A Note on Single-Server QPIR from One-Way Functions

Prabhanjan Ananth* UCSB
Divyanshu Bhardwaj† UCSB
Aditya Gulati‡ UCSB

Abstract

We observe that there exists a single-server quantum private information retrieval with polylogarithmic communication assuming post-quantum one-way functions. Our observation follows immediately from the compilation technique of [Kerenidis-Wolf STOC'03] when combined with distributed point functions by [Gilboa-Ishai EUROCRYPT'14].

1 Introduction

Classical PIR. Private information retrieval (PIR), introduced by Chor, Goldreich, Kushilevitz, and Sudan [5], allows a client to retrieve a designated bit $D[i]$ from a database $D \in \{0,1\}^N$ without revealing the index i to the server. In the standard single-server setting, information-theoretic privacy requires linear communication, and hence any nontrivial communication saving necessarily relies on computational assumptions. Such protocols are referred to as *computational private information retrieval (CPIR)* protocols. Kushilevitz and Ostrovsky gave the first sublinear-communication single-server CPIR protocol under the quadratic residuosity assumption [9], and Cachin, Micali, and Stadler subsequently obtained polylogarithmic communication under the Φ-hiding assumption [4]. Döttling, Garg, Ishai, Malavolta, Mour and Ostrovsky designed CPIR protocols with polylogarithmic communication from decisional Diffie Hellman assumption or the quadratic residousity assumption.

Brakerski and Vaikuntanathan later constructed an asymptotically efficient single-server CPIR protocol from the Learning With Errors (LWE) assumption [11]. More recently, Abram, Malavolta, and Roy gave the first nontrivial CPIR construction from a code-based assumption, relying on the quasi-polynomial hardness of low-noise Learning Parity with Noise (LPN) [12]. Beimel, Ishai, Kushilevitz, and Malkin showed that one-way functions are necessary for any nontrivial communication saving in single-server PIR [3], while Ishai, Kushilevitz, and Ostrovsky showed that every one-round single-server CPIR protocol with sub-linear communication implies collision-resistant hash functions [14]. Thus, in the standard no-preprocessing model, single-server CPIR is known from number-theoretic, lattice-based, and code-based assumptions, whereas one-way functions constitute a necessary lower bound on the required assumptions.

Quantum PIR. A few works have explored the possibility of using quantum communication to achieve private information retrieval. Kerenidis and de Wolf used quantum interference to convert two classical queries into one quantum query and derived applications to locally decodable codes and PIR [7]; they also constructed quantum symmetrically private PIR protocols [8]. For a single honest server, Le Gall gave an information-theoretically private QPIR protocol with $O(\sqrt{N})$ qubits of communication [10]. This positive result is sensitive to the adversarial model: Baumeler and Broadbent proved that single-server QPIR requires linear communication against specious servers, the quantum analogue of honest-but-curious adversaries [2]. Aharonov, Brakerski, Chung, Green, Lai, and Sattath further isolated the input-purification

*prabhanjan@cs.ucsb.edu
†dbhardwaj@ucsb.edu
‡adityagulati@ucsb.edu

attack underlying this barrier and studied an anchored privacy model in which the database input is fixed to be classical [1]. In this work, we allow the adversarial server to be an arbitrary (potentially malicious) quantum polynomial-time algorithm.

Main Theorem. The main theorem is as follows:

Theorem 1.1. *There exists a single-server private information retrieval protocol:*

- *assuming post-quantum one-way functions,*
- *with one round of quantum communication,*
- *with communication complexity that is polylogarithmic in N, where N is the size of the database.*

In the above theorem, we consider databases where each entry is a single bit. Although the central quantum idea is not new, its combination with modern computational PIR constructions does not appear to have been stated explicitly in the literature.

2 Preliminaries

Notation. For a positive integer N, let $[N] = \{1, \ldots, N\}$. We write $x \xleftarrow{\$} \mathcal{X}$ for a uniformly random element of a finite set $\mathcal{X}$. Classical strings are represented by computational-basis states. Two ensembles are computationally indistinguishable, denoted by $\approx_c$, if every non-uniform quantum polynomial-time distinguisher has negligible distinguishing advantage. All computational assumptions and classical primitives in this note are required to be secure against quantum polynomial-time algorithms.

