

Quantum Locally Recoverable Codes*

Louis Golowich
UC Berkeley
lgolowich@berkeley.edu

Venkatesan Guruswami
UC Berkeley
venkatg@berkeley.edu

November 14, 2023

Abstract

Classical locally recoverable codes, which permit highly efficient recovery from localized errors as well as global recovery from larger errors, provide some of the most useful codes for distributed data storage in practice. In this paper, we initiate the study of quantum locally recoverable codes (qLRCs). In the long term, like their classical counterparts, such qLRCs may be used for large-scale quantum data storage. Furthermore, our results have concrete implications for quantum LDPC codes, which are widely applicable to near-term quantum error-correction, as local recoverability is a weakening of the LDPC property.

After defining quantum local recoverability, we provide an explicit construction of qLRCs based on the classical LRCs of Tamo and Barg (2014), which we show have (1) a close-to-optimal rate-distance tradeoff (i.e. near the Singleton bound), (2) an efficient decoder, and (3) permit good spatial locality in a physical implementation. The analysis for both the distance and the efficient decoding of these quantum Tamo-Barg (qTB) codes is significantly more involved than in the classical case. Nevertheless, we obtain close-to-optimal parameters by introducing a “folded” version of these qTB codes, which we then analyze using a combination of algebraic techniques. We furthermore present and analyze two additional constructions using more basic techniques, namely random qLRCs, and qLRCs from AEL distance amplification. Each of these constructions has some advantages, but neither achieves all 3 properties of our folded qTB codes described above.

We complement these constructions with Singleton-like bounds that show our qLRC constructions achieve close-to-optimal parameters. We also apply these results to obtain Singleton-like bounds for qLDPC codes, which to the best of our knowledge are novel. We then show that even the weakest form of a stronger locality property called local correctability, which permits more robust local recovery and is achieved by certain classical codes, is impossible quantumly.

*Research supported in part by a Simons Investigator award, and a UC Noyce initiative award. L. Golowich is supported by a National Science Foundation Graduate Research Fellowship under Grant No. DGE 2146752.

Contents

1	Introduction	1
1.1	Our Contributions	2
1.1.1	Definition of Quantum Local Recoverability	2
1.1.2	Explicit Construction of qLRCs	2
1.1.3	Singleton-Like Bound for qLRCs, with Implications for qLDCP Codes	3
1.1.4	Comparison to Basic Constructions	4
1.1.5	Efficient Decoding Algorithms	5
1.1.6	Impossibility of Quantum Locally Correctable Codes	7
1.2	Overview of (Folded) Quantum Tamo-Barg codes	8
1.2.1	Background and Classical Tamo-Barg Codes	8
1.2.2	Quantum Tamo-Barg codes	8
1.2.3	Folded Quantum Tamo-Barg Codes	11
1.3	Open Questions	11
2	Preliminaries	12
2.1	Notation	12
2.2	Classical Codes	13
2.3	Polynomial Evaluation Codes	14
2.4	Quantum Codes	15
3	Singleton-Like Bounds for qLRCs	18
4	Basic Constructions Using Known Techniques	22
4.1	Random qLRCs	22
4.2	Explicit qLRCs from AEL Distance Amplification	25
4.2.1	Review of AEL Technique	25
4.2.2	Application to qLRCs	26
5	Explicit Construction of qLRCs: Quantum Tamo-Barg Codes	29
5.1	Classical Tamo-Barg Codes	29
5.2	Quantum Tamo-Barg Codes	29
5.3	Folded Quantum Tamo-Barg Codes	32
6	Bounding the Distance	33
6.1	Distance of Unfolded Quantum Tamo-Barg Codes	34
6.2	Distance of Folded Quantum Tamo-Barg Codes	35
7	Efficient Decoding Algorithm	40

7.1	Unfolded Quantum Tamo-Barg Codes	41
7.2	Folded Quantum Tamo-Barg Codes	44
8	Impossibility of Quantum Locally Correctable Codes	48
9	Acknowledgments	50
A	Technical Lemmas	53
B	Omitted Proofs	55

1 Introduction

Classical locally recoverable codes (LRCs) provide one of the most important coding theoretic tools for distributed data storage. Such codes are defined to permit highly efficient recovery from common localized errors, as well as larger-scale recovery from rarer but more “catastrophic” global errors.

In this paper, we initiate the study of quantum locally recoverable codes (qLRCs). In particular, we define qLRCs, present and analyze constructions, and also prove fundamental limitations on the achievable parameters and properties. While our constructions can be viewed as quantum generalizations of classical constructions, the analysis becomes surprisingly intricate, and requires new ideas that were not needed classically. Our results may also shed light on the study of locality in quantum coding theory, for instance as it pertains to quantum LDPC codes.

Classically, the properties of a LRC are well suited for the needs of large datacenters, which can cost billions of dollars to build and maintain, and must often account for localized server failures while also handling occasional more global failures. Indeed, companies such as Microsoft [HSX⁺12] and Facebook [MLR⁺14] have implemented LRCs to obtain improved performance for data storage.

Currently, experimental quantum computers remain at a vastly smaller scale than that of the classical datacenters in which LRCs are often used in practice. However, it is not implausible that quantum computing technology eventually follows its classical counterpart by growing to the scale where codes such as qLRCs become an integral part of quantum data storage.

Furthermore, our study of qLRCs reveals the potential for more broad and near-term implications as well. Indeed, locality properties in quantum codes, such as the ability to decode using local measurements (i.e. LDPC codes), are of particular importance for quantum error correction. Yet such locality is notoriously difficult to achieve in the quantum setting. Indeed, the first linear-distance quantum LDPC (qLDPC) codes were only recently constructed [PK22, LZ22, DHLV23], and good quantum codes with stronger properties such as local testability have yet to be constructed. This difficulty of achieving locality is in contrast to the classical setting, where good LDPC codes have been known for decades (e.g. [SS96]), good locally testable codes were recently constructed [PK22, DEL⁺22], and other strong locality properties such as local correctability exist in linear-distance, albeit low-rate codes.

From this perspective, our study of qLRCs provides a new angle to investigate locality properties in quantum codes. Indeed, classical local recoverability requires each code component to participate in one low-weight parity check, while quantum local recoverability requires each code component to participate in two low-weight stabilizers. Thus local recoverability can be viewed as a weakening of the LDPC property, in which each code component participates in many low-weight parity checks/stabilizers. Our study of qLRCs can therefore be viewed as progress towards understanding stronger locality properties possessed by qLDPC codes. One concrete example of this connection is provided in Section 3, where we show that qLRCs, and therefore also qLDPC codes, of constant locality $r = O(1)$ must have relative distance bounded away from $1/2$; to the best of our knowledge such a bound for qLDPC codes had not been previously shown.

1.1 Our Contributions

In this section we present the contributions of our paper. For details on notation or basic definitions, the reader is referred to Section 2.

1.1.1 Definition of Quantum Local Recoverability

To begin, we define qLRCs. Recall that a classical LRC (cLRC) is a classical code C such that for every $c \in C$ and every component i , the value of c_i can be recovered by looking at the restriction of c to just $r - 1$ other components.

Definition 1 (Informal statement of Definition 32). A **quantum locally recoverable code (qLRC) of locality r** is a quantum code \mathcal{C} such that if any single qudit of a code state $|\psi\rangle \in \mathcal{C}$ is erased (i.e. it experiences a completely depolarizing channel), the original code state $|\psi\rangle$ can be recovered by applying a recovery channel that accesses only $r - 1$ other code state qudits.

For intuition, recall that a classical linear code is a cLRC of locality r if each component takes part in a parity-check of weight $\leq r$. Similarly, we show that a quantum CSS code $\mathcal{C} = \text{CSS}(C_X, C_Z)$ is a qLRC if both C_X, C_Z are cLRCs, so that every qudit takes part in a low weight X -parity-check and a low-weight Z -parity-check (see Corollary 34).

We remark that classically, there is also a notion of message locally recoverable codes (mLRCs), which require that every message (instead of codeword) symbol can be recovered from $r - 1$ codeword symbols. As any linear classical code has a systematic encoding, meaning that the first k codeword symbols equal the message, classical mLRCs are strictly weaker than LRCs. However, the local indistinguishability property of quantum codes (see Lemma 26) implies that local queries to quantum codes cannot reveal anything about the message. Thus mLRCs do not exist quantumly, at least in the regime where the locality is less than the distance.

1.1.2 Explicit Construction of qLRCs

One of our principal technical contributions is the following explicit construction of qLRCs.

Theorem 2 (Folded quantum Tamo-Barg codes; informal statement of Corollary 64 combined with Lemma 56). *For every prime number r and every $0 < R < 1$, there exists an infinite explicit family of qLRCs of locality r , rate $\geq R$, relative distance*

$$\delta \geq \frac{1 - R}{2} - O\left(\frac{1}{\sqrt{r}}\right),$$

and alphabet size $n^{O(r^2)}$, where n denotes the block length.

In Theorem 2 (and in future informal result statements in Section 1), for readability we state slightly looser bounds than the formal result statements. For instance, our actual distance bound in Corollary 64 is stronger than than stated in Theorem 2 for high rates, and in particular shows that $\delta \geq \Omega(1/r)$ for all $R \leq 1 - 10/r$.

We remark that while the alphabet size $n^{O(r^2)}$ may seem large, in the LRC literature one typically thinks of each code component as being a fairly large entity, so such a polynomial alphabet

size is not unreasonable. Classically each code component could for instance be a hard drive, while quantumly each component would likely be itself a fault-tolerant quantum memory.

We prove Theorem 2 by introducing a quantum CSS version of the classical LRCs of Tamo and Barg [TB14], which achieve the optimal classical rate-distance-locality tradeoff [GHSY12]. The classical Tamo-Barg (TB) codes are constructed as subcodes of Reed-Solomon codes by carefully inserting low-weight parity checks.

Multiple complications arise when converting TB codes into CSS codes. Specifically, to ensure the CSS orthogonality relations are satisfied, we define a quantum Tamo-Barg (qTB) code to be a CSS code $\mathcal{C} = \text{CSS}(C, C)$ consisting of two copies of a classical code C that contains a TB code as a subcode, but also contains some added low-weight codewords. These added low-weight codewords also lie in C^\perp , so they do not necessarily degrade the distance of \mathcal{C} , which equals the minimum weight of an element of $C \setminus C^\perp$. However, these added low-weight codewords make the distance analysis significantly more challenging, and we are only able to show a “Johnson-like” bound on the distance of qTB codes:

Theorem 3 (Quantum Tamo-Barg codes; informal statement of Theorem 62). *For every prime number r and every $0 < R < 1$, there exists an infinite explicit family of qLRCs of locality r , rate $\geq R$, relative distance*

$$\delta \geq 1 - \sqrt{\frac{1+R}{2}} - O\left(\frac{1}{r}\right),$$

and alphabet size $n + 1$, where n denotes the block length.

To obtain the improved bound $\delta \geq 1 - (1+R)/2 - O(1/\sqrt{r})$ in Theorem 2, we “fold” together code components, thereby reblocking the symbols into larger components. We then use a combination of algebraic techniques to bound the distance of these folded qTB (fqTB) codes; the two main tools are a root detection method involving a determinant polynomial, and an uncertainty principle over finite fields.

In Section 1.2 below, we describe the formal construction of our (f)qTB codes, and outline the proofs of Theorem 3 and Theorem 2. A more detailed description of these codes is provided in Section 5, and the distance bounds are formally proven in Section 6.

1.1.3 Singleton-Like Bound for qLRCs, with Implications for qLDCP Codes

Our fqTB codes in Theorem 2 achieve an optimal rate-distance tradeoff as the locality r grows large. In particular, we prove the following fundamental limitation on any qLRC.

Theorem 4 (Singleton-like bound; informal statement of Theorem 35). *If \mathcal{C} is a qLRC of locality r and rate R , then \mathcal{C} has relative distance*

$$\delta \leq \frac{1-R}{2} - \Omega\left(\frac{1}{r}\right).$$

The reader is referred to Theorem 35 for the specific value of the constant hidden in the the $\Omega(1/r)$ term above.

For comparison, recall that the ordinary quantum Singleton bound states that every quantum code of rate R has relative distance $\delta \leq (1-R)/2 + O(1/n)$, where n denotes the block length. Thus

for a fixed rate, Theorem 4 shows that imposing local recoverability with locality r decreases the optimal relative distance of a quantum code by at least $\Omega(1/r)$. Meanwhile, our explicit construction in Theorem 64 has relative distance $O(1/\sqrt{r})$ below that of the Singleton bound.

Theorem 4 implies a fundamental difference between the classical and quantum cases for LRCs of low rate. Classically, the ordinary Singleton bound says that a code of rate R has relative distance $\delta \leq 1 - R + O(1/n)$, while for cLRCs of locality r this bound becomes $\delta \leq 1 - R \cdot r / (r - 1) + O(1/n)$. The TB codes achieve this latter bound, proving its tightness. Therefore in particular, for fixed locality r , by letting the rate $R \rightarrow 0$, we see that there exist classical LRCs of relative distance $\delta \rightarrow 1$. In contrast, whereas there exist quantum codes of relative distance approaching $1/2$, Theorem 4 shows every qLRC has relative distance at most $1/2 - \Omega(1/r)$, which is bounded away from $1/2$ for fixed r .

As we mentioned previously, because every (q)LRC is by definition a (q)LDPC code, it follows that qLDPC codes of locality r have relative distance $\leq 1/2 - \Omega(1/r)$ for arbitrarily large alphabets. To the best of our knowledge, such a bound has not been previously shown in the literature. This result is in again in contrast to the classical case, where for instance the q -ary Hadamard code is an LDPC code of locality 3 with relative distance approaching 1 as q grows large.

The discussion above raises the interesting question as to how tight Theorem 4 is for qLDPC codes. That is, what is the additional cost to the optimal rate-distance tradeoff of requiring a quantum code be LDPC, compared to just being an LRC?

1.1.4 Comparison to Basic Constructions

While our explicit fqTB codes in Theorem 2 are $O(1/\sqrt{r})$ below the Singleton-like bound in Theorem 4, we show that a randomized construction improves this gap to $O(1/r)$, at the cost of explicitness and efficiency.

Proposition 5 (Random qLRCs; informal statement of Proposition 40). *For every $r \geq 3$ and $\delta > 0$, there exists a randomized construction that with high probability gives a qLRC of locality r , rate R , relative distance*

$$\delta \geq \frac{1 - R}{2} - O\left(\frac{1}{r}\right),$$

and alphabet size $2^{O(r)}$.

The constant hidden in the $O(1/r)$ term in Proposition 5 is larger than the constant in the $\Omega(1/r)$ term in Theorem 4, so our bound on the randomized construction is still $O(1/r)$ below that of our Singleton-like bound.

Thus the randomized construction in Proposition 40 achieves relative distance $O(1/r)$ below our Singleton-like bound with has alphabet size $2^{O(r)}$. These parameters improve upon our fqTB construction, which has relative distance $O(1/\sqrt{r})$ below the Singleton-like bound, and has alphabet size $n^{O(r^2)}$.

However, the main disadvantage of the randomized construction is its non-explicitness. As a result, we have no efficient algorithm to certify that a randomly sampled qLRC has good distance, and we have no efficient decoding algorithm for errors in unknown locations. In contrast, our fqTB codes are explicit, so their distance bound is guaranteed. Furthermore, we show that they have an efficient decoding algorithm (see Section 1.1.5 below).

One way of derandomizing the random qLRCs in Proposition 5 is to use them as in inner code in the concatenation and distance amplification scheme of Alon, Edmonds, and Luby (AEL) [AEL95]. This technique has been used extensively in classical coding theory [GI01, GI02, GI03, GR08, HW18, KMRZS16, GKO⁺18, HRZW20], but has only recently been considered in the quantum setting [BGG22, WLH23]. We show that applying AEL using a random qLRC as an “inner code,” which is small enough to be found efficiently via brute force, yields the following result.

Proposition 6 (qLRCs from AEL; informal statement of Proposition 49). *For every fixed $0 < R < 1$, it holds for all sufficiently large $r \in \mathbb{N}$ that there exists an infinite family of efficiently constructable qLRCs of locality r , rate R , relative distance*

$$\delta \geq \frac{1 - R}{2} - O_R\left(\frac{1}{r^{1/4}}\right),$$

and alphabet size $q = 2^{O(r)}$, where the O_R above hides a constant depending on R .

As described in Remark 50, the codes in Proposition 6 are technically only efficiently constructable by a randomized algorithm with high probability, but can be made truly explicit by using a slightly more complicated construction, with slightly worse parameters. Also, like our fqTB codes, we show that these codes from AEL have efficient decoders from errors in unknown locations.

The alphabet size $q = 2^{O(r)}$ in Proposition 6 is smaller than the $q = n^{O(r^2)}$ of our fqTB codes in Theorem 2. However, qLRCs from AEL have worse rate-distance-locality tradeoff, as their relative distance is $O_R(1/r^{1/4})$ below the Singleton-like bound, compared to only $O(1/\sqrt{r})$ for our fqTB codes.

Our (f)qTB codes may have additional practical advantages over the qLRCs from AEL. For instance, our (f)qTB codes can achieve locality as small as $r = 3$, whereas the the minimum possible locality in the qLRCs from AEL is $r = 9$. Furthermore, as one step in the AEL construction redistributes code symbols according to the edges of an expander graph, the resulting size- r sets of code components used for local recovery form an r -uniform hypergraph with a complex expanding structure. In contrast, the recovery sets of our (f)qTB codes form a partition of the code components, which is ideally suited for a physical implementation with good spatial locality. That is, our (f)qTB codes can easily be implemented in 1, 2, or 3-dimensional space such that each local recovery operation only involves code components that are close together; such a spatially local implementation would be much less feasible for a qLRC from AEL.

1.1.5 Efficient Decoding Algorithms

This section presents our results on the decodability of our qLRCs. As we discussed previously, in the classical setting, LRCs are typically used for data storage where errors correspond to events such as server failures that are detectable. Such errors occur in known locations, so they can be treated as erasures. Because every linear code can be efficiently decoded from a number of erasures up to the distance using Gaussian elimination, the efficiency of decoding is often not a primary concern for cLRCs. Note that all of our qLRC constructions are stabilizer (and in fact CSS) codes, which can similarly be decoded efficiently from erasures.

However, in the quantum setting there are additional potential applications of efficient decoding from errors in unknown locations. Due to the inherently noisy nature of quantum states, even if

a qLRC is used to store a large quantum state where each code component is itself stored in a fault-tolerant memory, these individual fault-tolerant memories may eventually accumulate errors. As such, it may be beneficial to occasionally perform a global decoding procedure to reduce the overall error rate in the long term.

Furthermore, as we have previously discussed, qLRCs can be viewed as a stepping stone towards better understanding stronger locality properties, such as the LDPC property, which are important for near-term quantum error correction. As efficient decoding is critical for these error correction applications, it is desirable that the qLRCs we study are also efficiently decodable.

Below we present a polynomial-time decoding algorithm for the fqTB codes in Theorem 2.

Theorem 7 (Decoding fqTB codes; informal statement of Corollary 72). *The fqTB codes in Theorem 2 of block length n , prime locality parameter r , and rate R can be decoded from errors acting on an unknown*

$$\frac{1}{2} \left(\frac{1-R}{2} - O\left(\frac{1}{\sqrt{r}}\right) \right)$$

fraction of the code components in time $n^{O_r(1)}$, provided the alphabet size is increased to some sufficiently large $q = n^{O_r(1)}$ with respect to r . Here $O_r(1)$ denotes a sufficiently large constant depending only on r .

Theorem 7 provides a decoding algorithm for the fqTB codes with decoding radius up to half our distance bound in Theorem 2, which is therefore optimal barring an improvement in the distance bound. While the algorithm runs in polynomial time $n^{O_r(1)}$ for fixed locality parameter r , this algorithm is inefficient for growing r . We address this issue by providing the following decoding algorithm for unfolded qTB codes, which therefore also applies to folded qTB codes, and whose running time is a polynomial independent of r .

Theorem 8 (Decoding (f)qTB codes; informal statement of Theorem 69). *An (unfolded or folded) qTB code of block length n , prime locality parameter r , and rate R can be decoded from errors acting on an unknown*

$$\frac{1}{2} \left(1 - \sqrt{\frac{1+R}{2}} - O\left(\frac{1}{r}\right) \right)$$

fraction of the code components in time $n^{O(1)}$.

The decoder in Theorem 8 simply performs $r-1$ calls to a classical Reed-Solomon (list) decoder, and then performs some postprocessing on the resulting outputs (see Algorithm 1). Therefore this algorithm should be efficient in practice, as Reed-Solomon decoders have been optimized for practical use.

The decoding radius in Theorem 8 is half our distance bound in Theorem 3. Therefore this decoding radius is optimal among decoders for qTB codes, barring an improvement to our distance bound in Theorem 3. As

$$\frac{1-R}{4} \leq 1 - \sqrt{\frac{1+R}{2}} \leq \frac{1-R}{2},$$

the unfolded qTB codes achieve distance and decoding radius within roughly a factor of 2 of the optimal values as dictated by our Singleton-like bound in Theorem 4.

We prove both Theorem 8 and Theorem 7 using similar techniques as in the proof of Theorem 3. Specifically, let $\mathcal{C} = \text{CSS}(C, C)$ denote the qTB code. The decoding task for \mathcal{C} can be reduced to the following problem: given a corrupted codeword a of C , we want an efficient algorithm that recovers some $c \in C$ that is close to a in Hamming weight. For this purpose, we show how every $c \in C$ can be modified to obtain some c' that is a codeword of a Reed-Solomon code. Applying the same modification to a , we obtain a corrupted Reed-Solomon codeword a' . We then apply a Reed-Solomon (list) decoder to recover c' , which we can then map back to the desired codeword c .

To decode fqTB codes in Theorem 7, we use the same algorithm, except now we apply a folded Reed-Solomon list decoder [GR08] to recover c' from a' . As folded Reed-Solomon codes have a larger list-decoding radius than ordinary Reed-Solomon codes, and our fqTB distance bound in Theorem 2 is better than our qTB distance bound in Theorem 3, we obtain a larger decoding radius in Theorem 7 than in Theorem 8.

