\nu=0}^j \sum_{k \in S} 8 \cdot p_{j-\nu-1} \cdot h_{r-\ell,S \setminus \{k\},k}^{\nu,j-\nu} \right) \quad (\text{by Fact D.3}) \end{aligned}$$

$$\begin{aligned}
&\leq 8 \cdot \sum_{S'} \sum_{k \notin S'} \sum_{j=0}^B \sum_{\nu=0}^j P_{\ell+1}^{B-\nu} \cdot h_{r-\ell, S', k}^{\nu, j-\nu} && \text{(by Definition C.7)} \\
&= 8 \cdot \sum_{S'} \sum_{\nu=0}^B P_{\ell+1}^{B-\nu} \cdot \left(\sum_{k \notin S'} \sum_{d=0}^{B-\nu} h_{r-\ell, S', k}^{\nu, d} \right). && \text{(by } d = j - \nu)
\end{aligned}$$

Here $|\mathcal{S}| = r - \ell$ and $|\mathcal{S}'| = r - \ell - 1$.

Note that $P_{\ell+1}^{B-\nu}$ is non-increasing on ν by the monotonicity of the transition probability p_t , and [Fact D.4](#) gives the inequality of two sums of arrays. Thus following [Lemma D.1](#), for every S' ($|\mathcal{S}'| = r - \ell - 1$), we can continue our calculation as:

$$\begin{aligned}
\sum_{\nu=0}^B P_{\ell+1}^{B-\nu} \cdot \left(\sum_{k \notin S'} \sum_{d=0}^{B-\nu} h_{r-\ell, S', k}^{\nu, d} \right) &\leq \sum_{\nu=0}^B P_{\ell+1}^{B-\nu} \cdot \left(\sum_{k \notin S'} \sum_{d \in \mathbb{N}} h_{r-\ell, S', k}^{\nu, d} \right) \\
&\leq \sum_{\nu=0}^B P_{\ell+1}^{B-\nu} \cdot g_{r-\ell-1, S'}^{\nu}. && \text{(by Lemma D.1)}
\end{aligned}$$

Therefore

$$\sum_S \sum_{j=0}^B P_{\ell}^{B-j} \cdot g_{r-\ell, S}^j \leq 8 \cdot \sum_{S'} \sum_{\nu=0}^B P_{\ell+1}^{B-\nu} \cdot g_{r-\ell-1, S'}^{\nu}$$

giving us the induction step from ℓ to $\ell + 1$, with $|\mathcal{S}| = r$, $|\mathcal{S}'| = r - \ell$, and $|\mathcal{S}''| = r - \ell - 1$:

$$\begin{aligned}
\sum_S \sum_{j=0}^B g_{r, S}^j &\leq 8^{\ell} \cdot \sum_{S'} \sum_{j=0}^B P_{\ell}^{B-j} \cdot g_{r-\ell, S'}^j \\
&\leq 8^{\ell+1} \cdot \sum_{S''} \sum_{j=0}^B P_{\ell+1}^{B-j} \cdot g_{r-\ell-1, S''}^j.
\end{aligned}$$

Conclusion. Therefore, by induction, [Claim C.8](#) holds for all $\ell \in [r]$ (for all fixed $r \in [\kappa]$). \square

D.3 Proof of [Fact D.3](#)

Proof of [Fact D.3](#). In the following computation, we will use $\text{cO} = \text{CPhsO}$.

According to the definition of $g_{r, S, k}^{j, \nu}$ in [\(22\)](#) and $h_{r, S', k}^{j, \nu}$ in [\(23\)](#), we can write them as the following, as we define $S' = \mathcal{S} \setminus \{k\}$:

$$\begin{aligned}
g_{r, S, k}^{j, \nu} &= \|Q_{j+\nu} Q^{\nu} \Lambda_S \cdot \text{CPhsO} \cdot (\mathbb{I} - \Lambda_k) \tau_k \cdot U_{z_r} |\psi_{\mathbf{z}, S'}^{z_r-1}\rangle\|^2, \\
&= \|Q_{j+\nu} Q^{\nu} \Lambda_S \cdot \text{CPhsO} \cdot \tilde{Q}_{j+\nu}^{\nu} \cdot \Lambda_{S'} (\mathbb{I} - \Lambda_k) \cdot \tau_k \cdot U_{z_r} |\psi_{\mathbf{z}, S'}^{z_r-1}\rangle\|^2, \\
h_{r, S', k}^{j, \nu} &= \|\tilde{Q}_{j+\nu}^{\nu} \cdot \Lambda_{S'} (\mathbb{I} - \Lambda_k) \cdot \tau_k \cdot U_{z_r} |\psi_{\mathbf{z}, S'}^{z_r-1}\rangle\|^2.
\end{aligned}$$