Coherent evaluation and phase kickback. A deterministic polynomial-time function $f : \{0,1\}^m \to \{0,1\}$ can be implemented reversibly as

$$|x\rangle|y\rangle \longmapsto |x\rangle|y \oplus f(x)\rangle.$$

If the second register is initialized to

$$|-\rangle = \frac{|0\rangle - |1\rangle}{\sqrt{2}},$$

then this transformation acts as

$$|x\rangle|-\rangle \longmapsto (-1)^{f(x)}|x\rangle|-\rangle.$$

Discarding the unchanged $|-\rangle$ register implements the phase oracle $|x\rangle \mapsto (-1)^{f(x)}|x\rangle$.

2.1 Quantum Private Information Retrieval

A one-round single-server quantum CPIR protocol consists of the following quantum polynomial time algorithms:

- $\mathsf{Share}(1^\lambda, i)$: on input a security parameter λ and an index $i \in [N]$, it outputs a quantum state $\rho_{\mathbf{QE}}$.
- $\mathsf{Eval}(D, \sigma_{\mathbf{Q}})$: on input a database $D \in \{0,1\}^N$ and a quantum state $\sigma_{\mathbf{Q}}$, it outputs a quantum state $\xi_{\mathbf{A}}$.
- $\mathsf{Reconstruct}$: on input the registers $\mathbf{E}$ and $\mathbf{A}$, it outputs y.

The client on input an index $i \in [N]$, computes $\rho_{\mathbf{QE}} \leftarrow \mathsf{Share}(1^\lambda, i)$. It sends the register $\mathbf{Q}$ to the server. The server computes $\xi_{\mathbf{A}} \leftarrow \mathsf{Eval}(D, \sigma_{\mathbf{Q}})$, where $\sigma_{\mathbf{Q}}$ is the reduced density matrix on $\mathbf{Q}$, and sends this to the client. Finally, the client computes $\mathsf{Reconstruct}$ on the registers $\mathbf{A}$ and $\mathbf{E}$ to obtain the decoded answer $D[i]$.

Correctness. The correctness property stipulates that the output of Reconstruct on input registers $\mathbb{E}$ (client's private state) and $\mathbb{A}$ (server's output) is $D[i]$ if: (a) $\rho_{\mathbb{QE}} \leftarrow \mathsf{Share}(1^\lambda, i)$ and, (b) $\xi_{\mathbb{A}} \leftarrow \mathsf{Eval}(D, \sigma_{\mathbb{Q}})$.

Communication Cost. The total number of qubits exchanged between the server and the client is polylogarithmic in N.

Indistinguishability Security. For every $i, i' \in [N]$, the following holds:

$$\sigma_{\mathbb{Q}}^{(i)} \approx_c \sigma_{\mathbb{Q}}^{(i')},$$

where:

- $\rho_{\mathbb{QE}} \leftarrow \mathsf{Share}(1^\lambda, i)$ and $\sigma_{\mathbb{Q}}^{(i)}$ is obtained by tracing out the register $\mathbb{E}$.
- $\rho_{\mathbb{QE}} \leftarrow \mathsf{Share}(1^\lambda, i')$ and $\sigma_{\mathbb{Q}}^{(i')}$ is obtained by tracing out the register $\mathbb{E}$.

2.2 2-Server CPIR with Additive Reconstruction

A 2-server CPIR protocol consists of the following probabilistic polynomial time algorithms:

- $\mathsf{Share}(1^\lambda, i)$: on input a security parameter λ and an index $i \in [N]$, it outputs two strings (a_0, a_1) along with a private state st_i. We assume without loss of generality that $|a_0| = |a_1|$.
- $\mathsf{Eval}(b, D, a_b)$: on input a bit b, a database $D \in \{0,1\}^N$ and a string a_b, it outputs a bit z_b. This algorithm is a deterministic polynomial time algorithm.
- $\mathsf{Reconstruct}(\mathsf{st}_i, z_0, z_1)$: on input a private state st_i, answers z_0, z_1 and outputs the answer y.

The client, on input an index $i \in [N]$, computes $((a_0, a_1), \mathsf{st}) \leftarrow \mathsf{Share}(1^\lambda, i)$. It sends a_0 to the first server and a_1 to the second server. The first server computes $z_0 \leftarrow \mathsf{Eval}(0, D, a_0)$ and the second server computes $z_1 \leftarrow \mathsf{Eval}(1, D, a_1)$. The servers respectively send z_0 and z_1 to the client. The client then computes $\mathsf{Reconstruct}(\mathsf{st}, z_0, z_1)$ to obtain $D[i]$.

