

Information-theoretic two-server PIR requires $(6 - o(1)) \log n$ bits of communication

Keewoo Lee

Ethereum Foundation
keewoo.lee@ethereum.org

Abstract. We prove that every information-theoretic two-server private information retrieval scheme for n -bit databases requires $(6 - o(1)) \log n$ bits of communication. This improves on the 5 of Wehner and de Wolf (ICALP 2005), who had raised the 4.4 of Kerenidis and de Wolf (STOC 2003), who in turn had raised Mann’s original 4 (M.Sc. thesis, 1998). Our proof follows the quantum route of the earlier bounds: encode the database in a quantum state, recover an entry from enough copies of it, and apply Nayak’s bound (FOCS 1999) on quantum random access codes. Previous proofs read that entry as a binary outcome with a small bias toward the correct answer, and pay the inverse square of that bias to amplify it. We instead allow a real-valued outcome whose mean is the correct answer, and pay only its second moment. The same readout strategy improves the other lower bounds of Wehner and de Wolf, for smooth codes and locally decodable codes. In particular, it drops the linearity assumption in the lower bounds of Goldreich, Karloff, Schulman, and Trevisan (CCC 2002) almost for free.

1 Introduction

A private information retrieval (PIR) scheme [CGKS95] lets a client retrieve the i -th entry x_i of a database $x \in \{0, 1\}^n$ replicated at non-communicating servers, each of which individually learns nothing about i . Since there is a trivial solution in which the client downloads the entire database, the question is how far below n bits the communication can be pushed. In this work, we focus on the information-theoretic setting, and within it on two servers, the smallest number for which sublinear communication is possible.

On the construction side, the very first PIR paper [CGKS95] already presented a scheme with $O(n^{1/3})$ bits, and a long line of work [Yek10] culminated in the $2^{\tilde{O}(\sqrt{\log n})}$ -bit scheme of Dvir and Gopi [DG15]. On the other hand, we know embarrassingly little about lower bounds. Mann [Man98] proved that $(4 - o(1)) \log n$ bits are necessary, Kerenidis and de Wolf [KdW03] raised the constant to 4.4, and Wehner and de Wolf [WdW05] to 5, where the record has stood since 2005. The state of the art is thus $2^{\tilde{O}(\sqrt{\log n})}$ from above and $(5 - o(1)) \log n$ from below, and closing this huge gap is one of the central open problems in information-theoretic cryptography.

1.1 Our Contribution

In this work, we prove that every information-theoretic two-server PIR scheme for n -bit databases requires $(6 - o(1)) \log n$ bits of communication (Theorem 4). While this is still a very weak lower bound, it is the first improvement on the 5 of Wehner and de Wolf [WdW05] in two decades. The gain of 1 is also the largest single step since Mann [Man98] proved the first nontrivial bound.

Our proof follows the template established in [KdW03, WdW05]: prove a lower bound by a quantum argument, and combine it with the bound of Katz and Trevisan [KT00] (Theorem 1). The quantum argument encodes the database in a quantum state built from the PIR scheme’s queries and answers, recovers x_i from enough copies of that state, and bounds the total number of qubits by Nayak’s bound (Theorem 2) on quantum random access codes (QRACs). Our improvement comes from one simple observation: the readout need not output a bit. Earlier proofs recover x_i as a binary outcome with a small bias, $\Theta(2^{-\ell})$ for ℓ -bit answers in [WdW05], so amplifying it takes $\Theta(2^{2\ell})$ copies. We instead measure an observable whose outcome is real-valued, with mean exactly $(-1)^{x_i}$ and second moment 2^ℓ (Lemma 1), so $\Theta(2^\ell)$ copies suffice (Lemma 2).

The same observation improves the other lower bounds of Wehner and de Wolf [WdW05]. We obtain stronger bounds for smooth codes (Theorem 7), for locally decodable codes (Corollary 1), and for PIR with imperfect privacy and correctness (Corollary 2). All three also cover the case where only a few bits of each answer are used. In particular, we drop the linearity assumption of [GKST02] almost for free. They proved bounds for locally decodable codes and for PIR schemes when the answers are linear functions of the database, and our bounds (Corollaries 1 and 2) essentially match theirs with no assumption on the code or the scheme (Remark 2).

1.2 Organization

Section 2 collects background on PIR and QRACs, and Section 3 proves the two core lemmas that serve as a template for the proofs that follow. Our main theorem is then proved twice. Section 4 gives the first proof, which stays entirely within the setting of PIR with perfect privacy and perfect correctness. A reader interested only in that setting can stop after Section 4.

The second proof detours through the concept of smooth codes. It is longer, but it proves more than the main theorem. Section 5 recalls smooth codes, locally decodable codes, and PIR with imperfect privacy and correctness. Section 6 then runs the argument of Section 4 for smooth codes. The argument also covers codes and schemes that use only a few bits of each answer. The detour buys lower bounds for smooth codes (Section 6.1) and for locally decodable codes (Section 6.2), and a generalized version of the main theorem (Section 6.3). A reader who wants these results can skip Section 4, though it is the shortest place to see the whole argument and makes a good warm-up.

2 Preliminaries

2.1 Private Information Retrieval

Definition 1 (Information-Theoretic Two-Server PIR). *An information-theoretic two-server PIR scheme is a protocol between a client and two servers, indexed by $s \in \{0, 1\}$. The servers hold the same database $x = x_1 \cdots x_n \in \{0, 1\}^n$, the client holds an index $i \in [n]$, and the protocol proceeds as follows.*

- **Query:** *The client samples a uniformly random string $\rho \in \{0, 1\}^r$ and sends to each server s the t -bit query $q_s = Q_s(i, \rho)$, with Q_0, Q_1 the query functions.*
- **Answer:** *Server s returns the ℓ -bit answer $a_s = A_s(x, q_s)$, with A_0, A_1 the answer functions.*
- **Extract:** *The client outputs $\varphi(i, \rho, a_0, a_1)$, with φ the extraction function.*

The functions Q_s, A_s , and φ are deterministic, so all randomness is in ρ . Denote the total communication by $C = 2(t + \ell)$. We restrict attention to non-trivial schemes with $t, \ell \geq 1$. The scheme must satisfy the following.