1.1.6 Impossibility of Quantum Locally Correctable Codes

Our results on qLRCs described above indicate that while optimal local recoverability is more nuanced and difficult to achieve quantumly than classically, there do exist constructions of qLRCs approaching the optimal parameters. It is therefore natural to consider quantum analogues of stronger forms of locality that exist classically.

Local correctability and local decodability provide particularly notable examples of such properties. Recall that an LRC of block length n and locality r has the property that any single code component $i \in [n]$ is erased, the value of a codeword at that component can be recovered from $\leq r - 1$ unerased components. A *locally correctable code (LCC)* has the stronger property that after a linear number $\Omega(n)$ of code components are erased, the value of a codeword at each erased component i can be recovered from $\leq r - 1$ unerased components.

Equivalently, an LRC requires the value of a codeword at each component $i \in [n]$ to be recoverable from some recovery set of $\leq r - 1$ other components. In contrast, an LCC requires each component $i \in [n]$ to have $\Omega(n)$ disjoint recovery sets, each of which contains $\leq r - 1$ components that can be used to recover the codeword's value at i .

Local decodable codes (LDCs) are defined similarly, except they only need to support local recovery for message components, rather than for codeword components.

The Hadamard and Reed-Muller codes provide examples of classical LCCs and LDCs. It is therefore natural to ask whether there are quantum versions of these types of codes.

As local access to a quantum code of large distance cannot reveal any information about the message, local decodability seems to make little sense quantumly, at least in the regime of large distance and small locality. However, quantum local correctability may seem to be a reasonable strengthening of quantum local recoverability.

We show below that quantum local correctability is impossible in a strong sense: if a quantum code has even two disjoint recovery sets for a single qudit, then that qudit is unentangled with the remainder of the code state, and contains no information about the message state. Recall that LCCs are typically required to have linearly many such disjoint recovery sets for each code component. Thus qLRCs, which have a single recovery set for each qudit, are in some sense the limit of what is possible for quantum local correctability.

Theorem 9 (Impossibility of quantum LCCs; informal statement of Theorem 74). *Let \mathcal{C} be a*

quantum code of block length n such that for some qudit $i \in [n]$, there exist subsets $I_i^1, I_i^2 \subseteq [n]$ satisfying $I_i^1 \cap I_i^2 = \{i\}$ such that the following holds for each $b = 1, 2$: if qudit i of a code state $\psi \in \mathcal{C}$ is erased (i.e. it experiences a completely depolarizing channel), the original code state ψ can be recovered by applying a recovery channel that only accesses qudits in I_i^b . Then there exists a 1-qudit state α such that every $\psi \in \mathcal{C}$ can be decomposed as $\psi = \alpha_i \otimes \psi_{[n] \setminus \{i\}}$.

While we prove Theorem 9 for general quantum codes $\mathcal{C} \subseteq (\mathbb{C}^d)^{\otimes n}$, it is illustrative to consider CSS or stabilizer codes. Recall from Section 1.1.1 that for a CSS code $\mathcal{C} = \text{CSS}(C_X, C_Z)$, a set $I_i \ni i$ can be used to recover qudit i if both C_X and C_Z have parity checks whose support contains i and lies inside I_i . Thus if I_i^1, I_i^2 are both recovery sets for component i with $I_i^1 \cap I_i^2 = \{i\}$, then there are parity checks $c'_X \in C_X^\perp$ and $c'_Z \in C_Z^\perp$ such that $i \in \text{supp}(c'_X) \subseteq I_i^1$ and $i \in \text{supp}(c'_Z) \subseteq I_i^2$. But then $\text{supp}(c'_X) \cap \text{supp}(c'_Z) = \{i\}$, so $c'_X \cdot c'_Z = (c'_X)_i (c'_Z)_i \neq 0$, which contradicts the orthogonality condition $C_X^\perp \subseteq C_Z$ required for \mathcal{C} to be a well defined CSS code. Thus Theorem 9 holds for CSS codes. A similar proof holds for stabilizer codes; In Theorem 74 we prove the more general result for arbitrary quantum codes.

1.2 Overview of (Folded) Quantum Tamo-Barg codes

In this section, we provide more details on our (folded) quantum Tamo-Barg codes described in Section 1.1.2, and we overview the techniques we use to prove the distance bounds in Theorem 2 and Theorem 3. The construction and analysis of these codes comprise one of the main technical contributions of our paper.

1.2.1 Background and Classical Tamo-Barg Codes

Our qTB codes are CSS codes whose associated classical codes are polynomial evaluation codes that are closely related to the classical LRCs of Tamo and Barg [TB14]. To describe these codes, we use the following notation. For a subset $S \subseteq \mathbb{Z}_{\geq 0}$, let $\mathbb{F}_q[X]^S = \{\sum_{i \in S} a_i X^i : a \in \mathbb{F}_q^S\}$ denote the space of polynomials for which only monomials X^i for $i \in S$ can have nonzero coefficients. Then let $\text{ev} : \mathbb{F}_q[X]^S \rightarrow \mathbb{F}_q^{*\mathbb{F}_q^*} \cong \mathbb{F}_q^{q-1}$ denote the evaluation map on nonzero points in \mathbb{F}_q , so that $\text{ev}(f) = (f(x))_{x \in \mathbb{F}_q^*}$.

With this notation, a classical Tamo-Barg (TB) code [TB14] specified by a prime power q , a locality parameter $r|(q-1)$, and an integer $\ell \in [q]$ is given by $\text{ev}(\mathbb{F}_q[X]^S)$ for

$$S = \{i \in [\ell] : i \not\equiv r-1 \pmod{r}\}.$$

[TB14] showed that this code has alphabet size q , block length $q-1$, dimension $\ell - \lfloor \ell/r \rfloor$, distance $\geq q - \ell$, and is locally recoverable with locality r .

1.2.2 Quantum Tamo-Barg codes

In this section, we define our quantum analogue of TB codes, and give an overview of the analysis. For more details, the reader is referred to Section 5.2 and Section 6.1.

As described in Section 1.1.2, to define a qTB code, we must first modify the classical TB code by adding some low-weight codewords in order to construct a well-defined associated quantum

CSS code. These added low-weight codewords can be interpreted as *piecewise linear functions*, as described below.

Definition 10 (Quantum Tamo-Barg codes; restatement of Definition 53). For a prime power q , a locality parameter $r|(q-1)$ with $r \geq 3$, and an integer $\ell \in [q]$, we define the **quantum Tamo-Barg (qTB) code** to be the CSS code $\mathcal{C} = \text{CSS}(C, C)$ with $C = \text{ev}(\mathbb{F}_q[X]^S)$ for

$$S = \{i \in [\ell] : i \not\equiv r-1 \pmod{r}\} \cup \{i \in [q-1] : i \equiv 1 \pmod{r}\}. \quad (1)$$

By some basic algebraic manipulations (see Lemma 55), we show that $C^\perp \subseteq C$, so the qTB code $\mathcal{C} = \text{CSS}(C, C)$ is indeed a well-defined quantum CSS code. Note that by construction \mathcal{C} has alphabet size q and block length $q-1$. A straightforward calculation (see Lemma 56) shows that \mathcal{C} has dimension

$$\begin{aligned} k &= 1 + |\{q-\ell \leq i \leq \ell-1 : i \not\equiv \pm 1 \pmod{r}\}| \\ &= (2\ell - q) \cdot \left(1 - \frac{2}{r}\right) + \epsilon \end{aligned}$$

for some $\epsilon \in [-2, 2]$.

The code C in the definition above is by definition the sum of a classical TB code $\text{ev}(\mathbb{F}_q[X]^{[\ell] \setminus (-1+r\mathbb{Z})})$ with the code¹ $P := \text{ev}(\mathbb{F}_q[X]^{[q-1] \cap (1+r\mathbb{Z})})$. We call this latter code P the space of **piecewise linear functions** due to the following lemma.

Lemma 11 (Restatement of Lemma 57). *Let $\Omega_r = \{x \in \mathbb{F}_q^* : x^r = 1\}$ denote the r th roots of unity. Then $P = \text{ev}(\mathbb{F}_q[X]^{[q-1] \cap (1+r\mathbb{Z})})$ consists of all functions $f : \mathbb{F}_q^* \rightarrow \mathbb{F}_q$ that can be expressed in the form $f(x) = \beta_{x\Omega_r} \cdot x$ for some $\beta \in \mathbb{F}_q^{*\Omega_r}$.*

Lemma 11 follows directly from the fact that the evaluation of a polynomial $h(X)$ on inputs in $x\Omega_r$ equals the evaluation of $h(X) \pmod{X^r - x^r}$ on such inputs. The lemma implies that every function in P equals a linear function on the restriction to inputs in each coset $x\Omega_r$.

We show in Lemma 55 that $P \subseteq C^\perp$, from which we obtain the local recoverability of \mathcal{C} :

Corollary 12 (Restatement of Corollary 58). *The qTB code $\mathcal{C} = \text{CSS}(C, C)$ given in Definition 10 is locally recoverable with locality r .*

Proof. As described in Section 1.1.1 and formalized in Corollary 34, it suffices to show that for each $\alpha \in \mathbb{F}_q^*$, there exists some $f \in C^\perp$ such that $\alpha \in \text{supp}(f)$ and $|\text{supp}(f)| \leq r$. But because $P \subseteq C^\perp$ as described above, the piecewise linear function $f : \mathbb{F}_q^* \rightarrow \mathbb{F}_q$ given by $f(x) = x$ for $x \in \alpha\Omega_r$ and $f(x) = 0$ for $x \notin \alpha\Omega_r$ satisfies these criteria, as desired. \square

It only remains to bound the distance of the qTB code \mathcal{C} . For this purpose, we show the following result, which directly implies Theorem 3.

Theorem 13 (Restatement of Theorem 62). *The qTB code \mathcal{C} in Definition 10 with a prime locality parameter r has distance at least*

$$d = (q-1) \left(1 - \frac{1}{2r} - \sqrt{\frac{1}{4r^2} + \frac{r-1}{r} \cdot \frac{\ell-1}{q-1}}\right). \quad (2)$$

¹In Section 5.3 we denote the code $\text{ev}(\mathbb{F}_q[X]^{[q-1] \cap (1+r\mathbb{Z})})$ by B^\perp , but for simplicity of notation in this section we denote it by P .

As $r \rightarrow \infty$, the bound in Theorem 13 gives

$$\frac{d}{q-1} \rightarrow 1 - \sqrt{\frac{\ell-1}{q-1}}.$$

As was mentioned in Section 1.1.2, this bound is reminiscent of the Johnson bound (see e.g. Theorem 7.3.1 of [GRS22]), which in particular implies that a classical code of dimension ℓ and block length $q-1$ whose distance approaches the Singleton bound is *list-decodable* from at least a fraction of errors approaching $1 - \sqrt{\ell/(q-1)}$ as the alphabet size and block length grow large.

It is unclear if there is a deeper reason for this similarity between the Johnson bound and Theorem 62. However, it is interesting that we are able to improve the distance beyond this Johnson-like bound by folding our qTB codes, just as [GR08] introduced folded RS codes to improve beyond the Johnson bound for the list-decoding radius of RS codes.

Proof sketch of Theorem 13. Recall that the distance of $\mathcal{C} = \text{CSS}(C, C)$ equals the minimum weight of any element of $C \setminus C^\perp$. As $P \subseteq C^\perp$, it therefore suffices show that every element of $C \setminus P$ has weight at least the value d given in (2).

For this purpose, consider any $\text{ev}(f) \in C \setminus P$. At a high level, the proof will proceed as follows. We define a polynomial G such that whenever f has multiple roots in a coset $x\Omega_r \subseteq \mathbb{F}_q^*$, G also has many roots in that coset. We can then bound the number of roots of f by bounding the number of roots of G , which we in turn bound by the degree of G . By constructing G to have low degree relative to the number of roots of f , we obtain the desired result.

We now formally define G . By definition we may decompose $f(X) = g(X) + h(X)$, where $g(X) \in \mathbb{F}_q[X]^{[\ell] \setminus (\pm 1 + r\mathbb{Z})}$ evaluates to a nonzero classical TB codeword, and $h(X) \in \mathbb{F}_q[X]^{[q-1] \cap (1+r\mathbb{Z})}$ evaluates to a piecewise linear function. Define the polynomial $G(X) \in \mathbb{F}_q[X]$ by

$$G(X) = \prod_{i=1}^{r-1} (\omega_r^{-i} g(\omega_r^i X) - g(X)). \quad (3)$$

By definition, the degree of G satisfies

$$\deg(G) \leq (r-1)(\ell-1). \quad (4)$$

Meanwhile, we bound the number of roots of G as follows. For every $x \in \mathbb{F}_q^*$ and every r th root of unity $\omega_r^i \in \Omega_r \setminus \{1\}$ such that $f(x) = 0$ and $f(\omega_r^i x) = 0$, then

$$\omega_r^{-i} g(\omega_r^i x) - g(x) = -\omega_r^{-i} h(\omega_r^i x) + h(x) = -h(x) + h(x) = 0,$$

where the second equality above holds because h is piecewise linear. Thus G has a root for every ordered pair $(x, \omega_r^i x)$ of roots of f whose ratio is an r th root of unity ω_r^i . Summing over all such pairs of roots of f (see Section 6.1 for details), we find that the total number of roots of G is at least

$$\frac{r}{q-1} (q-1 - |\text{ev}(f)|)^2 - (q-1 - |\text{ev}(f)|).$$

This expression must be bounded above by the RHS of (4); rearranging terms in the resulting inequality yields the desired bound $|\text{ev}(f)| \geq d$ for d given in (2). \square

1.2.3 Folded Quantum Tamo-Barg Codes

As described in Section 1.1.2 and Section 1.2.2, we were only able to prove a “Johnson-like” bound on the distance of qTB codes, which does not approach our Singleton-like bounds for qLRCs described in Section 1.1.3. We address this issue by introducing a “folded” version of qTB codes, for which we show the distance does approach the Singleton bound for large localities r . The proof of this distance bound for folded qTB codes (Theorem 2) is quite involved, so in this section we simply provide a brief description of the main ideas involved. The reader is referred to Section 6.2 for the full proof.

We define folded qTB codes by grouping together symbols of qTB codes as described below.

Definition 14 (Restatement of Definition 59). As in Definition 10, let $\mathcal{C} = \text{CSS}(C, C)$ be the qTB code with parameters q, r, ℓ . Given an additional folding parameter $s|(q-1)/r$, we define the **folded quantum Tamo-Barg (fqTB) code** $\tilde{\mathcal{C}}$ to be the quantum code of alphabet size q^s and block length $(q-1)/s$ obtained as follows. Fix a generator ω_{q-1} for \mathbb{F}_q^* , and then for every $i \in [(q-1)/s]$, we block together the s components (each of alphabet size q) at positions $\{\omega_{q-1}^{si}, \omega_{q-1}^{si+1}, \dots, \omega_{q-1}^{si+s-1}\}$ in \mathcal{C} into a single component (of alphabet size q^s) of the folded code $\tilde{\mathcal{C}}$.

A folded qTB code by construction inherits the rate and local recoverability of the underlying unfolded qTB code. Thus to prove Theorem 2, we simply need to bound the distance of the fqTB code. For this purpose, define S as in (1), so that the underlying unfolded qTB code $\mathcal{C} = \text{CSS}(C, C)$ has $C = \text{ev}(\mathbb{F}_q[X]^S)$. The distance of $\tilde{\mathcal{C}}$ by definition equals the minimum over all codewords $\text{ev}(f) \in C \setminus C^\perp$ of the number of distinct blocks $\{\omega_{q-1}^{si}, \omega_{q-1}^{si+1}, \dots, \omega_{q-1}^{si+s-1}\}$ (see Definition 14) in which f takes at least one nonzero value.

At a high level, we bound this minimum distance by considering two types of message polynomials $f(X) = \sum_{i \in S} f_i X^i \in \mathbb{F}_q[X]^S$ separately. For the first type of message polynomial, we use a similar argument as described in the proof sketch of Theorem 13 in Section 1.2.2. We are able to leverage the folding to replace the polynomial $G(X)$ in (3) with another polynomial, which consists of the composition of f with a determinant polynomial. This alternative choice of $G(X)$ detects roots of f more efficiently relative to its degree, and hence yields a better distance bound (see Claim 68). However, this method breaks down when the coefficients f_i of f are supported in a small number of distinct values $i \pmod{r}$, that is, when there are $\ll r$ distinct values $i \pmod{r}$ among all i with $f_i \neq 0$. The reason for the breakdown on these “bad” polynomials f stems from the fact that we need $G(X)$ to be a nonzero polynomial, which becomes more difficult when f has fewer nonzero coefficients.

Fortunately, we can consider such “bad” message polynomials f separately, and instead apply an uncertainty principle (Proposition 66) to bound the weight of the associated codeword (Claim 67). Intuitively, this uncertainty principle implies that if the coefficients f_i are zero for most values of $i \pmod{r}$, then most of the evaluation points in $\text{ev}(f)$ must be nonzero, so $\text{ev}(f)$ has large weight.

Combining our bounds from the two types of message polynomials described above, we obtain a bound on the distance of the fqTB code \mathcal{C} , which yields Theorem 2.

1.3 Open Questions

Our work leads to the following open questions:

- Can explicit qLRCs be constructed with a better rate-distance-locality tradeoff than given by Theorem 2? Two approaches here are to improve our distance bound for fqTB codes in Theorem 2, or to introduce new constructions of qLRCs for which stronger bounds can be shown. In general, the goal is to close the gap to the Singleton-like bounds in Section 3.
- How does the optimal rate-distance-locality tradeoff of qLRCs differ from that of qLDPC codes? As described in Section 1.1.3, our bounds on qLRCs imply bounds on qLDPC codes, but it is an interesting question whether the additional structure in qLDPC codes can be used to show stronger bounds. This question also points towards the more general line of inquiry into relationships between notions of locality in quantum codes, of which local recoverability and LDPC are two examples of interest.

2 Preliminaries

In this section, we introduce some notation, and then present preliminary definitions and results pertaining to classical codes, quantum codes, and local recoverability.

2.1 Notation

For $n \in \mathbb{N}$, we let $[n] = \{0, \dots, n-1\}$.

For a prime power q , let \mathbb{F}_q denote the finite field of order q . For $n \in \mathbb{N}$ and for $x, y \in \mathbb{F}_q^n$, let $x \cdot y = \sum_{i \in [n]} x_i y_i$ denote the standard dot product.

Let $\mathbb{F}_q^* = \mathbb{F}_q \setminus \{0\}$ denote the multiplicative group of \mathbb{F}_q . For $r | (q-1)$, we let $\omega_r \in \mathbb{F}_q^*$ denote a primitive r th root of unity, and we let $\Omega_r = \{1, \omega_r, \dots, \omega_r^{r-1}\}$ denote the group of all r th roots of unity. While the choice of ω_r is not unique, we typically will think of fixing some $(q-1)$ st root of unity $\omega_{q-1} \in \mathbb{F}_q^*$, and then let $\omega_r = \omega_{q-1}^{(q-1)/r}$.

For a finite alphabet A , we denote the Hamming distance between two strings $x, y \in A^n$ by $\text{dis}(x, y) = |\{i \in [n] : x_i \neq y_i\}|$. For two sets $X, Y \subseteq A^n$, the minimum Hamming distance is denoted $\text{dis}(X, Y) = \min_{x \in X, y \in Y} \text{dis}(x, y)$. If A is an abelian group (e.g. $A = \mathbb{F}_q$, or $A = \mathbb{F}_q^s$) we denote the Hamming weight of a string $x \in A^n$ by $|x| = \text{dis}(0, x) = |\{i \in [n] : x_i \neq 0\}|$, so that $\text{dis}(x, y) = |x - y|$.

We denote by $\mathcal{M}(A)$ the set of density matrices on a set A of qudits, where the local dimension a of the qudits will be clear from context. Thus $\rho \in \mathcal{M}(A)$ if ρ is a positive semi-definite Hermitian operator $\rho : (\mathbb{C}^a)^{\otimes A} \rightarrow (\mathbb{C}^a)^{\otimes A}$ of trace 1. For a pure state $|\psi\rangle \in (\mathbb{C}^a)^{\otimes A}$, we let $\psi = |\psi\rangle\langle\psi| \in \mathcal{M}(A)$ denote the associated density matrix. For a density matrix $\rho \in \mathcal{M}(A)$, if $B \subseteq A$ is a subset of the qudits, we denote the reduced density matrix $\rho_B = \text{Tr}_{A \setminus B} \rho \in \mathcal{M}(B)$.

A quantum channel \mathcal{K} from qudits A to qudits B is a completely positive, trace preserving map $\mathcal{K} : \mathcal{M}(A) \rightarrow \mathcal{M}(B)$. Equivalently, a quantum channel is a map $\mathcal{K} : \mathcal{M}(A) \rightarrow \mathcal{M}(B)$ that can be expressed in the form $\mathcal{K}(\rho) = \sum_{\nu} K_{\nu} \rho K_{\nu}^{\dagger}$ for a set of operators $\{K_{\nu} : (\mathbb{C}^a)^{\otimes A} \rightarrow (\mathbb{C}^a)^{\otimes B}\}_{\nu}$, called Kraus operators, which satisfy $\sum_{\nu} K_{\nu}^{\dagger} K_{\nu} = I$.

We let the weight of a quantum channel $\mathcal{K} : \mathcal{M}(A) \rightarrow \mathcal{M}(A)$ refer to the number of distinct qudits on which the channel acts nontrivially. Therefore if \mathcal{K} has weight w , then there exists a subset $W \subseteq A$ of size $|W| = w$ such that each Kraus operator K_{ν} of \mathcal{K} can be decomposed as $K_{\nu} = K_{\nu}^W \otimes I_{A \setminus W}$ for $K_{\nu}^W : (\mathbb{C}^a)^{\otimes W} \rightarrow (\mathbb{C}^a)^{\otimes W}$.

2.2 Classical Codes

Below we define the notion of a classical code.

Definition 15. A **classical code** of **block length** n , **dimension** k , and **alphabet size** a is a subset $C \subseteq [a]^n$ of size $|C| = a^k$.

The **distance** d of C is defined as the minimum Hamming distance $d = \min_{c \neq c' \in C} \text{dis}(c, c')$ between any two distinct codewords $c, c' \in C$.