Additive Reconstruction Property. For every $D \in \{0,1\}^N$, every $i \in [N]$, and every $((a_0, a_1), \mathsf{st})$ in the support of $\mathsf{Share}(1^\lambda, i)$,

$$\mathsf{Reconstruct}(\mathsf{st}, z_0, z_1) := \mathsf{Eval}(0, D, a_0) \oplus \mathsf{Eval}(1, D, a_1) = D[i].$$

We can consider the relaxation of the above condition, where the correctness only holds with probability negligibly close to 1.

Communication Complexity. The total number of bits exchanged between the server and the client is polylogarithmic in N.

Simulation Security. There exists a probabilistic polynomial-time simulator Sim such that, for every $i \in [N]$ and $b \in \{0,1\}$,

$$\left\{\mathsf{Sim}(1^\lambda, 1^{\lceil \log(N) \rceil}, b)\right\}_\lambda \approx_c \left\{a_b : (a_0, a_1) \leftarrow \mathsf{Share}(1^\lambda, i)\right\}_\lambda.$$

The indistinguishability requirement is against quantum polynomial-time distinguishers. In particular, the simulator's output distribution is independent of i.

Theorem 2.1 (Gilboa–Ishai [6]; post-quantum instantiation)**.** *There exists a two-server (classical) computational PIR protocol:*

- *assuming post-quantum one-way functions,*
- *with one round of classical communication,*
- *with additive reconstruction (over $\mathbb{F}_2$),*
- *with 1-bit answers and with communication complexity polylogarithmic in N.*

Justification of Theorem 2.1. Gilboa and Ishai construct a distributed point function (DPF): on input i, its key-generation algorithm outputs two short keys (a_0, a_1) whose pointwise evaluations xor to the unit vector e_i, while either key individually hides i. Taking the inner product of each expanded DPF share with D gives answer bits z_0, z_1 satisfying $z_0 \oplus z_1 = D[i]$. Their construction is based on pseudorandom generators. Instantiating it with a quantum-secure pseudorandom generator gives security against quantum polynomial-time distinguishers; such pseudorandomness can be based on post-quantum one-way functions via standard black-box transformations.

3 Construction

We present the construction below. The main ingredient is a two-server classical CPIR protocol satisfying additive reconstruction. The phase-interference idea is the standard two-classical-query-to-one-quantum-query transformation of Kerenidis and de Wolf [7]; related quantum PIR constructions appear in [8].

- The client receives as input an index $i \in [N]$, and the server receives as input a database $D \in \{0,1\}^N$.
- The client computes $(a_0, a_1) \leftarrow \mathsf{Share}(1^\lambda, i)$. It also samples $z \xleftarrow{\$} \{0,1\}$. The client sends
$$|\psi^{(z)}_{a_0,a_1}\rangle := \frac{1}{\sqrt{2}}\left(|0\rangle|a_0\rangle + (-1)^z|1\rangle|a_1\rangle\right)$$
to the server, while retaining the classical values (z, a_0, a_1).
- The server implements the controlled phase transformation
$$|b\rangle|a\rangle \longmapsto (-1)^{\mathsf{Eval}(b,D,a)}|b\rangle|a\rangle.$$
Equivalently, this can be implemented by reversible evaluation into a $|-\rangle$ ancilla and phase kickback (see Section 2). Writing $z_b = \mathsf{Eval}(b, D, a_b)$, the server prepares:
$$\frac{1}{\sqrt{2}}\left((-1)^{z_0}|0\rangle|a_0\rangle + (-1)^{z\oplus z_1}|1\rangle|a_1\rangle\right) = \frac{(-1)^{z_0}}{\sqrt{2}}\left(|0\rangle|a_0\rangle + (-1)^{z\oplus z_0\oplus z_1}|1\rangle|a_1\rangle\right).$$
Since $(-1)^{z_0}$ is a global phase, it can be ignored. The server sends the resulting state to the client.
- Define
$$|\psi^{(0)}_{a_0,a_1}\rangle = \frac{1}{\sqrt{2}}\left(|0\rangle|a_0\rangle + |1\rangle|a_1\rangle\right) \quad \text{and} \quad |\psi^{(1)}_{a_0,a_1}\rangle = \frac{1}{\sqrt{2}}\left(|0\rangle|a_0\rangle - |1\rangle|a_1\rangle\right).$$
The client measures using the three-outcome projective measurement:
$$\left\{|\psi^{(0)}_{a_0,a_1}\rangle\langle\psi^{(0)}_{a_0,a_1}|, |\psi^{(1)}_{a_0,a_1}\rangle\langle\psi^{(1)}_{a_0,a_1}|, \mathsf{I} - \sum_{c\in\{0,1\}}|\psi^{(c)}_{a_0,a_1}\rangle\langle\psi^{(c)}_{a_0,a_1}|\right\}.$$
If the first outcome occurs, it outputs z. If the second outcome occurs, it outputs $\overline{z} = z \oplus 1$. Otherwise, it aborts.