- **Perfect Correctness:** *For every (x, i, ρ) , it holds that $\varphi(i, \rho, a_0, a_1) = x_i$, where $a_s = A_s(x, Q_s(i, \rho))$.*
- **Perfect Privacy:** *For each s , the distribution of the query $q_s = Q_s(i, \rho)$ over ρ is independent of i . That is, there are distributions p_0, p_1 on $\{0, 1\}^t$ such that $p_s(q) = \Pr_\rho[Q_s(i, \rho) = q]$ for all $i \in [n]$ and $q \in \{0, 1\}^t$.*

Theorem 1 ([KT00], as stated in [WdW05, Thm. 9]). *Every information-theoretic two-server PIR scheme (Definition 1) for n -bit databases with t -bit queries and ℓ -bit answers satisfies*

$$t \geq 2 \log \frac{n}{\ell} - O(1).$$

While we use Theorem 1 only as a black box, Appendix A gives a direct proof of it from Definition 1.

2.2 Quantum Random Access Code

Definition 2 (Quantum Random Access Code [ANTV99]). *For integers $n, m \geq 1$ and $1/2 < p \leq 1$, an (n, m, p) -quantum random access code (QRAC) is an encoding $x \mapsto \rho_x$ of $\{0, 1\}^n$ into m qubits with the following property: for every $i \in [n]$ there is a measurement that, for every $x = x_1 \cdots x_n \in \{0, 1\}^n$, recovers x_i from ρ_x with probability at least p .*

Theorem 2 (Nayak’s Bound [Nay99]). *Every (n, m, p) -QRAC satisfies the following, where H is the binary entropy function:*

$$m \geq (1 - H(p))n.$$

3 A Second-Moment Template for Lower Bounds

In this section, we prove the two core lemmas that Sections 4 and 6 share. Both sections build an encoding of $\{0, 1\}^n$ into qubits and, for each $i \in [n]$, a measurement that produces a state of the form $(|0, a_0\rangle + |1, a_1\rangle)/\sqrt{2}$ for two ℓ -bit strings a_0, a_1 . Lemma 1 reads any Boolean function f of a_0 and a_1 out of that state, with mean exactly $(-1)^{f(a_0, a_1)}$ and second moment 2^ℓ . Lemma 2 turns any such readout into a bound on n via Nayak's bound (Theorem 2).

Lemma 1. *For every $f : \{0, 1\}^\ell \times \{0, 1\}^\ell \rightarrow \{0, 1\}$ there is an observable M_f on $\ell + 1$ qubits such that, for all $a_0, a_1 \in \{0, 1\}^\ell$, measuring M_f on*

$$\frac{|0, a_0\rangle + |1, a_1\rangle}{\sqrt{2}}$$

yields an outcome Z with $\mathbb{E}[Z] = (-1)^{f(a_0, a_1)}$ and $\mathbb{E}[Z^2] = 2^\ell$.

Proof. Let $|\psi\rangle = (|0, a_0\rangle + |1, a_1\rangle)/\sqrt{2}$ and define the Hermitian operator

$$M_f = \sum_{u, v \in \{0, 1\}^\ell} (-1)^{f(u, v)} (|0, u\rangle\langle 1, v| + |1, v\rangle\langle 0, u|).$$

By construction M_f flips the first qubit:

$$M_f|0, a_0\rangle = \sum_{v \in \{0, 1\}^\ell} (-1)^{f(a_0, v)} |1, v\rangle, \quad M_f|1, a_1\rangle = \sum_{u \in \{0, 1\}^\ell} (-1)^{f(u, a_1)} |0, u\rangle.$$

Hence $\langle 0, a_0 | M_f | 0, a_0 \rangle = \langle 1, a_1 | M_f | 1, a_1 \rangle = 0$. The two images are supported on different values of the first qubit, hence orthogonal. Expanding $|\psi\rangle$ gives

$$\begin{aligned} \mathbb{E}[Z] &= \langle \psi | M_f | \psi \rangle = \frac{1}{2} \left(\langle 1, a_1 | M_f | 0, a_0 \rangle + \langle 0, a_0 | M_f | 1, a_1 \rangle \right) = (-1)^{f(a_0, a_1)}, \\ \mathbb{E}[Z^2] &= \|M_f |\psi\rangle\|^2 = \frac{1}{2} \left(\|M_f |0, a_0\rangle\|^2 + \|M_f |1, a_1\rangle\|^2 \right) = 2^\ell. \end{aligned}$$

The last equality holds since $M_f|0, a_0\rangle$ and $M_f|1, a_1\rangle$ have 2^ℓ entries of ± 1 . $\square$

Lemma 2. *Let $x \mapsto \rho_x$ be an encoding of $\{0, 1\}^n$ into L qubits. Suppose that for some $\beta, \sigma > 0$ and every $i \in [n]$ there is a measurement such that, for every $x \in \{0, 1\}^n$, applying it to ρ_x yields a real-valued outcome Z with*

$$(-1)^{x_i} \mathbb{E}[Z] \geq \beta \quad \text{and} \quad \mathbb{E}[Z^2] \leq \sigma^2.$$

Then the encoding must satisfy $n = O((\sigma^2/\beta^2) \cdot L)$. For pure-state encodings, say $\rho_x = |\Psi_x\rangle\langle\Psi_x|$, this bound can be improved to

$$n = O\left(\frac{\sigma^2}{\beta^2} \log\left(2 + \frac{\beta^2}{\sigma^2} \cdot 2^L\right)\right).$$