Here Q^{ν} is defined in [\(22\)](#); it is the projection that projects every superposition $|k, x, u, z\rangle |D\rangle$ to the states such that the size of the restricted database $D|_k$ on salt k is ν . Here $\tilde{Q}_{j+\nu}^{\nu}$ is defined in [\(23\)](#) as the projection onto the state that every superposition can contribute to falling in $Q_{j+\nu} Q^{\nu}$

after one oracle query CPhsO. That is, $Q_{j+\nu}Q^\nu \cdot \text{CPhsO} \cdot (\mathbb{I} - \tilde{Q}_{j+\nu}^\nu) |\psi\rangle = 0$ for all possible $|\psi\rangle$, and $Q_{j+\nu}Q^\nu \cdot \text{CPhsO} \cdot \tilde{Q}_{j+\nu}^\nu \neq 0$.

We define that state with the norm of $h_{r,S',k}^{j,\nu}$ as $|\phi\rangle$, and it can be written as

$$\begin{aligned} |\phi\rangle &= \tilde{Q}_{j+\nu}^\nu \cdot \Lambda_{S'}(\mathbb{I} - \Lambda_k) \cdot \tau_k \cdot U_{z_r} |\psi_{\mathbf{z},S'}^{z_r-1}\rangle \\ &= \sum_{x,u,z,D} \alpha_{x,u,z,D} |k, x, u, z\rangle \otimes |D\rangle + \sum_{x,u,z,D,y} \beta_{x,u,z,D,y} |k, x, u, z\rangle \otimes |D \cup ((k, x), y)\rangle. \end{aligned}$$

Here the sum is over D such that $(k, x) \notin D$. And we also define $|\phi_\alpha\rangle$ and $|\phi_\beta\rangle$ as

$$\begin{aligned} |\phi_\alpha\rangle &= \sum_{x,u,z,D} \alpha_{x,u,z,D} |k, x, u, z\rangle \otimes |D\rangle, \\ |\phi_\beta\rangle &= \sum_{x,u,z,D,y} \beta_{x,u,z,D,y} |k, x, u, z\rangle \otimes |D \cup ((k, x), y)\rangle \\ &= \sum_{x,u,z,D} |k, x, u, z\rangle \otimes \left(\sum_y \beta_{x,u,z,D,y} |D \cup ((k, x), y)\rangle \right). \end{aligned}$$

By definition, $|\phi_\alpha\rangle$ and $|\phi_\beta\rangle$ are orthogonal.

Now we analyze how the norm decreases for $|\phi_\alpha\rangle$ and $|\phi_\beta\rangle$ when CPhsO and the projector $Q_{j+\nu}Q^\nu\Lambda_k$ are applied.

Case of $|\phi_\alpha\rangle$. For $|\phi_\alpha\rangle$ of which superposition is over D such that $(k, x) \notin D$,

$$\begin{aligned} \text{CPhsO} |\phi_\alpha\rangle &= \sum_{x,u,z,D} \alpha_{x,u,z,D} \cdot \left(\text{StdDecomp}_{(k,x)} \cdot \text{CPhsO}' \cdot \text{StdDecomp}_{(k,x)} |k, x, u, z\rangle \otimes |D\rangle \right) \\ &= \sum_{x,u,z,D} \alpha_{x,u,z,D} |k, x, u, z\rangle \otimes \left(\frac{1}{\sqrt{N}} \sum_y \omega_N^{uy} |D \cup ((k, x), y)\rangle \right). \end{aligned}$$