We say that C is **linear** if $a = q$ is a prime power so that $[a] \cong \mathbb{F}_q$ and $C \subseteq \mathbb{F}_q^n$ is a linear subspace. Similarly, C is **\mathbb{F}_q -linear** if $a = q^s$ for some prime power q so that $[a] \cong \mathbb{F}_q^s$ and $C \subseteq (\mathbb{F}_q^s)^n \cong \mathbb{F}_q^{sn}$ is linear subspace over the field \mathbb{F}_q . The **dual** of a linear code is $C^\perp = \{c' \in \mathbb{F}_q^n : c' \cdot c = 0 \forall c \in C\}$. Nonzero elements of the dual code C^\perp are sometimes referred to as **parity checks** of C .

All classical codes satisfy that following fundamental limitation on the tradeoff between distance and dimension, whose proof is a simple application of the pigeonhole principle.

Proposition 16 (Singleton bound). *If C is a classical code of block length n , dimension $k > 0$, and distance d , then*

$$k \leq n - (d - 1).$$

A classical (linear) code can be expressed as the image of a (linear) encoding map $\text{Enc} : \mathbb{F}_q^k \rightarrow \mathbb{F}_q^n$ that maps messages to codewords encoding the messages. However, we will mostly consider the space of codewords without reference to the encoding map.

Throughout this paper we restrict attention to linear or \mathbb{F}_q -linear codes, as will be clear from context. In particular, \mathbb{F}_q -linear codes will arise from linear codes over \mathbb{F}_q by “folding”, that is, by grouping together symbols of the linear code into larger symbols.

Definition 17. A **classical locally recoverable code (cLRC)** of **locality** r , block length n , and alphabet size a is a classical code $C \subseteq [a]^n$ together with a set of **local recovery maps** Rec_i for every $i \in [n]$ satisfying the following properties:

1. There exists a set $I_i \subseteq [n]$ of size $|I_i| \leq r$ with $i \in I_i$ such that Rec_i is a function $\text{Rec}_i : [a]^{I_i \setminus \{i\}} \rightarrow [a]$.
2. It holds for every codeword $c \in C$ that

$$\text{Rec}_i(c_{|I_i \setminus \{i\}}) = c_i.$$

We specify that $i \in I_i$ for notational convenience due to the following well known form of linear cLRCs. However, we emphasize that an LRC with locality r can recover any given codeword symbol from $r - 1$ other symbols.

Lemma 18 (Well known). *If $C \subseteq \mathbb{F}_q^n$ is a linear code such that every $i \in [n]$ lies in the support of some parity check of Hamming weight $\leq r$, then C is locally recoverable with locality r .*

Proof. Fix $i \in [n]$, and let $c' \in C^\perp$ be a parity check of weight $\leq r$ with $i \in I_i := \text{supp}(c')$. Then by definition every $c \in C$ satisfies $c_i = -\sum_{j \in I_i \setminus \{i\}} (c'_j/c'_i)c_j$, which gives the desired recovery function for index i . \square

2.3 Polynomial Evaluation Codes

Below we describe a particularly useful type of classical linear code given by evaluations of polynomials. Here we let $\mathbb{F}_q[X_1, \dots, X_m]$ denote the ring of m -variate polynomials over \mathbb{F}_q .

Definition 19. For a subset $S \subseteq \mathbb{Z}_{\geq 0}$, let

$$\mathbb{F}_q[X]^S = \left\{ \sum_{i \in S} a_i X^i : a \in \mathbb{F}_q^S \right\}$$

denote the $|S|$ -dimensional subspace of $\mathbb{F}_q[X]$ consisting of polynomials with zero coefficients for all monomials X^i with $i \notin S$. Define the **polynomial evaluation map**

$$\text{ev} : \mathbb{F}_q[X]^S \rightarrow \mathbb{F}_q^{\mathbb{F}_q^*} \cong \mathbb{F}_q^{q-1}$$

by $\text{ev}(f) = (f(x))_{x \in \mathbb{F}_q^*}$, so that $\text{ev}(f)$ outputs the list of evaluations of f at all points in \mathbb{F}_q^* . The image $\text{ev}(\mathbb{F}_q[X]^S)$ of ev is a linear code, called a **polynomial evaluation code**, with alphabet size q , block length $q - 1$, and dimension $|S|$.

Below, we present the well-known Reed-Solomon codes, along with their folded counterpart.

Definition 20. For a prime power q and an integer $\ell \in [q]$, the **Reed-Solomon (RS) code** is the polynomial evaluation code $C = \text{ev}(\mathbb{F}_q[X]^{\ell})$.

Given an additional folding parameter $s|(q-1)$, the **folded Reed-Solomon (fRS) code** \tilde{C} is the \mathbb{F}_q -linear code of alphabet size q^s and block length $(q-1)/s$ obtained as follows. Fix a generator ω_{q-1} for \mathbb{F}_q^* , and then for every $i \in [(q-1)/s]$, block together the s components (each an element of \mathbb{F}_q) at positions $\{\omega_{q-1}^{si}, \omega_{q-1}^{si+1}, \dots, \omega_{q-1}^{si+s-1}\}$ in C into a single component (which is an element of \mathbb{F}_q^s) of the folded code \tilde{C} .

By definition Reed-Solomon codes have block length $q-1$, dimension ℓ , and distance $d = q - \ell$. Specifically, it holds that $d \geq (q-1) - (\ell-1) = q - \ell$ because every polynomial of degree $< \ell$ has $< \ell$ roots, and $d \leq q - \ell$ by the Singleton bound. Folded Reed-Solomon codes similarly have block length $(q-1)/s$, dimension ℓ/s , and distance $(q-\ell)/s$.

Our decoding algorithm for the quantum LRCs we construct will rely on the efficient list-decodability of the (folded) Reed-Solomon codes, as stated in the known results below. However, we first must define list decoding.

Definition 21. Let C be a code of block length n over alphabet A , and let $\delta \geq 0$. An **e -list-decoding** algorithm for C is an algorithm that takes as input a corrupted codeword $b \in A^n$, and outputs the list of every codeword $c \in C$ such that $\text{dis}(b, c) \leq e$.

In the statements below, recall that the RS code has block length $q-1$ and rate $R = \ell/(q-1)$, while the fRS code has block length $(q-1)/s$ and rate $R = \ell/(q-1)$.

Theorem 22 ([GS98]). *The RS code with parameters q, ℓ has an e -list-decoding algorithm that runs in time $q^{O(1)}$ for*

$$e = (q-1)(1 - \sqrt{R})$$

Theorem 23 ([GR08]). *The fRS code with parameters q, ℓ, s has an e -list decoding algorithm that runs in time $q^{O(\sqrt{s})}$, where*

$$e = \frac{q-1}{s} \left(1 - \left(1 + \frac{2}{\sqrt{s}} \right) R^{1-1/\sqrt{s}} \right) - 2,$$

assuming that $s \geq s_0$ and $q \geq q_0 = q_0(s)$ for sufficiently large constants $s_0, q_0(s)$.

Note that the running times of the above algorithms implicitly give bounds on the lengths of the lists they output.

2.4 Quantum Codes

In this section, we describe relevant background on quantum codes.

Definition 24. A quantum code of **block length** n , **dimension** k , and **local dimension (that is, alphabet size)** a is a a^k -dimensional linear subspace $\mathcal{C} \subseteq (\mathbb{C}^a)^{\otimes n}$.

The definition above of a quantum code as a linear subspace of Hilbert space assumes a unitary encoding map, as we may express $\mathcal{C} = \{ |\text{Enc } \phi\rangle |0^{n-k}\rangle : |\phi\rangle \in (\mathbb{C}^a)^{\otimes k} \}$ for a unitary encoding map $\text{Enc} : (\mathbb{C}^a)^{\otimes k} \rightarrow (\mathbb{C}^a)^{\otimes n}$. In this paper we restrict attention to codes with such unitary encodings unless explicitly stated otherwise. However, there are more general coding schemes with arbitrary channel encodings; we leave it as an open question whether such non-unitary encodings could be beneficial for local recovery.

Below we define the distance of a quantum code as one plus the maximum number of erasures that the code can correct.

Definition 25. Let \mathcal{C} be a quantum code of block length n . Given a subset $S \subseteq [n]$, we say that \mathcal{C} can **decode from erasures in S** if there exists a decoding channel $\text{Dec}^S : \mathcal{M}([n] \setminus S) \rightarrow \mathcal{M}([n])$ that satisfies

$$\text{Dec}^S(\psi_{[n] \setminus S}) = \psi.$$

for every $\psi \in \mathcal{C}$.

The **distance** of \mathcal{C} is then defined as the maximum value $d \in \mathbb{N}$ such that \mathcal{C} can decode from erasures in every $S \subseteq [n]$ of size $|S| < d$.

The well-known fact below (see e.g. Fact 2 of [AN22] for a proof) shows that for a quantum code of distance d , any $d-1$ codeword symbols contain no information about the encoded message. Note that there is no classical analogue to this fact, as classically every linear code has a systematic encoding, meaning that the first k codeword symbols equal the message.

Lemma 26 (Local indistinguishability). *If \mathcal{C} is a quantum code of block length n that can decode from erasures in some set $S \subseteq [n]$, then the reduced density matrix ψ_S is the same for all $\psi \in \mathcal{C}$.*

The following well known quantum analogue of Proposition 16 presents a tradeoff between distance and dimension for quantum codes. See for instance Section 12.4.3 of [NC10] for a proof.

Proposition 27 (Singleton bound). *If \mathcal{C} is a quantum code of block length n , dimension $k > 0$, and distance d , then*

$$k \leq n - 2(d - 1).$$

We next define stabilizer codes, a well-known class of quantum codes that are defined as the simultaneous +1 eigenspace of a set of Pauli operators, called *stabilizers*. We will first need to define the Pauli operators.

Definition 28. Let \mathbb{F}_q be a finite field of characteristic p . For $\alpha \in \mathbb{F}_q$, define the q -ary **Pauli operators** $X^\alpha, Z^\alpha \in \mathbb{C}^{\mathbb{F}_q \times \mathbb{F}_q} \cong \mathbb{C}^{q \times q}$ so that for $x \in \mathbb{F}_q$,

$$\begin{aligned} X^\alpha |x\rangle &= |x + \alpha\rangle \\ Z^\alpha |x\rangle &= e^{(2\pi i/p) \operatorname{tr}_{\mathbb{F}_q/\mathbb{F}_p}(\alpha x)} |x\rangle. \end{aligned}$$

Observe that Paulis always commute up to a phase, and specifically $e^{(2\pi i/p) \operatorname{tr}_{\mathbb{F}_q/\mathbb{F}_p}(\alpha\beta)} X^\alpha Z^\beta = Z^\beta X^\alpha$.

For $n \in \mathbb{N}$, an n -**qudit Pauli** is an operator of the form $X^\alpha Z^\beta = \bigotimes_{i=1}^n X^{\alpha_i} Z^{\beta_i}$ for $\alpha, \beta \in \mathbb{F}_q^n$. We take \mathcal{P}_q^n denote the group of n -qudit Paulis modulo the global phase, so that $\mathcal{P}_q^n \cong \mathbb{F}_q^{2n}$. The **support** $\operatorname{supp}(P)$ of an n -qudit Pauli $P \in \mathcal{P}_q^n$ is defined as the set $i \in [n]$ of qudits on which P acts nontrivially. That is, $\operatorname{supp}(X^\alpha Z^\beta) = \{i \in [n] : (\alpha_i, \beta_i) \neq (0, 0)\}$. The **Hamming weight** of P is $|P| = |\operatorname{supp}(P)|$.

As \mathcal{P}_q^n is isomorphic to \mathbb{F}_q^{2n} , it is a vector space where the action of $\gamma \in \mathbb{F}_q$ on $X^\alpha Z^\beta \in \mathcal{P}_q^n$ gives $X^{\gamma\alpha} Z^{\gamma\beta}$. We therefore let a *subspace*, or (by slight abuse of terminology) *subgroup*, of \mathcal{P}_q^n be any subset that is a \mathbb{F}_q -subspace under the isomorphism $\mathcal{P}_q^n \cong \mathbb{F}_q^{2n}$. Therefore in particular we require a “subgroup” of \mathcal{P}_q^n to be closed under the action of \mathbb{F}_q . Furthermore, although we defined \mathcal{P}_q^n to be the group of Paulis modulo a global phase, we abuse notation and call a subgroup of \mathcal{P}_q^n *abelian* if all the operators in the subgroup commute with each other, taking into account the global phase. We also denote $(\mathcal{P}_q^s)^n = \mathcal{P}_q^{sn}$, which we interpret as the group of length- n strings of elements in \mathcal{P}_q^s ; this group will naturally arise when we consider folded stabilizer codes in Section 5.3.

Definition 29. For a vector space \mathbb{F}_q^s over a finite field, a **stabilizer code of block length n over the alphabet \mathbb{F}_q^s** is a subgroup of the Hilbert space $(\mathbb{C}^{\mathbb{F}_q^s})^{\otimes n}$ that is specified as the simultaneous +1 eigenspace of some abelian subgroup of $(\mathcal{P}_q^s)^n$. This abelian subgroup is called the **stabilizer group** of the code, and its elements are the code’s **stabilizers**.

We will mostly be concerned with a specific type of stabilizer code called a CSS code, which is specified by two classical codes satisfying an orthogonality condition.

Definition 30. Given classical codes $C_X, C_Z \subseteq (\mathbb{F}_q^s)^n$ such that $C_X^\perp \subseteq C_Z$, the associated **CSS code** $\mathcal{C} = \operatorname{CSS}(C_X, C_Z)$ is given by $\mathcal{C} = \operatorname{span}\{\sum_{y \in C_X^\perp} |x + y\rangle : x \in C_Z\} \subseteq (\mathbb{C}^{\mathbb{F}_q^s})^{\otimes n}$. Equivalently, \mathcal{C} is the stabilizer code with stabilizer group $\{X^\alpha Z^\beta : \alpha \in C_X^\perp, \beta \in C_Z^\perp\}$.

It is well known that to decode a CSS code $\mathcal{C} = \operatorname{CSS}(C_X, C_Z)$, it is sufficient to have classical decoders for C_X and C_Z , as stated below:

Proposition 31 (Well known). *Let $\mathcal{C} = \operatorname{CSS}(C_X, C_Z)$ be a CSS code of block length n over the alphabet $A = \mathbb{F}_q^s$ of size $a = |A|$. Let $e \geq 0$ be an integer such that for each permutation (α, β) of (X, Z) , there exists a (classical) decoding algorithm $\operatorname{Dec}_\alpha$ that takes as input a (classical) corrupted codeword $c + b$ for some $c \in C_\alpha$ and some corruption $b \in A^n$ of Hamming weight $|b| \leq e$, and outputs some $c' \in C_\alpha$ such that $c' - c \in C_\beta^\perp$.*

Then \mathcal{C} has a decoding algorithm Dec that recovers from errors of weight e , so \mathcal{C} has distance $d \geq 2e + 1$. Furthermore, if each Dec_α has running time $T_\alpha(n, a)$, then Dec has running time $T_X(n, a) + T_Z(n, a) + O(n^3 \text{poly log } a)$.

Proof sketch. The result is standard, so we just provide a brief outline. Let $C_X = \ker H_X$ and $C_Z = \ker H_Z$. Given a corrupted code state $\rho = \mathcal{E}(\psi)$ for a code state $\psi \in \mathcal{C}$ and an error channel \mathcal{E} acting on $\leq e$ of the codeword qudits, the decoder $\text{Dec}(\rho)$ first performs syndrome measurements for H_X, H_Z to collapse ρ to $E\psi E^\dagger$, so that the error on ψ is collapsed to some Pauli $E = X^{b_X} Z^{b_Z}$ of weight $\leq e$, where $s_X = H_X b_X$ and $s_Z = H_Z b_Z$ are the outputs of the syndrome measurements. Then for each permutation (α, β) of (X, Z) , the decoder performs Gaussian elimination to find some a_α such that $H_\alpha a_\alpha = s_\alpha$, so that $a_\alpha = c_\alpha + b_\alpha$ for some (currently unknown) $c_\alpha \in C_\alpha$. The decoder then runs $\text{Dec}_\alpha(a_\alpha)$, which outputs $c_\alpha + p_\alpha$ for some $p_\alpha \in C_\beta^\perp$. Letting $b'_\alpha = a_\alpha - (c_\alpha + p_\alpha) = b_\alpha - p_\alpha$ and $E' = X^{b'_X} Z^{b'_Z}$, the decoder then applies E'^\dagger to output $E'^\dagger E\psi E^\dagger E' = \psi$. This final equality holds because by assumption X^{p_X} and Z^{p_Z} are stabilizers of \mathcal{C} , so they preserve ψ .

The running time of Dec defined above is $T_X(n, a) + T_Z(n, a) + O(n^3 \text{poly log } a)$ because outside of the calls to Dec_X and Dec_Z , all operations run in time $O(n^2 \text{poly log } a)$, except the Gaussian elimination which runs in time $O(n^3 \text{poly log } a)$. \square

Definition 32. A **quantum locally recoverable code (qLRC)** of **locality** r , **block length** n , and **local dimension** (i.e. alphabet size) a is a quantum code $\mathcal{C} \subseteq (\mathbb{C}^a)^{\otimes n}$ together with a set of **local recovery channels** Rec_i for every $i \in [n]$ satisfying the following properties:

1. There exists a set $I_i \subseteq [n]$ of size $|I_i| \leq r$ with $i \in I_i$ such that Rec_i is a quantum channel $\text{Rec}_i : \mathcal{M}(I_i \setminus \{i\}) \rightarrow \mathcal{M}(I_i)$ with input qudits $I_i \setminus \{i\}$ and output qudits I_i .
2. It holds for every code state $\psi \in \mathcal{C}$ that

$$\text{Rec}_i \otimes I_{[n] \setminus I_i}(\psi_{[n] \setminus \{i\}}) = \psi.$$

That is, a qLRC permits local recovery against erasure of a single qudit, or equivalently, against a corruption in a single known location.

Below, we present a quantum analogue of Lemma 18 for stabilizer codes.

Proposition 33. *Let \mathcal{C} be a stabilizer code with stabilizer group $S \subseteq \mathcal{P}_q^n$. Assume that for every $i \in [n]$, there exist stabilizers $P = X^\alpha Z^\beta, Q = X^\gamma Z^\delta \in S$ with $|\text{supp}(P) \cup \text{supp}(Q)| \leq r$ such that $(\alpha_i, \beta_i) = (1, 0)$ and $(\gamma_i, \delta_i) = (0, 1)$. Then \mathcal{C} is locally recoverable with locality r .*

Proof sketch. The result is a standard application of the well-known fact that a stabilizer code detects any Pauli error that anticommutes with one of the code's stabilizers. We will therefore just provide a brief proof sketch. For any given $i \in [n]$, let $P = X^\alpha Z^\beta, Q = X^\gamma Z^\delta \in S$ be the Paulis given by the proposition statement, and let $I_i = \text{supp}(P) \cup \text{supp}(Q)$. We will construct a recovery channel to revert the action of an arbitrary error channel (such as a totally depolarizing channel) acting on qudit i . The recovery channel first performs syndrome measurements for P and Q , specifically by performing measurements on the corrupted code state for the operators P^η and Q^η for all $\eta \in \mathbb{F}_q$. These measurements will collapse the error on qudit i to a simultaneous eigenspace of these operators P^η and Q^η for $\eta \in \mathbb{F}_q$, and the measurement outcomes (i.e. the syndrome) give the eigenvalues of the projected state for each of these operators. That is, the error on qudit i is

projected to some single-qudit Pauli $E = X^{e_X} Z^{e_Z}$ for $e_X, e_Z \in \mathbb{F}_q^n$ both supported in component i . The syndrome measurement provides the phases $(EP^\eta)(P^\eta E)^\dagger$ and $(EQ^\eta)(Q^\eta E)^\dagger$ for all $\eta \in \mathbb{F}_q$. These phases can be used to compute the symplectic inner products $s_P = e_X \cdot \beta - e_Z \cdot \alpha$ and $s_Q = e_X \cdot \delta - e_Z \cdot \gamma$. By the assumption that $(\alpha_i, \beta_i) = (1, 0)$ and $(\gamma_i, \delta_i) = (0, 1)$, it follows that $s_P = -(e_Z)_i$ and $s_Q = (e_X)_i$. That is, s_P and s_Q give the Z and X errors that occurred on qudit i , respectively. The recovery channel then simply applies $X^{-e_X} Z^{-e_Z}$ to revert the error and recover the original code state. \square

Proposition 33 immediately implies that a CSS code is locally recoverable if its two classical codes are locally recoverable linear codes in the sense of Lemma 18:

Corollary 34. *A CSS code $\mathcal{C} = \text{CSS}(C_X, C_Z)$ is locally recoverable with locality r if for every $i \in [n]$, there exist parity checks $c'_X \in C_X^\perp, c'_Z \in C_Z^\perp$ such that $i \in \text{supp}(c'_X) \cap \text{supp}(c'_Z)$ and $|\text{supp}(c'_X) \cup \text{supp}(c'_Z)| \leq r$.*

3 Singleton-Like Bounds for qLRCs

This section presents Singleton-like bounds for qLRCs, by combining techniques for proving quantum Singleton bounds (see for example Section 12.4.3 of [NC10]) and classical LRC Singleton-like bounds [GHSY12].

Below, for a qLRC \mathcal{C} of block length n and locality i , we let $I_i \subseteq [n]$ denote the set of size r with $i \in I_i$ such that the local recovery map Rec_i for \mathcal{C} recovers the i th code component from the components in $I_i \setminus \{i\}$.

Theorem 35. *Let \mathcal{C} be a qLRC of block length n , dimension $k > 0$, distance d , and locality r . Then*

$$k \leq n - 2(d - 1) - \left\lfloor \frac{n - (d - 1)}{r} \right\rfloor - \left\lfloor \frac{n - 2(d - 1) - \left\lfloor \frac{n - (d - 1)}{r} \right\rfloor}{r} \right\rfloor,$$

Omitting the last term in the bound in Theorem 35, it follows that $k \leq n - 2(d - 1) - (n - (d - 1))/r + 1$, or equivalently,

$$d - 1 \leq \frac{(1 - \frac{1}{r})n - k + 1}{2 + \frac{1}{r}}.$$

Thus for fixed locality r , even as the rate $k/n \rightarrow 0$, the relative distance of a qLRC satisfies $d/n \leq 1/2 - \Omega(1/r)$. In contrast, general (non-LRC) quantum codes can have $d/n \rightarrow 1/2$ as $k/n \rightarrow 0$. Thus for codes of low rate, there is a fundamental cost to imposing a constant locality r in quantum codes.

