

LOCALLY DECODABLE QUANTUM CODES

JOP BRIËT¹ AND RONALD DE WOLF¹

¹ CWI. Kruislaan 413. 1098SJ Amsterdam. The Netherlands.

E-mail address: {jop.briet, rdewolf}@cwi.nl

ABSTRACT. We study a quantum analogue of locally decodable error-correcting codes. A q -query *locally decodable quantum code* encodes n classical bits in an m -qubit state, in such a way that each of the encoded bits can be recovered with high probability by a measurement on at most q qubits of the quantum code, even if a constant fraction of its qubits have been corrupted adversarially. We show that such a quantum code can be transformed into a *classical* q -query locally decodable code of the same length that can be decoded well on average (albeit with smaller success probability and noise-tolerance). This shows, roughly speaking, that q -query quantum codes are not significantly better than q -query classical codes, at least for constant or small q .

1. Introduction

Locally decodable codes (LDCs) have received much attention in the last decade. They are error-correcting codes that encode n bits into m bits, with the usual error-correcting properties, and the additional feature that any one of the n encoded bits can be recovered (with high probability) by a randomized decoder that queries at most q bits in the codeword, for some small q . In other words, to decode small parts of the encoded data, we only need to look at a small part of the codeword instead of “unpacking” the whole thing. Precise definitions will be given in the next sections. Such codes are potentially useful in their own right (think of decoding small pieces from a large encoded library), and also have a variety of applications in complexity theory and cryptography. For instance, it is well known that they can be turned into private information retrieval schemes and vice versa. For further details about such connections, we refer to Trevisan’s survey [Tre04] and the references therein.

The most interesting question about LDCs is the tradeoff between their length m and the number of queries q . The former measures the space efficiency of the code, while the latter measures the efficiency of decoding. The larger we make q , the smaller we can

1998 ACM Subject Classification: E1, E4, F1.

Key words and phrases: Data structures, locally decodable codes, quantum computing.

Jop Briët is partially supported by a Vici grant from the Netherlands Organization for Scientific Research (NWO). Ronald de Wolf is partially supported by Veni and Vidi grants from the Netherlands Organization for Scientific Research (NWO). Both authors are partially supported by the European Commission under the Integrated Project Qubit Applications (QAP) funded by the IST directorate as Contract Number 015848.

make m . On one extreme, if we allow $q = \text{polylog}(n)$ queries, the codelength m can be made polynomial in n [BFLS91]. On the other extreme, for $q = 1$ and sufficiently large n , LDCs do not exist at all [KT00]. For $q = 2$ they do exist but need exponential length, $m = \exp(n)$ [KW04]. Between these two extremes, interesting but hard questions persist. In particular, we know little about the length of LDCs with constant $q > 2$. The best upper bounds are based on Yekhanin's construction [Yek07]. He gives 3-query LDCs with length $m = \exp(n^{1/t})$ for every *Mersenne prime* $p = 2^t - 1$. (The largest known Mersenne prime has $t = 32582657$, but it has been conjectured that there are infinitely many.) Recently, Efremenko [Efr08] used Yekhanin's basic underlying combinatorial structure to obtain, for integer $r \geq 2$, 2^r -query LDCs with $m = \exp(\exp(O(\log n \log \log^{r-1} n)^{1/r}))$, and 3-query LDCs with $m = \exp(\exp(O(\sqrt{\log n \log \log n}))$). On the lower bound side, the best we know for $q > 2$ is $m = \Omega((n/\log n)^{1+1/(\lceil q/2 \rceil - 1)})$ (for fixed success probability and noise rate) [KT00, KW04, Woo06]. For $q = 3$ and $q = 4$, this is slightly less than n^2 .

Interestingly, the best known lower bounds were obtained using tools from quantum information theory. It is thus a natural question to consider also the potential *positive* effects of quantum: can we construct shorter q -query LDCs by somehow harnessing the power of quantum states and quantum algorithms? There are two natural ways to generalize locally decodable codes to the quantum world:

- We can keep the code classical, but allow q quantum queries. This means we can query positions of the codeword in quantum superposition, and process the results using quantum circuits. This approach was investigated in [KW04]. A q -query quantum decoder can simulate a $2q$ -query classical decoder with high success probability, and this simulation can be made exact if the classical decoder took the parity of its $2q$ bits. This implies for instance that Efremenko's 3-query LDC can be decoded by only 2 quantum queries. In contrast, we know that every 2-query LDC needs length $\exp(n)$. Allowing quantum queries thus results in very large savings in m when we consider a fixed number of queries q .
- We can also make the code itself quantum: instead of encoding an n -bit x into an m -bit string $C(x)$, we could encode it into an m -qubit state $Q(x)$. A q -query decoder for such a code would select up to q qubits of the state $Q(x)$, and make a 2-outcome measurement on those qubits to determine its output. In this case our notion of noise also needs to be generalized: instead of up to δm bitflip-errors, we allow any set of up to δm qubits of $Q(x)$ to be arbitrarily changed.¹

¹While a classical LDC can be reused as often as we want, a quantum code has the problem that a measurement made to predict one bit changes the state, so predicting another bit based on the changed state may give the wrong results. However, if the error probability is small then the changes incurred by each measurement will be small as well, and we can reuse the code many times with reasonable confidence. Another issue is that more general decoders could be allowed. For instance, we could consider allowing any quantum measurement on the m -qubit state that can be written as a linear combination of m -qubit Pauli-matrices that have support on at most q positions. This is potentially stronger than what we do now (it is an interesting open question whether it is really stronger). However, we feel this is a somewhat unnatural formalization of the idea that a measurement should be localized to at most q qubits. Our current set-up, where we classically select up to q positions and then apply an arbitrary quantum measurement to those q qubits, seems more natural. Similarly, more general noise operators can be defined, but we will not consider these here.

