Electronic Colloquium on Computational Complexity, Report No. 59 (2002)

Exponential Lower Bound for 2-Query Locally Decodable Codes

Iordanis Kerenidis* Ronald de Wolf†

Abstract

We prove exponential lower bounds on the length of 2-query locally decodable codes. Goldreich et al. recently proved such bounds for the special case of *linear* locally decodable codes. Our proof shows that a 2-query locally decodable code can be decoded with only 1 *quantum* query, and then proves an exponential lower bound for such 1-query locally quantum-decodable codes. We also exhibit q-query locally quantum-decodable codes that are much shorter than the best known q-query classical codes. Finally, we give some new lower bounds for (not necessarily linear) private information retrieval systems.

Keywords: Locally decodable codes, error correction, lower bounds, private information retrieval, quantum computing.

1 Introduction

1.1 Setting

Error-correcting codes allow one to encode an n-bit string x into an m-bit codeword $C(x)$, in such a way that x can still be recovered even if the codeword is corrupted in a number of places. For example, codewords of length $m = O(n)$ already suffice to recover from errors in a constant fraction of the bitpositions of the codeword (even in linear time [20]). One disadvantage of such "standard" error-correction, is that one usually needs to consider all or most of the (corrupted) codeword to recover anything about x. If one is only interested in recovering one or a few of the bits of x, then more efficient schemes are possible, so-called locally decodable codes (LDCs). LDCs allow us to extract small parts of encoded information from a corrupted codeword, while looking at ("querying") only a few positions of that word. They have found various applications in complexity theory and cryptography, such as self-correcting computations, PCPs, worst-case to average-case reductions, and private information retrieval. Informally, LDCs are described as follows:

> A (q, δ, ε)-*locally decodable code* encodes n-bit strings x into m-bit codewords $C(x)$, such that for each i, the bit x_i can be recovered with probability $1/2 + \varepsilon$ making only q queries, even if the codeword is corrupted in δm of the bits.

For example, the Hadamard code is a locally decodable code where *two* queries are sufficient in order to predict any bit with constant advantage, even with a constant fraction of errors. The code has $m = 2^n$ and $C(x)_j = j \cdot x \bmod 2$ for all $j \in \{0,1\}^n$. Recovery from a corrupted codeword y is possible by picking a random $j \in \{0,1\}^n$, querying y_j and $y_{j\oplus e_i}$, and outputting the XOR of those two bits. If neither bit has been corrupted, then we output $y_j \oplus y_{j\oplus e_i} = j\cdot x \oplus (j \oplus e_i)\cdot x = e_i \cdot x = x_i$,

*UC Berkeley, `jkeren@cs.berkeley.edu`. Supported by DARPA under agreement number F 30602–01-2–0524.
†UC Berkeley, `rdewolf@cs.berkeley.edu`. Supported by Talent grant S 62–565 from the Netherlands Organization for Scientific Research (NWO).

ISSN 1433-8092

as we should. If $C(x)$ has been corrupted in at most δm positions, then a fraction of at least $1 - 2\delta$ of all $(j, j \oplus e_i)$ pairs of indices is uncorrupted, so the recovery probability is at least $1 - 2\delta$. This is $> 1/2$ as long as $\delta < 1/4$. The main drawback of the Hadamard code is its exponential length.

Clearly, we would like both the codeword length m and the number of queries q to be small. The main complexity question about LDCs is how large m needs to be, as a function of n, q, δ, and ε. For $q = polylog(n)$, Babai et al. [2] showed how to achieve $m = n^{1+o(1)}$, for some fixed δ, ε. For constant q, the best known upper bounds are of the form $m = 2^{O(n^{1/(q-1)})}$(see e.g. [3]).

The study of *lower* bounds on m was initiated by Katz and Trevisan [13]. They proved that for $q = 1$, LDCs do not exist if n is larger than some constant depending on δ and ε. For $q \geq 2$, they proved a bound of $m = \Omega(n^{q/(q-1)})$ if the q queries are made non-adaptively; this bound was generalized to the adaptive case by Deshpande et al. [10]. This establishes superlinear but at most quadratic lower bounds on the length of LDCs with a constant number of queries. There is still a large gap between the best known upper and lower bounds. In particular, it is open whether $m = poly(n)$ is achievable with constant q. Recently, Goldreich et al. [11] examined the case $q = 2$, and showed that $m \geq 2^{\delta\varepsilon n/8}$ if C is a *linear* code. Obata [18] subsequently strengthened the dependence on ε to $m \geq 2^{\Omega(\delta n/(1-2\varepsilon))}$, which is essentially optimal.

1.2 Our results

The main result of this paper is an exponential lower bound for general 2-query LDCs:

A $(2, \delta, \varepsilon)$-locally decodable code requires length $m \geq 2^{cn-1}$,

for $c = 1 - H(1/2 + 3\delta\varepsilon/14)$, where $H(\cdot)$ is the binary entropy function. This is the first super-polynomial lower bound on general LDCs with more than 1 query. Our constant c in the exponent is somewhat worse than the ones of Goldreich et al. and of Obata, but our proof establishes the exponential lower bound for *all* LDCs, not just linear ones. Goldreich et al. also give extensions of their result for codewords over larger alphabets, but consider the result for the binary alphabet their "main result". We focus only on the binary case in this paper (though see Section 3.4).

Our proof introduces one radically new ingredient: *quantum* computing. We show that if 2 classical queries can recover x_i with probability $1/2 + \varepsilon$, then x_i can also be recovered with probability $1/2 + 4\varepsilon/7$ using only 1 quantum query. In other words, a $(2, \delta, \varepsilon)$-locally decodable code is a $(1, \delta, 4\varepsilon/7)$-locally *quantum*-decodable code. We then prove an exponential lower bound for 1-query LQDCs by showing, roughly speaking, that a 1-query LQDC of length m induces a *quantum random access code* for x of length $\log m$. Nayak's [16] linear lower bound on such codes finishes off the proof (for the sake of completeness, we include a proof of his result in Appendix A).