Theorem 3.1. *If the underlying two-server CPIR protocol is correct, per-server query cost θ and simulation secure against quantum polynomial-time adversaries, then the above construction is correct, has total communication cost $2\theta + 2$ qubits and satisfies indistinguishability security.*

Proof. We prove correctness and security separately.

Correctness. Let $z_b = \mathsf{Eval}(b, D, a_b)$. Ignoring the global phase $(-1)^{z_0}$, the state returned by the honest server is

$$|\psi^{(z\oplus z_0\oplus z_1)}_{a_0,a_1}\rangle.$$

By additive reconstruction,

$$z_0 \oplus z_1 = D[i].$$

Consequently, the client's measurement returns:

$$c = z \oplus D[i]$$

with probability one. If $c = 0$, the client outputs z, and the equality $0 = z \oplus D[i]$ implies $z = D[i]$. If $c = 1$, the client outputs $z \oplus 1$, and the equality $1 = z \oplus D[i]$ implies $z \oplus 1 = D[i]$. Thus the client always outputs $D[i]$ in an honest execution.

Security. We first prove the following simple claim.

Claim 3.2. *For fixed classical strings a_0, a_1,*

$$\mathbb{E}_{z\xleftarrow{\$}\{0,1\}}\left[|\psi^{(z)}_{a_0,a_1}\rangle\langle\psi^{(z)}_{a_0,a_1}|\right] = \sum_{b\in\{0,1\}}\frac{1}{2}|b\rangle\langle b|\otimes|a_b\rangle\langle a_b|.$$

Proof. Expanding the density operator gives

$$\begin{aligned}|\psi^{(z)}_{a_0,a_1}\rangle\langle\psi^{(z)}_{a_0,a_1}| = \frac{1}{2}|0\rangle\langle 0|\otimes|a_0\rangle\langle a_0| + \frac{1}{2}|1\rangle\langle 1|\otimes|a_1\rangle\langle a_1| \\ + \frac{(-1)^z}{2}|0\rangle\langle 1|\otimes|a_0\rangle\langle a_1| + \frac{(-1)^z}{2}|1\rangle\langle 0|\otimes|a_1\rangle\langle a_0|.\end{aligned}$$

Since $\mathbb{E}_z[(-1)^z] = 0$, averaging over z removes the two off-diagonal terms and proves the claim. □

Fix $i, i' \in [N]$. Consider the following hybrids:

$\underline{\mathsf{Hybrid}_1}$: This is the real experiment for index i. The client samples

$$(a_0, a_1) \leftarrow \mathsf{Share}(1^\lambda, i), \qquad z \xleftarrow{\$} \{0,1\},$$

and sends $|\psi^{(z)}_{a_0,a_1}\rangle$.

$\underline{\mathsf{Hybrid}_2}$: The client samples $(a_0, a_1) \leftarrow \mathsf{Share}(1^\lambda, i)$ and $b \xleftarrow{\$} \{0,1\}$, and sends the computational-basis state $|b\rangle|a_b\rangle$.

By Claim 3.2, after averaging over the client's private phase bit, the density operators in Hybrid_1 and Hybrid_2 are identical.

$\underline{\mathsf{Hybrid}_3}$: The client samples $b \xleftarrow{\$} \{0,1\}$ and $\widetilde{a}_b \leftarrow \mathsf{Sim}(1^\lambda, 1^{\lceil\log(N)\rceil}, b)$, and sends $|b\rangle|\widetilde{a}_b\rangle$.

The ensembles in Hybrid_2 and Hybrid_3 are computationally indistinguishable assuming the simulation security of classical two-server CPIR.