Our results. In this paper we investigate the second kind of code, which we call a “ q -query locally decodable quantum code”, or q -query LDQC. The question is whether the ability to encode our n bits into a *quantum* state enables us to make codes much shorter. There are some small examples where quantum encodings achieve things that are impossible for classical encodings. Ambainis et al. [ANTV02] give an encoding of 2 classical bits into 1 qubit, such that each of the bits—though not both simultaneously—can be recovered from the qubit with success probability 0.85. They even cite an example due to Chuang where 3 bits are encoded into 1 qubit, and each bit can be recovered with success probability 0.78. However, they also show that asymptotically large savings are not possible in their setting (explained in Section 3.4 below). Their setting, however, considers neither noise nor local decodability, and hence does not answer our question about locally decodable codes: can LDCs be made significantly shorter if we allow quantum encodings?

Our main result is a negative answer to this question: essentially it says that q -query LDQCs can be turned into classical q -query LDCs of the same length, with some deterioration in their other parameters. The precise statement of this result (Corollaries 5.3 and 5.4) is a little bit dirty. We obtain a cleaner statement for so-called “smooth (quantum) codes”, which have the property that they query the codewords fairly uniformly. These smooth (quantum) codes can be converted into LD(Q)Cs and vice versa. For these, the precise statement is as follows (Theorem 5.1). Suppose we are given a smooth quantum code of m qubits from which we can recover (with success probability at least $1/2 + \epsilon$) each bit x_i of the encoded n -bit string x , while only looking at q qubits of the state. Let μ be a distribution on the n -bit inputs. Then we can construct a *randomized* classical code R of the same length (for each x , the “codeword” $R(x)$ is a *distribution* over m -bit strings) from which we can recover each x_i with μ -average success probability at least $1/2 + \epsilon/4^{q+1}$, while only looking at q bits of the codeword. Thus a q -query quantum code is turned into a q -query classical code of the same length, at the expense of reducing the advantage of the algorithm (over random guessing) by a factor of roughly 4^q .²

For those who do not like the idea of encoding x into a *distribution* $R(x)$, we can turn the randomized code R into a *deterministic* code C , where $C(x)$ is a fixed m -bit codeword instead of a distribution, at the expense of correctly decoding only a constant fraction of all indices i instead of all n of them (Corollary 5.2). Since all known lower bounds on LDCs also apply to randomized classical codes that work well under a uniform distribution on the n -bit strings, those lower bounds carry over to LDQCs. In particular we obtain as corollaries of our result:

- For sufficiently large n , 1-query LDQCs do not exist for any m (from [KT00]).³
- 2-query LDQCs need $m = \exp(n)$ (from [KW04]).
- For fixed q , q -query LDQCs need $m = \Omega\left(\left(\frac{n}{\log n}\right)^{1+1/(\lceil q/2 \rceil - 1)}\right)$ (from [KW04]).

Techniques. Our main technique is to apply to the m qubits of the quantum code a randomly selected sequence of m Pauli measurements. The randomized “codeword” $R(x)$ will be the probability distribution on m -bit outcomes resulting from such a measurement on the quantum state $Q(x)$. The meat of our proof is to show that there is a choice of measurements that roughly preserves correct decodability for all i .

²Oded Regev showed us how to improve this to 3^q , but we won’t give the details of his improvement here.

³Actually, this result can more easily be shown directly, by combining Katz and Trevisan’s proof for classical codes with the quantum random access code lower bound mentioned below in Section 3.4.

2. Preliminaries

We write $[n]$ for the set $\{1, \dots, n\}$. We use $\mathcal{P}(S)$ to denote the set of all probability distributions (or random variables) on set S . If z is distributed according to the distribution of a random variable Z , we write $z \sim Z$. Probabilities and expectations with a subscript ' $i \in S$ ' should be read as taken over a uniformly random $i \in S$. We give a brief overview of quantum mechanics here, see [NC00] for more.

Quantum states. In quantum mechanics, a physical system is mathematically represented by a complex Hilbert space. A d -dimensional complex Hilbert space consists of all d -dimensional vectors with complex entries, endowed with the standard inner product. The *state* of a physical system is in turn represented by a density operator (a positive semidefinite linear operator with trace 1) acting on a Hilbert space. We use $\mathcal{B}_+^1(\mathcal{H}_d)$ to denote the set of all density operators on a d -dimensional complex Hilbert space. Two-dimensional Hilbert spaces are called *qubits*. An n -qubit state is a density operator on the tensor product of n qubits (a 2^n -dimensional Hilbert space).

Measurements. The most general k -outcome *measurement* on a physical system is defined as a set $\{A_1, \dots, A_k\}$ of k positive semidefinite matrices that satisfy $\sum_{i=1}^k A_i = I$. The probability that the measurement of a system in a state ρ yields the i 'th outcome is $\text{Tr}(A_i \rho)$. Hence, the measurement yields a random variable $A(\rho)$ with $\Pr[A(\rho) = i] = \text{Tr}(A_i \rho)$. With a measurement that has outcomes $+1$ and -1 (and corresponding operators A^+ and A^-) we associate an operator $A = A^+ - A^-$. The expected value of this measurement on a state ρ is then $\text{Tr}(A\rho)$. Note that this equals the difference between the probabilities of outcomes $+1$ and -1 , respectively.

Pauli matrices. The one-qubit Pauli operators are given by $I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$, $X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, $Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}$ and $Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$. For integer $k \geq 1$, the set of k -qubit Pauli operators is $\mathcal{P}_k := \{I, X, Y, Z\}^{\otimes k}$. These 4^k matrices form an orthonormal basis for the space of all $2^k \times 2^k$ complex matrices endowed with the inner product $\langle A, B \rangle = \frac{1}{2^k} \text{Tr}(A^\dagger B)$. Each Pauli operator $S \in \mathcal{P}_k$ has a unique decomposition $S = S^+ - S^-$, with S^+ and S^- orthogonal projectors that satisfy $S^+ + S^- = I$. For this reason we associate a unique two-outcome measurement $\{S^+, S^-\}$ with each such S . A Pauli measurement $S \in \mathcal{P}_k$ of a k -qubit state ρ yields a ± 1 -valued random variable $S(\rho)$ with expected value $\text{Tr}(S\rho)$. However, we can also view $S \in \mathcal{P}_k$ as k separate one-qubit Pauli measurements, to be applied to the k qubits of the state, respectively. When viewed in this way, the result of measuring ρ is a k -bit random variable, i.e., a probability distribution on $\{\pm 1\}^k$. The product of those k bits equals the ± 1 -valued random variable $S(\rho)$ mentioned before.