This lower bound for classical LDCs is one of the very few examples where tools from quantum computing enable one to prove *new* results in *classical* computer science. The only other example of this that we know, are the lower bounds on the set membership data structure of Radhakrishnan et al. [19]. Their lower bounds are proved for quantum computers (hence also apply to classical computers), but are in fact stronger than the previous classical lower bounds of Buhrman et al. [6].[1]

We also observe that our construction implies the existence of 1-query quantum-decodable codes for all n. The Hadamard code is an example of this. Here the codewords are still classical, but the decoding algorithm is quantum. As mentioned before, if we only allow one *classical* query, then LDCs do not exist for n larger than some constant depending on δ and ε [13]. For larger q we show

[1]The quantum lower bound on the communication complexity of the inner product function of Cleve et al. [8] provides new insight in a classical result, but does not establish a *new* result for classical CS.

that the best known $(2q, \delta, \varepsilon)$-LDCs (which have length $m = 2^{O(n^{1/(2q-1)})}$) are actually (q, δ, ε)-LQDCs. Hence for fixed number of queries q, we obtain LQDCs that are significantly shorter than the best known LDCs. We summarize the situation in the following table, where our contributions are indicated by boldface.

Queries	Length of LDC	Length of LQDC
$q=1$	don't exist	$\mathbf{2^{\Theta(n)}}$
$q=2$	$\mathbf{2^{\Theta(n)}}$	$\mathbf{2^{O(n^{1/3})}}$
$q=3$	$2^{O(n^{1/2})}$	$\mathbf{2^{O(n^{1/5})}}$

Table 1: Best known bounds on the length of LDC and QLDC with q queries

Katz and Trevisan, and Goldreich et al. established a close connection between locally decodable codes and *private information retrieval (PIR)* schemes. These schemes allow a user to extract a bit x_i from an n-bit database x that is replicated over one or more servers, without the server(s) learning *which* i the user wants. The complexity of such schemes is measured by the total number of bits communicated. Our techniques allow us to reduce classical 2-server PIR schemes with 1-bit answers to quantum 1-server PIRs. Since Nayak [16] established a linear lower bound for the latter (see Appendix B), we obtain a linear lower bound on the communication complexity for all classical 2-server PIRs with 1-bit answers. Previously, such a bound was known only for PIRs where the answer bits are *linear combinations* of the bits of x (this was first proven in [7, Section 5.2] and extended to linear PIRs with constant-length answers in [11]).

2 Preliminaries

2.1 Quantum

Below we give more precise definitions of locally decodable codes and related notions, but we first briefly explain the standard notation of quantum computing. We refer to Nielsen and Chuang [17] for more details. A *qubit* is a linear combination of the basis states $|0\rangle$ and $|1\rangle$, also viewed as a 2-dimensional complex vector:

$$\alpha_0|0\rangle + \alpha_1|1\rangle = \begin{pmatrix} \alpha_0 \\ \alpha_1 \end{pmatrix},$$

where α_0, α_1 are complex *amplitudes*, and $|\alpha_0|^2 + |\alpha_1|^2 = 1$.

The 2^m basis states of an m-qubit system are the m-fold tensor products of the states $|0\rangle$ and $|1\rangle$. For example, the basis states of a 2-qubit system are the four 4-dimensional unit vectors $|0\rangle \otimes |0\rangle$, $|0\rangle \otimes |1\rangle$, $|1\rangle \otimes |0\rangle$, and $|1\rangle \otimes |1\rangle$. We abbreviate, e.g., $|1\rangle \otimes |0\rangle$ to $|0\rangle|1\rangle$, or $|1, 0\rangle$, or $|10\rangle$, or even $|2\rangle$ (since 2 is 10 in binary). With these basis states, an m-qubit state $|\phi\rangle$ is a 2^m-dimensional complex unit vector

$$|\phi\rangle = \sum_{i \in \{0,1\}^m} \alpha_i |i\rangle.$$

We use $\langle\phi| = |\phi\rangle^*$ to denote the conjugate transpose of the vector $|\phi\rangle$, and $\langle\phi|\psi\rangle = \langle\phi| \cdot |\psi\rangle$ for the inner product between states $|\phi\rangle$ and $|\psi\rangle$. These two states are *orthogonal* if $\langle\phi|\psi\rangle = 0$. The *density matrix* corresponding to $|\phi\rangle$ is the outer product $|\phi\rangle\langle\phi|$. The density matrix corresponding to a *mixed state*, which is in pure state $|\phi_i\rangle$ with probability p_i, is $\rho = \sum_i p_i |\phi_i\rangle\langle\phi_i|$. If a 2-register quantum state has the form $|\phi\rangle = \sum_i \sqrt{p_i}|i\rangle|\phi_i\rangle$, then the state of a system holding only the second register of $|\phi\rangle$ is described by the (reduced) density matrix $\sum_i p_i |\phi_i\rangle\langle\phi_i|$.

The most general measurement allowed by quantum mechanics is a so-called *positive operator-valued measurement (POVM)*. A k-outcome POVM is specified by positive operators $E_i = M_i^* M_i$, $1 \leq i \leq k$, subject to the condition that $\sum_i E_i = I$. Given a state ρ, the probability of getting the ith outcome is $p_i = \mathrm{Tr}(M_i \rho)$. If the outcome is indeed i, then the resulting state is $M_i \rho M_i^* / \mathrm{Tr}(M_i \rho M_i^*)$. In particular, if $\rho = |\phi\rangle\langle\phi|$, then $p_i = \langle\phi|E_i|\phi\rangle = \| M_i|\phi\rangle \|^2$, and the resulting state is $M_i|\phi\rangle / \| M_i|\phi\rangle \|$. A special case is where $k = 2^m$ and $B = \{|\psi_i\rangle\}$ forms an orthonormal basis of the m-qubit space. "Measuring in the B-basis" means that we apply the POVM given by $E_i = M_i = |\psi_i\rangle\langle\psi_i|$. Applying this to a pure state $|\phi\rangle$ gives resulting state $|\psi_i\rangle$ with probability $p_i = |\langle\phi|\psi_i\rangle|^2$. Apart from measurements, the basic operations that quantum mechanics allows us to do, are *unitary* (i.e., linear norm-preserving) transformations of the vector of amplitudes.