This fact differs from the classical case, as the classical Singleton-like bound [GHSY12] states that every cLRC of block length n , dimension $k > 0$, distance d , and locality r satisfies

$$d \leq n - k - \left\lceil \frac{k}{r} \right\rceil + 2.$$

Tamo and Barg [TB14] provided an explicit construction of cLRCs that exactly meet this bound. Therefore there exist cLRCs with $k/n \rightarrow 0$ and $d/n \rightarrow 1$, which approaches the classical Singleton bound for general (non-locally-recoverable) codes.

Because all quantum LDPC stabilizer codes are by definition quantum LRCs, the above observation in particular implies that there is a cost (in terms of rate-distance tradeoff) to requiring a quantum code be LDPC with some fixed locality. As for LRCs, letting the rate $k/n \rightarrow 0$ gives a disconnect between the quantum case and the classical case. Specifically, Theorem 35 implies that all qLDPC codes with constant locality (i.e. check weight) r have relative distance $d/n \leq 1/2 - \Omega(1/r)$. In contrast, classically there are LDPC codes with constant locality r and distance approaching 1, such as the repetition code ($r = 2$) and the Hadamard code ($r = 3$) over a growing field size.

It is an interesting question to determine the optimal tradeoff between rate, distance, and locality for qLDPC codes in addition to for qLRCs, and to classify the additional cost of requiring a quantum code to be LDPC compared to just requiring local recoverability.

Proof of Theorem 35. We partition the code symbols $Q = [n]$ into five disjoint subsets S_1, V_1, S_2, V_2, W . Here as a shorthand we denote $A \sqcup B$ by AB , so we will have $Q = S_1V_1S_2V_2W$. To construct this partition, we first construct S_1, V_1 as follows:

1. Initialize sets $S_1 = \emptyset$ and $T_1 = \emptyset$.
2. Repeat the following step

$$s_1 = \left\lfloor \frac{n - (d - 1)}{r} \right\rfloor$$

times: choose some $i \in Q \setminus S_1T_1$, add i to S_1 , and add all elements of $I_i \setminus (\{i\} \cup S_1)$ to T_1 . Note that by construction $S_1 \cap T_1 = \emptyset$.

3. Set V_1 to be any $(d - 1)$ -element subset of $Q \setminus S_1T_1$.

The above procedure is guaranteed to successfully output S_1, V_1 with $|S_1| = s_1$ and $|V_1| = d - 1$, as each iteration in step 2 adds 1 element to S_1 and adds at most r elements to S_1T_1 , so after all s_1 iterations it still holds that $|Q \setminus S_1T_1| \geq n - rs_1 \geq d - 1$.

Let $Q_1 = Q \setminus S_1V_1$. For a code state $\rho \in \mathcal{C}$, given access to the restriction $\rho_{Q_1} = \text{Tr}_{Q \setminus Q_1} \rho$ to components in Q_1 , we may by construction pass through all $i \in S_1$ one-by-one in the same order S_1 was constructed, and apply Rec_i at each step, to recover $\rho_{Q_1S_1}$. Then as $Q \setminus Q_1S_1 = V_1$, which has size $d - 1$, the global decoding algorithm of \mathcal{C} recovers ρ from $\rho_{Q_1S_1}$. Thus the code components in Q_1 can be used to completely recover the entire codeword. It follows by Lemma 26 that the reduced density matrix $\rho_{Q \setminus Q_1} = \rho_{S_1V_1}$ is the same for all code states $\rho \in \mathcal{C}$.

Now if $|Q_1| < d - 1$, then the global decoding algorithm for \mathcal{C} can recover a code state from components in $Q \setminus Q_1 = S_1V_1$. But as we showed above that a code state can also be recovered from components in Q_1 , the code must have dimension $k = 0$, as otherwise we would be able to clone a code state by breaking it into the parts S_1V_1 and Q_1 , and recovering the entire state from each part. But we assumed $k > 0$, a contradiction, so it must be that $|Q_1| \geq d - 1$.

We then construct S_2, V_2 using a similar procedure as above, but on the restriction to components in Q_1 :

1. Initialize sets $S_2 = \emptyset$ and $T_2 = \emptyset$.
2. Repeat the following step

$$s_2 = \left\lfloor \frac{|Q_1| - (d - 1)}{r} \right\rfloor = \left\lfloor \frac{n - 2(d - 1) - s_1}{r} \right\rfloor = \left\lfloor \frac{n - 2(d - 1) - \left\lfloor \frac{n - (d - 1)}{r} \right\rfloor}{r} \right\rfloor$$

times: choose some $i \in Q_1 \setminus S_2T_2$, add i to S_2 , and add all elements of $I_i \setminus (\{i\} \cup S_2)$ to T_2 .

3. Set V_2 to be any $(d-1)$ -element subset of $Q_1 \setminus S_2T_2$.

Note that the procedure above is guaranteed to successfully output S_2, V_2 with $|S_2| = s_2$ and $|V_2| = d-1$ because $|Q_1| \geq d-1$ as shown above, and each iteration of step (2) adds 1 element to S_2 and adds at most r elements to S_2T_2 , so after all s_2 iterations $|S_2| = s_2$ and $|Q_1 \setminus S_2T_2| \geq |Q_1| - |S_2T_2| \geq |Q_1| - rs_2 \geq d-1$. Letting $Q_2 = Q \setminus S_2V_2$, then by similar reasoning as above for Q_1 , the code components in S_2V_2 can be recovered from components in Q_2 , so all code states $\rho \in \mathcal{C}$ have the same reduced density matrix $\rho_{Q \setminus Q_2} = \rho_{S_2V_2}$.

Now let $W = Q \setminus S_1V_1S_2V_2$, so that we have our desired partition $Q = S_1V_1S_2V_2W$. Let q be the local dimension of \mathcal{C} , and let A be a set of k additional qudits of local dimension q . Define

$$|\psi\rangle_{AQ} = \frac{1}{q^k} \sum_{x \in [q]^k} |x\rangle_A \otimes (\text{Enc } |x\rangle)_Q \quad (5)$$

to be the state obtained by applying the encoding map Enc of \mathcal{C} to one register from a maximally entangled pair of message states $\sum_{x \in [q]^k} |x\rangle \otimes |x\rangle$. Let $\psi = |\psi\rangle \langle \psi|$ be the associated density matrix.

For $B \subseteq AQ$, let $S(B) = S(\psi_B) = \text{Tr}(-\psi_B \log_q \psi_B)$ denote the von Neumann entropy of qudits B in the state $|\psi\rangle$. Then

$$\begin{aligned} S(A) + S(S_1V_1) &= S(AS_1V_1) = S(S_2V_2W) \leq S(S_2V_2) + S(W) \\ S(A) + S(S_2V_2) &= S(AS_2V_2) = S(S_1V_1W) \leq S(S_1V_1) + S(W). \end{aligned} \quad (6)$$

In both lines above, the inequality holds by the subadditivity of entropy. Meanwhile, the first equality in both lines holds because for $i = 1, 2$, we showed above that all code states $\rho \in \mathcal{C}$ have the same reduced density matrix $\rho_{S_iV_i}$, which must equal $\psi_{S_iV_i}$. Therefore $\psi_{AS_iV_i} = q^{-k} \sum_{x \in [q]^k} |x\rangle \langle x| \otimes \rho_{S_iV_i} = \psi_A \otimes \psi_{S_iV_i}$, and thus $S(\psi_{AS_iV_i}) = S(\psi_A) + S(\psi_{S_iV_i})$ as entropy is additive over tensor products.

Adding the two inequalities in (6) gives that $S(A) \leq S(W)$. Now $\psi_A = q^{-k} \sum_{x \in [q]^k} |x\rangle \langle x|$, so $S(A) = k$. Meanwhile, $S(W) \leq |W| = n - |S_1V_2S_wV_2| = n - 2(d-1) - s_1 - s_2$. Thus

$$k = S(A) \leq S(W) = n - 2(d-1) - \left\lfloor \frac{n - (d-1)}{r} \right\rfloor - \left\lfloor \frac{n - 2(d-1) - \left\lfloor \frac{n - (d-1)}{r} \right\rfloor}{r} \right\rfloor,$$

as desired. \square

While Theorem 35 applies for arbitrary qLRCs, our explicit construction of qLRCs in Section 5 has a specific structure: the recovery sets I_i form a partition of the code components $[n]$, that is, for all $i, j \in [n]$ either $I_i = I_j$ or $I_i \cap I_j = \emptyset$.

Below, we show that if we assume the recovery sets I_i have this partition structure, then a stronger Singleton-like bound than in Theorem 35 holds. We will assume for simplicity that all recovery sets I_i have size exactly $|I_i| = r$, as is the case for our explicit constructions in Section 5; a similar result also holds in the more general case where $|I_i| \leq r$.

Theorem 36. *Let \mathcal{C} be a qLRC of block length n , dimension $k > 0$, distance d , and locality r . Assume that for all $i \in [n]$, $|I_i| = r$, and for all $i, j \in [n]$, either $I_i = I_j$ or $I_i \cap I_j = \emptyset$. Then*

$$k \leq \left(1 - \frac{2}{r}\right)n - 2 \left(d - 1 - \left\lceil \frac{d-1}{r-1} \right\rceil\right).$$

Proof. The proof is similar to that of Theorem 35, except we choose the partition $S_1V_1S_2V_2W$ of the qudits more carefully. Specifically, again let $Q = [n]$ denote the set of code qudits, and as a shorthand denote $AB = A \sqcup B$ for disjoint $A, B \subseteq Q$.

Because $\{I_i : i \in Q\}$ forms a partition of $Q = [n]$ with all $|I_i| = r$, we must have $r|n$. Denote by $J_1, \dots, J_{n/r}$ the distinct partition elements, so $\{I_i : i \in [n]\} = \{J_j : j \in [n/r]\}$ and therefore $[n] = J_1 \sqcup \dots \sqcup J_{n/r}$.

We define a partition $Q = S_1V_1S_2V_2W$ of the qudits as follows. Let $a = \lceil (d-1)/(r-1) \rceil$, and fix an arbitrary partition $[n/r] = A_1 \sqcup A_2 \sqcup B$ such that $|A_1| = |A_2| = a$. Choose $S_1, S_2 \subseteq [n]$ to be a pair of disjoint subsets such that S_1 consists of a single element of J_j for each $j \in A_1 \sqcup B$, and S_2 consists of a single element of J_j for each $j \in A_2 \sqcup B$. Note that because $r \geq 2$ by definition, we can indeed choose such S_1, S_2 that are disjoint.

Then define V_1 to be any set of $d-1$ elements in $\sqcup_{j \in A_2} J_j \setminus S_2$, and define V_2 to be any set of $d-1$ elements in $\sqcup_{j \in A_1} J_j \setminus S_2$. Note that there exist such sets V_1, V_2 because by construction $|\sqcup_{j \in A_2} J_j \setminus S_2| = |\sqcup_{j \in A_1} J_j \setminus S_2| = (r-1)a \geq d-1$.

Finally define $W = Q \setminus S_1V_1S_2V_2$, so that by construction $Q = S_1V_1S_2V_2W$ forms a partition of $Q = [n]$.

For $b = 1, 2$, if we are given access to components in $Q \setminus S_bV_b$ of a code state $\rho \in \mathcal{C}$, then as V_b is by construction disjoint from $\sqcup_{i \in S_b} I_i$, we can apply the local recovery maps Rec_i for $i \in S_b$ to recover the components of ρ in S_b . Then we have all components of ρ in $Q \setminus V_b$, so as $|V_b| = d-1$, we can apply the global decoding map for \mathcal{C} to recover ρ . It follows that the components in $Q \setminus S_bV_b$ completely determine the code state ρ , so the reduced density matrix $\rho_{S_bV_b}$ must be the same for all $\rho \in \mathcal{C}$.

We now proceed as in the proof of Theorem 35, though with different values for $|S_1|, |S_2|$. Define $|\psi\rangle$ as in (5), so that (6) again holds here. Summing these inequalities in (6), we conclude as in the proof of Theorem 35 that

$$\begin{aligned} k &\leq |W| \\ &= n - 2(d-1) - |S_1| - |S_2| \\ &= n - 2(d-1) - 2(a + (n/r - 2a)) \\ &= \left(1 - \frac{2}{r}\right)n - 2 \left(d - 1 - \left\lceil \frac{d-1}{r-1} \right\rceil\right), \end{aligned}$$

as desired. □

The separation between the bounds in Theorem 35 and Theorem 36 raises the possibility that qLRCs could benefit from having recovery sets I_i that do not simply form a partition of $[n]$. Such a situation would be in contrast to the classical case. Specifically, the classical LRCs of [TB14] have recovery sets I_i forming a partition of $[n]$, yet they exactly meet the Singleton-like bound of [GHSY12], which applies to LRCs with arbitrary recovery structure (see also Chapter 19 of

[GRS22]). Note that the cLRCs of [TB14] have linear-sized alphabets; the case of smaller alphabets is considered in [GXY19].

4 Basic Constructions Using Known Techniques

In this section, we describe how known techniques yield some basic qLRCs, and highlight the issues with such techniques that our construction addresses.

In Section 4.1, we show how to randomly sample a qLRC whose rate is within $O(1/r)$ of the Singleton-like bounds in Section 3 over an alphabet of size $q = 2^{O(r)}$. However, this construction is non-explicit, and as a result we have no efficient algorithm to certify the distance bound, and no efficient algorithm to decode from errors in unknown locations. In contrast, our fqTB codes presented in Section 5 below only have rate within $O(1/\sqrt{r})$ of the Singleton-like bound, and have alphabet size growing polynomially in the block length. However, our fqTB codes are explicit, so their distance bound is guaranteed, and we provide an efficient decoding algorithm against errors in unknown locations in Section 7.

In Section 4.2, we then show how to explicitly construct an efficiently-decodable qLRC using Alon-Edmunds-Luby (AEL) [AEL95] distance amplification and alphabet reduction. This construction concatenates a small random qLRC, which can be found efficiently via brute force, with a large quantum code of high rate and constant relative distance. The key step of the AEL construction is to then permute the symbols according to the edges of an expander graph, which amplifies the relative distance to almost that of the inner code. The resulting construction (see Proposition 49 and the subsequent discussion) yields qLRCs whose rate is within $O(1/r^{1/4})$ of the Singleton-like bounds in Section 3 over an alphabet of size $q = 2^{O(r)}$. In contrast, as mentioned above, our fqTB codes in Section 5 have a better rate-distance-locality tradeoff, as their rate is within $O(1/\sqrt{r})$ of the Singleton-like bound.

4.1 Random qLRCs

In this section, we analyze the following natural random construction of a qLRC.

Definition 37 (Random qLRC). Given a block length n , a locality parameter $r|n$, an integer $\ell \in [n/2 - n/r]$, and an alphabet \mathbb{F}_q , we define a random qLRC to be a CSS code $\mathcal{C} = \text{CSS}(C_X, C_Z)$ that is sampled as follows. Initialize parity check matrices H_X, H_Z for C_X, C_Z respectively to be $n/r \times n$ matrices such that for $i \in [n]$ and $j \in [n/r]$,

$$(H_X)_{i,j} = \begin{cases} 1, & j \in \{ri, \dots, ri + r - 1\} \\ 0, & \text{otherwise} \end{cases}$$

$$(H_Z)_{i,j} = \begin{cases} -(r-1), & j = ri \\ 1, & j \in \{ri + 1, \dots, ri + r - 1\} \\ 0, & \text{otherwise.} \end{cases}$$

That is, H_X, H_Z are initialized so that $\text{CSS}(\ker H_X, \ker H_Z)$ is a well-defined CSS code that is a qLRC of maximal possible rate, whose recovery sets form a partition of the code components $[n]$. We then sequentially add ℓ random rows to each of H_X and H_Z subject to the orthogonality relations, as follows:

1. Repeat ℓ times: sample a uniformly random vector in $\text{row-span}(H_Z)^\perp \setminus \text{row-span}(H_X)$ and add it as a new row in H_X .
2. Repeat ℓ times: sample a uniformly random vector in $\text{row-span}(H_X)^\perp \setminus \text{row-span}(H_Z)$ and add it as a new row in H_Z .

Thus we have sampled matrices $H_X, H_Z \in \mathbb{F}_q^{(n/r+\ell) \times n}$ with orthogonal row-spans. Thus we may let $C_X = \ker H_X, C_Z = \ker H_Z$, obtain our desired well-defined CSS code $\mathcal{C} = \text{CSS}(C_X, C_Z)$.

The following lemma is immediate from Definition 37.

Lemma 38. *The random qLRC \mathcal{C} in Definition 37 is a qLRC of locality r and dimension $k = n - 2(n/r + \ell)$. Furthermore, each $i \in [n]$ has recovery set $I_i = \{ar, \dots, ar + r - 1\}$ for $a = \lfloor i/r \rfloor$.*

We now analyze this distance of these random qLRCs. We will need the following definition.

Definition 39. The q -ary entropy function $H_q : [0, 1] \rightarrow [0, 1]$ is defined by

$$H_q(x) = x \log_q(q-1) - x \log_q(x) - (1-x) \log_q(1-x).$$

The following proposition bounds the distance of random qLRCs. The proof is similar to that of the well-known Gilbert-Varshamov (GV) bound for truly random codes.

Proposition 40. *Given any $\delta, \epsilon > 0$, the distance $d(\mathcal{C})$ of a random qLRC \mathcal{C} from Definition 37 with parameters n, r , and $\ell \geq (H_q(\delta) + \epsilon)n$ over the alphabet \mathbb{F}_q satisfies*

$$\Pr[d(\mathcal{C}) \geq \delta n] > 1 - 2q^{-\epsilon n}. \quad (7)$$

In particular, if $q \geq 2^{2/\epsilon}$, then it holds for all $\ell \geq (\delta + \epsilon)n$ that

$$\Pr[d(\mathcal{C}) \geq \delta n] > 1 - 2q^{-\epsilon n/2}. \quad (8)$$

Proof. We first prove (7). Consider any fixed nonzero $y \in \mathbb{F}_q^n$. We will first compute $\Pr[y \in C_Z \setminus C_X^\perp]$, so we can then union bound over all low-weight y . After step 1 in the sampling procedure in Definition 37, then H_X is fully constructed and $C_X^\perp = \text{row-span}(H_X)$. If $y \in C_X^\perp$ then $y \notin C_Z \setminus C_X^\perp$, so we will condition on $y \notin C_X^\perp$. For $0 \leq i \leq \ell$, let $H_Z^{(i)}$ denote the matrix H_Z after the i th iteration of step 2 of the sampling procedure.

Now if $y \in C_Z \setminus C_X^\perp$, then for every $0 \leq i \leq \ell - 1$, it must hold that $y \in \ker H_Z^{(i)}$, and that $y \in \ker H_Z^{(i+1)}$, where $H_Z^{(i+1)}$ is obtained from $H_Z^{(i)}$ by adding a row given by a uniformly random $h \in C_X \setminus \text{row-span}(H_Z^{(i)})$. Because $y \notin C_X^\perp$ by assumption, exactly $1/q$ -fraction of the vectors in C_X are orthogonal to y . Conditioning on the event that $y \in \ker H_Z^{(i)}$, then all vectors in $\text{row-span}(H_Z^{(i)})$ are orthogonal to y . Thus less than $1/q$ -fraction of all $h \in C_X \setminus \text{row-span}(H_Z^{(i)})$ are orthogonal to y , that is,

$$\Pr\left[y \in \ker H_Z^{(i+1)} \setminus C_X^\perp \mid y \in \ker H_Z^{(i)} \setminus C_X^\perp\right] = \Pr_{h \sim \text{Unif}(C_X \setminus \text{row-span}(H_Z^{(i)}))} [y \cdot h = 0] < \frac{1}{q}.$$

Therefore

$$\begin{aligned} \Pr\left[y \in C_Z \setminus C_X^\perp\right] &= \Pr\left[y \notin C_X^\perp\right] \cdot \Pr\left[y \in \ker H_Z^{(0)} \mid y \notin C_X^\perp\right] \\ &\quad \cdot \prod_{i=0}^{\ell-1} \Pr\left[y \in \ker H_Z^{(i+1)} \setminus C_X^\perp \mid y \in \ker H_Z^{(i)} \setminus C_X^\perp\right] \\ &< \frac{1}{q^\ell}. \end{aligned}$$

Union bounding over all $y \in \mathbb{F}_q^n$ of Hamming weight $|y| \leq \delta n$ gives

$$\begin{aligned} \Pr\left[\exists y \in C_Z \setminus C_X^\perp \text{ s.t. } |y| \leq \delta n\right] &< |\{y \in \mathbb{F}_q^n : |y| \leq \delta n\}| \cdot \frac{q}{q^\ell} \\ &\leq q^{H_q(\delta)n - \ell} \\ &= q^{-\epsilon n}, \end{aligned}$$

where the second inequality above follows from the well-known fact that a radius- δn Hamming ball in \mathbb{F}_q^n has volume $\leq q^{H_q(\delta)n}$ (see for instance Proposition 3.3.3 in [GRS22]), and the final equality follows from the assumption that $\ell \geq (H_q(\delta) + \epsilon)n$.