Super-operators. A super-operator is a mathematical representation of the most general transformation of a quantum state allowed by the laws of quantum mechanics. A super-operator \mathcal{E} can be defined by a finite set $\{E_1, \dots, E_k\}$ of linear operators (known as Kraus operators) that satisfy $\sum_{i=1}^k E_i^\dagger E_i = I$. The corresponding operation on a state ρ yields another density operator, $\mathcal{E}(\rho) = \sum_{i=1}^k E_i \rho E_i^\dagger$. This $\mathcal{E}(\rho)$ may act on a Hilbert space of a possibly different dimension, though we will not need that here. We say that \mathcal{E} "acts trivially" on, say, the first qubit of the state if all its Kraus operators have the form $E_i = I \otimes E'_i$ for some E'_i acting on all but the first qubit.

3. Codes

3.1. Classical codes

For convenience, we write bits as ± 1 instead of 0/1. This way, if random variable $A \in \{\pm 1\}$ predicts bit x_i , we can write the bias of this prediction as an expectation: $\mathbb{E}[A \cdot x_i] = \Pr[A = x_i] - \Pr[A \neq x_i]$. Note that $\Pr[A = x_i] \geq 1/2 + \varepsilon$ iff $\mathbb{E}[A \cdot x_i] \geq 2\varepsilon$.

We start with classical codes. The formal definition of a locally decodable code involves a decoder \mathcal{A} that receives input $i \in [n]$ and oracle access to a string $y \in \{\pm 1\}^m$, usually written as a superscript to \mathcal{A} . This y will be a codeword $C(x) \in \{\pm 1\}^m$ corrupted by some “error string” $E \in \{\pm 1\}^m$, which negates some of the bits of $C(x)$ (below, $C(x) \circ E$ denotes the entry-wise product of the two m -bit vectors $C(x)$ and E). The oracle “queries index $j \in [m]$ ” if it reads the j ’th bit of y .

Definition 3.1 (LDC). A function $C : \{\pm 1\}^n \rightarrow \{\pm 1\}^m$ is a (q, δ, ε) -LDC if there exists a probabilistic oracle algorithm \mathcal{A} such that

- (1) For every $x \in \{\pm 1\}^n$, every $i \in [n]$, and every $E \in \{\pm 1\}^m$ with at most δm -1 ’s, we have $\Pr[\mathcal{A}^{C(x) \circ E}(i) = x_i] \geq 1/2 + \varepsilon$, where the probability is taken over the internal coin tosses of \mathcal{A} .
- (2) \mathcal{A} queries at most q indices of y . Queries are made non-adaptively, meaning that the indices to be queried are all selected before the querying starts.

An \mathcal{A} satisfying the above is called a (q, δ, ε) -local decoder for C .

Since any δm indices can be corrupted, a local decoder must query the indices fairly uniformly. Otherwise, an adversary could choose to corrupt the most queried part of the code and ruin the decoder’s success probability. Motivated by this property, Katz and Trevisan [KT00] defined a variation of LDCs called a *smooth codes*, defined only for uncorrupted codewords.

Definition 3.2 (Smooth code). A function $C : \{\pm 1\}^n \rightarrow \{\pm 1\}^m$ is a (q, c, ε) -smooth code if there exists a probabilistic oracle algorithm \mathcal{A} such that:

- (1) For every $x \in \{\pm 1\}^n$ and $i \in [n]$, we have $\Pr[\mathcal{A}^{C(x)}(i) = x_i] \geq 1/2 + \varepsilon$.
- (2) For every $i \in [n]$ and $j \in [m]$, we have $\Pr[\mathcal{A}^{(\cdot)}(i) \text{ queries index } j] \leq c/m$.
- (3) \mathcal{A} queries at most q indices (non-adaptively).

An \mathcal{A} satisfying the above is called a (q, c, ε) -smooth decoder for C .

Katz and Trevisan showed that LDCs and smooth codes are essentially equivalent, in the sense that a decoder for one can be transformed into a decoder for the other. We prove the same for quantum codes in Section 3.3, using essentially their proof.

3.2. Randomized codes

Here we define our first generalization, incorporating randomness into the definition of the code. A *randomized locally decodable code* maps $\{\pm 1\}^n$ to *random variables* over $\{\pm 1\}^m$ (rather than fixed codewords), such that any x_i can be decoded well using a constant number of queries, even if up to δm indices are corrupted.

Definition 3.3 (Randomized LDC). A function $R : \{\pm 1\}^n \rightarrow \mathcal{P}(\{\pm 1\}^m)$ is a (q, c, ε) -randomized LDC if there exists a probabilistic oracle algorithm \mathcal{A} such that:

- (1) For every $x \in \{\pm 1\}^n$, every $i \in [n]$, and every $E \in \{\pm 1\}^m$ with at most δm -1 's, we have $\Pr[\mathcal{A}^{R(x) \circ E}(i) = x_i] \geq 1/2 + \varepsilon$, where the probability is taken over the internal coin tosses of \mathcal{A} as well as the distribution $R(x)$.
- (2) \mathcal{A} queries at most q indices (non-adaptively).