Finally, a word about quantum *queries*. A query to an m-bit string y is commonly formalized as the following unitary transformation, where $j \in [m]$, and $b \in \{0,1\}$ is called the *target bit*:

$$|j\rangle|b\rangle \mapsto |j\rangle|b \oplus y_j\rangle.$$

A quantum computer may apply this to any superposition. An equivalent formalization that we will be using here, is:

$$|c\rangle|j\rangle \mapsto (-1)^{c \cdot y_j}|c\rangle|j\rangle.$$

Here c is a *control bit* that controls whether the phase $(-1)^{y_j}$ is added or not. Given some extra workspace, one query of either type can be simulated exactly by one query of the other type.

2.2 Codes

Below, by a 'decoding algorithm' we mean an algorithm (quantum or classical depending on context) with oracle access to the bits of some (possibly corrupted) codeword y for x. The algorithm gets input i and is supposed to recover x_i while making only few queries to y.

Definition 1 *$C : \{0,1\}^n \rightarrow \{0,1\}^m$ is a (q, δ, ε)-*locally decodable code *(LDC) if there is a classical randomized decoding algorithm A such that*

1. *A makes at most q non-adaptive queries*
2. *For all x and i, and all $y \in \{0,1\}^m$ with Hamming distance $d(C(x), y) \leq \delta m$ we have $\Pr[A(y,i) = x_i] \geq 1/2 + \varepsilon$.*

The LDC is called linear *if C is a linear function over $GF(2)$ (i.e., $C(x+y) = C(x) + C(y)$).*

By allowing A to be a quantum computer and to make queries in superposition, we can similarly define (q, δ, ε)-locally quantum-*decodable codes (LQDCs).*

It will be convenient to work with *non-adaptive* queries, as used in the above definition, so the distribution on the queries that A makes is independent of y. However, our main lower bound also holds for adaptive queries, see the first remark at the end of Section 3.3.

2.3 Private information retrieval

Next we formally define private information retrieval schemes.

Definition 2 *A one-round, $(1-\delta)$-secure, k-server private information retrieval (PIR) scheme with recovery probability $1/2 + \varepsilon$, query size t, and answer size a, consists of a randomized algorithm representing the user, and k deterministic algorithms $S_1, \ldots, S_k$ (the servers), such that*

1. *On input $i \in [n]$, the user produces k t-bit queries $q_1, \ldots, q_k$ and sends these to the respective servers. The jth server sends back an a-bit string $a_j = S_j(x, q_j)$. The user outputs a bit b depending on $i, a_1, \ldots, a_k$, and his randomness.*

2. *For all x and i, the probability (over the user's randomness) that $b = x_i$ is at least $1/2 + \varepsilon$.*

3. *For all x and j, the distributions on q_j (over the user's randomness) are δ-close (in total variation distance) for different i.*

The scheme is called linear *if, for every j and q_j, the jth server's answer $S_j(x, q_j)$ is a linear combination (over $GF(2)$) of the bits of x.*

All known upper bounds on PIR have one round, $\varepsilon = 1/2$ (perfect recovery) and $\delta = 0$ (the servers get no information whatsoever about i). Below we will assume one round and $\delta = 0$ without mentioning this further. We can generalize these definitions to *quantum* PIR. For the $\delta = 0$-case this generalization is straightforward (the server's state after the query should be independent of i), and that is the only case we will need here.

The main complexity measure of a PIR scheme is its communication complexity, i.e., the sum of the lengths of the queries that the user sends to each server, and the length of the servers' answers. If there is only one server ($k = 1$), then privacy can be maintained by letting the server send the whole n-bit database to the user. This takes n bits of communication and is optimal. If the database is replicated over $k \geq 2$ servers, then smarter protocols are possible. Chor et al. [7] exhibited a 2-server PIR with communication complexity $O(n^{1/3})$ and with $O(n^{1/k})$ for $k > 2$. Ambainis [1] improved the latter to $O(n^{1/(2k-1)})$, and some more recent references are [3, 4]. No general lower bounds better than $\Omega(\log n)$ are known for PIRs with $k \geq 2$ servers. Goldreich et al. [11] proved that *linear* 2-server PIRs with t-bit queries, and a-bit answers where the user looks only at k predetermined positions in each answer, require $t = \Omega(n/a^k)$.

3 Lower Bound for Locally Decodable Codes with Two Queries

Here we will show that the linearity constraint is not needed for the exponential lower bound on 2-query LDCs. The proof consists of two parts, both of which have a clear intuition but require quite a few technicalities:

1. A 2-query LDC gives a 1-query LQDC, because 1 quantum query can compute the same Boolean functions as 2 classical queries (albeit with slightly worse error probability).

2. The length m of a 1-query LQDC must be exponential, because it induces a $\log m$-qubit quantum random access code for x, for which a linear lower bound is already known [16].

3.1 From 2 classical queries to 1 quantum query

The key to the first step is the following lemma:

Lemma 1 *Let $f : \{0,1\}^2 \to \{0,1\}$ and suppose we can make queries to the bits of some input string $a = a_1a_2 \in \{0,1\}^2$. There exists a quantum algorithm that makes only one query (one that is independent of f) and outputs $f(a)$ with probability* exactly $11/14$, *and outputs $1 - f(a)$ otherwise.*

Proof. The quantum algorithm makes the following query:

$$\frac{1}{\sqrt{3}}\left(|0\rangle|1\rangle + |1\rangle|1\rangle + |1\rangle|2\rangle\right),$$

where the first bit is the control bit, and the appropriate phase $(-1)^{a_j}$ is added if the control bit is 1. The result of the query is the state

$$|\phi\rangle = \frac{1}{\sqrt{3}}\left(|0\rangle|1\rangle + (-1)^{a_1}|1\rangle|1\rangle + (-1)^{a_2}|1\rangle|2\rangle\right).$$

The algorithm then measures this state in a basis containing the following 4 states ($b \in \{0,1\}^2$):

$$|\psi_b\rangle = \frac{1}{2}\left(|0\rangle|1\rangle + (-1)^{b_1}|1\rangle|1\rangle + (-1)^{b_2}|1\rangle|2\rangle + (-1)^{b_1+b_2}|0\rangle|2\rangle\right).$$

The probability of getting outcome a is $|\langle\phi|\psi_a\rangle|^2 = 3/4$, and each of the other 3 outcomes has probability $1/12$. The algorithm determines its output based on f and on the measurement outcome b. We distinguish 3 cases for f:

1. $|f(1)^{-1}| = 1$ (the case $|f(1)^{-1}| = 3$ is completely analogous, with 0 and 1 reversed). If $f(b) = 1$, then the algorithm outputs 1 with probability 1. If $f(b) = 0$ then it outputs 0 with probability $6/7$ and 1 with probability $1/7$. Accordingly, if $f(a) = 1$, then the probability of outputting 1 is $\Pr[f(b) = 1] \cdot 1 + \Pr[f(b) = 0] \cdot 1/7 = 3/4 + 1/28 = 11/14$. If $f(a) = 0$, then the probability of outputting 0 is $\Pr[f(b) = 0] \cdot 6/7 = (11/12) \cdot (6/7) = 11/14$.
2. $|f(1)^{-1}| = 2$. Then $\Pr[f(a) = f(b)] = 3/4 + 1/12 = 5/6$. If the algorithm outputs $f(b)$ with probability $13/14$ and outputs $1 - f(b)$ with probability $1/14$, then its probability of outputting $f(a)$ is exactly $11/14$.
3. f is constant. In that case the algorithm just outputs that value with probability $11/14$.

□

Theorem 1 *A $(2, \delta, \varepsilon)$-locally decodable code is a $(1, \delta, 4\varepsilon/7)$-locally quantum-decodable code.*

Proof. Consider some i, x, and y such that $d(C(x), y) \leq \delta m$. Let us fix the randomness of the 2-query classical decoder. This determines two indices $j, k \in [m]$ and an $f : \{0,1\}^2 \to \{0,1\}$ such that

$$\Pr[f(y_j, y_k) = x_i] = p \geq 1/2 + \varepsilon,$$

where the probability is taken over the decoder's randomness. We now use Lemma 1 to obtain a 1-query quantum decoder that outputs some value o such that

$$\Pr[o = f(y_j, y_k)] = 11/14.$$

The success probability of this quantum decoder is:

$$\begin{aligned}
\Pr[o = x_i] &= \Pr[o = f(y_j, y_k)] \cdot \Pr[f(y_j, y_k) = x_i] + \Pr[o \neq f(y_j, y_k)] \cdot \Pr[f(y_j, y_k) \neq x_i] \\
&= \frac{11}{14}p + \frac{3}{14}(1-p) \;=\; \frac{3}{14} + \frac{4}{7}p \;\geq\; \frac{1}{2} + \frac{4\varepsilon}{7},
\end{aligned}$$

as claimed. (Here we use the 'exactly' part of Lemma 1. If our quantum algorithm would have success probability $11/14$ for f where $f(y_j, y_k) = x_i$ but success probability 1 for f where $f(y_j, y_k) \neq x_i$, then we could actually end up with overall recovery probability less than $1/2$.) □

3.2 Exponential lower bound for 1-query LQDCs

A quantum *random access code* is an encoding $x \mapsto \rho_x$ of n-bit strings x into m-qubit states ρ_x, such that any bit x_i can be recovered with some probability $p \geq 1/2+\varepsilon$ from ρ_x. That is, for each i there is a 2-outcome POVM E_i, $I - E_i$, such that $\mathrm{Tr}(E_i\rho_x) \geq 1/2 + \varepsilon$ if $x_i = 1$ and $\mathrm{Tr}(E_i\rho_x) \leq 1/2 - \varepsilon$ if $x_i = 0$. The following lower bound is known on the length of such quantum codes [16] (see Appendix A for a proof).

Theorem 2 (Nayak) *An encoding $x \mapsto \rho_x$ of n-bit strings into m-qubit states with recovery probability at least p, has $m \geq (1 - H(p))n$.*

This allows us to prove an exponential lower bound for 1-query LQDC:

Theorem 3 *If $C : \{0,1\}^n \to \{0,1\}^m$ is a $(1, \delta, \varepsilon)$-locally quantum-decodable code, then*

$$m \geq 2^{cn-1},$$

for $c = 1 - H(1/2 + \delta\varepsilon/4)$.

Proof. We fix i. Let $|Q\rangle = \sum_{c\in\{0,1\},j\in[m]} \alpha_{cj}|c\rangle|j\rangle$ be the query that the quantum decoder makes to recover x_i. Without loss of generality, we assume that all α_{cj} are non-negative reals (complex phases and entanglement with its workspace can always be added by the decoder after the query). Let D and $I - D$ be the two POVM operators that the decoder uses on the state $|R\rangle$ returned by the query, corresponding to outcomes 1 and 0, respectively. Its probability of outputting 1 on $|R\rangle$ is $p(R) = \langle R|D|R\rangle = \|\sqrt{D}|R\rangle\|^2$.

Since C is a LQDC, the decoder can recover x_i with probability $1/2 + \varepsilon$ from the state

$$\sum_{c\in\{0,1\},j\in[m]} \alpha_{cj}(-1)^{c\cdot y_j}|c\rangle|j\rangle$$

for every y such that $d(C(x), y) \leq \delta m$. Our goal below is to show that we can also recover x_i with probability $1/2 + \delta\varepsilon/4$ from the uniform state

$$|U(x)\rangle = \frac{1}{\sqrt{2m}} \sum_{c\in\{0,1\},j\in[m]} (-1)^{c\cdot C(x)_j}|c\rangle|j\rangle.$$

Since $|U(x)\rangle$ is independent of i, it forms a $(\log(m) + 1)$-qubit random access code for x. The theorem then follows from Theorem 2.