Proof. First, we claim that $k = \lceil 3\sigma^2/\beta^2 \rceil$ copies of ρ_x suffice to recover each bit of x with probability at least $2/3$. Indeed, fix $i \in [n]$, apply its measurement independently to each of the k copies, and output the sign of the mean $\bar{Z}$ of the k outcomes. The sign of $\bar{Z}$ differs from $(-1)^{x_i}$ only if $|\bar{Z} - \mathbb{E}[Z]| \geq \beta$. Since $\mathbb{E}[\bar{Z}] = \mathbb{E}[Z]$ and $\text{Var}(\bar{Z}) \leq \mathbb{E}[Z^2]/k \leq \sigma^2/k$, Chebyshev's inequality bounds the probability of the latter by $\sigma^2/(k\beta^2) \leq 1/3$ for every $x \in \{0, 1\}^n$.

The k copies occupy kL qubits, so $x \mapsto \rho_x^{\otimes k}$ is an $(n, kL, 2/3)$ -QRAC (Definition 2) and Nayak's bound (Theorem 2) gives $n = O(kL)$. Since $\beta^2 \leq \mathbb{E}[Z]^2 \leq \mathbb{E}[Z^2] \leq \sigma^2$ we have $\sigma^2/\beta^2 \geq 1$ and hence $k \leq 4\sigma^2/\beta^2$, so $n = O((\sigma^2/\beta^2) \cdot L)$.

For a pure-state encoding $\rho_x = |\Psi_x\rangle\langle\Psi_x|$, the states $|\Psi_x\rangle^{\otimes k}$ all lie in the same symmetric subspace of $(\mathbb{C}^{2^L})^{\otimes k}$, of dimension $\binom{2^L+k-1}{k}$, so a fixed isometry can compress them losslessly into $m = \lceil \log \binom{2^L+k-1}{k} \rceil$ qubits. As the compression is lossless, undoing it and then measuring as above recovers each bit of x with probability at least $2/3$, so the compressed encoding is an $(n, m, 2/3)$ -QRAC and Theorem 2 again gives $n = O(m)$. Finally

$$m \leq k \log \frac{e(2^L + k - 1)}{k} + 1 = O\left(\frac{\sigma^2}{\beta^2} \log\left(2 + \frac{\beta^2}{\sigma^2} \cdot 2^L\right)\right),$$

where the inequality uses $\binom{N}{k} \leq (eN/k)^k$ with $N = 2^L + k - 1$, and the last step again uses $k \leq 4\sigma^2/\beta^2$. $\square$

4 A Direct Proof of the Main Theorem

In this section, we prove the main theorem (Theorem 4) directly from Definition 1. Theorem 3 obtains a lower bound by following the template of Section 3, and Theorem 4 combines it with the bound of Katz and Trevisan (Theorem 1).

Theorem 3. *Every information-theoretic two-server PIR scheme (Definition 1) for n -bit databases with t -bit queries and ℓ -bit answers satisfies*

$$t = \Omega\left(\frac{n}{2^\ell}\right).$$

Proof. With the notation of Definition 1, consider the following encoding of a database $x \in \{0, 1\}^n$ into $L = t + \ell + 1$ qubits:

$$|\Psi_x\rangle = \sum_{q \in \{0, 1\}^t} \sum_{s \in \{0, 1\}} \sqrt{\frac{p_s(q)}{2}} |q\rangle |s, A_s(x, q)\rangle.$$

Each term holds a query, the server that receives it, and the answer that server gives. The two servers are weighted equally, and each query by the probability $p_s(q)$ that the client sends it. Perfect privacy is what makes these weights well defined, and no index appears in the encoding.

Fix $i \in [n]$. To apply Lemma 2, we now build a measurement of $|\Psi_x\rangle$ whose outcome has mean $(-1)^{x_i}$ and second moment 2^ℓ . We first invert the query

functions at i : each query becomes a uniform superposition over the randomness that produces it. Define a linear map, not depending on x , by

$$|q\rangle|s, a\rangle \mapsto \frac{1}{\sqrt{p_s(q) \cdot 2^r}} \sum_{\rho: Q_s(i, \rho)=q} |\rho\rangle|s, a\rangle$$

for $s \in \{0, 1\}$, $a \in \{0, 1\}^\ell$, and $q \in \{0, 1\}^t$ with $p_s(q) > 0$. The images are unit vectors, since $|\{\rho \in \{0, 1\}^r : Q_s(i, \rho) = q\}| = p_s(q) \cdot 2^r$ by perfect privacy, and they are pairwise orthogonal: two with different $|s, a\rangle$ are orthogonal, and two with the same $|s, a\rangle$ but different q are supported on disjoint sets of ρ . So the map is isometric. We may assume $r \geq t$ without loss of generality, and then it extends to an isometry V_i from L qubits to $r + \ell + 1$ qubits. Applying it to $|\Psi_x\rangle$,

$$V_i|\Psi_x\rangle = \frac{1}{\sqrt{2^r}} \sum_{\rho \in \{0, 1\}^r} |\rho\rangle \left(\frac{|0, a_0(x, i, \rho)\rangle + |1, a_1(x, i, \rho)\rangle}{\sqrt{2}} \right),$$

where $a_s(x, i, \rho) = A_s(x, Q_s(i, \rho))$.