Similarly, we define a *randomized smooth code*:

Definition 3.4 (Randomized smooth code). A function $R : \{\pm 1\}^n \rightarrow \mathcal{P}(\{\pm 1\}^m)$ is a (q, c, ε) -*randomized smooth code* if there exists a probabilistic oracle algorithm \mathcal{A} such that:

- (1) For every $x \in \{\pm 1\}^n$ and every $i \in [n]$, $\Pr[\mathcal{A}^{R(x)}(i) = x_i] \geq 1/2 + \varepsilon$.
- (2) For every $i \in [n]$, and every $j \in [m]$, $\Pr[\mathcal{A}^{(\cdot)}(i)$ queries index $j] \leq c/m$.
- (3) \mathcal{A} queries at most q indices (non-adaptively).

It will be convenient to also have a version of these codes that are only required to work well on average, instead of for all x :

Definition 3.5 (μ -average codes). Let μ be a distribution on $\{\pm 1\}^n$. A function $C : \{\pm 1\}^n \rightarrow \{\pm 1\}^m$ is a μ -*average* (q, δ, ε) -*LDC* if Definition 3.1 holds with the first clause replaced by:

- (1) For every $i \in [n]$ and $E \in \{\pm 1\}^m$ with at most δm -1 's, $\Pr_{x \sim \mu}[\mathcal{A}^{C(x) \circ E}(i) = x_i] \geq \frac{1}{2} + \varepsilon$.

Analogously, we define μ -average versions of smooth codes, randomized LDCs, and randomized smooth codes. For these codes, we assume without loss of generality that for each i and queried set $r \subseteq [m]$, the decoder \mathcal{A} always uses the same function $f_{i,r} : \{\pm 1\}^q \rightarrow \{\pm 1\}$ to determine its output.

A μ -average randomized smooth code can actually be “derandomized” to a μ -average smooth code on a smaller number of bits:

Lemma 3.6. *Let $R : \{\pm 1\}^n \rightarrow \mathcal{P}(\{\pm 1\}^m)$ be a μ -average (q, c, ε) -randomized smooth code. Then there exists a μ -average $(q, c, \varepsilon/2)$ -smooth code $C : \{\pm 1\}^n \rightarrow \{\pm 1\}^m$ for at least εn of the indices i (that is, a smooth code with μ -success probability at least $1/2 + \varepsilon/2$ for at least εn of the n indices).*

Proof. As a first step we will view R as a function to strings: there exists a random variable W (over some possibly infinite set \mathcal{W}) and a function $R : \{\pm 1\}^n \times \mathcal{W} \rightarrow \{\pm 1\}^m$ such that for every $x \in \{\pm 1\}^n$, the random variables $R(x, W)$ and $R(x)$ are the same. A decoder \mathcal{A} for R also works for $R(\cdot, W)$, so we have bias $\mathbb{E}_{x \sim \mu, w \sim W}[\mathcal{A}^{R(x,w)}(i) \cdot x_i] \geq 2\varepsilon$ for every $i \in [n]$. For every $i \in [n]$ and $w \in \mathcal{W}$, define variables $X_{i,w} \in \{0, 1\}$, with

$$X_{i,w} = 1 \iff \mathbb{E}_{x \sim \mu}[\mathcal{A}^{R(x,w)}(i) \cdot x_i] \geq \varepsilon,$$

and $X_w := \sum_{i=1}^n X_{i,w}$. Using the definition of a μ -average randomized smooth code:

$$\begin{aligned} 2\varepsilon n &\leq \sum_{i=1}^n \mathbb{E}_{x \sim \mu, w \sim W}[\mathcal{A}^{R(x,w)}(i) \cdot x_i] = \mathbb{E}_{w \sim W} \left[\sum_{i=1}^n \mathbb{E}_{x \sim \mu}[\mathcal{A}^{R(x,w)}(i) \cdot x_i] \right] \\ &< \mathbb{E}_{w \sim W}[X_w + (n - X_w)\varepsilon] = \varepsilon n + (1 - \varepsilon)\mathbb{E}_{w \sim W}[X_w]. \end{aligned}$$

Hence $\mathbb{E}_{w \sim W}[X_w] \geq \varepsilon n$. Thus there exists a $w \in \mathcal{W}$ such that for at least εn of the n indices i , we have $\mathbb{E}_{x \sim \mu}[\mathcal{A}^{R(x,w)}(i) \cdot x_i] \geq \varepsilon$, equivalently, $\mathbb{E}_{x \sim \mu}[\Pr[\mathcal{A}^{R(x,w)}(i) = x_i]] \geq 1/2 + \varepsilon/2$. Defining the code $C(\cdot) := R(\cdot, w)$ gives the lemma. \blacksquare

3.3. Quantum codes

Our second level of generalization brings quantum mechanics into the picture: now our code maps classical n -bit strings to m -qubit quantum states. Below, a “quantum oracle algorithm” is an algorithm \mathcal{A} with oracle access to an m -qubit state ρ . This ρ could be a corrupted version of an m -qubit “codeword” $Q(x)$, obtained by applying some super-operator \mathcal{E} to $Q(x)$. This \mathcal{E} should only affect a δ -fraction of the m qubits. This error model generalizes the classical case: a classical error pattern $E \in \{\pm 1\}^m$ corresponds to a super-operator \mathcal{E} that applies an X to the qubits at positions where E has a -1 , and I to the positions where E has a $+1$. On input $i \in [n]$, the algorithm probabilistically selects a set $r \subseteq [m]$ of at most q indices of qubits of ρ , and applies a two-outcome measurement to the selected qubits with operators $A_{i,r}^+$ and $A_{i,r}^-$. As before, we will use “ $\mathcal{A}^\rho(i)$ ” to denote the ± 1 -valued random variable that is the output. We say that “ \mathcal{A} queries r ”, and “ \mathcal{A} queries index j ” if j is in r . Note that such algorithms are non-adaptive by definition: the set of qubits r is selected before it is measured. We now define a *locally decodable quantum code* (LDQC) as follows:

Definition 3.7 (LDQC). A function $Q : \{\pm 1\}^n \rightarrow \mathcal{B}_+^1(\mathcal{H}_{2^m})$ is a (q, δ, ε) -LDQC if there is a quantum oracle algorithm \mathcal{A} s.t.:

- (1) For every $x \in \{\pm 1\}^n$, every $i \in [n]$, and every super-operator \mathcal{E} that acts non-trivially on at most δm qubits, we have $\Pr[\mathcal{A}^{\mathcal{E}(Q(x))}(i) = x_i] \geq 1/2 + \varepsilon$, where the probability is taken over the coin tosses and measurements in \mathcal{A} .
- (2) \mathcal{A} queries at most q indices (non-adaptively).