Inspired by the "smoothing" technique of [13], we split the amplitudes of the query $|Q\rangle$ into small and large ones: $A = \{cj : \alpha_{cj} \leq \sqrt{1/\delta m}\}$ and $B = \{cj : \alpha_{cj} > \sqrt{1/\delta m}\}$. We can assume that α_{0j} is the same for all j, so $\alpha_{0j} \leq 1/\sqrt{m} \leq 1/\sqrt{\delta m}$ and hence $0j \in A$. Let $a = \sqrt{\sum_{cj\in A}\alpha_{cj}^2}$ be the norm of the "small-amplitude" part of the state. Since $\sum_{cj\in B}\alpha_{cj}^2 \leq 1$, we have $|B| < \delta m$. Define non-normalized states

$$\begin{aligned}
|A(x)\rangle &= \sum_{cj\in A} (-1)^{c\cdot C(x)_j}\alpha_{cj}|c\rangle|j\rangle \\
|B\rangle &= \sum_{cj\in B} \alpha_{cj}|c\rangle|j\rangle
\end{aligned}$$

The states $|A(x)\rangle + |B\rangle$ and $|A(x)\rangle - |B\rangle$ each correspond to a $y \in \{0,1\}^m$ that is corrupted (compared to $C(x)$) in at most $|B| \leq \delta m$ positions, so the decoder can recover x_i from each of these states. If x has $x_i = 1$, then

$$p(A(x) + B) \geq 1/2 + \varepsilon \quad\text{and}\quad p(A(x) - B) \geq 1/2 + \varepsilon.$$

Since $p(A\pm B)=p(A)+p(B)\pm(\langle A|D|B\rangle+\langle B|D|A\rangle)$, averaging the previous two inequalities gives

$$p(A(x))+p(B)\geq 1/2+\varepsilon.$$

Similarly, if x' has $x'_i=0$, then

$$p(A(x'))+p(B)\leq 1/2-\varepsilon.$$

Hence, for the normalized states $|A(x)\rangle/a$ and $|A(x')\rangle/a$ we have

$$p(A(x)/a)-p(A(x')/a)\geq 2\varepsilon/a^2.$$

Since this holds for every x,x' with $x_i=1$ and $x'_i=0$, there are constants $q_1,q_0\in[0,1]$, $q_1-q_0\geq 2\varepsilon/a^2$, such that $p(A(x)/a)\geq q_1$ whenever $x_i=1$ and $p(A(x)/a)\leq q_0$ whenever $x_i=0$.

If we had a copy of the state $|A(x)\rangle/a$, then we could run the following procedure, where for simplicity we assume $q_1\geq 1/2+\varepsilon/a^2$ (if not, then we must have $q_0\leq 1/2-\varepsilon/a^2$ and we can use the same argument with 0 and 1 reversed):

> Output 0 with probability $q=1-1/(q_1+q_0)$,
> and otherwise output the result of running the decoder's POVM on $|A(x)\rangle/a$.

If $x_i=1$, then the probability that this procedure outputs 1 is

$$(1-q)p(A(x)/a)\geq(1-q)q_1=\frac{q_1}{q_1+q_0}=\frac{1}{2}+\frac{q_1-q_0}{2(q_1+q_0)}\geq\frac{1}{2}+\frac{\varepsilon}{2a^2}.$$

If $x_i=0$, then the probability that it outputs 0 is

$$q+(1-q)(1-p(A(x)/a))\geq q+(1-q)(1-q_0)=1-\frac{q_0}{q_1+q_0}=\frac{q_1}{q_1+q_0}\geq\frac{1}{2}+\frac{\varepsilon}{2a^2}.$$

Thus, we can recover x_i with good probability if we had the state $|A(x)\rangle/a$.

It remains to show how we can obtain $|A(x)\rangle/a$ from $|U(x)\rangle$ with reasonable probability. This we do by applying a POVM with operators $M^\dagger M$ and $I-M^\dagger M$ to $|U(x)\rangle$, where $M=\sqrt{\delta m}\sum_{cj\in A}\alpha_{cj}|cj\rangle\langle cj|$. Note that both $M^\dagger M$ and $I-M^\dagger M$ are positive operators (as is required for a POVM) because $0\leq\sqrt{\delta m}\alpha_{cj}\leq 1$ for all $cj\in A$. The measurement gives the first outcome with probability

$$\langle U(x)|M^\dagger M|U(x)\rangle=\frac{\delta m}{2m}\sum_{cj\in A}\alpha_{cj}^2=\delta a^2/2.$$

In this case we have obtained the normalized version of $M|U(x)\rangle$, which is $|A(x)\rangle/a$, so then we can run the above procedure to recover x_i. If the measurement gives the second outcome, then we just output a fair coin flip. Thus we recover x_i from $|U(x)\rangle$ with probability at least

$$(\delta a^2/2)(1/2+\varepsilon/2a^2)+(1-\delta a^2/2)1/2=1/2+\delta\varepsilon/4,$$

which concludes the proof (note that the user can do all of the above without knowing x). □

3.3 Exponential lower bound for 2-query LDCs

Theorem 4 *If $C : \{0,1\}^n \to \{0,1\}^m$ is a $(2, \delta, \varepsilon)$-locally decodable code, then*

$$m \geq 2^{cn-1},$$

for $c = 1 - H(1/2 + 3\delta\varepsilon/14)$.

Proof. The theorem follows by combining Theorems 1 and 3. Straightforwardly, this would give a constant of $1 - H(1/2 + \delta\varepsilon/7)$. We get the better constant claimed here by observing that the 1-query LQDC derived from the 2-query LDC actually has 1/3 of the overall squared amplitude on queries where the control bit c is zero (and all those α_{0j} are in A). Hence in the proof of Theorem 3, we can redefine "small amplitude" to $\alpha_{cj} \leq \sqrt{2/3\delta m}$, and still B will have at most δm elements because $\sum_{cj \in B} \alpha_{cj}^2 \leq 2/3$. This in turns allows us to make M a factor $\sqrt{3/2}$ larger, which improves the probability of getting $|A(x)\rangle/a$ from $|U(x)\rangle$ to $3\delta a^2/4$ and the recovery probability to $1/2 + 3\delta\varepsilon/8$. Combining that with the first step (which makes ε a factor 4/7 smaller) gives $c = 1 - H(1/2 + 3\delta\varepsilon/14)$, as claimed. □

Remarks:

(1) Note that a $(2, \delta, \varepsilon)$-LDC with *adaptive* queries gives a $(2, \delta, \varepsilon/2)$-LDC with non-adaptive queries: if query q_1 would be followed by query q_2^0 or q_2^1 depending on the outcome of q_1, then we can just guess in advance whether to query q_1 and q_2^0, or q_1 and q_2^1. With probability 1/2, the second query will be the one we would have made in the adaptive case and we're fine, in the other case we just flip a coin, giving overall recovery probability $1/2(1/2 + \varepsilon) + 1/2(1/2) = 1/2 + \varepsilon/2$. Thus we also get slightly weaker but still exponential lower bounds for *adaptive* 2-query LDCs.