$\underline{\mathsf{Hybrid}_4}$: The client samples $(a'_0, a'_1) \leftarrow \mathsf{Share}(1^\lambda, i')$ and $b \xleftarrow{\$} \{0,1\}$, and sends $|b\rangle|a'_b\rangle$.

The ensembles in Hybrid_3 and Hybrid_4 are computationally indistinguishable by the same simulation-security argument, now applied to shares generated for i'.

$\underline{\mathsf{Hybrid}_5}$: This is the real experiment for index i'. The client samples

$$(a'_0, a'_1) \leftarrow \mathsf{Share}(1^\lambda, i'), \qquad z \xleftarrow{\$} \{0,1\},$$

and sends $|\psi^{(z)}_{a'_0,a'_1}\rangle$.

By Claim 3.2, the density operators in Hybrid_4 and Hybrid_5 are identical.

This shows that the hybrids Hybrid_1 and Hybrid_5 are computationally indistinguishable.

□

By combining Theorem 3.1 and Theorem 2.1, the proof of Theorem 1.1 follows.

Acknowledgements

The authors acknowledge the use of OpenAI's ChatGPT as an interactive research and writing assistant and that discussions helped identify the consequence, while the authors independently verified all statements and proofs. The authors wrote the proofs, and take full responsibility for the correctness of the results and any remaining errors.

References

[1] Dorit Aharonov, Zvika Brakerski, Kai-Min Chung, Ayal Green, Ching-Yi Lai, and Or Sattath. On quantum advantage in information-theoretic single-server PIR. In *Advances in Cryptology–EUROCRYPT 2019, Part III*, pages 219–246, 2019.

[2] Ämin Baumeler and Anne Broadbent. Quantum private information retrieval has linear communication complexity. *Journal of Cryptology*, 28(1):161–175, 2015.

[3] Amos Beimel, Yuval Ishai, Eyal Kushilevitz, and Tal Malkin. One-way functions are essential for single-server private information retrieval. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing*, pages 89–98, 1999.

[4] Christian Cachin, Silvio Micali, and Markus Stadler. Computationally private information retrieval with polylogarithmic communication. In *Advances in Cryptology–EUROCRYPT '99*, pages 402–414, 1999.

[5] Benny Chor, Oded Goldreich, Eyal Kushilevitz, and Madhu Sudan. Private information retrieval. *Journal of the ACM*, 45(6):965–981, 1998.

[6] Niv Gilboa and Yuval Ishai. Distributed point functions and their applications. In *Advances in Cryptology–EUROCRYPT 2014*, pages 640–658, 2014.

[7] Iordanis Kerenidis and Ronald de Wolf. Exponential lower bound for 2-query locally decodable codes via a quantum argument. *Journal of Computer and System Sciences*, 69(3):395–420, 2004.

[8] Iordanis Kerenidis and Ronald de Wolf. Quantum symmetrically-private information retrieval. *Information Processing Letters*, 90(3):109–114, 2004.

[9] Eyal Kushilevitz and Rafail Ostrovsky. Replication is not needed: Single database, computationally-private information retrieval. In *Proceedings of the 38th Annual IEEE Symposium on Foundations of Computer Science*, pages 364–373, 1997.

[10] François Le Gall. Quantum private information retrieval with sublinear communication complexity. *Theory of Computing*, 8(16):369–374, 2012.

[11] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 97–106, 2011.

[12] Damiano Abram, Giulio Malavolta, and Lawrence Roy. Slightly sublinear trapdoor hash functions and PIR from low-noise LPN. In *Theory of Cryptography Conference (TCC), Part I*, pages 280–300, 2025.

[13] Nico Döttling, Sanjam Garg, Yuval Ishai, Giulio Malavolta, Tamer Mour, and Rafail Ostrovsky. Trapdoor Hash Functions and Their Applications. In *Advances in Cryptology – CRYPTO 2019, Part III*, Lecture Notes in Computer Science, vol. 11694, pp. 3–32. Springer, 2019.

[14] Yuval Ishai, Eyal Kushilevitz, and Rafail Ostrovsky. Sufficient Conditions for Collision-Resistant Hashing. In *Theory of Cryptography Conference (TCC)*, Lecture Notes in Computer Science, vol. 3378, pp. 445–456. Springer, 2005.

[15] Caicai Chen, Yuval Ishai, Tamer Mour, and Alon Rosen. Secret-Key PIR from Random Linear Codes. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC)*. ACM, 2026.