An \mathcal{A} satisfying the above is called a (q, δ, ε) -local quantum decoder for Q .

LDQCs generalize randomized LDCs, because probability distributions are diagonal density operators. We can also establish a smoothness property for quantum codes:

Definition 3.8 (Smooth quantum code). A function $Q : \{\pm 1\}^n \rightarrow \mathcal{B}_+^1(\mathcal{H}_{2^m})$ is a (q, c, ε) -smooth quantum code if there exists a quantum oracle algorithm \mathcal{A} such that:

- (1) For every $x \in \{\pm 1\}^n$ and every $i \in [n]$, we have $\Pr[\mathcal{A}^{Q(x)}(i) = x_i] \geq 1/2 + \varepsilon$.
- (2) For every $i \in [n]$ and every $j \in [m]$, we have $\Pr[\mathcal{A}^{(\cdot)}(i) \text{ queries index } j] \leq c/m$.
- (3) \mathcal{A} queries at most q indices (non-adaptively).

An \mathcal{A} satisfying the above is called a (q, c, ε) -smooth quantum decoder for Q .

As Katz and Trevisan [KT00] did for classical LDCs, we can establish a strong connection between LDQCs and smooth quantum codes. Either one can be used as the other, as the next theorems show. Analogues of these theorems also hold between randomized LDCs and randomized smooth codes, and between the μ -average versions of these codes.

Theorem 3.9. *Let $Q : \{\pm 1\}^n \rightarrow \mathcal{B}_+^1(\mathcal{H}_{2^m})$ be a (q, c, ε) -smooth quantum code. Then, as long as $\delta \leq \varepsilon/c$, Q is also a $(q, \delta, \varepsilon - \delta c)$ -locally decodable quantum code.*

Proof. Let \mathcal{A} be a (q, c, ε) -smooth quantum decoder for Q . Suppose we run it on $\mathcal{E}(Q(x))$ with at most δm corrupted qubits. The probability that \mathcal{A} queries a specific qubit is at most c/m . Then by the union bound, the probability that \mathcal{A} queries any of the corrupted qubits is at most $\delta mc/m = \delta c$. Hence \mathcal{A} itself is also a $(q, \delta, \varepsilon - \delta c)$ -local quantum decoder for Q . ■

Theorem 3.10. *Let $Q : \{\pm 1\}^n \rightarrow \mathcal{B}_+^1(\mathcal{H}_{2^m})$ be a (q, δ, ε) -locally decodable quantum code. Then Q is also a $(q, q/\delta, \varepsilon)$ -smooth quantum code.*

Proof. Let \mathcal{A} be a (q, δ, ε) -local quantum decoder for Q . For each $i \in [n]$, let $p_i(j)$ be the probability that on input i , \mathcal{A} queries qubit j . Let $H_i = \{j \mid p_i(j) > q/(\delta m)\}$. Then $|H_i| \leq \delta m$, because \mathcal{A} queries no more than q indices. Let \mathcal{B} be the quantum decoder that simulates \mathcal{A} , except that on input i it does not query qubits in H_i , but instead acts as if those qubits are in a completely mixed state. Then \mathcal{B} does not measure any qubit j with probability greater than $q/(\delta m)$. Also, \mathcal{B} 's behavior on input i and $Q(x)$ is the same as \mathcal{A} 's behavior on input i and $\mathcal{E}(Q(x))$ that is obtained by replacing all qubits in H_i by completely mixed states. Since \mathcal{E} acts non-trivially on at most $|H_i| \leq \delta m$ qubits, we have $\Pr[\mathcal{B}^{Q(x)}(i) = x_i] = \Pr[\mathcal{A}^{\mathcal{E}(Q(x))}(i) = x_i] \geq 1/2 + \varepsilon$. ■

3.4. A weak lower bound from random access codes

We can immediately establish a weak lower bound on the length of LDQCs and smooth quantum codes by considering a *quantum random access code* (QRAC), introduced by Ambainis et al. [ANTV02].

Definition 3.11 (QRAC). A function $Q : \{\pm 1\}^n \rightarrow \mathcal{B}_+^1(\mathcal{H}_{2^m})$ is an (n, m, ε) -QRAC if there exists a quantum oracle algorithm \mathcal{A} such that for every $x \in \{\pm 1\}^n$ and $i \in [n]$, $\Pr[\mathcal{A}^{Q(x)}(i) = x_i] \geq 1/2 + \varepsilon$.

LDQCs and smooth quantum codes are QRACs with some additional properties, such as constraints on the way the qubits of the codeword are accessed. Hence the following well-known lower bound on the length of QRACs also holds for them.