(2) For a $(2, \delta, \varepsilon)$-LDC where the decoder's output is the XOR of its two queries, we can give a better reduction than in Theorem 1. In this case, the quantum decoder can apply his query to

$$\frac{1}{\sqrt{2}} (|1\rangle|1\rangle + |1\rangle|2\rangle),$$

giving

$$\frac{1}{\sqrt{2}} ((-1)^{a_1}|1\rangle|1\rangle + (-1)^{a_2}|1\rangle|2\rangle) = (-1)^{a_1} \frac{1}{\sqrt{2}} (|1\rangle|1\rangle + (-1)^{a_1 \oplus a_2}|1\rangle|2\rangle),$$

and extract $a_1 \oplus a_2$ from this with certainty. Thus the recovery probability remains $1/2 + \varepsilon$ instead of going down to $1/2 + 4\varepsilon/7$. Accordingly, we also get slightly better lower bounds for 2-query LDCs where the output is the XOR of the two queried bits, namely $c = 1 - H(1/2 + 3\delta\varepsilon/8)$.

Kenji Obata (unpublished, personal communication) also has proved exponential lower bounds for the length of (not necessarily linear) LDCs with this XOR-property.

3.4 Larger alphabets

We emphasize that our exponential lower bound only applies to the case where the codewords are over the *binary* alphabet. If the codewords are over a larger alphabet Σ, then a query to a symbol in the codeword can give more than one bit of information about x, and the length of the codewords (measured in number of Σ-symbols) may be smaller. Goldreich et al. [11] were able to extend their lower bound for linear binary PIRs to larger alphabets. Below we sketch the extension of the second step of our proof (the lower bound for 1-query LQDCs) to larger alphabets. Unfortunately, so far we have been unable to generalize the first step of our proof (the 2-classical-to-1-quantum-query reduction).

The extension uses the Bernstein-Vazirani algorithm [5] to reduce codes over Σ to codes over the binary alphabet. Very briefly, that algorithm does the following: given access to the Hadamard code for a string $a \in \{0,1\}^k$, it queries a uniform superposition, giving

$$\frac{1}{\sqrt{2^k}} \sum_{i \in \{0,1\}^k} (-1)^{i \cdot a} |i\rangle,$$

and then applies a Hadamard transform to turn this into $|a\rangle$. Accordingly, a quantum query to a can be replaced by a binary query to the Hadamard code for a.

Now consider some (q, δ, ε)-LQDC $C : \{0,1\}^n \to \Sigma^m$, and let $\ell = 2^{\lceil \log |\Sigma| \rceil}$. Define a binary code C' by replacing each $C(x)_j$ (viewed as a $\log \ell$-bit string) by its Hadamard code, which has length ℓ bits. We obtain a q-query decoder for C' from the decoder for C, by replacing each of its q queries to $C(x)_j$ by the Bernstein-Vazirani algorithm. If $d(C'(x), y') \leq (\delta/\ell)m'$, then at most δm of the positions in the corresponding $C(x)$ are corrupted, so the decoder for $C'(x)$ will output x_i with probability $1/2 + \varepsilon$. This gives:

Theorem 5 *Let $\ell = 2^{\lceil \log |\Sigma| \rceil}$. If there exists a (q, δ, ε)-LQDC $C : \{0,1\}^n \to \Sigma^m$, then there exists a $(q, \delta/\ell, \varepsilon)$-LQDC $C' : \{0,1\} \to \{0,1\}^{m'}$, where $m' = m \cdot \ell$.*

Combining with our lower bound for the binary alphabet (Theorem 3):

Corollary 1 *If $C : \{0,1\}^n \to \Sigma^m$ is a $(1, \delta, \varepsilon)$-locally quantum-decodable code, then*

$$m \geq 2^{cn-1}/\ell,$$

for $\ell = 2^{\lceil \log |\Sigma| \rceil}$ and $c = 1 - H(1/2 + \delta\varepsilon/4\ell)$.

Reductions with smaller loss in δ are possible by observing that a small number of errors in the Hadamard code for some $a \in \Sigma$ will give the Bernstein-Vazirani algorithm only a negligibly small error probability.

4 Locally Quantum-Decodable Codes with Few Queries

The second remark of Section 3.3 immediately generalizes to:

Theorem 6 *A $(2q, \delta, \epsilon)$-LDC where the decoder's output is the XOR of the $2q$ queried bits, is a (q, δ, ε)-LQDC.*

Reasonably good $(2q, \delta, \epsilon)$-LDCs with this XOR property can be obtained from known $2q$-server PIR schemes with 1-bit answers [3, Theorem 6.8]. For every $k \geq 2$ there exist k-server PIRs with $\delta = 0$, $\varepsilon = 1/2$, query length $t = O(k \log(k) \cdot n^{1/(k-1)})$ and answer length 1, where the user's output is the XOR of the k answer bits. By concatenating all 2^t answers for all k servers, we obtain a k-query LDC of length $m = k \cdot 2^t$ where the recovery algorithm makes one query in each of the k blocks, and outputs the XOR of the queried bits. Within each block, all 2^t positions are equally likely to be queried (though knowing the query-positions in $k-1$ of the blocks determines which position will be queried in the k-th block, so queries are not quite independent). Without any errors in the codeword, the recovery probability would be 1. The worst-case corruption is if all δm errors occur in one of the k blocks. In this case we still have recovery probability at least $1 - \delta m/2^t = 1 - \delta k$, so $\varepsilon = 1/2 - \delta k$. Of course, this only makes sense if $\delta < 1/2k$. Plugging in $k = 2q$ and combining with Theorem 6 gives:

Corollary 2 *For all n, $q \geq 1$ and $\delta < 1/4q$, there exists a $(q, \delta, 1/2 - \delta 2q)$-LQDC of length $m = 2q \cdot 2^{O(q \log(2q) \cdot n^{1/(2q-1)})}$.*

For example, for every n, the Hadamard code is a $(1, \delta, 1/2 - 2\delta)$-LQDC of exponential length (which is optimal by Theorem 3). For $q = 2$ it suffices to have length $m = 2^{O(n^{1/3})}$, for $q = 3$ it suffices to have $m = 2^{O(n^{1/5})}$, etc.