By definition, every superposition of $|\phi_\alpha\rangle$ is in the image of $\tilde{Q}_{j+\nu}^\nu$, and we do not care the part that will lie in the image of $\mathbb{I} - Q_{j+\nu}Q^\nu\Lambda_k$ after one oracle query. Thus we can assume that the size of $D \cup ((k, x), y)|_k$ is ν , and the size of $D|_k$ is $\nu - 1$, and

$$\begin{aligned} \|Q_{j+\nu}Q^\nu\Lambda_k\text{CPhsO} |\phi_\alpha\rangle\| &\leq \|\Lambda_k\text{CPhsO} |\phi_\alpha\rangle\| \\ &= \left\| \sum_{x,u,z,D} \alpha_{x,u,z,D} |k, x, u, z\rangle \otimes \left(\frac{1}{\sqrt{N}} \sum_{\substack{y \text{ s.t.} \\ D \cup ((k, x), y) \in P}} \omega_N^{uy} |D \cup ((k, x), y)\rangle \right) \right\| \\ &\leq \sqrt{p_{\nu-1}} \cdot \left\| \sum_{x,u,z,D} \alpha_{x,u,z,D} |k, x, u, z\rangle \otimes \left(\frac{1}{\sqrt{N}} \sum_y \omega_N^{uy} |D \cup ((k, x), y)\rangle \right) \right\| \\ &= \sqrt{p_{\nu-1}} \cdot \|\text{CPhsO} |\phi_\alpha\rangle\|. \end{aligned}$$

Note that for all $y \in [N]$, the number of y such that $D \cup \{(k, x), y\}|_k \in P$ while $D|_k \in \bar{P}$ is no more than $N \cdot p_{\nu-1}$, with the size of $D|_k$ no more than $\nu - 1$. This is by the definition of the transition

probability $p_{\nu-1}$. Thus the norm of every superposition of $|k, x, u, z\rangle$ in $|\phi_\alpha\rangle$ is reduced to at most $\sqrt{p_{\nu-1}}$ times its original norm, which means,

$$\|Q_{j+\nu}Q^\nu\Lambda_k\text{CPhsO}|\phi_\alpha\rangle\|^2 \leq p_{\nu-1} \cdot \|\phi_\alpha\|^2. \quad (24)$$

Case of $|\phi_\beta\rangle$. For $|\phi_\beta\rangle$ of which superposition is over D such that $(k, x) \notin D$,

$$\begin{aligned} & \text{CPhsO}|\phi_\beta\rangle \\ &= \sum_{x,u,z,D,y} \beta_{x,u,z,D,y} \cdot \left(\text{StdDecomp}_{(k,x)} \cdot \text{CPhsO}' \cdot \text{StdDecomp}_{(k,x)} |k, x, u, z\rangle \otimes |D \cup ((k, x), y)\rangle \right) \\ &= \sum_{x,u,z,D,y} \beta_{x,u,z,D,y} |k, x, u, z\rangle \otimes \\ & \quad \left(\omega_N^{uy} |D \cup ((k, x), y)\rangle + (1 + \omega_N^{uy}) |D\rangle - \frac{1}{N} \sum_{y'} (\omega_N^{uy} + \omega_N^{yy'}) |D \cup ((k, x), y')\rangle \right). \end{aligned}$$

By definition, every superposition of $|\phi_\beta\rangle$ is in the image of $\tilde{Q}_{j+\nu}^\nu$, and we do not care the part that will lie in the image of $\mathbb{I} - Q_{j+\nu}Q^\nu\Lambda_k$ after one oracle query. Since $D|_k \in \bar{P}$, the size of $D \cup ((k, x), y)|_k$ is ν , and as $D \cup ((k, x), y)|_k \in \bar{P}$,

$$\begin{aligned} & \|Q_{j+\nu}Q^\nu\Lambda_k\text{CPhsO}|\phi_\beta\rangle\| \leq \|\Lambda_k\text{CPhsO}|\phi_\beta\rangle\| \\ &= \left\| \sum_{x,u,z,D,y} \beta_{x,u,z,D,y} |k, x, u, z\rangle \otimes \left(\frac{1}{N} \sum_{\substack{y' \text{ s.t.} \\ D \cup ((k, x), y') \in P}} (\omega_N^{uy} + \omega_N^{yy'}) |D \cup ((k, x), y')\rangle \right) \right\| \\ &= \left\| \sum_{x,u,z,D} |k, x, u, z\rangle \otimes \left(\sum_{\substack{y' \text{ s.t.} \\ D \cup ((k, x), y') \in P}} \sum_y \frac{\beta_{x,u,z,D,y}}{N} \cdot (\omega_N^{uy} + \omega_N^{yy'}) |D \cup ((k, x), y')\rangle \right) \right\|. \quad (25) \end{aligned}$$