Theorem 3.12 ([ANTV02, Nay99]). *Every (n, m, ε) -QRAC has $m \geq (1 - H(1/2 + \varepsilon))n$.*

4. Pauli decoding from disjoint subsets

In this section we consider a (q, c, ε) -smooth quantum code Q . Fix a distribution μ on $\{\pm 1\}^n$. We will show that there exists a sequence $S^* \in \mathcal{P}_m$ such that if the m qubits of $Q(x)$ are measured by the m Pauli measurements in S^* , then each x_i can be retrieved by querying only q bits of the m -bit measurement outcome $S^*(Q(x))$, in a very structured way. Specifically, we prove:

Theorem 4.1. *Let $Q : \{\pm 1\}^n \rightarrow \mathcal{B}_+^1(\mathcal{H}_{2^m})$ be a (q, c, ε) -smooth quantum code and μ be a distribution on $\{\pm 1\}^n$. Then, for sufficiently large n , there exists a sequence $S^* \in \mathcal{P}_m$, and for every $i \in [n]$ a set M_i of at least $\varepsilon m/(qc)$ disjoint sets $r \subseteq [m]$ (each of size at most q) with associated signs $a_{i,r} \in \{\pm 1\}$, such that*

$$\mathbb{E}_{x \sim \mu} \left[\frac{1}{|M_i|} \sum_{r \in M_i} \Pr[a_{i,r} \prod_{j \in r} S_j^*(Q(x)) = x_i] \right] \geq \frac{1}{2} + \frac{\varepsilon}{4q+1}.$$

The proof consists of two parts. We start by constructing the sets M_i and then we show that decoding Q can be done by using only Pauli measurements. Putting these two observations together enables us to prove Theorem 4.1.

4.1. Decoding from disjoint subsets

First we obtain the sets M_i of disjoint q -sets that allow reasonable prediction of x_i .

Theorem 4.2 (modified from Lemma 4 in [KT00]). *Let $Q : \{\pm 1\}^n \rightarrow \mathcal{B}_+^1(\mathcal{H}_{2^m})$ be a (q, c, ε) -smooth quantum code with decoder \mathcal{A} , and μ a distribution on $\{\pm 1\}^n$. Then for every $i \in [n]$ there exists a set M_i of at least $\varepsilon m / (qc)$ disjoint sets $r \subseteq [m]$ (each of size at most q) satisfying $\Pr_{x \sim \mu}[\mathcal{A}^{Q(x)}(i) = x_i \mid \mathcal{A}^{(\cdot)}(i) \text{ queries } r] \geq 1/2 + \varepsilon/2$.*

Proof. Call a set $r \subseteq [m]$ “good for i ” if it satisfies the above inequality. Define for every $i \in [n]$ a hypergraph $H_i = (V, E_i)$ with vertex-set $V = [m]$ and a set of hyperedges $E_i := \{e \mid e \text{ is good for } i\}$. A smooth quantum decoder \mathcal{A} for Q “queries E_i ” if \mathcal{A} queries an $e \in E_i$. Let $p(e) := \Pr[\mathcal{A}^{(\cdot)}(i) \text{ queries } e]$. Then the probability that this decoder queries E_i is $p(E_i) := \sum_{e \in E_i} p(e)$. For all $e \notin E_i$ we have

$$\Pr[\mathcal{A}^{Q(x)}(i) = x_i \mid \mathcal{A}^{(\cdot)}(i) \text{ queries } e] < \frac{1}{2} + \frac{\varepsilon}{2}.$$

But since for every x and i , \mathcal{A} decodes bit x_i with probability at least $1/2 + \varepsilon$, we have

$$\frac{1}{2} + \varepsilon \leq \Pr[\mathcal{A}^{Q(x)}(i) = x_i] < p(E_i) + (1 - p(E_i))\left(\frac{1}{2} + \frac{\varepsilon}{2}\right) = \frac{1}{2} + \frac{\varepsilon}{2} + p(E_i)\left(\frac{1}{2} - \frac{\varepsilon}{2}\right).$$

Hence $p(E_i) > \varepsilon / (1 - \varepsilon) \geq \varepsilon$. Since Q is *smooth*, we know that the probability that \mathcal{A} queries an index j is $\sum_{e \in E_i \mid j \in e} p(e) = \Pr[\mathcal{A}^{(\cdot)}(i) \text{ queries } j] \leq c/m$.

Let M_i be a maximal set of disjoint hyperedges in H_i , and define the vertex set $T = \cup_{e \in M_i} e$. Note that T intersects each $e \in E_i$ (since otherwise M_i would not be maximal), and has at most $q|M_i|$ elements. We can now lower bound $|M_i|$ as follows:

$$\varepsilon < p(E_i) = \sum_{e \in E_i} p(e) \stackrel{(*)}{\leq} \sum_{j \in T} \sum_{e \in E_i \mid j \in e} p(e) \leq \frac{c|T|}{m} \leq \frac{cq|M_i|}{m},$$

where $(*)$ holds because each $e \in E_i$ is counted exactly once on the left-hand side, and at least once on the right-hand side (since T intersects each $e \in E_i$). \blacksquare

4.2. Pauli decoding

In the second part of the proof of Theorem 4.1, we find the appropriate Pauli measurements. Recall that to decode x_i , a smooth quantum decoder first selects a set $r \subseteq [m]$ of at most q indices, and then applies some measurement with operators $A_{i,r}^+, A_{i,r}^-$ to determine its output. Let $A_{i,r} = A_{i,r}^+ - A_{i,r}^-$. Strictly speaking these operators act only on the qubits indexed by r , but we can view them as acting on the m -qubit state $Q(x)$ by tensoring them with $m - |r|$ identities. The difference between the probabilities of obtaining outcomes $+1$ and -1 is $\text{Tr}(A_{i,r} \cdot Q(x))$. For every $i \in [n]$ and $r \in M_i$ we define the following bias:

$$B(i, r) := \mathbb{E}_{x \sim \mu}[\text{Tr}(A_{i,r} \cdot Q(x)) \cdot x_i].$$

This measures how well the measurement outcome is correlated with x_i (with x weighted according to μ). By Theorem 4.2 we have $B(i, r) \geq \varepsilon$ for every $i \in [n]$ and each $r \in M_i$.