Now measure $|\Psi_x\rangle$ as follows: apply V_i , measure the register holding ρ , and then measure the observable $M_{\varphi(i, \rho, \cdot, \cdot)}$ of Lemma 1 on the $\ell + 1$ qubits carrying the two answers. Let Z be the outcome of this last measurement. The measured ρ is uniform on $\{0, 1\}^r$, and the remaining $\ell + 1$ qubits are left in the state $(|0, a_0\rangle + |1, a_1\rangle)/\sqrt{2}$, where $a_s = a_s(x, i, \rho)$. Conditioned on ρ , Lemma 1 gives $\mathbb{E}[Z] = (-1)^{\varphi(i, \rho, a_0, a_1)}$ and $\mathbb{E}[Z^2] = 2^\ell$. Perfect correctness makes the former $(-1)^{x_i}$. Neither depends on ρ , so $\mathbb{E}[Z] = (-1)^{x_i}$ and $\mathbb{E}[Z^2] = 2^\ell$ unconditionally.

Lemma 2 therefore applies to the pure-state encoding $x \mapsto |\Psi_x\rangle$ with $L = t + \ell + 1$, $\beta = 1$, and $\sigma^2 = 2^\ell$, giving $n = O(2^\ell \log(2 + 2^{t+1})) = O(2^\ell t)$. $\square$

Theorem 4 (Main Theorem). *Every information-theoretic two-server PIR scheme (Definition 1) for n -bit databases requires total communication*

$$C \geq (6 - o(1)) \log n.$$

Proof. Following Definition 1, let the scheme have t -bit queries and ℓ -bit answers, so that its total communication is $C = 2(t + \ell)$. We may assume $C < 6 \log n$, as the claim is immediate otherwise, and then $t, \ell < 3 \log n$. By Theorem 3 we have $2^\ell = \Omega(n/t)$, so $\ell \geq \log n - \log t - O(1) \geq \log n - \log \log n - O(1)$. Meanwhile Theorem 1 gives $t \geq 2 \log n - 2 \log \ell - O(1) \geq 2 \log n - 2 \log \log n - O(1)$. Adding, $C = 2(t + \ell) \geq 6 \log n - 6 \log \log n - O(1)$, which is $(6 - o(1)) \log n$. $\square$

5 More Preliminaries

In this section, we recall the definitions and imported results that Section 6 needs. A *code* is a map $C : \{0, 1\}^n \rightarrow (\{0, 1\}^\ell)^m$, and we write $C_j(x) \in \{0, 1\}^\ell$ for the j -th symbol of $C(x)$. Smooth codes (Definition 3) are the common intermediate: locally decodable codes (Definition 4) are smooth by Lemma 3, and PIR schemes (Definition 5) yield smooth codes by Lemma 4.

5.1 Smooth Code

Definition 3 (Smooth Code [KT00]). A code $C : \{0, 1\}^n \rightarrow (\{0, 1\}^\ell)^m$ is a (k, c, ε) -smooth code if there is a decoder that, on input $i \in [n]$ and with oracle access to a word $y = y_1 \cdots y_m \in (\{0, 1\}^\ell)^m$, proceeds as follows.

- **Oracle Query:** The decoder samples a uniformly random string $\rho \in \{0, 1\}^r$ and queries the k distinct positions $j_s = J_s(i, \rho) \in [m]$, receiving the symbols $a_s = y_{j_s} \in \{0, 1\}^\ell$, with $J_0, \dots, J_{k-1}$ the position functions.
- **Extract:** The decoder outputs $\varphi(i, \rho, a_0, \dots, a_{k-1})$, with φ the extraction function.

The functions J_s and φ are deterministic, so all randomness is in ρ . The decoder must satisfy the following.

- **ε -Correctness:** For every $x = x_1 \cdots x_n \in \{0, 1\}^n$ and $i \in [n]$, it holds that $\Pr_\rho[\varphi(i, \rho, a_0, \dots, a_{k-1}) = x_i] \geq \frac{1}{2} + \varepsilon$, where $a_s = C_{j_s}(x)$ and $j_s = J_s(i, \rho)$.
- **c -Smoothness:** For every $i \in [n]$ and $j \in [m]$, the decoder queries position j with probability at most c/m , i.e., $\Pr_\rho[j \in \{J_0(i, \rho), \dots, J_{k-1}(i, \rho)\}] \leq c/m$.

We say that C uses b bits of each answer if the decoder can be taken to depend on each a_s only through b of its bits, whose positions depend only on i and ρ .

Theorem 5 (Implicit in [KT00, Thm. 6]). For every integer $k \geq 2$ and all $c, \varepsilon > 0$, every (k, c, ε) -smooth code $C : \{0, 1\}^n \rightarrow (\{0, 1\}^\ell)^m$ (Definition 3) satisfies

$$\log m \geq \frac{k}{k-1} \log \frac{n}{\ell} - O_{k,c,\varepsilon}(1).$$

5.2 Locally Decodable Code

Definition 4 (Locally Decodable Code [KT00]). A code $C : \{0, 1\}^n \rightarrow (\{0, 1\}^\ell)^m$ is a (k, δ, ε) -locally decodable code if there is a decoder of the form in Definition 3 such that, for every $x = x_1 \cdots x_n \in \{0, 1\}^n$, every $i \in [n]$, and every word $y \in (\{0, 1\}^\ell)^m$ with Hamming distance $d(y, C(x)) \leq \delta m$,

$$\Pr_\rho[\varphi(i, \rho, a_0, \dots, a_{k-1}) = x_i] \geq \frac{1}{2} + \varepsilon,$$

where $a_s = y_{j_s}$ and $j_s = J_s(i, \rho)$.

Lemma 3 ([KT00], as stated in [WdW05, Thm. 1]). Every (k, δ, ε) -locally decodable code (Definition 4) that uses b bits of each answer is a $(k, k/\delta, \varepsilon)$ -smooth code (Definition 3) that uses b bits of each answer.¹

¹ The definitions in [KT00] let the decoder read at most k positions, not necessarily distinct, whereas our Definitions 3 and 4 follow [WdW05] in requiring k distinct positions. The statement still holds if, in the proof, the decoder reads a replacement and ignores it rather than skipping a position it must avoid. We omit the details.