Accordingly, for even k, the best known (k, δ, ε)-LDCs just happen to be $(k/2, \delta, \varepsilon)$-LQDCs, because they happen to output the XOR of their k queries. For more general LDCs we can do something nearly as good, using van Dam's result that a k-bit oracle can be recovered with probability nearly 1 using $k/2 + O(\sqrt{k})$ quantum queries [9]:

Theorem 7 *A (k, δ, ϵ)-LDC is a $(k/2 + O(\sqrt{k}), \delta, \varepsilon/2)$-LQDC.*

5 Private Information Retrieval

As mentioned, there is a close connection between locally decodable codes and private information retrieval. Our techniques also allow us to give new lower bounds for 2-server PIRs, but only for PIRs having answer length 1. Again we give a 2-step proof: a reduction of 2 classical servers to 1 quantum server, combined with a lower bound for quantum 1-server PIR.

Theorem 8 *If there exists a classical 2-server PIR scheme with t-bit queries, 1-bit answers, and recovery probability $1/2 + \varepsilon$, then there exists a quantum 1-server PIR scheme with $(t + 2)$-qubit queries, $(t + 2)$-qubit answers, and recovery probability $1/2 + 4\varepsilon/7$.*

Proof. The proof is analogous to the proof for locally decodable codes. If we fix the classical user's randomness, the problem boils down to computing some $f(a_1, a_2)$, where a_1 is the first server's 1-bit answer to query q_1, and a_2 is the second server's 1-bit answer to query q_2. However, in addition we now have to hide i from the quantum server. This we do by making the quantum user set up the $(4 + t)$-qubit state

$$\frac{1}{\sqrt{3}}\left(|0, 0, 0^t\rangle + |1, 1, q_1\rangle + |2, 2, q_2\rangle\right),$$

where '0^t' is a string of t 0s. The user sends everything but the first 2 qubits to the server. The state of the server is now a uniform mixture of $|0, 0^t\rangle$, $|1, q_1\rangle$, and $|2, q_2\rangle$. By the security of the classical protocol, $|1, q_1\rangle$ contains no information about i (averaged over the user's randomness), and the same holds for $|2, q_2\rangle$. Therefore the uniform mixture of $|0, 0^t\rangle$, $|1, q_1\rangle$, and $|2, q_2\rangle$ contains no information about i.

The quantum server then puts $(-1)^{a_s}$ in front of $|s, q_s\rangle$ ($s \in \{1, 2\}$), leaves $|0, 0^t\rangle$ alone, and sends everything back. Note that we need to supply the name of the classical server $s \in \{1, 2\}$ to tell the server in superposition whether it should play the role of server 1 or 2. The user now has

$$\frac{1}{\sqrt{3}}\left(|0, 0, 0^t\rangle + (-1)^{a_1}|1, 1, q_1\rangle + (-1)^{a_2}|2, 2, q_2\rangle\right).$$

From this we can compute $f(a_1, a_2)$ with success probability exactly $11/14$, giving overall recovery probability $1/2 + 4\varepsilon/7$ as before. □

Nayak [16] proved (see Appendix B; a proof may also be found in [12]):

Theorem 9 (Nayak) *A quantum 1-server PIR scheme with recovery probability p has communication complexity at least $(1 - H(p))n$.*

Combining the above two theorems, we obtain the first linear lower bound that holds for all 1-bit-answer 2-server PIRs, not just for linear ones.

Theorem 10 *A classical 2-server PIR scheme with t-bit queries, 1-bit answers, and recovery probability $1/2 + \varepsilon$, has $t \geq \frac{1}{2}(1 - H(1/2 + 4\varepsilon/7))n - 2$.*

Acknowledgments

We thank Harry Buhrman, Ashwin Nayak, Kenji Obata, Ashish Thapliyal, and Luca Trevisan for helpful discussions.

References

[1] A. Ambainis. Upper bound on communication complexity of private information retrieval. In *Proceedings of the 24th ICALP*, volume 1256 of *Lecture Notes in Computer Science*, pages 401–407, 1997.

[2] L. Babai, L. Fortnow, L. Levin, and M. Szegedy. Checking computations in polylogarithmic time. In *Proceedings of 23rd ACM STOC*, pages 21–31, 1991.

[3] A. Beimel and Y. Ishai. Information-theoretic private information retrieval: A unified construction. In *Proceedings of 28th ICALP*, pages 912–926, 2001. References are to the longer version on ECCC.

[4] A. Beimel, Y. Ishai, E. Kushilevitz, and J. Raymond. Breaking the $O(n^{1/(2k-1)})$ barrier for information-theoretic Private Information Retrieval. In *Proceedings of 43rd IEEE FOCS*, 2002. To appear.

[5] E. Bernstein and U. Vazirani. Quantum complexity theory. *SIAM Journal on Computing*, 26(5):1411–1473, 1997. Earlier version in STOC'93.

[6] H. Buhrman, P. B. Miltersen, J. Radhakrishnan, and S. Venkatesh. Are bitvectors optimal? In *Proceedings of 32nd ACM STOC*, pages 449–458, 2000.

[7] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan. Private information retrieval. *Journal of the ACM*, 45(6):965–981, 1998. Earlier version in FOCS 95.

[8] R. Cleve, W. van Dam, M. Nielsen, and A. Tapp. Quantum entanglement and the communication complexity of the inner product function. In *Proceedings of 1st NASA QCQC conference*, volume 1509 of *Lecture Notes in Computer Science*, pages 61–74. Springer, 1998. quant-ph/9708019.

[9] W. van Dam. Quantum oracle interrogation: Getting all information for almost half the price. In *Proceedings of 39th IEEE FOCS*, pages 362–367, 1998. quant-ph/9805006.

[10] A. Deshpande, R. Jain, T. Kavita, S. Lokam, and J. Radhakrishnan. Better lower bounds for locally decodable codes. In *Proceedings of 17th IEEE Conference on Computational Complexity*, pages 184–193, 2002.

[11] O. Goldreich, H. Karloff, L. Schulman, and L. Trevisan. Lower bounds for linear locally decodable codes and private information retrieval. In *Proceedings of 17th IEEE Conference on Computational Complexity*, pages 175–183, 2002. Also on ECCC.