Since \mathcal{P}_q is a basis for all $2^q \times 2^q$ complex matrices we can write

$$A_{i,r} = \sum_{S \in \mathcal{P}_q} \widehat{A_{i,r}}(S) S,$$

with $\widehat{A}_{i,r}(S) := \langle A_{i,r}, S \rangle = \frac{1}{2^q} \text{Tr}(A_{i,r} \cdot S) \in [-1, 1]$. We now have:

$$\varepsilon \leq B(i, r) = \sum_{S \in \mathcal{P}_q} \widehat{A}_{i,r}(S) \mathbb{E}_{x \sim \mu} [\text{Tr}(S \cdot Q(x)) \cdot x_i] \leq \sum_{S \in \mathcal{P}_q} |\mathbb{E}_{x \sim \mu} [\text{Tr}(S \cdot Q(x)) \cdot x_i]|. \quad (4.1)$$

Suppose we measure the r -qubits of $Q(x)$ with some $S \in \mathcal{P}_q$ and get outcome $b \in \{\pm 1\}$. The quantity $\mathbb{E}_{x \sim \mu} [\text{Tr}(S \cdot Q(x)) \cdot x_i]$ is the difference between $\Pr_{x \sim \mu}[b = x_i]$ and $\Pr_{x \sim \mu}[b \neq x_i]$. If we output b if this difference is nonnegative, and $-b$ otherwise, then we would predict x_i with bias

$$B'(i, S, r) := |\mathbb{E}_{x \sim \mu} [\text{Tr}(S \cdot Q(x)) \cdot x_i]|.$$

From Equation (4.1) we know that this bias is at least $\varepsilon/4^q$ for at least one “good” $S \in \mathcal{P}_q$. Hence, with some loss in success probability, we can decode Q by only using Pauli measurements. We now use a probabilistic argument to prove that a good sequence S^* of Pauli measurements exists, which is simultaneously good, for every $i \in [n]$, for most of the elements $r \in M_i$.

Proof (of Theorem 4.1). Suppose we let $\mathbf{S} \in \mathcal{P}_q$ be a random variable uniformly distributed over \mathcal{P}_q , and we use it to predict x_i as above. Then $B'(i, \mathbf{S}, r)$ is a random variable in the interval $[0, 1]$, with expectation

$$\mathbb{E}_{S \in \mathcal{P}_q} [B'(i, S, r)] = \frac{1}{4^q} \sum_{S \in \mathcal{P}_q} |\mathbb{E}_{x \sim \mu} [\text{Tr}(S \cdot Q(x)) \cdot x_i]| \geq \frac{\varepsilon}{4^q}.$$

Now we consider m -qubit Pauli measurements and replace all elements not in r with I 's: for $S \in \mathcal{P}_m$ and $r \subseteq [m]$, let $S_{(r)}$ denote S with all its $m - |r|$ elements outside of r replaced by I . If we let \mathbf{S} be uniform over \mathcal{P}_m , we get biases $B'(i, \mathbf{S}_{(r)}, r)$ for each $r \in M_i$, each in $[0, 1]$ and with expectation at least $\varepsilon/4^q$ (over the choice of $S_{(r)}$). But note that the random variables $B'(i, \mathbf{S}_{(r)}, r)$ are independent for different $r \in M_i$, since the elements of M_i are disjoint. Hence the average bias over all $r \in M_i$,

$$B'(\mathbf{S}, i) := \frac{1}{|M_i|} \sum_{r \in M_i} B'(i, \mathbf{S}_{(r)}, r),$$

is the average of $|M_i|$ independent random variables, each in $[0, 1]$ and with expectation at least $\varepsilon/4^q$. By a Chernoff bound⁴ the probability that $B'(\mathbf{S}, i)$ is much smaller than its expectation, is small:

$$\Pr_{S \in \mathcal{P}_m} \left[B'(S, i) < \frac{1}{2} \frac{\varepsilon}{4^q} \right] \leq \Pr_{S \in \mathcal{P}_m} \left[B'(S, i) < \frac{1}{2} \mathbb{E}[B'(S, i)] \right] \leq \exp \left(-\frac{|M_i| \varepsilon}{8 \cdot 4^q} \right).$$

By Theorems 3.12 and 4.2 we may assume $|M_i| > 8 \cdot 4^q \log(n)/\varepsilon$. It follows that for each $i \in [n]$, the above probability is less than $1/n$. Hence, the union bound gives

$$\Pr_{S \in \mathcal{P}_m} \left[\exists i \text{ s.t. } B'(S, i) < \frac{1}{2} \frac{\varepsilon}{4^q} \right] \leq \sum_{i=1}^n \Pr_{S \in \mathcal{P}_m} \left[B'(S, i) < \frac{1}{2} \frac{\varepsilon}{4^q} \right] < 1.$$

We can thus conclude that there exists an $S^* \in \mathcal{P}_m$ such that for every $i \in [n]$ we have

$$\frac{1}{|M_i|} \sum_{r \in M_i} B'(i, S_{(r)}^*, r) \geq \frac{1}{2} \frac{\varepsilon}{4^q}.$$

⁴See Equation (7) in [HR90]. A small modification of their proof shows that this bound not only holds for independent 0/1-variables, but also for independent variables in the interval $[0, 1]$.