Now observe that the final distribution over codes $\mathcal{C} = \text{CSS}(C_X, C_Z)$ remains the same if we swap the order of steps 1,2 in the sampling procedure in Definition 37, as in both cases by symmetry we obtain a uniform distribution over pairs of matrices $H_X, H_Z \in \mathbb{F}_q^{(n/r+\ell) \times n}$ satisfying $H_X H_Z^\top = 0$ such that the first n/r rows of each matrix are as given in Definition 37. Thus swapping C_X and C_Z in the above argument gives that

$$\Pr\left[\exists y \in C_X \setminus C_Z^\perp \text{ s.t. } |y| \leq \delta n\right] < q^{-\epsilon n}.$$

Then because $d(\mathcal{C}) = \min_{y \in (C_Z \setminus C_X^\perp) \cup (C_X \setminus C_Z^\perp)} |y|$, we conclude that

$$\Pr[d(\mathcal{C}) \leq \delta n] \leq 2q^{-\epsilon n}.$$

It remains to prove the second claim in the proposition statement. That is, assuming that $q \geq 2^{2/\epsilon}$ and $\ell \geq (\delta + \epsilon)n$, we want to prove that (8) holds. But by definition $H_q(\delta) \leq \delta + 1/\log_2 q$, so for $q \geq 2^{2/\epsilon}$ then $H_q(\delta) \leq \delta + \epsilon/2$, which means that $\ell \geq (\delta + \epsilon)n \geq (H_q(\delta) + \epsilon/2)n$. Then applying (7) with $\epsilon/2$ replacing ϵ gives

$$\Pr[d(\mathcal{C}) \geq \delta n] > 1 - 2q^{-\epsilon n/2},$$

as desired. \square

Letting the alphabet size $q \rightarrow \infty$, we immediately obtain the following corollary.

Corollary 41. *For every fixed $\delta > 0$ and for every alphabet size q , there exists $\epsilon = \epsilon(q) > 0$ with $\epsilon(q) \rightarrow 0$ as $q \rightarrow \infty$ such that a random qLRC \mathcal{C} from Definition 37 with parameters $n, r, \ell = \lceil (H_q(\delta) + \epsilon/4)n \rceil$ has dimension*

$$k \geq \left(1 - \frac{2}{r}\right)n - 2d - \epsilon n,$$

and with high probability has distance $d = \delta n$.

Thus for large q , the dimension of these random qLRCs is within $2d/r$ of the Singleton-like upper bound in Theorem 36.

Remark 42. For simplicity and consistency with the rest of the paper, we have chosen to consider random CSS qLRCs. By not restricting to CSS codes and instead taking random stabilizer qLRCs, one can obtain a slightly better dependence on alphabet size. That is, a smaller alphabet size q would be needed to achieve any given rate-distance tradeoff dictated by Lemma 38 and Proposition 40. However, in the limit of large q , which is our primary concern, the rate-distance tradeoff is the same whether or not we restrict to CSS codes. A similar phenomenon occurs with ordinary random quantum codes (with no locality constraints), and is the reason that random stabilizer codes (i.e. the quantum GV bound) have better distance than random CSS codes over small alphabets.

4.2 Explicit qLRCs from AEL Distance Amplification

In this section, we show how to construct explicit qLRCs using Alon-Edmunds-Luby (AEL) distance amplification and alphabet reduction [AEL95]. This distance amplification technique has seen extensive use in classical coding theory [GI01, GI02, GI03, GR08, HW18, KMRZS16, GKO⁺18, HRZW20], but has only recently been studied in the quantum setting [BGG22, WLH23]. Here we apply the AEL technique with an inner code given by a random qLRC from Section 4.1.

4.2.1 Review of AEL Technique

Here review the AEL distance amplification and alphabet reduction technique of [AEL95]. We will subsequently specify how we parameterize it to obtain qLRCs. The construction will use the following notion of an expander.

Definition 43. A Δ -regular bipartite graph $G = (L, R, E)$ with $|L| = |R| = n$ is λ -**pseudorandom** if it holds for every $S \subseteq L$ and $T \subseteq R$ that

$$\left| |E(S, T)| - \frac{\Delta|S||T|}{n} \right| \leq \lambda\Delta\sqrt{|S||T|}.$$

Recall that a λ -*spectral expander* is a graph for which the second largest absolute value of an eigenvalue of the random walk matrix is $\leq \lambda$. By the well-known expander mixing lemma (see for instance Lemma 4.15 in [Vad12]), the double cover of a λ -spectral expander is λ -pseudorandom, so the explicit near-Ramanujan graphs of [MOP21] prove the following.

Proposition 44 ([MOP21]). *For every integer $\Delta \geq 3$, there exists an infinite explicit family of Δ -regular λ -pseudorandom graphs for $\lambda = 2/\sqrt{\Delta}$.*

We now describe the AEL construction. Let \mathcal{C}_{out} and \mathcal{C}_{in} be $[[n_{\text{out}}, k_{\text{out}} = R_{\text{out}}n_{\text{out}}, d_{\text{out}} = \delta_{\text{out}}n_{\text{out}}]]_{q_{\text{out}}}$ and $[[n_{\text{in}}, k_{\text{in}} = R_{\text{in}}n_{\text{in}}, d_{\text{in}} = \delta_{\text{in}}n_{\text{in}}]]_{q_{\text{in}}}$ quantum codes respectively such that $q_{\text{out}} = q_{\text{in}}^{k_{\text{in}}}$. Let \mathcal{C}_{\diamond} denote the $[[n_{\text{out}}n_{\text{in}}, k_{\text{out}}, k_{\text{in}}]]_{q_{\text{in}}}$ concatenated code, so that its encoder

$$\text{Enc}_{\diamond} = \text{Enc}_{\text{in}}^{\otimes n_{\text{out}}} \circ \text{Enc}_{\text{out}}$$

first applies the encoder for \mathcal{C}_{out} , and then applies the encoder for \mathcal{C}_{in} to each symbol of the resulting codeword of \mathcal{C}_{out} , where we use the assumption that $q_{\text{out}} = q_{\text{in}}^{k_{\text{in}}}$ to view each such symbol as a message for \mathcal{C}_{in} .

Now for some $\Delta|n_{\text{in}}$, we partition each inner code block of size n_{in} into n_{in}/Δ subsets each of size Δ , and we fold together the Δ components in each subset to obtain a folded concatenated code $\tilde{\mathcal{C}}_\diamond$, which is a $[[n = n_{\text{out}}n_{\text{in}}/\Delta, k = k_{\text{out}}k_{\text{in}}/\Delta]]_q$ code of alphabet size $q = q_{\text{in}}^\Delta$.

To construct our final $[[n, k]]_q$ code \mathcal{C} , we permute the symbols of $\tilde{\mathcal{C}}_\diamond$ according to a Δ -regular λ -pseudorandom bipartite graph $G = (L, R, E)$ with $|L| = |R| = \tilde{n}_\diamond$. Formally, let $\pi_G : [n] \times [\Delta] \rightarrow [n] \times [\Delta]$ be the permutation that maps $(i, j) \in [n] \times [\Delta]$ to the unique $(i', j') \in [n] \times [\Delta]$ such that the j th edge incident to vertex $i \in L$ equals the j' th edge incident to vertex $i' \in R$. Note that here we have fixed some ordering of the edges incident to each vertex. We then let \mathcal{C} have encoder

$$\text{Enc} = \pi_G \circ \text{Enc}_\diamond$$

that simply encodes the message into \mathcal{C}_\diamond , and then permutes the symbols according to π_G . That is, for $j \in [\Delta]$, the j th symbol of component $i \in [n]$ in $\tilde{\mathcal{C}}_\diamond$ is mapped to the j' th symbol of component i' in \mathcal{C} for $(i', j') = \pi_G(i, j)$. Thus for the composition here we view π_G as an invertible map $\mathbb{F}_q^{n \times \Delta} \rightarrow \mathbb{F}_q^{n \times \Delta}$ in the natural way.

The following lemma is immediate from the definition of the code \mathcal{C}

Lemma 45. *\mathcal{C} has rate $R = R_{\text{out}}R_{\text{in}}$.*

The distance of \mathcal{C} is given by the following proposition. The statement and proof of this AEL distance bound follows from [AEL95], and was previously shown in the quantum setting by [BGG22]; we repeat the proof in Appendix B for completeness in our specific setting. Below, we use the definitions and notation from the AEL construction above. We also say a code has *decoding radius* a if it has a (possibly inefficient) decoding algorithm that corrects against errors in up to a unknown locations.

Proposition 46. *If \mathcal{C}_{out} and \mathcal{C}_{in} have respective decoding radii $\alpha_{\text{out}}n_{\text{out}}$ and $\alpha_{\text{in}}n_{\text{in}}$, then \mathcal{C} has decoding radius at least αn for $\alpha = \alpha_{\text{in}} - \lambda\sqrt{\alpha_{\text{in}}/\alpha_{\text{out}}}$. Furthermore, if the decoding algorithms for \mathcal{C}_{out} and \mathcal{C}_{in} run in time $\text{poly}(n \log q)$, so does the decoding algorithm for \mathcal{C} .*

Because a code of distance d may be (possibly inefficiently) decoded from e errors for every $e < d/2$, and a code with decoding radius e has distance $> 2e$, the following corollary is immediate.

Corollary 47. *If \mathcal{C}_{out} and \mathcal{C}_{in} have respective distances $\delta_{\text{out}}n_{\text{out}}$ and $\delta_{\text{in}}n_{\text{in}}$, then \mathcal{C} has distance at least δn for $\delta = \delta_{\text{in}} - 2\lambda\sqrt{\delta_{\text{in}}/\delta_{\text{out}}}$.*

4.2.2 Application to qLRCs

We now apply the AEL described above to qLRCs. In this section we again continue to use the definitions and notation for the AEL construction from Section 4.2.1.

Lemma 48. *Let the inner code \mathcal{C}_{in} be a qLRC of locality $r_{\text{in}}|\Delta$, whose recovery sets $\{I_i : i \in [n_{\text{in}}]\}$ all have size r_{in} and form a partition of $[n_{\text{in}}]$. The AEL construction then yields a qLRC \mathcal{C} of locality Δr_{in} .*

Proof. We may perform the folding step of the AEL construction described in Section 4.2.1 such that all elements of each I_i are folded into the same length- Δ component in $\tilde{\mathcal{C}}_\diamond$. The permutation π_G

then sends these r elements of I_i to distinct components in \mathcal{C} , assuming we choose a λ -pseudorandom graph G arising from a 2-cover of a λ -spectral expander with no self-loops; the expanders of [MOP21] described in Proposition 44 indeed have no self-loops. Thus each of the Δ symbols of a component of \mathcal{C} can be recovered from r_{in} symbols in other components, so \mathcal{C} is a qLRC of locality $r_{\text{in}}\Delta$. \square

In summary, if we apply AEL with a qLRC \mathcal{C}_{in} of locality r_{in} , we obtain a qLRC of locality $r = \Delta r_{\text{in}}$ that has rate $R = R_{\text{out}}R_{\text{in}}$ and relative distance at least $\delta_{\text{in}} - 2\lambda\sqrt{\delta_{\text{in}}/\delta_{\text{out}}}$, where we can take $\lambda = 2/\sqrt{\Delta}$.

To optimize the parameters for the resulting qLRC \mathcal{C} , we can take \mathcal{C}_{out} to be a quantum Reed-Solomon code, which is the CSS code formed by two (generalized) Reed-Solomon codes². This code has alphabet size $q_{\text{out}} = O(n_{\text{out}})$ and lies on the quantum Singleton bound so that $R_{\text{out}} \geq 1 - 2\delta_{\text{out}}$. Then taking a random qLRC from Section 4.1 as the inner code \mathcal{C}_{in} , we obtain the following.

Proposition 49. *For every $r_{\text{in}}, \Delta \geq 3$ such that $r_{\text{in}}|\Delta$, and for every $\delta_{\text{in}}, \delta_{\text{out}}, \epsilon > 0$, there exists an infinite family of qLRCs with locality $r = r_{\text{in}}\Delta$, alphabet size $q = 2^{O(\Delta/\epsilon)}$, rate*

$$R = (1 - 2\delta_{\text{out}}) \left(1 - 2\delta_{\text{in}} - \frac{2}{r_{\text{in}}} - 2\epsilon \right),$$

and relative distance at least

$$\delta = \delta_{\text{in}} - 4\sqrt{\frac{\delta_{\text{in}}}{\delta_{\text{out}}\Delta}}.$$

Furthermore, a qLRC of block length n in this family can be constructed by a $\text{poly}(n \log q)$ -time randomized algorithm with high probability, and can be decoded by a $\text{poly}(n \log q)$ -time algorithm from up to $\delta n/2$ errors.

Proof. As described above, we instantiate the AEL construction with \mathcal{C}_{out} to be a (generalized) qRS code of relative distance at least δ_{out} , rate $R_{\text{out}} = 1 - 2\delta_{\text{out}}$, and alphabet size $q_{\text{out}} = O(n_{\text{out}})$, and we instantiate \mathcal{C}_{in} to be a random qLRC from Proposition 40 of relative distance at least δ_{in} , rate $R_{\text{in}} = 1 - 2\delta_{\text{in}} - 2/r_{\text{in}} - 2\epsilon$, locality r_{in} , and alphabet size $q_{\text{in}} = 2^{O(1/\epsilon)}$. We take the expander G to be given by Proposition 44. Taking \mathcal{C} to be the resulting code from applying the AEL construction, then the bounds on the parameters of \mathcal{C} in the proposition statement follow from Lemma 48, Proposition 40, Lemma 38, Lemma 45, and Corollary 47.

To efficiently construct \mathcal{C} , we must simply find an inner code \mathcal{C}_{in} satisfying the distance bound in Proposition 40, as the code \mathcal{C}_{out} and the expander G are by definition explicit. But because $q_{\text{out}} = O(n_{\text{out}})$, in $\text{poly}(n \log q)$ time we may repeatedly generate a random qLRC $\mathcal{C}_{\text{in}} = \text{CSS}(C_X, C_Z)$ as described in Section 4.1, and via brute force check every codeword in $C_X \setminus C_Z$ and $C_Z \setminus C_X$ to compute the distance of \mathcal{C}_{in} . By Proposition 40, with high probability a random such \mathcal{C}_{in} will have the desired distance δ_{in} , so we successfully construct \mathcal{C}_{in} and therefore \mathcal{C} with high probability. The

²If C is a length- $(q-1)$ classical RS code as defined in Definition 20, then indeed $\text{CSS}(C, C)$ is a well-defined CSS code of alphabet size q that lies on the quantum Singleton bound. This construction can be extended to general block lengths n (such that $n+1$ is not a prime power) by using *generalized* RS codes. In a generalized RS code, the messages are still low-degree polynomials, but the encoding map only evaluates the message polynomial at some points in \mathbb{F}_q^* , and multiplies these evaluations by some fixed coefficients in \mathbb{F}_q^* . By choosing appropriate such coefficients for generalized RS codes C_X and C_Z , we can ensure that the CSS orthogonality relations are satisfied, so that $\text{CSS}(C_X, C_Z)$ is a well-defined CSS code on the quantum Singleton bound. See for instance Section 3.1 of [BGG22] for more details.

efficiency of the decoding algorithm follows from Proposition 46 along with the fact that RS (and therefore also qRS) codes are known to have efficient decoders from errors of weight up to half the distance (e.g. Theorem 22). \square

Remark 50. The codes in Proposition 49 are efficiently constructable, but not technically explicit. This distinction arises because $q_{\text{out}} = \Theta(n_{\text{out}})$, so while the number of possible inner codewords $q_{\text{out}} = |\mathbb{F}_{q_{\text{in}}}^{k_{\text{in}}}|$ is polynomial in $n \log q$, the number of possible inner codes $\mathcal{C}_{\text{in}} = \text{CSS}(C_X, C_Z)$ is superpolynomial in $n \log q$. We can resolve this issue to obtain explicit qLRCs \mathcal{C} with AEL by using an outer code \mathcal{C}_{out} with a smaller alphabet size than a quantum RS code. For instance, the alphabet size of \mathcal{C}_{out} can be reduced with another application of AEL, at the cost of a small loss in parameters.

The qLRCs of locality r and relative distance δ in Proposition 49 have rate

$$R \leq (1 - 2\delta_{\text{out}}) \left(1 - 2\delta - 8\sqrt{\frac{\delta}{\delta_{\text{out}}\Delta}} - \frac{2\Delta}{r} - 2\epsilon \right)$$

for any $\epsilon > 0$ that determines the alphabet size $q = 2^{O(\Delta/\epsilon)}$. The RHS above is maximized when Δ is chosen to minimize the expression $8\sqrt{\delta/\delta_{\text{out}}\Delta} + 2\Delta/r$; solving this minimization gives $\Delta = (2r)^{2/3}(\delta/\delta_{\text{out}})^{1/3}$, which in turn gives

$$\begin{aligned} R &\leq (1 - 2\delta_{\text{out}}) \left(1 - 2\delta - 6 \left(\frac{4\delta}{\delta_{\text{out}}r} \right)^{1/3} - 2\epsilon \right) \\ &\leq 1 - 2\delta - \Omega \left(\frac{\delta}{r} \right)^{1/3}. \end{aligned}$$

Thus the qLRCs constructed with AEL in Proposition 49 have rate at least $\Omega(\delta/r)^{1/3}$ below the Singleton-like bound in Theorem 35.

Furthermore, if we let $0 < \delta < 1/2$ be a fixed constant and let r grow large, then the above inequality implies that

$$R \leq 1 - 2\delta - \Omega \left(\delta_{\text{out}} + \frac{1}{(\delta_{\text{out}}r)^{1/3}} \right).$$

The RHS above is maximized at $\delta_{\text{out}} = \Theta(1/r^{1/4})$, giving

$$R \leq 1 - 2\delta - \Omega \left(\frac{1}{r} \right)^{1/4}. \tag{9}$$

Thus for fixed δ as r grows large, the qLRCs in Proposition 49 have rate at least $\Omega(r)^{1/4}$ below the Singleton-like bound in Theorem 35.

The bound in (9) is tight, in the sense that if we fix $0 < \delta < 1/2$ in Proposition 49 and then take $\delta_{\text{out}} = 1/r^{1/4}$, $\Delta = r^{3/4}$, and $\epsilon = 1/r^{1/4}$, Proposition 49 gives qLRCs of relative distance δ and rate

$$R \geq 1 - 2\delta - O \left(\frac{1}{r} \right)^{1/4}$$

over an alphabet of size $q = 2^{O(r)}$.

5 Explicit Construction of qLRCs: Quantum Tamo-Barg Codes

In this section, we present our construction of qLRCs, and show that they are indeed locally recoverable. The construction is essentially a quantum CSS version of the classical construction of Tamo and Barg [TB14], though the quantum case is more involved due to issues surrounding orthogonality and degeneracy.

5.1 Classical Tamo-Barg Codes

We begin by describing the classical LRC construction of Tamo and Barg [TB14]. Recall below that for $n \in \mathbb{N}$, we denote $[n] = \{0, \dots, n-1\}$. We will also use the notation from Definition 19. In particular, recall that the map $\text{ev} : \mathbb{F}_q[X] \rightarrow \mathbb{F}_q^*$ maps a polynomial f to the list of its evaluations on all inputs in \mathbb{F}_q^* .

Definition 51 ([TB14]). For a prime power q , a locality parameter $r|(q-1)$ with $r \geq 3$, and an integer $\ell \in [q]$, the **Tamo-Barg (TB) code** is the polynomial evaluation code $\text{ev}(\mathbb{F}_q[X]^S)$ for

$$S = \{i \in [\ell] : i \not\equiv r-1 \pmod{r}\}.$$

The following proposition of [TB14] states that the TB code is a cLRC achieving the Singleton-like bound.

Proposition 52 ([TB14]). *The TB code in Definition 51 has alphabet size q , block length $q-1$, dimension $\ell - \lfloor \ell/r \rfloor$, distance $\geq q - \ell$, and is locally recoverable with locality r .*

Below we present a quantum analogue of TB codes that we show are also locally recoverable (Corollary 58). The proof of local recoverability in both the classical and quantum cases simply rely on some polynomial manipulations (Lemma 55 and Lemma 57 below).

5.2 Quantum Tamo-Barg Codes

We now introduce our quantum analogue of Tamo-Barg codes.

Definition 53. For a prime power q , a locality parameter $r|(q-1)$ with $r \geq 3$, and an integer $\ell \in [q]$, we define the **quantum Tamo-Barg (qTB) code** to be the CSS code $\mathcal{C} = \text{CSS}(C_X, C_Z)$ with $C_X = C_Z = \text{ev}(\mathbb{F}_q[X]^S)$ for

$$S = \{i \in [\ell] : i \not\equiv r-1 \pmod{r}\} \cup \{i \in [q-1] : i \equiv 1 \pmod{r}\}. \quad (10)$$

By construction \mathcal{C} has local dimension q and block length $q-1$.

To understand the definition of a qTB code, and to see that it gives a well defined CSS code, consider the following general fact, which presents a certain kind of “intersection” of CSS codes.

Lemma 54. *For CSS codes $\mathcal{A} = \text{CSS}(A_X, A_Z)$ and $\mathcal{B} = \text{CSS}(B_X, B_Z)$ of block length n over \mathbb{F}_q , there exists a CSS code $\mathcal{C} = \text{CSS}(C_X, C_Z)$ given by*

$$\begin{aligned} C_X &= (A_X \cap B_X) + B_Z^\perp \\ C_Z &= (A_Z \cap B_Z) + B_X^\perp, \end{aligned}$$

so that

$$\begin{aligned} C_X^\perp &= (A_X^\perp \cap B_Z) + B_X^\perp \\ C_Z^\perp &= (A_Z^\perp \cap B_X) + B_Z^\perp. \end{aligned}$$

Proof. By definition

$$\begin{aligned} C_X^\perp &= ((A_X \cap B_X) + B_Z^\perp)^\perp \\ &= (A_X \cap B_X)^\perp \cap B_Z \\ &= (A_X^\perp + B_X^\perp) \cap B_Z \\ &= (A_X^\perp \cap B_Z) + B_X^\perp, \end{aligned}$$

and by the same reasoning $C_Z^\perp = (A_Z^\perp \cap B_X) + B_Z^\perp$. Then because $A_X^\perp \subseteq A_Z$, it follows that $C_X^\perp \subseteq C_Z$, so \mathcal{C} is a well defined CSS code. \square

We then obtain the following.