[12] R. Jain, J. Radhakrishnan, and P. Sen. Privacy and interaction in quantum communication complexity and a theorem about the relative entropy of quantum states. In *Proceedings of 43rd IEEE FOCS*, 2002. To appear.

[13] J. Katz and L. Trevisan. On the efficiency of local decoding procedures for error-correcting codes. In *Proceedings of 32nd ACM STOC*, pages 80–86, 2000.

[14] I. Kremer. Quantum communication. Master's thesis, Hebrew University, Computer Science Department, 1995.

[15] H-K. Lo. Insecurity of quantum secure computations. *Physical Review A*, 56:1154, 1997. quant-ph/9611031.

[16] A. Nayak. Optimal lower bounds for quantum automata and random access codes. In *Proceedings of 40th IEEE FOCS*, pages 369–376, 1999. quant-ph/9904093.

[17] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information.* Cambridge University Press, 2000.

[18] K. Obata. Optimal lower bounds for 2-query locally decodable linear codes. In *Proceedings of 6th RANDOM*, 2002. To appear.

[19] J. Radhakrishnan, P. Sen, and S. Venkatesh. The quantum complexity of set membership. In *Proceedings of 41st IEEE FOCS*, pages 554–562, 2000. quant-ph/0007021.

[20] M. Sipser and D. A. Spielman. Expander codes. *IEEE Transactions on Information Theory*, 42:1710–1722, 1996. Earlier version in FOCS'94.

[21] A. C-C. Yao. Quantum circuit complexity. In *Proceedings of 34th IEEE FOCS*, pages 352–360, 1993.

A Lower Bound for Quantum Random Access Codes

As mentioned before, a quantum random access code is an encoding $x \mapsto \rho_x$, such that any bit x_i can be recovered with some probability $p \geq 1/2 + \varepsilon$ from ρ_x. Below we reprove Nayak's [16] linear lower bound on the length m of such encodings.

We assume familiarity with the following notions from quantum information theory, referring to [17, Chapters 11 and 12] for more details. Very briefly, if we have a bipartite quantum system AB (given by some density matrix), then we use A and B to denote the states (reduced density matrices) of the individual systems; $S(A) = -\mathrm{Tr}(A \log A)$ is the *(Von Neumann) entropy* of A; $S(A|B) = S(AB) - S(B)$ is the *conditional entropy* of A given B; and $S(A:B) = S(A) + S(B) - S(AB) = S(A) - S(A|B)$ is the *mutual information* between A and B.

We define an $n+m$-qubit state XM as follows:

$$\frac{1}{2^n} \sum_{x \in \{0,1\}^n} |x\rangle\langle x| \otimes \rho_x.$$

We use X to denote the first subsystem, X_i for its individual bits, and M for the second subsystem. By [17, Theorem 11.8.4] we have

$$S(XM) = n + \frac{1}{2^n}\sum_x S(\rho_x) \geq n = S(X).$$

Since M has m qubits we have $S(M) \leq m$, hence

$$S(X:M) = S(X) + S(M) - S(XM) \leq S(M) \leq m.$$

Using a chain rule for relative entropy we get

$$S(X|M) = \sum_{i=1}^{n} S(X_i|X_1 \dots X_{i-1}M) \leq \sum_{i=1}^{n} S(X_i|M).$$

Since we can predict X_i from M with success probability p, Fano's inequality implies

$$H(p) \geq S(X_i|M).$$

Putting the above equations together we obtain

$$(1 - H(p))n \leq S(X) - \sum_{i=1}^{n} S(X_i|M) \leq S(X) - S(X|M) = S(X:M) \leq m.$$

B Lower Bound for Quantum PIR

Here we give Nayak's lower bound on quantum 1-server PIR schemes that have good recovery probability, following [16, Section 4.4] with one additional ingredient from [15]. Though we will only need the result for *1-round* quantum PIR schemes, the proof actually applies equally well to multi-round protocols.

Consider a quantum 1-server PIR scheme with recovery probability p. Without loss of generality, we assume the only measurement in the protocol is the user's measurement on his part of the final state to recover x_i. If the communication is m qubits, then the final state on inputs x and i can be written as [21, 14]:

$$|\phi_{xi}\rangle = \sum_{k\in\{0,1\}^m} \underbrace{|a_k(x)\rangle}_{\text{server}} \underbrace{|b_k(i)\rangle}_{\text{user}},$$

where $|a_k(x)\rangle$ is a vector (not necessarily normalized) that depends on x but not on i, and similarly for $|b_k(i)\rangle$. Note that the user's part of the state $|\phi_{xi}\rangle$ lives in a 2^m-dimensional space that is independent of x. From his part of the state, the user can recover x_i with success probability p.

Now suppose the server starts with a uniform superposition over all x. Then the final state can be written as

$$|\phi_i\rangle = \frac{1}{\sqrt{2^n}} \sum_{x\in\{0,1\}^n} |x\rangle|\phi_{xi}\rangle,$$

where we assumed that the server keeps a copy of x around. The privacy constraint on the protocol means that the server's part of each $|\phi_{xi}\rangle$ is independent of i, therefore the server's part of $|\phi_i\rangle$ is independent of i as well. But then there exists a unitary U_{ij} on the user's part of the state such that $(I \otimes U_{ij})|\phi_i\rangle = |\phi_j\rangle$. This $I \otimes U_{ij}$ must then actually map $|x\rangle|\phi_{xi}\rangle$ to $|x\rangle|\phi_{xj}\rangle$ for all x. Defining ρ_x as the user's part of $|\phi_{x1}\rangle$, we obtain a quantum random access code with recovery probability p: for any j, the user can apply U_{1j} to ρ_x to obtain a state from which he can recover x_j with probability p. All the ρ_x lie in the same 2^m-dimensional space $span\{|b_k(1)\rangle : k \in \{0,1\}^m\}$, therefore we just need m qubits to represent them. Applying the lower bound on random access codes concludes the proof.

ECCC ISSN 1433-8092
http://www.eccc.uni-trier.de/eccc
ftp://ftp.eccc.uni-trier.de/pub/eccc
ftpmail@ftp.eccc.uni-trier.de, subject 'help eccc'