This implies the statement of the theorem. \blacksquare

5. Classical codes from quantum codes

Theorem 4.1 implies that if we measure all m indices of a smooth quantum quantum code Q with the elements of S^* , then we get distributions on $\{\pm 1\}^m$ that can be massaged to “codewords” $R(x)$ of a randomized smooth code:

Theorem 5.1. *Let $Q : \{\pm 1\}^n \rightarrow \mathcal{B}_+^1(\mathcal{H}_{2^m})$ be a (q, c, ε) -smooth quantum code. Then, for sufficiently large n , for every input distribution μ on $\{\pm 1\}^n$, there exists a μ -average $(q, qc/\varepsilon, \varepsilon/4^{q+1})$ -randomized smooth code $R : \{\pm 1\}^n \rightarrow \mathcal{P}(\{\pm 1\}^m)$.*

Proof. We use Theorem 4.1. Let $R(x)$ be the distribution on $\{\pm 1\}^m$ obtained by measuring $Q(x)$ with S^* . We define a decoder \mathcal{A} for R as follows: on input $i \in [m]$ and oracle $y \in \{\pm 1\}^m$, pick a set r from the set M_i uniformly at random, and return $a_{i,r} \prod_{j \in r} y_j$. It is straightforward to check that \mathcal{A} is a μ -average $(q, qc/\varepsilon, \varepsilon/4^{q+1})$ decoder for R ; in particular, since \mathcal{A} picks r uniformly from a set of at least $\varepsilon m/(qc)$ disjoint sets, each index $j \in [m]$ has probability at most $qc/(\varepsilon m)$ of being queried. \blacksquare

Together, Lemma 3.6 and Theorem 5.1, give the following “derandomization”:

Corollary 5.2. *Let $Q : \{\pm 1\}^n \rightarrow \mathcal{B}_+^1(\mathcal{H}_{2^m})$ be a (q, c, ε) -smooth quantum code. Then, for sufficiently large n , for every distribution μ on $\{\pm 1\}^n$, there exists a $C : \{\pm 1\}^n \rightarrow \{\pm 1\}^m$ which is a μ -average $(q, qc/\varepsilon, \varepsilon/(2 \cdot 4^{q+1}))$ -smooth code for at least $\varepsilon n/4^{q+1}$ of the n indices.*

Following the path through Theorems 3.10, 5.1, and the μ -average version of Theorem 3.9, we can turn an LDQC into a μ -average randomized LDC:

Corollary 5.3. *Let $Q : \{\pm 1\}^n \rightarrow \mathcal{B}_+^1(\mathcal{H}_{2^m})$ be a (q, δ, ε) -LDQC. Then, as long as $\delta' \leq \delta \varepsilon^2/(q^2 4^{q+1})$ and n is sufficiently large, for every distribution μ over $\{\pm 1\}^n$, there exists an $R : \{\pm 1\}^n \rightarrow \mathcal{P}(\{\pm 1\}^m)$ which is a μ -average $(q, \delta', \varepsilon/4^{q+1} - \delta' q^2/(\delta \varepsilon))$ -randomized LDC.*

Going through Theorem 3.10, Corollary 5.2, and the μ -average version of Theorem 3.9 instead, we can also turn an LDQC into a μ -average LDC:

Corollary 5.4. *Let $Q : \{\pm 1\}^n \rightarrow \mathcal{B}_+^1(\mathcal{H}_{2^m})$ be a (q, δ, ε) -LDQC. Then, as long as $\delta' \leq \delta \varepsilon^2/(2q^2 4^{q+1})$ and n is sufficiently large, for every distribution μ over $\{\pm 1\}^n$, there exists a $C : \{\pm 1\}^n \rightarrow \{\pm 1\}^m$ which is a μ -average $(q, \delta', \varepsilon/(2 \cdot 4^{q+1}) - \delta' q^2/(\delta \varepsilon))$ -LDC for at least $\varepsilon n/4^{q+1}$ of the n indices.*

6. Conclusion and open problems

We defined quantum generalizations of q -query LDCs in which q queries correspond to a measurement on q qubits of the m -qubit codeword. By a reduction to (classical) randomized smooth codes through a special sequence of Pauli measurements on an LDQC, we showed that the use of quantum systems for this type of encoding cannot provide much advantage in terms of length, at least for small q . An obvious open problem is reducing the gap between upper and lower bound on the length m of LDCs for fixed small number of queries q . Our results show that an upper bound for LDQCs would carry over to (μ -average) LDCs. This might perhaps be a way to improve the best known classical upper bounds on m .

Acknowledgments

We thank Harry Buhrman, Peter Høyer, Oded Regev, Falk Unger and Stephanie Wehner for useful discussions. JB is especially indebted to Peter Høyer for suggesting to turn an LDQC into an LDC via measurements.

References

- [ANTV02] A. Ambainis, A. Nayak, A. Ta-Shma, and U. Vazirani. Dense quantum coding and quantum finite automata. *J. of ACM*, 49(4):496–511, 2002.
- [BFLS91] L. Babai, L. Fortnow, L. Levin, and M. Szegedy. Checking computations in polylogarithmic time. In *Proc. of 23rd ACM STOC*, pages 21–31, 1991.
- [CGKS98] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan. Private information retrieval. *J. of ACM*, 45(6):965–981, 1998.
- [Efr08] K. Efremenko. 3-Query Locally Decodable Codes of Subexponential Length. Technical report, ECCS TR08–069, 2008.
- [HR90] T. Hagerup and C. Rüb. A guided tour of Chernoff bounds. *Information Processing Letters*, 33(6):305–308, 1990.
- [KT00] J. Katz and L. Trevisan. On the efficiency of local decoding procedures for error-correcting codes. In *Proc. of 32nd ACM STOC*, pages 80–86, 2000.
- [KW04] I. Kerenidis and R. de Wolf. Exponential lower bound for 2-query locally decodable codes via a quantum argument. *J. of Computer and System Sciences*, 69(3):395–420, 2004.
- [Nay99] A. Nayak. Optimal lower bounds for quantum automata and random access codes. In *Proc. of 40th IEEE FOCS*, pages 369–376, 1999.
- [NC00] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [Tre04] L. Trevisan. Some applications of coding theory in computational complexity. *Quaderni di Matematica*, 13:347–424, 2004.
- [Woo06] D. Woodruff. New lower bounds for general locally decodable codes. Technical report, ECCS TR07–006, 2006.
- [Yao77] A. C-C. Yao. Probabilistic computations: Toward a unified measure of complexity. In *Proc. of 18th IEEE FOCS*, pages 222–227, 1977.
- [Yek07] S. Yekhanin. Towards 3-query locally decodable codes of subexponential length. In *Proc. of 39th ACM STOC*, pages 266–274, 2007.