Lemma 55. For $r|(q-1)$ with $r \geq 3$ and $\ell \in [q]$, let

$$\begin{aligned} A &= \text{ev}(\mathbb{F}_q[X]^{[\ell]}) \\ B &= \text{ev}(\mathbb{F}_q[X]^{[q-1] \setminus (-1+r\mathbb{Z})}). \end{aligned}$$

Then $A \cap B$ is a TB code. Furthermore,

$$\begin{aligned} A^\perp &= \text{ev}(\mathbb{F}_q[X]^{[q-\ell] \setminus \{0\}}) \\ B^\perp &= \text{ev}(\mathbb{F}_q[X]^{[q-1] \cap (1+r\mathbb{Z})}) \subseteq B. \end{aligned} \tag{11}$$

If $\ell \geq q/2$, then $A^\perp \subseteq A$, so letting $C = (A \cap B) + B^\perp$, then $\text{CSS}(C, C)$ is a q TB code with

$$C^\perp = (A^\perp \cap B) + B^\perp = \text{ev}(\mathbb{F}_q[X]^T) \subseteq C.$$

for

$$T = \{i \in [q-\ell] \setminus \{0\} : i \not\equiv r-1 \pmod{r}\} \cup \{i \in [q-1] : i \equiv 1 \pmod{r}\}. \tag{12}$$

Proof. It suffices to show the equalities in (11), as then the remaining claims in the lemma statement follow from the definitions of TB and q TB codes along with Lemma 54.

The first equality in (11) simply states the well known fact that the dual of a Reed-Solomon code is a (generalized) Reed-Solomon code. Formally, by dimension counting it suffices to show that $\text{ev}(\mathbb{F}_q[X]^{[q-\ell] \setminus \{0\}}) \subseteq A^\perp$. By using monomials as a basis for $\mathbb{F}_q[X]^{[\ell]}$ and for $\mathbb{F}_q[X]^{[q-\ell] \setminus \{0\}}$, it then suffices to show that for every $i \in [\ell]$ and $j \in [q-\ell] \setminus \{0\}$, we have $\text{ev}(X^i) \cdot \text{ev}(X^j) = 0$. But

$$\text{ev}(X^i) \cdot \text{ev}(X^j) = \sum_{\alpha \in \mathbb{F}_q^*} \alpha^{i+j}, \tag{13}$$

which indeed equals 0 for $i+j \not\equiv 0 \pmod{q-1}$.

We now prove the second equality in (11). By dimension counting, it suffices to show that $\text{ev}(\mathbb{F}_q[X]^{[q-1] \cap (1+r\mathbb{Z})}) \subseteq B^\perp$. By using monomials as a basis for $\mathbb{F}_q[X]^{[q-1] \setminus (-1+r\mathbb{Z})}$ and for $\mathbb{F}_q[X]^{[q-1] \cap (1+r\mathbb{Z})}$, it then suffices to show that for every $i \in [q-1] \setminus (-1+r\mathbb{Z})$ and $j \in [q-1] \cap (1+r\mathbb{Z})$, we have $\text{ev}(X^i) \cdot \text{ev}(X^j) = 0$. But then (13) again holds, and $i+j \not\equiv 0 \pmod{q-1}$, so $\text{ev}(X^i) \cdot \text{ev}(X^j) = 0$ as desired. \square

Lemma 55 implies that the qTB code $\mathcal{C} = \text{CSS}(C, C)$ satisfies $C^\perp \subseteq C$, so it is indeed a well defined CSS code. Below we compute the dimension of this code.

Lemma 56. *The qTB code $\mathcal{C} = \text{CSS}(C, C)$ with parameters q, r, ℓ has dimension*

$$\begin{aligned} k &= 1 + |\{q - \ell \leq i \leq \ell - 1 : i \not\equiv \pm 1 \pmod{r}\}| \\ &= (2\ell - q) \cdot \left(1 - \frac{2}{r}\right) + \epsilon \end{aligned}$$

for some $\epsilon \in [-2, 2]$.

Proof. Define $S \subseteq [q - 1]$ as in (10) and $T \subseteq S$ as in (12), so that $C = \mathbb{F}_q[X]^S$ and $C^\perp = \mathbb{F}_q[X]^T$. Then by definition \mathcal{C} has dimension

$$\begin{aligned} k &= \dim(C) - \dim(C^\perp) \\ &= |S \setminus T| \\ &= |\{0\} \cup \{q - \ell \leq i \leq \ell - 1 : i \not\equiv \pm 1 \pmod{r}\}|, \end{aligned}$$

which proves the first equality in the lemma statement. The second equality in the lemma statement then follows because $|\{q - \ell \leq i \leq \ell - 1 : i \not\equiv \pm 1 \pmod{r}\}|$ can differ from $(2\ell - q)(1 - 2/r)$ by at most 2. \square

We now show that the qTB code $\mathcal{C} = \text{CSS}(C, C)$ is locally recoverable. By Lemma 34, it suffices to show that C^\perp contains low-weight parity checks whose supports cover all n code components. As $B^\perp \subseteq C^\perp$ (see Lemma 55), it in fact suffices to show that B^\perp contains such low-weight parity checks, as is shown below.

Lemma 57. *The code $B^\perp = \text{ev}(\mathbb{F}_q[X]^{[q-1] \cap (1+r\mathbb{Z})})$ as defined in Lemma 55 may be equivalently characterized as follows. Let $\Omega_r = \{x \in \mathbb{F}_q^* : x^r = 1\} \subseteq \mathbb{F}_q^*$ denote the r th roots of unity. Then B^\perp consists of all functions $f : \mathbb{F}_q^* \rightarrow \mathbb{F}_q$ that can be expressed in the form $f(x) = \beta_{x\Omega_r} \cdot x$ for some $\beta \in \mathbb{F}_q^{*\mathbb{F}_q^*/\Omega_r}$.*

We refer to elements of B^\perp as **piecewise linear functions**, as Lemma 57 states that $f \in B^\perp$ if and only if for each coset $x\Omega_r$ of Ω_r , the function $f(x)$ agrees with the linear function $\beta_{x\Omega_r} \cdot x'$ on the restriction to inputs $x' \in x\Omega_r$. Note that these linear functions $x' \mapsto \beta_{x\Omega_r} \cdot x'$ are truly linear maps (not just affine linear).

Proof of Lemma 57. By dimension counting, it suffices to show that every $\text{ev}(f(X)) \in B^\perp = \text{ev}(\mathbb{F}_q[X]^{[q-1] \cap (1+r\mathbb{Z})})$ is piecewise linear, that is, that $f(x) = \beta_{x\Omega_r} \cdot x$ for some $\beta \in \mathbb{F}_q^{*\mathbb{F}_q^*/\Omega_r}$.

Fix some $f(X) = \sum_{i \in [q-1]} f_i X^i$ with $\text{ev}(f) \in B^\perp$, so that

$$f_i = 0 \quad \forall i \not\equiv 1 \pmod{r}. \tag{14}$$

For every coset $\alpha\Omega_r \in \mathbb{F}_q^*/\Omega_r$, then $X^r - \alpha^r$ vanishes on all $X \in \alpha\Omega_r$, so $g(X) := f(X) \pmod{X^r - \alpha^r}$ agrees with $f(X)$ on all $X \in \alpha\Omega_r$. But by definition $g(X) = \sum_{i \in [r]} g_i X^i$ with $g_i = \sum_{j \in [(q-1)/r]} f_{i+rj} \alpha^{rj}$, so (14) implies that $g_i = 0$ for all $i \neq 1$. Thus for all $X \in \alpha\Omega_r$ we have $f(X) = g(X) = g_1 X$, so f is indeed piecewise linear, as desired. \square

Lemma 57 immediately yields the result of [TB14] that TB codes are cLRCs (Proposition 52), and similarly shows that qTB codes are qLRCs:

Corollary 58. *The qTB code given in Definition 53 is locally recoverable with locality r .*

Proof. By Lemma 34, it suffices to show that the qTB code $\mathcal{C} = \text{CSS}(C, C)$ has C (classically) locally recoverable with locality r . Now by Lemma 55, C^\perp contains $B^\perp = \text{ev}(\mathbb{F}_q[X]^{[q-1] \cap (1+r\mathbb{Z})})$. For every $\alpha \in \mathbb{F}_q^*$, Lemma 57 implies that B^\perp contains the function

$$f_\alpha(x) = \begin{cases} x, & x \in \alpha\Omega_r \\ 0, & x \notin \alpha\Omega_r, \end{cases}$$

which has $\text{supp}(f) = \alpha\Omega_r$, that is, f_α is supported on r components including α . Thus Lemma 18 implies that C is a cLRC with locality r , so \mathcal{C} is a qLRC with locality r . \square

5.3 Folded Quantum Tamo-Barg Codes

We are not able to show that the qTB codes in Definition 53 have distance approaching the Singleton-like bound. Rather, we will only obtain such near-optimal distance for the *folded* version of these codes, defined below.

Definition 59. As in Definition 53, let $\mathcal{C} = \text{CSS}(C_X, C_Z)$ be the qTB code with parameters q, r, ℓ , so that $C_X = C_Z = \text{ev}(\mathbb{F}_q[X]^S)$ for $S \subseteq [q-1]$ given by (10). Given an additional folding parameter $s|(q-1)/r$, we define the **folded quantum Tamo-Barg (fqTB) code** $\tilde{\mathcal{C}}$ to be the quantum code of local dimension q^s and block length $(q-1)/s$ obtained as follows. Fix a generator ω_{q-1} for \mathbb{F}_q^* , and then for every $i \in [(q-1)/s]$, we block together the s components (each of local dimension q) at positions $\{\omega_{q-1}^{si}, \omega_{q-1}^{si+1}, \dots, \omega_{q-1}^{si+s-1}\}$ in \mathcal{C} into a single component (of local dimension q^s) of the folded code $\tilde{\mathcal{C}}$. Let \tilde{C}_X, \tilde{C}_Z denote the \mathbb{F}_q -linear codes obtained by similarly folding C_X, C_Z , so that $\tilde{\mathcal{C}} = \text{CSS}(\tilde{C}_X, \tilde{C}_Z)$.

As a point of notation, we will almost universally use $\tilde{\alpha}$ to denote the s -folded version of an object α , where s will be clear from context. Specifically, we denote $\tilde{\mathbb{F}}_q^* = \{\{\omega_{q-1}^{si}, \dots, \omega_{q-1}^{si+s-1}\} : i \in [(q-1)/s]\}$ to be the partition of \mathbb{F}_q^* into the $(q-1)/s$ blocks of size s into which we folded the qTB code in Definition 59. For a function $a : \mathbb{F}_q^* \rightarrow \mathbb{F}_q$, we denote $\tilde{a} : \tilde{\mathbb{F}}_q^* \rightarrow \mathbb{F}_q^s$ the associated folded function given by $\tilde{a}(S) = (a(x))_{x \in S}$ for $S \in \tilde{\mathbb{F}}_q^*$. For instance, if a is a codeword of the classical code C associated to a qTB code $\mathcal{C} = \text{CSS}(C, C)$, then \tilde{a} is a codeword of the folded code $\tilde{\mathcal{C}}$ associated to the fqTB code $\tilde{\mathcal{C}} = \text{CSS}(\tilde{C}, \tilde{C})$. We also extend this notation to subspaces $V \subseteq \mathbb{F}_q^{\mathbb{F}_q^*}$, so that $\tilde{V} = \{\tilde{a} : a \in V\}$; thus indeed \tilde{C} refers to the folded version of a code C , meaning that \tilde{C} is obtained by replacing all codewords $a \in C$ with their folded version \tilde{a} . We furthermore define $\tilde{\text{ev}} : \mathbb{F}_q[X] \rightarrow (\mathbb{F}_q^s)^{\tilde{\mathbb{F}}_q^*}$ by $\tilde{\text{ev}}(f) = \widetilde{\text{ev}(f)}$, meaning that $\tilde{\text{ev}}$ for applies the evaluation map ev , and then folds the output.

For a given $x = \omega_{q-1}^b \in \mathbb{F}_q^*$, we let $F_x \in \tilde{\mathbb{F}}_q^*$ denote the unique element of $\tilde{\mathbb{F}}_q^*$ that contains x , so that

$$F_x = \{\omega_{q-1}^{\lfloor b/s \rfloor \cdot s + j} : j \in [s]\}.$$

For a polynomial $f(X) \in \mathbb{F}_q[X]$, we will also let $\tilde{f}(F_x) = \tilde{ev}(f)(F_x)$. Therefore in particular if $\tilde{ev}(f)$ is a folded codeword, then $\tilde{f}(F_x) \in \mathbb{F}_q^s$ denotes the component of this folded codeword associated to the partition element in $\tilde{\mathbb{F}}_q^*$ that contains $x \in \mathbb{F}_q^*$.

Folding a code by definition preserves the rate, so a folded qTB code has the same rate as the associated unfolded qTB code, which we computed in Lemma 56. Below, we use our proof in Lemma 57 and Corollary 58 that qTB codes are qLRCs to show that folded qTB codes are also qLRCs.

Corollary 60. *The fqTB code give in Definition 59 is locally recoverable with locality r .*

Proof. Let $\mathcal{C} = \text{CSS}(C, C)$ denote the qTB code, and let $\tilde{\mathcal{C}} = \text{CSS}(\tilde{C}, \tilde{C})$ denote the associated folded qTB code. Our proof of Lemma 57 and Corollary 58 implies that for every $x \in \mathbb{F}_q^*$, component x of \mathcal{C} can be recovered from the components at positions in $x(\Omega_r \setminus \{1\})$.

Now for every $F_x \in \tilde{\mathbb{F}}_q^*$, component F_x of $\tilde{\mathcal{C}}$ consists of the components x' of \mathcal{C} for all $x' \in F_x$. As stated above, each such component x' of \mathcal{C} can be recovered from the components in $x'(\Omega_r \setminus \{1\})$ of \mathcal{C} . But by definition the union of all these recovery sets $x'(\Omega_r \setminus \{1\})$ for $x' \in F_x$ equals the union of the $r - 1$ partition elements $F_{x\omega_r^i}$ for $\omega_r^i \in \Omega_r \setminus \{1\}$, that is,

$$\bigsqcup_{x' \in F_x} x'(\Omega_r \setminus \{1\}) = F_x(\Omega_r \setminus \{1\}).$$

Therefore component F_x of $\tilde{\mathcal{C}}$ can be recovered from these $r - 1$ components $F_{x\omega_r^i}$ for $\omega_r^i \in \Omega_r \setminus \{1\}$ by simply separately recovering the symbol at each position $x' \in F_x \subseteq \mathbb{F}_q^*$ from the symbols at positions in $x'(\Omega_r \setminus \{1\})$ using the recovery maps for the unfolded qTB code \mathcal{C} . Therefore $\tilde{\mathcal{C}}$ is a qLRC of locality r , as desired. \square

6 Bounding the Distance

In this section, we construct explicit qLRCs approaching the Singleton-like bound. We first bound the distance of the qTB codes introduced in Definition 53 in Section 5. Interestingly, we are only able to get a distance bound for these codes that is reminiscent of the Johnson bound, so that this bound in particular only approaches the Singleton bound for codes of small rate. However, we then show that the folded version of these qTB codes introduced in Defintion 59 have distance approaching the Singleton bound for all rates, for sufficiently large locality parameters r . It remains an open question whether the unfolded qTB codes have distance approaching the Singleton bound.

Note that classical Tamo-Barg codes are subcodes of Reed-Solomon codes, so their distance bound immediately follows from the distance of RS codes. That is, the distance of a classical TB code with parameters q, r, ℓ is at least the distance $d = q - \ell$ of a classical RS code with parameters q, ℓ . However, our quantum TB codes are not subcodes of RS codes due to the orthogonality condition for CSS codes which necessitates extra high degree terms in the polynomials. As a result our distance analysis in the quantum case is much more involved and delicate. Indeed, the following corollary of Theorem 36 shows that the distance of a qTB code with parameters q, r, ℓ must be strictly worse asymptotically than the distance $d = q - \ell$ of a qRS code with parameters q, ℓ .

Corollary 61. *The distance d of the qTB code $\mathcal{C} = \text{CSS}(C, C)$ with parameters q, r, ℓ satisfies*

$$d \leq \left(1 - \frac{1}{r}\right)(q - \ell) + 5.$$

Proof. Plugging the expression for the dimension k of \mathcal{C} from Lemma 56 into the bound in Theorem 36 gives that

$$(2\ell - q) \left(1 - \frac{2}{r}\right) - 2 \leq \left(1 - \frac{2}{r}\right)(q - 1) - 2 \left(1 - \frac{1}{r-1}\right)(d - 1) + 2.$$

Rearranging terms in the above expression then yields the desired result. \square

Similar reasoning as in Corollary 61 also implies that the folded qTB code with parameters q, r, ℓ, s has asymptotically worse distance than the folded qRS code with parameters q, ℓ, s , as the folded qRS code lies on the quantum Singleton bound like its unfolded counterpart.

6.1 Distance of Unfolded Quantum Tamo-Barg Codes

We now show the following bound on the distance of the qTB codes from Definition 53. Recall that the following result was stated informally as Theorem 3 in Section 1.1.2.

Theorem 62. *The qTB code \mathcal{C} in Definition 53 with a prime locality parameter r has distance at least*

$$d = (q - 1) \left(1 - \frac{1}{2r} - \sqrt{\frac{1}{4r^2} + \frac{r-1}{r} \cdot \frac{\ell-1}{q-1}}\right). \quad (15)$$

The reader is referred to Section 1.2.2 for more context surrounding the above bound and its apparent similarity to the Johnson bound.

Proof of Theorem 62. Let $\mathcal{C} = \text{CSS}(C, C)$ be a qTB code with parameters q, r, ℓ . Fix an arbitrary $\text{ev}(f) \in C \setminus C^\perp$. Our goal is to show that $|\text{ev}(f)| \geq d$.

For a high-level proof overview, the reader is referred to the proof sketch in Section 1.2.2. In brief, we will define a polynomial G that has many roots associated to the roots of f , but also has low degree. We then obtain the desired result by comparing bounds on the number of roots and the degree of G .

We now give the formal proof. As $C = \text{ev}(\mathbb{F}_q[X]^S)$ for $S = ([\ell] \setminus (-1 + r\mathbb{Z})) \cup ([q-1] \cap (1 + r\mathbb{Z}))$, we may write $f(X) = g(X) + h(X)$, where $g(X) \in \mathbb{F}_q[X]^{[\ell] \setminus (\pm 1 + r\mathbb{Z})}$ and $h(X) \in \mathbb{F}_q[X]^{[q-1] \cap (1 + r\mathbb{Z})}$. Then $h(X)$ is piecewise linear by Lemma 57. As $\text{ev}(h) \in C^\perp$ and by assumption $\text{ev}(f) \notin C^\perp$, we must have $g \neq 0$.

Define a polynomial $G(X) \in \mathbb{F}_q[X]$ by

$$G(X) = \prod_{i=1}^{r-1} (\omega_r^{-i} g(\omega_r^i X) - g(X)). \quad (16)$$

Letting $g(X) = \sum_j g_j X^j$, then by definition the j th coefficient of $\omega_r^{-i} g(\omega_r^i X) - g(X)$ equals $(\omega_r^{i(j-1)} - 1)g_j$. Because $g_j = 0$ for all $j \equiv 1 \pmod{r}$, and r is prime so that $\omega_r^{i(j-1)} \neq 1$ for all $j \not\equiv 1 \pmod{r}$,

it follows that $(\omega_r^{i(j-1)} - 1)g_j = 0$ iff $g_j = 0$. Therefore in particular $\deg(\omega_r^{-i}g(\omega_r^i X) - g(X)) = \deg(g) \leq \ell - 1$, so $\deg(G) = (r - 1)\deg(g) \leq (r - 1)(\ell - 1)$. It also follows that $G \neq 0$.

We now bound the number of roots of G in terms of the number of roots of f . If $f(x) = 0$ and $f(\omega_r^i x) = 0$ for some $x \in \mathbb{F}_q^*$ and some $1 \leq i \leq r - 1$, then

$$\omega_r^{-i}g(\omega_r^i x) - g(x) = -\omega_r^{-i}h(\omega_r^i x) + h(x) = -h(x) + h(x) = 0,$$

where the second equality above holds because h is piecewise linear.

Therefore letting $\Omega_r = \{1, \omega_r, \dots, \omega_r^{r-1}\}$ denote the r th roots of unity in \mathbb{F}_q^* , then for every ordered pair of distinct elements $(x, y = \omega_r^i x)$ within some coset $x\Omega_r$ such that $f(x) = f(y) = 0$, the polynomial $G(X)$ has an associated root at $X = x$. Therefore $G(X)$ has $(r - |\text{ev}(f)|_{x\Omega_r})(r - |\text{ev}(f)|_{x\Omega_r} - 1)$ roots within a given coset $x\Omega_r$, where $|\text{ev}(f)|_{x\Omega_r}$ denotes the Hamming weight of the restriction of $\text{ev}(f)$ to $x\Omega_r$. Thus the total number of roots of G is at least

$$\begin{aligned} & \sum_{x\Omega_r \in \mathbb{F}_q^*/\Omega_r} (r - |\text{ev}(f)|_{x\Omega_r})(r - |\text{ev}(f)|_{x\Omega_r} - 1) \\ &= \frac{q-1}{r} \mathbb{E}_{x\Omega_r \sim \text{Unif}(\mathbb{F}_q^*/\Omega_r)} [(r - |\text{ev}(f)|_{x\Omega_r})^2] - (q-1 - |\text{ev}(f)|) \\ &\geq \frac{q-1}{r} \mathbb{E}_{x\Omega_r \sim \text{Unif}(\mathbb{F}_q^*/\Omega_r)} [(r - |\text{ev}(f)|_{x\Omega_r})]^2 - (q-1 - |\text{ev}(f)|) \\ &= \frac{r}{q-1} (q-1 - |\text{ev}(f)|)^2 - (q-1 - |\text{ev}(f)|). \end{aligned}$$

But we showed above that G is a nonzero polynomial of degree $\leq (r-1)(\ell-1)$, so we must have

$$\frac{r}{q-1} (q-1 - |\text{ev}(f)|)^2 - (q-1 - |\text{ev}(f)|) \leq (r-1)(\ell-1).$$

Solving this quadratic equation and rearranging terms gives that $|\text{ev}(f)| \geq d$ for d given in (15), as desired. \square

6.2 Distance of Folded Quantum Tamo-Barg Codes

In this section, we show the following bound on the distance of the fqTB codes from Definition 59.