For every superposition of $|k, x, u, z\rangle$ in $|\phi_\beta\rangle$, the norm squared is $\sum_y |\beta_{x,u,z,D,y}|^2$. Note that for all $y \in [N]$, the number of y such that $D \cup \{(k, x), y\}|_k \in P$ while $D|_k \in \bar{P}$ is no more than $N \cdot p_{\nu-1}$, with the size of $D|_k$ no more than $\nu - 1$. This is by the definition of the transition probability $p_{\nu-1}$. Therefore, for every superposition of $|k, x, u, z\rangle$ in (25), the norm squared is

$$\begin{aligned} & \sum_{\substack{y' \text{ s.t.} \\ D \cup ((k, x), y') \in P}} \left| \sum_y \frac{\beta_{x,u,z,D,y}}{N} \cdot (\omega_N^{uy} + \omega_N^{yy'}) \right|^2 \leq \frac{p_{\nu-1}}{N} \cdot \max_{y' \in [N]} \left| \sum_y \beta_{x,u,z,D,y} \cdot (\omega_N^{uy} + \omega_N^{yy'}) \right|^2 \\ & \leq \frac{p_{\nu-1}}{N} \cdot \left(\sum_y |\beta_{x,u,z,D,y}|^2 \right) \cdot \max_{y' \in [N]} \left(\sum_y |\omega_N^{uy} + \omega_N^{yy'}|^2 \right) \\ & \leq \frac{p_{\nu-1}}{N} \cdot \left(\sum_y |\beta_{x,u,z,D,y}|^2 \right) \cdot 4N \end{aligned}$$

$$= 4p_{\nu-1} \cdot \left(\sum_y |\beta_{x,u,z,D,y}|^2 \right).$$

The second line is by Cauchy–Schwarz inequality $|\sum_i a_i b_i|^2 \leq (\sum_i |a_i|^2)(\sum_i |b_i|^2)$. Therefore, we can conclude from (25) that

$$\|Q_{j+\nu} Q^\nu \Lambda_k \text{CPhsO} |\phi_\beta\rangle\|^2 \leq 4p_{\nu-1} \cdot \|\phi_\beta\|^2. \quad (26)$$

Conclusion. Note that for every $|\phi_\alpha\rangle$ and $|\phi_\beta\rangle, \forall k \in [\kappa], Q_{j+\nu} Q^\nu \Lambda_k \text{CPhsO}$ will not change their first register $|k\rangle$. Therefore, by (24), (26), and the orthogonality on the first register,

$$\begin{aligned} g_{r,S,k}^{j,\nu} &= \|Q_{j+\nu} Q^\nu \Lambda_S \text{CPhsO} |\phi\rangle\|^2 \\ &\leq \|Q_{j+\nu} Q^\nu \Lambda_k \text{CPhsO} |\phi_\alpha\rangle + Q_{j+\nu} Q^\nu \Lambda_k \text{CPhsO} |\phi_\beta\rangle\|^2 \\ &\leq 2 \cdot \left(\|Q_{j+\nu} Q^\nu \Lambda_k \text{CPhsO} |\phi_\alpha\rangle\|^2 + \|Q_{j+\nu} Q^\nu \Lambda_k \text{CPhsO} |\phi_\beta\rangle\|^2 \right) \\ &\hspace{15em} \text{(since } \|a+b\|^2 \leq 2(\|a\|^2 + \|b\|^2)\text{)} \\ &\leq 8 \cdot p_{\nu-1} \cdot (\|\phi_\alpha\|^2 + \|\phi_\beta\|^2) \hspace{10em} \text{(by (24) and (26))} \\ &= 8 \cdot p_{\nu-1} \cdot \|\phi\|^2 = 8 \cdot p_{\nu-1} \cdot h_{r,S',k}^{j,\nu}. \hspace{10em} \text{(by orthogonality)} \end{aligned}$$

Thus [Fact D.3](#) holds. \square

Actually, within the analysis to prove [Fact D.3](#), we can have the following corollary that shows the transition pattern.