5.3 PIR with Imperfect Privacy and Correctness

Definition 5 (PIR with Imperfect Privacy and Correctness). A two-server (ε, η) -PIR scheme is defined as in Definition 1, with the two conditions replaced by the following.

- **ε -Correctness:** For every $x = x_1 \cdots x_n \in \{0, 1\}^n$ and $i \in [n]$, it holds that $\Pr_\rho[\varphi(i, \rho, a_0, a_1) = x_i] \geq \frac{1}{2} + \varepsilon$, where $a_s = A_s(x, q_s)$ and $q_s = Q_s(i, \rho)$.
- **η -Privacy:** For each $s \in \{0, 1\}$ and all $i, i' \in [n]$, the distributions of the queries $Q_s(i, \rho)$ and $Q_s(i', \rho)$ over the randomness ρ are within total variation distance η .

We say that the scheme uses b bits of each answer if φ depends on each a_s only through b of its bits, whose positions depend only on i and ρ .

Lemma 4 ([GKST02], as stated in [WdW05, Lem. 5]). Every two-server (ε, η) -PIR scheme (Definition 5) for n -bit databases with t -bit queries and ℓ -bit answers yields a $(2, 3, \varepsilon - \eta)$ -smooth code $C : \{0, 1\}^n \rightarrow (\{0, 1\}^\ell)^m$ (Definition 3) with $m \leq 6 \cdot 2^t$. The code uses b bits of each answer if the scheme does.²

Theorem 6. For all $\varepsilon > \eta$, every two-server (ε, η) -PIR scheme (Definition 5) for n -bit databases with t -bit queries and ℓ -bit answers satisfies

$$t \geq 2 \log \frac{n}{\ell} - O_{\varepsilon, \eta}(1).$$

Proof. By Lemma 4 the scheme yields a $(2, 3, \varepsilon - \eta)$ -smooth code $C : \{0, 1\}^n \rightarrow (\{0, 1\}^\ell)^m$ with $m \leq 6 \cdot 2^t$, so $t \geq \log m - O(1)$, and Theorem 5 gives $\log m \geq 2 \log(n/\ell) - O_{\varepsilon, \eta}(1)$. $\square$

6 A Lower Bound for Smooth Codes and Applications

In this section, we repeat the argument of Section 4 for smooth codes, which gives Theorem 7. The reductions of Section 5 then carry it to locally decodable codes (Corollary 1) and to PIR (Corollary 2), and Theorem 8 generalizes the main theorem (Theorem 4).

6.1 A Lower Bound for Smooth Codes

Lemma 5. Let $C : \{0, 1\}^n \rightarrow (\{0, 1\}^\ell)^m$ be a $(2, c, \varepsilon)$ -smooth code (Definition 3) with a fixed decoder that uses b bits of each answer, and write $a_s(x, i, \rho)$ for those bits of $C_{J_s(i, \rho)}(x)$. With u the number of sets $T \subseteq [\ell]$ with $|T| \leq b$, there is an encoding $x \mapsto |\Psi_x\rangle$ of $\{0, 1\}^n$ into $\lceil \log(mu) \rceil$ qubits with the following property.

² The definition in [GKST02] lets the decoder read at most two positions, not necessarily distinct, whereas our Definition 3 follows [WdW05] in requiring two distinct positions. The statement still holds if, in the proof, the decoder reads a replacement and ignores it rather than skipping a query. We omit the details.

For every $i \in [n]$ there is a measurement, with outcomes $\{0, 1\}^r$ and **Fail**, such that for every $x \in \{0, 1\}^n$, applied to $|\Psi_x\rangle$ it returns each $\rho \in \{0, 1\}^r$ with probability $2^{b+1}/(cu \cdot 2^r)$ and leaves the state

$$\frac{|0, a_0(x, i, \rho)\rangle + |1, a_1(x, i, \rho)\rangle}{\sqrt{2}}.$$

Proof. For $T \subseteq [\ell]$ and $a = a_1 \cdots a_\ell \in \{0, 1\}^\ell$ let $\chi_T(a) = (-1)^{\sum_{k \in T} a_k}$. Consider the following encoding of $\{0, 1\}^n$ into $\lceil \log(mu) \rceil$ qubits:

$$|\Psi_x\rangle = \frac{1}{\sqrt{mu}} \sum_{j \in [m]} \sum_{|T| \leq b} \chi_T(C_j(x)) |j, T\rangle.$$

That is, we encode the low-degree parities of the symbols $C_j(x)$ into the phases.

Fix $i \in [n]$. We now build the measurement claimed for this i . Write $j_s^{(\rho)} = J_s(i, \rho)$ and let $S_s^{(\rho)} \subseteq [\ell]$ be the positions of the b bits used from the s -th answer, so that $a_s(x, i, \rho)$ is the restriction of $C_{j_s^{(\rho)}}(x)$ to $S_s^{(\rho)}$. For $j = j_s^{(\rho)}$ and $T \subseteq S_s^{(\rho)}$, the parity $\chi_T(C_j(x))$ depends only on the bits $a_s(x, i, \rho)$ that the decoder reads. For each $\rho \in \{0, 1\}^r$ we therefore project onto those terms, taking

$$\Pi_\rho = \sum_{s \in \{0, 1\}} \sum_{T \subseteq S_s^{(\rho)}} |j_s^{(\rho)}, T\rangle \langle j_s^{(\rho)}, T|, \quad K_\rho = \sqrt{\frac{m}{c \cdot 2^r}} \Pi_\rho.$$

Since $j_0^{(\rho)} \neq j_1^{(\rho)}$, each Π_ρ is an orthogonal projection, and a basis state $|j, T\rangle$ lies in its range only if the decoder queries j , which by c -smoothness happens for at most $c \cdot 2^r/m$ choices of ρ . Hence

$$\sum_\rho K_\rho^\dagger K_\rho = \frac{m}{c \cdot 2^r} \sum_\rho \Pi_\rho \leq I,$$

and the K_ρ can be completed to a measurement with one more outcome **Fail**.