Theorem 63. *Let \tilde{C} be the fqTB code in Definition 59 with parameters q, r, ℓ, s such that r is prime and such that the uncertainty principle in Proposition 66 holds for r over \mathbb{F}_q . Then \tilde{C} has distance at least*

$$d = \frac{q-1}{s} \left(1 - \frac{\ell-1}{q-1} - \epsilon \right). \quad (17)$$

for

$$\epsilon = \max_{1 \leq m \leq r} \min \left\{ \left(1 - \frac{\ell-1}{q-1} \right) \frac{m-1}{r}, \frac{1}{m} + \frac{m-1}{s} \right\}. \quad (18)$$

Applying the technical Lemma 75 proved in Appendix A to bound the expression ϵ in (18) yields us the following corollary.

Corollary 64. *Let $\tilde{\mathcal{C}}$ be the fqTB code in Definition 59 with parameters q, r, ℓ, s such that r is prime and such that the uncertainty principle in Proposition 66 holds for r over \mathbb{F}_q . If $s \geq 2r^2$, then $\tilde{\mathcal{C}}$ has distance at least*

$$d \geq \frac{q-1}{s} \left(1 - \frac{\ell-1}{q-1} - \left(1 + \frac{r^2}{s} \right) \sqrt{\frac{1}{r} \left(1 - \frac{\ell-1}{q-1} \right)} \right).$$

Recall that an informal statement of Corollary 64 was given in Theorem 2 in Section 1.1.2, which intuitively states the following. Because folding a code preserves its rate, Lemma 56 implies that as $q, r, \ell, s \rightarrow \infty$, Corollary 64 provides a bound on the relative distance of a fqTB code of rate R that approaches

$$\frac{d}{(q-1)/s} \rightarrow \frac{1-R}{2}.$$

This bound indeed approaches the ordinary Singleton bound (Proposition 27) for general quantum codes.

There is one subtlety with our bound from Theorem 63 and Corollary 64, namely, that we require the uncertainty principle in Proposition 66 to hold for r over \mathbb{F}_q . Proposition 66 shows that for every fixed prime r , this uncertainty principle can only fail to hold for finite fields whose characteristic lies in a finite set depending on r . Thus we must first fix the locality parameter r , and then let the block length grow arbitrarily large. Specifically, by letting q be an arbitrary power of an arbitrary prime outside of this bad finite set depending on r , we obtain our desired infinite family of qLRCs satisfying the distance bounds described above.

To prove Theorem 63, we will use the following well known fact about determinants. For completeness, a proof is provided in Appendix A.

Lemma 65 (Well known). *For every $m \in \mathbb{N}$, the determinant polynomial $\det \in \mathbb{F}_q[(X_{ij})_{i,j \in [m]}]$ has a root of multiplicity $m - t$ at every matrix $(x_{ij})_{i,j \in [m]}$ of rank t .*

We will also use the following uncertainty principle over finite fields. Below we use the following notation for a finite field \mathbb{F}_q and a prime $r|(q-1)$. For a polynomial $f(X) = \sum_{i \in [q]} f_i X^i \in \mathbb{F}_q[X]$, we let $|f|$ denote the number of $i \in [q]$ for which the coefficient f_i is nonzero, while we let $|\text{ev}(f)|_{\Omega_r}$ denote the number of distinct r th roots of unity $\omega_r^i \in \Omega_r \subseteq \mathbb{F}_q$ for which $f(\omega_r^i)$ is nonzero.

Proposition 66 ([GGI05]). *For every fixed prime r , the following holds for all but finitely many primes p : if \mathbb{F}_q is a field of characteristic p with $r|(q-1)$, then every nonzero $f \in \mathbb{F}_q[X]$ of degree $< r$ satisfies*

$$|f| + |\text{ev}(f)|_{\Omega_r} \geq r + 1.$$

Proposition 66 can be viewed as an uncertainty principle, as it says that if f has few nonzero coefficients, then it has many nonzero evaluation points on r th roots of unity.

This uncertainty principle is implied by the statement that all minors of the Vandermonde matrix $(\omega_r^{ij})_{i,j \in [r]}$ are nonzero (a proof of this implication is given by [Tao04]). Over the complex field, this Vandermonde matrix also has no vanishing minors, as was first proved by Chebotarëv; additional proofs have since been given, e.g. [EI76, Tao04, Fre04, GGI05]. The discussion in Section 6 of [GGI05] shows that for every fixed r , Chebotarëv's theorem over \mathbb{C} implies the same result over finite fields of all but finitely many characteristics. Thus Proposition 66 holds. We remark that a similar result is also shown in [Zha19].

Proof of Theorem 63. Let $\mathcal{C} = \text{CSS}(C, C)$ be a qTB code with parameters q, r, ℓ as in the theorem statement, and let $\tilde{\mathcal{C}} = \text{CSS}(\tilde{C}, \tilde{C})$ be the folding of \mathcal{C} for folding parameter $s|(q-1)/r$. Fix an arbitrary $f(X) = \sum_{i \in [q-1]} f_i X^i \in \mathbb{F}_q[X]$ such that $\text{ev}(f) \in C \setminus C^\perp$, with associated folded codeword $\tilde{\text{ev}}(f) \in \tilde{C} \setminus \tilde{C}^\perp$. Our goal is to show that $|\tilde{\text{ev}}(f)| \geq d$.

Let $M = \{i \in [r] : \exists j \equiv i \pmod{r} \text{ with } f_j \neq 0\}$ and $m = |M|$. In the two claims below, we show two upper bounds on $|\tilde{\text{ev}}(f)| \geq d$; the first bound is tighter when m is small, while the second is tighter when m is large.

Claim 67.

$$|\tilde{\text{ev}}(f)| \geq \frac{q-1}{s} \left(1 - \frac{\ell-1}{q-1}\right) \left(1 - \frac{m-1}{r}\right).$$

Proof. We will show this bound on $|\tilde{\text{ev}}(f)|$ by first bounding $|\text{ev}(f)|$, and then applying the fact that $|\tilde{\text{ev}}(f)| \geq |\text{ev}(f)|/s$. To bound $|\text{ev}(f)|$, we will use the uncertainty principle in Proposition 66. Recall that for every $x \in \mathbb{F}_q^*$, then on the restriction to inputs in the coset $x\Omega_r$, f agrees with $f \pmod{X^r - x^r}$, as $X^r - x^r$ vanishes on all $X \in x\Omega_r$. Now by definition $f \pmod{X^r - x^r}$ is a polynomial of degree $< r$ whose coefficients are supported within M . Therefore $|f \pmod{X^r - x^r}| \leq m$, so by Proposition 66, we have that either $\text{ev}(f)|_{x\Omega_r} = 0$ or $|\text{ev}(f)|_{x\Omega_r}| \geq r + 1 - m$.

We now bound the number of cosets $x\Omega_r \in \mathbb{F}_q^*/\Omega_r$ for which $\text{ev}(f)|_{x\Omega_r} = 0$. If $M \subseteq \{1\}$ then $\text{ev}(f) \in C^\perp$ by Lemma 55, contradicting the assumption that $\text{ev}(f) \in C \setminus C^\perp$, so there must be some $i \in M \setminus \{1\}$. Then $\text{ev}(f)|_{x\Omega_r} = 0$ iff $f \pmod{X^r - x^r} = 0$, which in turn can only occur if the i th coefficient of $f \pmod{X^r - x^r}$ equals 0. But this i th coefficient is precisely $\sum_{j \in [(q-1)/r]} f_{i+rj} (x^r)^j$. Now the polynomial $\sum_{j \in [(q-1)/r]} f_{i+rj} Y^j$ has degree $\leq (\ell-1)/r$ because $f_{i+rj} = 0$ if $i+rj \geq \ell$ by the definition of C . Thus this polynomial has $\leq (\ell-1)/r$ roots, so there are $\leq (\ell-1)/r$ values of $Y = x^r$ for which $\sum_{j \in [(q-1)/r]} f_{i+rj} (x^r)^j = 0$, and thus there are $\leq (\ell-1)/r$ cosets $x\Omega_r \in \mathbb{F}_q^*/\Omega_r$ for which $\text{ev}(f)|_{x\Omega_r} = 0$.

Because we showed above that $|\text{ev}(f)|_{x\Omega_r}| \geq r + 1 - m$ whenever $\text{ev}(f)|_{x\Omega_r} \neq 0$, we may sum over all cosets $x\Omega_r$ to conclude that

$$\begin{aligned} |\text{ev}(f)| &= \sum_{x\Omega_r \in \mathbb{F}_q^*/\Omega_r} |\text{ev}(f)|_{x\Omega_r}| \\ &\geq \left(\frac{q-1}{r} - \frac{\ell-1}{r}\right) (r + 1 - m) \\ &= (q-1) \left(1 - \frac{\ell-1}{q-1}\right) \left(1 - \frac{m-1}{r}\right). \end{aligned}$$

Then by definition

$$|\tilde{\text{ev}}(f)| \geq \frac{|\text{ev}(f)|}{s} \geq \frac{q-1}{s} \left(1 - \frac{\ell-1}{q-1}\right) \left(1 - \frac{m-1}{r}\right). \quad \square$$

Claim 68.

$$|\tilde{\text{ev}}(f)| \geq \frac{q-1}{s} \left(1 - \frac{\ell-1}{q-1} - \frac{1}{m} - \frac{m-1}{s}\right).$$

Proof. We show this bound on $|\tilde{\text{ev}}(f)|$ using a similar method as in the proof of Theorem 62, except we replace the polynomial $G(X)$ in (16) with (19) below. At a high level, we obtain our improved

distance bound here by leveraging the code folding to more efficiently detect roots of f using a certain determinant polynomial.

Our use of the determinant polynomial to detect roots of f is similar in spirit to analysis in [GK16] of subspace designs based on folded RS codes. However, the details are different, and we require a more involved argument to show that the polynomial G we obtain using the determinant is nonzero.

As $C = \text{ev}(\mathbb{F}_q[X]^S)$ for $S = ([\ell] \setminus (-1 + r\mathbb{Z})) \cup ([q-1] \cap (1 + r\mathbb{Z}))$, we have a decomposition

$$f(X) = g(X) + h(X)$$

for $g(X) = \sum_i g_i X^i \in \mathbb{F}_q[X]^{[\ell] \setminus (\pm 1 + r\mathbb{Z})}$ and $h(X) = \sum_i h_i X^i \in \mathbb{F}_q[X]^{[q-1] \cap (1 + r\mathbb{Z})}$. Because $\text{ev}(f) \notin C^\perp$ and $\text{ev}(h) \in C^\perp$ by assumption, we must have $g \neq 0$.

Let $\det \in \mathbb{F}_q[(X_{ij})_{i,j \in [m]}]$ denote the determinant polynomial, which takes as input an $m \times m$ matrix of variables $(X_{ij})_{i,j \in [m]}$ over \mathbb{F}_q , and outputs the determinant of the input matrix. Define $G(X) \in \mathbb{F}_q[X]$ by

$$G(X) = \det\left(\omega_r^{-i} \cdot g(\omega_r^i \omega_{q-1}^j \cdot X)\right)_{i,j \in [m]}. \quad (19)$$

As $\deg(\det) = m$, we have $\deg(G) \leq m \cdot \deg(g) \leq m(\ell - 1)$.

We next show that $G(X)$ is a nonzero polynomial. As $g(X) = \sum_{a \in [\ell] \setminus (\pm 1 + r\mathbb{Z})} g_a X^a$, we may decompose the matrix

$$\begin{aligned} A(X) &:= \left(\omega_r^{-i} \cdot g(\omega_r^i \omega_{q-1}^j \cdot X)\right)_{i,j \in [m]} \\ &= \sum_a g_a X^a \left(\omega_r^{(a-1)i} \omega_{q-1}^{aj}\right)_{i,j \in [m]} \\ &= \sum_a g_a X^a \begin{pmatrix} 1 \\ \omega_r^{a-1} \\ \omega_r^{2(a-1)} \\ \vdots \\ \omega_r^{(m-1)(a-1)} \end{pmatrix} \cdot \left(1 \quad \omega_{q-1}^a \quad \omega_{q-1}^{2a} \quad \cdots \quad \omega_{q-1}^{(m-1)a}\right) \\ &= \sum_{i \in M} \begin{pmatrix} 1 \\ \omega_r^{i-1} \\ \omega_r^{2(i-1)} \\ \vdots \\ \omega_r^{(m-1)(i-1)} \end{pmatrix} \cdot \sum_{a \equiv i \pmod{r}} g_a X^a \left(1 \quad \omega_{q-1}^a \quad \omega_{q-1}^{2a} \quad \cdots \quad \omega_{q-1}^{(m-1)a}\right). \end{aligned} \quad (20)$$

By definition $A(X)$ is an $m \times m$ matrix with entries in the ring $\mathbb{F}_q[X]$, and $G(X) = \det(A(X))$. To show that $G \neq 0$, it therefore suffices to show that A has full rank m , for which it in turn suffices to show that the sets of vectors

$$\left\{ \begin{pmatrix} 1 \\ \omega_r^{i-1} \\ \omega_r^{2(i-1)} \\ \vdots \\ \omega_r^{(m-1)(i-1)} \end{pmatrix} : i \in M \right\} \quad \text{and} \quad \left\{ \sum_{a \equiv i \pmod{r}} g_a X^a \begin{pmatrix} 1 \\ \omega_{q-1}^a \\ \omega_{q-1}^{2a} \\ \vdots \\ \omega_{q-1}^{(m-1)a} \end{pmatrix} : i \in M \right\} \quad (21)$$

from the rank-1 decomposition of A in (20) are linearly independent over $\mathbb{F}_q[X]$. Here it is equivalent to show linear independence over the field of rational functions $\mathbb{F}_q(X)$ and (as we may clear denominators in any linear dependency) over the ring of polynomials $\mathbb{F}_q[X]$, so we will show the latter.

Now the first set of vectors in (21) above form the columns of an $m \times m$ Vandermonde matrix, which is known to have full rank. Meanwhile, if there is some nontrivial $\mathbb{F}_q[X]$ -linear dependency among the second set of vectors in (21), then taking the highest-degree term of the associated polynomials over X gives a nontrivial \mathbb{F}_q -linear dependency among the vectors $(\omega_{q-1}^{aj})_{j \in [m]}$ for m distinct values of a . But these vectors again form the columns of an $m \times m$ Vandermonde matrix and thus cannot have any nontrivial linear dependencies. Therefore both sets of vectors in (21) are indeed linearly independent, so the decomposition in (20) expressing $A(X)$ as the sum of m rank-1 matrices implies that $A(X)$ has full rank m , and thus $G(X) = \det(A(X))$ is a nonzero polynomial.

We now bound the number of roots of $G(X)$ in terms of the Hamming weight of $\tilde{e}\tilde{v}(f)$. For a given $x = \omega_{q-1}^b \in \mathbb{F}_q^*$, recall that $F_x = \{\omega_{q-1}^{\lfloor b/s \rfloor \cdot s + j} : j \in [s]\} \in \tilde{\mathbb{F}}_q^*$ denotes the index of the folded component of $\tilde{\mathcal{C}}$ that contains the component x of \mathcal{C} .

If $b \pmod{s} \in \{0, 1, \dots, s-m\}$, then because $s|(q-1)/r$ by assumption, it holds for each $i \in [m]$ that

$$\{\omega_r^i \omega_{q-1}^j x : j \in [m]\} = \{\omega_{q-1}^{b+i(q-1)/r+j} x : j \in [m]\} \subseteq F_{\omega_r^i x}.$$

Thus the i th row of the matrix $(\omega_r^i \omega_{q-1}^j x)_{i,j \in [m]}$ consists of m elements all lying inside $F_{\omega_r^i x}$. Applying f to each component, it follows that that i th row of $(f(\omega_r^i \omega_{q-1}^j x))_{i,j \in [m]}$ consists of m out of the s components in the folded symbol $\tilde{e}\tilde{v}(f)_{F_{\omega_r^i x}}$.

Let $Z_x = \{i \in [m] : \tilde{e}\tilde{v}(f)_{F_{\omega_r^i x}} = 0\}$. If $b \pmod{s} \in \{0, 1, \dots, s-m\}$, then for every $i \in Z_x$, the i th row of $(f(\omega_r^i \omega_{q-1}^j x))_{i,j \in [m]}$ by definition consists of all zeros. In this case because $f(X) = g(X) + h(X)$, for every $i \in Z_x$ we have

$$(\omega_r^{-i} g(\omega_r^i \omega_{q-1}^j x))_{j \in [m]} = (-\omega_r^{-i} h(\omega_r^i \omega_{q-1}^j x))_{j \in [m]} = (-h(\omega_{q-1}^j x))_{j \in [m]},$$

where the second equality holds because $h(X)$ is by definition piecewise linear. Now the left hand side above equals the i th row of $A(x) \in \mathbb{F}_q^{m \times m}$, while the right hand side is a vector that does not depend on i . Thus we have shown that if $b \pmod{s} \in \{0, 1, \dots, s-m\}$, then the i th row of $A(x)$ is the same for all $i \in Z_x$. Therefore in this case $A(x)$ has rank $\leq m+1 - |Z_x|$, so Lemma 65 implies that the $m \times m$ determinant polynomial has a root of multiplicity $\geq |Z_x| - 1$ at $A(x)$. Thus $G(X) = \det(A(X))$ has a root of multiplicity $\geq |Z_x| - 1$ at $X = x$.

Summing over all $x = \omega_{q-1}^b \in \mathbb{F}_q^*$, it follows that the number of roots (including multiplicities) of $G(X)$ is at least

$$\begin{aligned} \sum_{b \in [q-1]: b \pmod{s} \in \{0, \dots, s-m\}} (|Z_{\omega_{q-1}^b}| - 1) &\geq \sum_{x \in \mathbb{F}_q^*} (|Z_x| - 1) - \frac{(q-1)m}{s} \cdot (m-1) \\ &= \sum_{x \in \mathbb{F}_q^*} |Z_x| - (q-1) \left(\frac{m(m-1)}{s} + 1 \right) \\ &= (q-1 - |\tilde{e}\tilde{v}(f)| \cdot s) \cdot m - (q-1) \left(\frac{m(m-1)}{s} + 1 \right) \end{aligned}$$

$$= (q-1)m \left(1 - \frac{|\tilde{e}\tilde{v}(f)| \cdot s}{q-1} - \frac{1}{m} - \frac{m-1}{s} \right)$$

where the second equality above holds because $|Z_x|$ equals the number of elements $x' \in x\Omega_r$ for which $\tilde{e}\tilde{v}(f)_{F_{x'}} = 0$, so each of the $q-1 - |\tilde{e}\tilde{v}(f)| \cdot s$ values $x' \in \mathbb{F}_q^*$ with $\tilde{e}\tilde{v}(f)_{F_{x'}} = 0$ contributes 1 to the sum $\sum_{x \in \mathbb{F}_q^*} |Z_x|$ a total of m times, once for each $x \in x'\Omega_r$.

But we showed above that G is a nonzero polynomial of degree $\leq (\ell-1)m$, so G has $\leq (\ell-1)m$ roots, and thus

$$(q-1)m \left(1 - \frac{|\tilde{e}\tilde{v}(f)| \cdot s}{q-1} - \frac{1}{m} - \frac{m-1}{s} \right) \leq (\ell-1)m.$$

Rearranging terms above then gives that

$$|\tilde{e}\tilde{v}(f)| \geq \frac{q-1}{s} \left(1 - \frac{\ell-1}{q-1} - \frac{1}{m} - \frac{m-1}{s} \right). \quad \square$$

Combining the bounds in the two claims above immediately implies that $|\tilde{e}\tilde{v}(f)| \geq d$ for d defined in (17), as desired. □

7 Efficient Decoding Algorithm

This section presents an efficient decoding algorithm for the qLRCs constructed in Section 5. Note that the local recovery algorithms Rec_i are by construction efficient, as they are a special case of erasure-decoding for a CSS code. Here we present an efficient algorithm for the more difficult task of globally decoding from errors at a number of unknown locations approaching half the code distance. As in Section 5, we consider both unfolded and folded quantum Tamo-Barg codes. For the unfolded codes our decoding algorithm is relatively simple but suboptimal. For the folded codes, our decoding algorithm approaches the optimal possible decoding radius as the locality r grows large.

Classical Tamo-Barg codes are subcodes of Reed-Solomon codes, so a classical Tamo-Barg code can be efficiently decoded by simply running an efficient Reed-Solomon decoder, which was known to exist. However, our (folded) quantum Tamo-Barg codes are not subcodes of quantum Reed-Solomon codes due to duality conditions for CSS codes. Just as this distinction between the quantum and classical cases made our distance analysis in Section 6 much more complicated than that of classical TB codes, our decoding algorithm for the quantum case is also more involved, though it does eventually reduce to Reed-Solomon decoding.

Note that in classical distributed storage applications of LRCs, errors often correspond to phenomena like a server being unresponsive, and therefore occur at known locations in the codeword. In such a case, it suffices to be able to decode from erasures, which can be done efficiently for any linear code using Gaussian elimination. Similarly, quantum stabilizer codes, of which our construction in Section 5 is an instance, can be efficiently decoded from erasures.

However, it is not implausible that decoding from errors in unknown locations would be useful in applications of quantum LRCs. For instance, even if a qLRC is using in a distributed setting where each code component is itself stored in a fault-tolerant memory, because quantum states are naturally more error-prone than classical data, it is not unreasonable that occasional global error

correction could be beneficial to maintaining high fidelity for the encoded state, even if most errors are caused by local failures in known locations.

7.1 Unfolded Quantum Tamo-Barg Codes

In this section we present an efficient decoding algorithm for the qTB codes in Definition 53. Our algorithm is motivated by the proof of Theorem 62, and ultimately reduces decoding the qTB codes to list decoding a classical Reed-Solomon code.

In particular, we apply Reed-Solomon list decoding up to the Johnson bound (Theorem 22) to obtain an efficient list decoder for our qTB codes, and then apply the qTB distance bound in Theorem 62 to show that when the error is not too large, the list decoder in fact serves as a unique decoder, that is, there are no extraneous list elements.

Algorithm 1: Classical decoding algorithm Dec_C for a qTB code $\mathcal{C} = \text{CSS}(C, C)$ with parameters q, r, ℓ . By Proposition 31, Dec_C can be used to obtain an efficient quantum decoding algorithm for \mathcal{C} from $\leq e$ errors in unknown locations. As a point of notation, $\text{ListDec}_{\text{RS}(q, \ell)}$ refers to the Reed-Solomon list decoder in Theorem 22, whose output is a list of message polynomials. Also, we let $B^\perp = \text{ev}(\mathbb{F}_q[X]^{[q-1] \cap (1+r\mathbb{Z})})$ as in (11) denote the space of piecewise linear functions (see Lemma 57).