Fact (Fact B.3 Restated). *Let P be a monotone property with transition probability p_t . For any joint state $|\psi\rangle$ of the algorithm's registers and the oracle registers, define $\Gamma_k^{\leq t}$ as the projection that projects every superposition $|i, x, u, z\rangle |D\rangle$ to the states such that $i = k$ and the size of the database $D|_k$ is less than t . Define Γ_k^P as the projection that projects to $D|_k \in P$. Then we can have the following result for the decrease of norm during the transition from $D|_k \in \bar{P}$ to $D|_k \in P$:*

$$\left\| \Gamma_k^P \cdot \text{CPhsO} \cdot (\mathbb{I} - \Gamma_k^P) \Gamma_k^{\leq t} |\psi\rangle \right\| \leq \sqrt{8 \cdot p_t} \left\| (\mathbb{I} - \Gamma_k^P) \Gamma_k^{\leq t} |\psi\rangle \right\|.$$

Proof Outline of Fact B.3. We similarly define $|\psi_{\alpha,k}\rangle$ to be the part where every superposition $|k, x, u, z\rangle |D\rangle$ satisfies $D(k, x) = \perp$, and define $|\psi_{\beta,k}\rangle$ to be the part $D(k, x) \neq \perp$. Note that every state in the image of $\Gamma_k^{\leq t}$ satisfies that the size of $D|_k$ is no more than t .

Now we can adapt the proof of (24) for every $(\mathbb{I} - \Gamma_k^P) \Gamma_k^{\leq t} |\psi_{\alpha,k}\rangle$, and obtain

$$\|\Gamma_k^P \cdot \text{CPhsO} \cdot (\mathbb{I} - \Gamma_k^P) \Gamma_k^{\leq t} |\psi_{\alpha,k}\rangle\|^2 \leq p_t \cdot \|(\mathbb{I} - \Gamma_k^P) \Gamma_k^{\leq t} |\psi_{\alpha,k}\rangle\|^2. \quad (27)$$

Similarly, we can adapt the proof of (26) for every $(\mathbb{I} - \Gamma_k^P) \Gamma_k^{\leq t} |\psi_{\beta,k}\rangle$ and obtain

$$\|\Gamma_k^P \cdot \text{CPhsO} \cdot (\mathbb{I} - \Gamma_k^P) \Gamma_k^{\leq t} |\psi_{\beta,k}\rangle\|^2 \leq 4p_t \cdot \|(\mathbb{I} - \Gamma_k^P) \Gamma_k^{\leq t} |\psi_{\beta,k}\rangle\|^2. \quad (28)$$

Since $\Gamma_k^{\leq t} (|\psi_{\alpha,k}\rangle + |\psi_{\beta,k}\rangle) = \Gamma_k^{\leq t} |\psi\rangle$, we can obtain

$$\|\Gamma_k^P \cdot \text{CPhsO} \cdot (\mathbb{I} - \Gamma_k^P) \Gamma_k^{\leq t} |\psi\rangle\|^2 = \|\Gamma_k^P \cdot \text{CPhsO} \cdot (\mathbb{I} - \Gamma_k^P) \Gamma_k^{\leq t} (|\psi_{\alpha,k}\rangle + |\psi_{\beta,k}\rangle)\|^2$$

$$\begin{aligned}
&\leq 2 \cdot (\text{LHS of (27)} + \text{LHS of (28)}) \\
&\hspace{15em} (\text{since } \|a + b\|^2 \leq 2 \cdot (\|a\|^2 + \|b\|^2)) \\
&\leq 8p_t \cdot \left(\|(\mathbb{I} - \Gamma_k^P)\Gamma_k^{\leq t} |\psi_{\alpha,k}\rangle\|^2 + \|(\mathbb{I} - \Gamma_k^P)\Gamma_k^{\leq t} |\psi_{\beta,k}\rangle\|^2 \right) \\
&\hspace{15em} (\text{by (27) and (28)}) \\
&= 8p_t \cdot \|(\mathbb{I} - \Gamma_k^P)\Gamma_k^{\leq t} |\psi\rangle\|^2 \hspace{10em} (\text{by orthogonality})
\end{aligned}$$

as desired. □