Write $a_s = a_s(x, i, \rho)$ and identify each $T \subseteq S_s^{(\rho)}$ with its indicator $t \in \{0, 1\}^b$ within $S_s^{(\rho)}$, so that $\chi_T(C_{j_s^{(\rho)}}(x)) = (-1)^{t \cdot a_s}$. On outcome ρ , the measurement relabels $|j_s^{(\rho)}, T\rangle$ as $|s, t\rangle$, an isometry on the range of Π_ρ . Since K_ρ keeps the terms of $|\Psi_x\rangle$ in the range of Π_ρ and scales them by $\sqrt{m/(c \cdot 2^r)}$, the relabeling turns $K_\rho |\Psi_x\rangle$ into

$$\frac{1}{\sqrt{cu \cdot 2^r}} \sum_{s \in \{0, 1\}} \sum_{t \in \{0, 1\}^b} (-1)^{t \cdot a_s} |s, t\rangle.$$

The measurement then applies the Hadamard transform $H^{\otimes b}$ to the last b qubits. As $H^{\otimes b} \sum_{t \in \{0, 1\}^b} (-1)^{t \cdot a_s} |t\rangle = 2^{b/2} |a_s\rangle$, this leaves

$$\sqrt{\frac{2^{b+1}}{cu \cdot 2^r}} \cdot \frac{|0, a_0\rangle + |1, a_1\rangle}{\sqrt{2}}.$$

The relabeling and the Hadamard transform preserve norms, and the second factor above is a unit vector. So the outcome ρ has probability $2^{b+1}/(cu \cdot 2^r)$ and leaves $(|0, a_0\rangle + |1, a_1\rangle)/\sqrt{2}$. $\square$

Remark 1 (Correspondence between Theorem 3 and Lemma 5). The position j in the proof of Lemma 5 plays the role of the pair (s, q) in the proof of Theorem 3, and the symbol $C_j(x)$ that of the answer $A_s(x, q)$. For $b = \ell$, identifying each $T \subseteq [\ell]$ with its indicator string, one has $\sum_T \chi_T(C_j(x))|T\rangle = 2^{\ell/2} H^{\otimes \ell} |C_j(x)\rangle$, so the encoding of Lemma 5 is $\frac{1}{\sqrt{m}} \sum_j |j\rangle \otimes H^{\otimes \ell} |C_j(x)\rangle$. This is the encoding in the proof of Theorem 3 with the pairs (s, q) weighted equally, so at $b = \ell$ the two agree up to the fixed unitary $H^{\otimes \ell}$ on the answer register. We use a different form in each proof in order to better suit its own setting.

Theorem 7. *For all $\varepsilon > 0$, every $(2, c, \varepsilon)$ -smooth code $C : \{0, 1\}^n \rightarrow (\{0, 1\}^\ell)^m$ (Definition 3) that uses b bits of each answer satisfies*

$$\log m = \Omega\left(\frac{\varepsilon^2 n}{cu}\right),$$

where u is the number of sets $T \subseteq [\ell]$ with $|T| \leq b$. In particular, for $b = \ell$ we have $\log m = \Omega(\varepsilon^2 n / (c \cdot 2^\ell))$.

Proof. Take the encoding $x \mapsto |\Psi_x\rangle$ of Lemma 5 and fix $i \in [n]$. To apply Lemma 2 we build a measurement whose outcome Y satisfies $(-1)^{x_i} \mathbb{E}[Y] \geq 2\varepsilon$ and $\mathbb{E}[Y^2] = cu/2$. We first apply the measurement of Lemma 5 for this i , which returns some $\rho \in \{0, 1\}^r$ with total probability $p = 2^{b+1}/(cu)$ and Fail otherwise. On Fail, return $Y = 0$. On outcome ρ , measure the observable $M_{\varphi(i, \rho, \cdot, \cdot)}$ of Lemma 1 and return $Y = Z/p$, where Z is the outcome of that measurement.

By Lemma 5, the measured ρ is uniform on $\{0, 1\}^r$ and leaves $(|0, a_0\rangle + |1, a_1\rangle)/\sqrt{2}$, where $a_s = a_s(x, i, \rho)$ is as in that lemma. Conditioned on ρ , Lemma 1 therefore gives $\mathbb{E}[Z] = (-1)^{\varphi(i, \rho, a_0, a_1)}$ and $\mathbb{E}[Z^2] = 2^b$. Hence

$$\mathbb{E}[Y] = \mathbb{E}_\rho[(-1)^{\varphi(i, \rho, a_0, a_1)}] = (-1)^{x_i} (2 \Pr_\rho[\varphi(i, \rho, a_0, a_1) = x_i] - 1),$$

so $(-1)^{x_i} \mathbb{E}[Y] \geq 2\varepsilon$ by ε -correctness, while $\mathbb{E}[Y^2] = 2^b/p = cu/2$.

Lemma 2 therefore applies to the pure-state encoding $x \mapsto |\Psi_x\rangle$ with $L = \lceil \log(mu) \rceil$, $\beta = 2\varepsilon$, and $\sigma^2 = cu/2$, giving

$$n = O\left(\frac{cu}{8\varepsilon^2} \log\left(2 + \frac{8\varepsilon^2}{cu} \cdot 2^L\right)\right).$$

By definition $c, m \geq 2$ and $\varepsilon \leq 1/2$. Hence $\log(2 + 2^L \cdot 8\varepsilon^2/(cu)) = O(\log m)$ and $n = O((cu/\varepsilon^2) \log m)$. $\square$

Remark 2 (Comparison with Previous Bounds). For $(2, c, \varepsilon)$ -smooth codes that use b bits of each answer, [GKST02, Cor. 4.2] gives $\log m \geq \Omega(\varepsilon n / (cu)) - \log u$ under the assumption that C is linear, with every bit of $C(x)$ a parity of bits

of x . Without linearity, [WdW05, Thm. 6] gives $\log m \geq \Omega(\varepsilon^2 n / (2^b c u)) - \log u$. Our bound $\log m = \Omega(\varepsilon^2 n / (c u))$ (Theorem 7) has neither the factor 2^b nor the additive term, so for constant ε it matches [GKST02] without assuming linearity. The factor 2^b is saved by the real-valued readout of Lemma 1, and the additive term by compressing the copies of the encoding into their symmetric subspace in Lemma 2.

6.2 Application to Locally Decodable Codes

Corollary 1. *For all $\delta, \varepsilon > 0$, every $(2, \delta, \varepsilon)$ -locally decodable code $C : \{0, 1\}^n \rightarrow (\{0, 1\}^\ell)^m$ (Definition 4) that uses b bits of each answer satisfies*

$$\log m = \Omega\left(\frac{\delta \varepsilon^2 n}{u}\right),$$

where u is the number of sets $T \subseteq [\ell]$ with $|T| \leq b$. In particular, for $b = \ell$ we have $\log m = \Omega(\delta \varepsilon^2 n / 2^\ell)$.

Proof. By Lemma 3 the code is a $(2, 2/\delta, \varepsilon)$ -smooth code that uses b bits of each answer. Apply Theorem 7 with $c = 2/\delta$. $\square$

6.3 Application to Private Information Retrieval

Corollary 2. *For all $\varepsilon > \eta$, every two-server (ε, η) -PIR scheme (Definition 5) for n -bit databases with t -bit queries and ℓ -bit answers that uses b bits of each answer satisfies*

$$t = \Omega\left(\frac{(\varepsilon - \eta)^2 n}{u}\right),$$

where u is the number of sets $T \subseteq [\ell]$ with $|T| \leq b$. In particular, for $b = \ell$ we have $t = \Omega((\varepsilon - \eta)^2 n / 2^\ell)$.

Proof. By Lemma 4 the scheme yields a $(2, 3, \varepsilon - \eta)$ -smooth code $C : \{0, 1\}^n \rightarrow (\{0, 1\}^\ell)^m$ that uses b bits of each answer, with $t = \Omega(\log m)$, and Theorem 7 gives $\log m = \Omega((\varepsilon - \eta)^2 n / u)$. $\square$

Theorem 8 (Main Theorem, Generalized). *For all $\varepsilon > \eta$, every two-server (ε, η) -PIR scheme (Definition 5) for n -bit databases requires total communication*

$$C \geq (6 - o_{\varepsilon, \eta}(1)) \log n.$$

Proof. Following Definition 5, let the scheme have t -bit queries and ℓ -bit answers, so that its total communication is $C = 2(t + \ell)$. We may assume $C < 6 \log n$, as the claim is immediate otherwise, and then $t, \ell < 3 \log n$. By Corollary 2 with $b = \ell$ we have $2^\ell = \Omega((\varepsilon - \eta)^2 n / t)$, so $\ell \geq \log n - \log t - O_{\varepsilon, \eta}(1) \geq \log n - \log \log n - O_{\varepsilon, \eta}(1)$. Meanwhile Theorem 6 gives $t \geq 2 \log n - 2 \log \ell - O_{\varepsilon, \eta}(1) \geq 2 \log n - 2 \log \log n - O_{\varepsilon, \eta}(1)$. Adding, $C = 2(t + \ell) \geq 6 \log n - 6 \log \log n - O_{\varepsilon, \eta}(1)$, which is $(6 - o_{\varepsilon, \eta}(1)) \log n$. $\square$

References

- ANTV99. Andris Ambainis, Ashwin Nayak, Amnon Ta-Shma, and Umesh V. Vazirani. Dense quantum coding and a lower bound for 1-way quantum automata. In *31st ACM STOC*, pages 376–383. ACM Press, May 1999. [doi:10.1145/301250.301347](#). 3
- CGKS95. Benny Chor, Oded Goldreich, Eyal Kushilevitz, and Madhu Sudan. Private information retrieval. In *36th FOCS*, pages 41–50. IEEE Computer Society Press, October 1995. [doi:10.1109/SFCS.1995.492461](#). 1
- DG15. Zeev Dvir and Sivakanth Gopi. 2-server PIR with sub-polynomial communication. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *47th ACM STOC*, pages 577–584. ACM Press, June 2015. [doi:10.1145/2746539.2746546](#). 1
- GKST02. Oded Goldreich, Howard J. Karloff, Leonard J. Schulman, and Luca Trevisan. Lower bounds for linear locally decodable codes and private information retrieval. In *Proceedings of the 17th Annual IEEE Conference on Computational Complexity, Montréal, Québec, Canada, May 21-24, 2002*, pages 175–183. IEEE Computer Society, 2002. [doi:10.1109/CCC.2002.1004353](#). 2, 8, 10, 11
- KdW03. Iordanis Kerenidis and Ronald de Wolf. Exponential lower bound for 2-query locally decodable codes via a quantum argument. In *35th ACM STOC*, pages 106–115. ACM Press, June 2003. [doi:10.1145/780542.780560](#). 1, 2
- KT00. Jonathan Katz and Luca Trevisan. On the efficiency of local decoding procedures for error-correcting codes. In *32nd ACM STOC*, pages 80–86. ACM Press, May 2000. [doi:10.1145/335305.335315](#). 2, 3, 7, 12
- Man98. Eran Mann. Private access to distributed information. Master’s thesis, Technion – Israel Institute of Technology, 1998. 1, 2, 12
- Nay99. Ashwin Nayak. Optimal lower bounds for quantum automata and random access codes. In *40th FOCS*, pages 369–377. IEEE Computer Society Press, October 1999. [doi:10.1109/SFCS.1999.814608](#). 3
- WdW05. Stephanie Wehner and Ronald de Wolf. Improved lower bounds for locally decodable codes and private information retrieval. In Luís Caires, Giuseppe F. Italiano, Luís Monteiro, Catuscia Palamidessi, and Moti Yung, editors, *ICALP 2005*, volume 3580 of *LNCS*, pages 1424–1436. Springer, Berlin, Heidelberg, July 2005. [doi:10.1007/11523468_115](#). 1, 2, 3, 7, 8, 11
- Yek10. Sergey Yekhanin. *Locally Decodable Codes and Private Information Retrieval Schemes*. ISC. Springer, 2010. [doi:10.1007/978-3-642-14358-8](#). 1

A A Direct Proof of Theorem 1

Katz and Trevisan [KT00] prove their lower bound for locally decodable codes by proving a bound for smooth codes (Theorem 5), which yields lower bounds for multi-server PIR, including Theorem 1. This appendix proves Theorem 1 directly from Definition 1, without passing through smooth codes. In the two-server setting with perfect privacy and perfect correctness, much of their argument simplifies, and we include the proof for its pedagogical value. We do not claim novelty for this proof. Its key idea is the same as in [KT00].³

³ It is presumably also close to Mann’s proof [Man98], which is not available online.

Theorem 1 (Restated). *Every information-theoretic two-server PIR scheme (Definition 1) for n -bit databases with t -bit queries and ℓ -bit answers satisfies*

$$t \geq 2 \log \frac{n}{\ell} - O(1).$$

Proof. We claim that it is enough to show that we can choose a small number of queries per server such that, for a constant fraction of the indices i , some random string lets the client retrieve x_i using only the chosen queries. That is, it suffices to find sets $S_0, S_1 \subseteq \{0, 1\}^t$ with $|S_0|, |S_1| \leq \lceil 2^{t/2} \rceil$ and a set $I \subseteq [n]$ with $|I| = \Omega(n)$ such that

$$\forall i \in I, \quad \exists \rho_i \in \{0, 1\}^r \text{ with } Q_0(i, \rho_i) \in S_0 \text{ and } Q_1(i, \rho_i) \in S_1. \quad (1)$$

Indeed, suppose such S_0, S_1, I exist. For each $i \in I$, fix one witness ρ_i satisfying (1), and write $q_0^{(i)} = Q_0(i, \rho_i) \in S_0$ and $q_1^{(i)} = Q_1(i, \rho_i) \in S_1$. By perfect correctness, $x_i = \varphi(i, \rho_i, A_0(x, q_0^{(i)}), A_1(x, q_1^{(i)}))$ is fully determined by $A_0(x, q_0^{(i)})$ and $A_1(x, q_1^{(i)})$ for every database x . Hence $(x_i)_{i \in I}$ is determined by the answers $A_0(x, q)$ for $q \in S_0$ and $A_1(x, q)$ for $q \in S_1$. These answers contain at most

$$(|S_0| + |S_1|)\ell \leq 2\lceil 2^{t/2} \rceil \ell \leq 2^{t/2+2}\ell$$

bits. On the other hand, $(x_i)_{i \in I}$ ranges over all $2^{|I|}$ possible strings, so necessarily

$$2^{t/2+2}\ell \geq |I| = \Omega(n),$$

and hence $t \geq 2 \log(n/\ell) - O(1)$, as claimed. It therefore remains only to prove the existence of S_0, S_1, I satisfying (1).

To this end, choose S_0 to be the $\lceil 2^{t/2} \rceil$ queries of largest probability under p_0 . Then $p_0(S_0) \geq 2^{-t/2}$. For each index $i \in [n]$, define

$$T_1^{(i)} = \{q_1 \in \{0, 1\}^t : \exists \rho \text{ such that } Q_0(i, \rho) \in S_0 \text{ and } Q_1(i, \rho) = q_1\}.$$

By definition, $Q_0(i, \rho) \in S_0$ implies $Q_1(i, \rho) \in T_1^{(i)}$. Therefore $\Pr_\rho[Q_1(i, \rho) \in T_1^{(i)}]$ is at least $\Pr_\rho[Q_0(i, \rho) \in S_0]$. Perfect privacy identifies these two marginal probabilities with $p_1(T_1^{(i)})$ and $p_0(S_0)$, respectively, so for every $i \in [n]$,

$$p_1(T_1^{(i)}) \geq p_0(S_0) \geq 2^{-t/2}.$$

Now sample $\lceil 2^{t/2} \rceil$ independent queries from p_1 , and let S_1 be the set of these samples, so that $|S_1| \leq \lceil 2^{t/2} \rceil$. For each i ,

$$\Pr_{S_1}[S_1 \cap T_1^{(i)} = \emptyset] \leq (1 - p_1(T_1^{(i)}))^{\lceil 2^{t/2} \rceil} \leq (1 - 2^{-t/2})^{\lceil 2^{t/2} \rceil} \leq e^{-1} < \frac{1}{2},$$

so by linearity of expectation the expected number of indices i with $S_1 \cap T_1^{(i)} \neq \emptyset$ exceeds $n/2$. Fix a choice of S_1 for which this number is at least $n/2$, and let I be the set of these indices. By the definition of $T_1^{(i)}$, every $i \in I$ satisfies (1). $\square$