Input : Corrupted codeword $a : \mathbb{F}_q^* \rightarrow \mathbb{F}_q$ with $\text{dis}(a, C) \leq e$

Output: $c' \in C$ such that $\text{dis}(c' - a, C^\perp) \leq e$

Function $\text{Dec}_C(a)$:

```

 $\mathcal{L} \leftarrow \emptyset$ 
for  $i \in \{1, \dots, r-1\}$  do
    Define  $a^{(i)} : \mathbb{F}_q^* \rightarrow \mathbb{F}_q$  by  $a^{(i)}(x) = \omega_r^{-i} a(\omega_r^i x) - a(x)$ 
     $\mathcal{L}_i \leftarrow \text{ListDec}_{\text{RS}(q, \ell)}(a^{(i)})$ 
    for  $g^{(i)}(X) = \sum_{j \in [\ell]} g_j^{(i)} X^j \in \mathcal{L}_i$  do
        if  $g_j^{(i)} = 0$  for every  $j \equiv \pm 1 \pmod{r}$  then
            Add  $\text{ev}\left(g(X) := \sum_{j \in [\ell]} (\omega_r^{(j-1)i} - 1)^{-1} g_j^{(i)} X^j\right)$  to  $\mathcal{L}$ 
return  $\arg \min_{\text{ev}(g) \in \mathcal{L}} \text{dis}(\text{ev}(g) - a, B^\perp)$ 

```

Theorem 69. *Let \mathcal{C} be the qTB code in Definition 53 with a prime locality parameter r . Then \mathcal{C} can be decoded from errors in up to*

$$e = (q-1) \cdot \frac{1}{2} \left(1 - \frac{1}{2r} - \sqrt{\frac{1}{4r^2} + \frac{r-1}{r} \cdot \frac{\ell}{q-1}} \right)$$

unknown locations in time $q^{O(1)}$.

Note that the error bound e in Theorem 69 above is just under half the distance bound d from Theorem 62; replacing the ℓ above with $\ell - 1$ gives the expression for $d/2$. Thus Theorem 69 shows that we can efficiently decode from adversarial errors up to nearly half the distance, which is optimal.

To show that Algorithm 1 runs in polynomial time, we need the following lemma, which shows that the distance computations $\text{dis}(\text{ev}(g) - a, B^\perp)$ in the final line of the algorithm can be performed efficiently.

Lemma 70. *Letting $B^\perp = \text{ev}(\mathbb{F}_q[X]^{[q-1] \cap (1+r\mathbb{Z})})$, there exists a $O(q \text{ poly log } q)$ -time algorithm that takes as input $b : \mathbb{F}_q^* \rightarrow \mathbb{F}_q$, and outputs $\text{dis}(b, B^\perp)$.*

Proof. To compute $\text{dis}(b, B^\perp) = \min_{b' \in B^\perp} |b - b'|$ efficiently, recall that B^\perp is the space of piecewise linear functions by Lemma 57. Therefore we may simply find the closest linear function to b within each coset $x\Omega_r \in \mathbb{F}_q^*/\Omega_r$ separately. Formally, for each coset $x\Omega_r \in \mathbb{F}_q^*/\Omega_r$, we may compute $\beta_{x\Omega_r} := \arg \max_{\beta \in \mathbb{F}_q} |\{x' \in x\Omega_r : b(x') = \beta \cdot x'\}|$, and then we set $b'(x') = \beta_{x\Omega_r} x'$ for all $x' \in x\Omega_r$. Then by construction $\text{dis}(b, B^\perp) = \text{dis}(b, b')$.

Note that above each $\beta_{x\Omega_r}$ can be computed in time $O(r \text{ poly log } q)$, as we may simply set $\beta_{x\Omega_r}$ equal to the mode value of $\beta = b(x')/x'$ across all $x' \in x\Omega_r$. Thus $\text{dis}(b, B^\perp)$ may be computed in time $O(q \text{ poly log } q)$. \square

Proof of Theorem 69. Let $\mathcal{C} = \text{CSS}(C, C)$ be a qTB code with parameters q, r, ℓ . By Proposition 31, it suffices to construct an algorithm $\text{Dec}_{\mathcal{C}}$ that takes as input a corrupted codeword $a = c + b$ for some $c \in C$ and some corruption $b \in \mathbb{F}_q^* \rightarrow \mathbb{F}_q$ of Hamming weight $|b| \leq e$, and outputs some $c' \in C$ such that $c' - c \in C^\perp$.

The desired algorithm is given in Algorithm 1. We first show it correctly decodes as described above, and then analyze the running time.

Consider a corrupted codeword $a = c + b$ for some codeword $c = \text{ev}(f) \in C$ and some corruption $b : \mathbb{F}_q^* \rightarrow \mathbb{F}_q$ of weight $|b| \leq e$. For a given $1 \leq i \leq r - 1$, by definition $a^{(i)}(x) = \omega_r^{-i} a(\omega_r^i x) - a(x)$ equals $c^{(i)}(x) = f^{(i)}(x) = \omega_r^{-i} f(\omega_r^i x) - f(x)$ at every point x for which $x, \omega_r^i x \notin \text{supp}(b)$. Meanwhile, within a given coset $x\Omega_r \in \mathbb{F}_q^*/\Omega_r$, the number of ordered pairs of distinct points $y, y' \in x\Omega_r$ such that $y, y' \notin \text{supp}(b)$ is precisely $(r - |b|_{x\Omega_r})(r - |b|_{x\Omega_r} - 1)$. Thus the sum over all $i \in \{1, \dots, r - 1\}$ of the number of points $x \in \mathbb{F}_q^*$ where $a^{(i)}(x) = f^{(i)}(x)$ satisfies

$$\begin{aligned}
\sum_{i=1}^{r-1} |\{x \in \mathbb{F}_q^* : a^{(i)}(x) = f^{(i)}(x)\}| &\geq \sum_{i=1}^{r-1} |\{x \in \mathbb{F}_q^* : x, \omega_r^i x \notin \text{supp}(b)\}| \\
&= \sum_{x\Omega_r \in \mathbb{F}_q^*/\Omega_r} |\{(y, y') \in (x\Omega_r)^2 : y \neq y' \text{ and } y, y' \notin \text{supp}(b)\}| \\
&= \sum_{x\Omega_r \in \mathbb{F}_q^*/\Omega_r} (r - |b|_{x\Omega_r})(r - |b|_{x\Omega_r} - 1) \\
&\geq \frac{q-1}{r} \mathbb{E}_{x\Omega_r \sim \text{Unif}(\mathbb{F}_q^*/\Omega_r)} [(r - |b|_{x\Omega_r})^2] - (q-1 - |b|) \\
&\geq \frac{q-1}{r} \mathbb{E}_{x\Omega_r \sim \text{Unif}(\mathbb{F}_q^*/\Omega_r)} [(r - |b|_{x\Omega_r})^2] - (q-1 - |b|) \\
&= \frac{r}{q-1} (q-1 - |b|)^2 - (q-1 - |b|) \\
&\geq \frac{r}{q-1} (q-1 - e)^2 - (q-1 - e),
\end{aligned}$$

where the final inequality above holds because the function $\beta \mapsto r(q-1-\beta)^2/(q-1) - (q-1-\beta)$ is decreasing for all $\beta \leq e \leq (q-1)/2$. Then averaging over all $i \in \{1, \dots, r-1\}$, there must be some such i for which

$$|\{x \in \mathbb{F}_q^* : a^{(i)}(x) = f^{(i)}(x)\}| \geq \frac{r(q-1-e)^2}{(r-1)(q-1)} - \frac{q-1-e}{r-1}. \quad (22)$$

Recall by Lemma 55, we may decompose $f(X) = g(X) + h(X)$ for $g(X) = \sum_j g_j X^j \in \mathbb{F}_q[X]^{[\ell] \setminus (\pm 1 + r\mathbb{Z})}$ and $h(X) \in \mathbb{F}_q[X]^{[q-1] \cap (1+r\mathbb{Z})}$. Define $g^{(i)}, h^{(i)}$ analogously to $f^{(i)}$, so that $f^{(i)} = g^{(i)} + h^{(i)}$. Then because

$$\omega_r^{-i}(\omega_r^i X)^j - X^j = (\omega_r^{i(j-1)} \pmod{r} - 1)X^j \quad (23)$$

equals 0 for all $j \equiv 1 \pmod{r}$, it follows that $h^{(i)} = 0$ and thus $f^{(i)} = g^{(i)}$. Therefore (22) is equivalent to

$$|\{x \in \mathbb{F}_q^* : a^{(i)}(x) = g^{(i)}(x)\}| \geq \frac{r(q-1-e)^2}{(r-1)(q-1)} - \frac{q-1-e}{r-1}. \quad (24)$$

But by (23), the coefficients of $g^{(i)}(X) = \sum_{j \in [\ell] \setminus (\pm 1 + r\mathbb{Z})} g_j^{(i)} X^j$ are given by $g_j^{(i)} = (\omega_r^{i(j-1)} \pmod{r} - 1)g_j$, and thus g and $g^{(i)}$ have coefficients of the same support, that is, $g_j = 0$ if and only if $g_j^{(i)} = 0$. In particular, it follows that $\deg(g^{(i)}) = \deg(g) < \ell$. In other words, $\text{ev}(g^{(i)}) \in \text{RS}(q, \ell)$ is a Reed-Solomon codeword, so (24) says that $a^{(i)} = \text{ev}(g^{(i)}) + b^{(i)}$ is a corrupted Reed-Solomon codeword with the corruption $b^{(i)}$ of weight

$$\begin{aligned} |b^{(i)}| &= q-1 - |\{x \in \mathbb{F}_q^* : a^{(i)}(x) = g^{(i)}(x)\}| \\ &\leq q-1 - \left(\frac{r(q-1-e)^2}{(r-1)(q-1)} - \frac{q-1-e}{r-1} \right) \\ &= (q-1) \left(1 - \left(\frac{r}{r-1} \left(1 - \frac{e}{q-1} \right)^2 - \frac{1}{r-1} \left(1 - \frac{e}{q-1} \right) \right) \right). \end{aligned}$$

By Theorem 22, the output of running the list decoder $\text{ListDec}_{\text{RS}(q, \ell)}(a^{(i)})$ is a list containing $\text{ev}(g^{(i)})$ as long as the right hand side above is at most $(q-1)(1 - \sqrt{\ell/(q-1)})$, which simplifies to needing that

$$e \leq (q-1) \left(1 - \frac{1}{2r} - \sqrt{\frac{1}{4r^2} + \frac{r-1}{r} \cdot \sqrt{\frac{\ell}{q-1}}} \right).$$

But the above holds by the definition of e along with Lemma 76, so the list \mathcal{L}_i in Algorithm 1 will contain $g^{(i)}$, and thus after the i th iteration of the for loop, the list \mathcal{L} will contain $\text{ev}(g(X)) = \text{ev}(\sum_j (\omega_r^{(j-1)i} - 1)^{-1} g_j^{(i)} X^j)$. Thus if Algorithm 1 outputs $\text{ev}(g') \in \mathcal{L}$, then $g' \in \mathbb{F}_q[X]^{[\ell] \setminus (\pm 1 + r\mathbb{Z})}$ and

$$\text{dis}(\text{ev}(g') - a, B^\perp) \leq \text{dis}(\text{ev}(g) - a, B^\perp) \leq |\text{ev}(g+h) - a| = |\text{ev}(f) - a| = |b| \leq e.$$

But as $B^\perp \subseteq C^\perp$, it follows that $\text{dis}(\text{ev}(g) - a, C^\perp) \leq e$ and $\text{dis}(\text{ev}(g') - a, C^\perp) \leq e$, so $\text{dis}(\text{ev}(g) - \text{ev}(g'), C^\perp) \leq 2e$. But by the definition of e along with Theorem 62, \mathcal{C} has distance $\min_{c \in C \setminus C^\perp} |c| >$

$2e$, so it follows that $\text{ev}(g) - \text{ev}(g') \in C^\perp$. Thus $\text{Dec}_C(a)$ outputs some $\text{ev}(g') \in \text{ev}(g) + C^\perp$, as desired.

It remains to show that Algorithm 1 runs in time $q^{O(1)}$. Theorem 22 implies that the calls to $\text{ListDec}_{\text{RS}(q,\ell)}$ run in $q^{O(1)}$ time, while Lemma 70 implies that the argmin computation in the final line of the algorithm runs in $O(|\mathcal{L}| \cdot q \text{ poly log } q) = q^{O(1)}$ time. The rest of the algorithm by definition also runs in $q^{O(1)}$ time, so the result follows. \square

7.2 Folded Quantum Tamo-Barg Codes

We now present an efficient decoding algorithm for the fqTB codes in Definition 59. The algorithm is similar to the algorithm presented in Section 7.1 for the unfolded qTB codes, and decodes up to an error fraction approaching half our distance bound in Theorem 63. Yet because the distance bound in Theorem 63 approaches the quantum Singleton bound as the locality parameter r grows large, our decoding algorithm for fqTB codes for large r tolerates error fractions approaching half the quantum Singleton, which is optimal.

We obtain this improved error tolerance for fqTB decoding compared to qTB decoding by leveraging the folding in two separate ways, namely, through the improved distance of fqTB codes (Theorem 63) compared to qTB codes (Theorem 62), and through the improved list-decodability of classical fRS codes (Theorem 23) compared to RS codes (Theorem 22). Combining these improvements for folded codes, but otherwise following the proof of Theorem 69, we obtain the following result.

Algorithm 2: Classical decoding algorithm $\text{Dec}_{\tilde{C}}$ for a fqTB code $\tilde{C} = \text{CSS}(\tilde{C}, \tilde{C})$ with parameters q, r, ℓ, s . By Proposition 31, $\text{Dec}_{\tilde{C}}$ can be used to obtain an efficient quantum decoding algorithm for \tilde{C} from $\leq e$ errors in unknown locations. As a point of notation, $\text{ListDec}_{\text{fRS}(q,\ell)}$ refers to the folded Reed-Solomon list decoder in Theorem 23, whose output is a list of message polynomials. Also, we let $\tilde{B}^\perp = \tilde{\text{ev}}(\mathbb{F}_q[X]^{[q-1] \cap (1+r\mathbb{Z})})$ as in (11) denote the folded space of piecewise linear functions (see Lemma 57).

Input : Corrupted codeword $\tilde{a} \in \widetilde{\mathbb{F}_q^*}$ with $\text{dis}(\tilde{a}, \tilde{C}) \leq e$

Output: $\tilde{c}' \in \tilde{C}$ such that $\text{dis}(\tilde{c}' - \tilde{a}, \tilde{C}^\perp) \leq e$

Function $\text{Dec}_{\tilde{C}}(\tilde{a})$:

```

 $\mathcal{L} \leftarrow \emptyset$ 
for  $i \in \{1, \dots, r-1\}$  do
    Define  $a^{(i)} : \mathbb{F}_q^* \rightarrow \mathbb{F}_q$  by  $a^{(i)}(x) = \omega_r^{-i} a(\omega_r^i x) - a(x)$ 
     $\mathcal{L}_i \leftarrow \text{ListDec}_{\text{fRS}(q,\ell,s)}(\tilde{a}^{(i)})$ 
    for  $g^{(i)}(X) = \sum_{j \in [\ell]} g_j^{(i)} X^j \in \mathcal{L}_i$  do
        if  $g_j^{(i)} = 0$  for every  $j \equiv \pm 1 \pmod{r}$  then
            Add  $\tilde{\text{ev}}(g)$  to  $\mathcal{L}$  for  $g(X) := \sum_{j \in [\ell]} (\omega_r^{(j-1)i} - 1)^{-1} g_j^{(i)} X^j$ 
return  $\arg \min_{\tilde{\text{ev}}(g) \in \mathcal{L}} \text{dis}(\tilde{\text{ev}}(g) - \tilde{a}, \tilde{B}^\perp)$ 

```

Theorem 71. Let $\tilde{C} = \text{CSS}(\tilde{C}, \tilde{C})$ be the fqTB code in Definition 59 with parameters q, r, ℓ, s such

that r is prime and such that the uncertainty principle in Proposition 66 holds for r over \mathbb{F}_q . Then \mathcal{C} can be decoded from errors in up to

$$e = \min \left\{ \frac{d}{2} - 1, e' \right\} \quad (25)$$

unknown locations in time $q^{O(\sqrt{s})}$, where d is defined as in (17) and e' is defined as in (28).

For sufficiently large folding parameter s , we can simplify the decoding radius in (25) to obtain the following bound, which approaches half the distance bound in Corollary 64.

Corollary 72. *For any $\gamma > 0$, let $\tilde{\mathcal{C}} = \text{CSS}(\tilde{C}, \tilde{C})$ be the fqTB code in Definition 59 with parameters q, r, ℓ, s such that r is a prime for which the uncertainty principle in Proposition 66 holds for r over \mathbb{F}_q . Furthermore assume that $s \geq s_0(\gamma, r)$ and $q \geq q_0(s)$, where $s_0(\gamma, r)$ and $q_0(s)$ are some sufficiently large numbers with respect to γ, r and s respectively. Then \mathcal{C} can be decoded from errors in up to*

$$e \geq \frac{q-1}{s} \cdot \frac{1}{2} \left(1 - \frac{\ell-1}{q-1} - \sqrt{\frac{1}{r} \left(1 - \frac{\ell-1}{q-1} \right)} - \gamma \right),$$

unknown locations in time $q^{O(\sqrt{s})}$.

Specifically, Corollary 72 follows from showing that our current bound on $d/2 - 1$ in Corollary 64 is smaller than e' , so our distance bound is the limiting constraint on decoding radius. Thus if we could show a tighter distance bound for fqTB codes than Theorem 63, we would immediately obtain efficient decoding up to a larger radius. For instance, if we could show that the fqTB codes have relative distance $d/((q-1)/s) \geq 1 - \ell/(q-1) - O(1/r)$, it would immediately follow that for all sufficiently large s, q , Algorithm 2 decodes fqTB codes for errors on up to $(1 - \ell/(q-1) - O(1/r))/2$ fraction of the qudits

Proof of Corollary 72. Letting $s \rightarrow \infty$ in the bound on the fqTB distance d in Corollary 64 and in the expression for e' in (28), we see that for all sufficiently large s relative to γ, r and sufficiently large q relative to s , then

$$\frac{d}{2} - 1 \geq \frac{q-1}{s} \cdot \frac{1}{2} \left(1 - \frac{\ell-1}{q-1} - \sqrt{\frac{1}{r} \left(1 - \frac{\ell-1}{q-1} \right)} - \gamma \right)$$

and

$$e' \geq \frac{q-1}{s} \cdot \left(1 - \frac{1}{2r} - \sqrt{\frac{1}{4r^2} + \frac{r-1}{r} \cdot \frac{\ell-1}{q-1}} - \frac{\gamma}{2} \right).$$

Thus to show the desired inequality in the corollary statement, because $e = \min\{d/2 - 1, e'\}$ by Theorem 71, it suffices to show that

$$\frac{q-1}{s} \cdot \frac{1}{2} \left(1 - \frac{\ell-1}{q-1} - \sqrt{\frac{1}{r} \left(1 - \frac{\ell-1}{q-1} \right)} - \gamma \right) \leq \frac{q-1}{s} \cdot \left(1 - \frac{1}{2r} - \sqrt{\frac{1}{4r^2} + \frac{r-1}{r} \cdot \frac{\ell-1}{q-1}} - \frac{\gamma}{2} \right),$$

or equivalently, that

$$\frac{1}{2} \left(1 - \frac{\ell - 1}{q - 1} - \sqrt{\frac{1}{r} \left(1 - \frac{\ell - 1}{q - 1} \right)} \right) \leq 1 - \frac{1}{2r} - \sqrt{\frac{1}{4r^2} + \frac{r - 1}{r} \cdot \frac{\ell - 1}{q - 1}}.$$

But the above inequality follows directly from Lemma 77. \square

To show that Algorithm 2 runs in polynomial time, we need the following analogue of Lemma 70 for the folded case, which shows that the distance computations $\text{dis}(\tilde{\text{ev}}(g) - \tilde{a}, \tilde{B}^\perp)$ in the final line of the algorithm can be performed efficiently.

Lemma 73. *Letting $\tilde{B}^\perp = \tilde{\text{ev}}(\mathbb{F}_q[X]^{[q-1] \cap (1+r\mathbb{Z})})$, there exists a $O(r \cdot q \text{ poly log } q)$ -time algorithm that takes as input $\tilde{b} \in \widetilde{\mathbb{F}_q^*}$, and outputs $\text{dis}(\tilde{b}, \tilde{B}^\perp)$.*

Proof. To compute $\text{dis}(\tilde{b}, \tilde{B}^\perp) = \min_{\tilde{b}' \in \tilde{B}^\perp} |\tilde{b} - \tilde{b}'|$ efficiently, recall that B^\perp is the space of piecewise linear functions by Lemma 57. Therefore we may consider each $F_x \Omega_r \in \tilde{\mathbb{F}_q^*}/\Omega_r$ separately, and for each $F_x \Omega_r$ we must compute the sequence of values $\beta = (\beta_{y \Omega_r})_{y \in F_x} \in \mathbb{F}_q^{F_x}$ that maximizes the number $w(\beta)$ of coset elements $x' \in x \Omega_r$ for which all $y' \in F_{x'}$ have $b(y') = \beta_{y' \Omega_r} \cdot y'$; note here that our notation suppresses the implicit dependence of β and $w(\beta)$ on the choice of coset $F_x \Omega_r$. Indeed, by definition we will then have $\text{dis}(\tilde{b}, \tilde{B}^\perp) = (q - 1)/s - \sum_{F_x \Omega_r \in \tilde{\mathbb{F}_q^*}/\Omega_r} w(\beta^{(F_x \Omega_r)})$, where $\beta^{(F_x \Omega_r)}$ denotes the optimal choice of β for a given $F_x \Omega_r$.

To compute the optimal $\beta = \beta^{(F_x \Omega_r)}$ for a given $F_x \Omega_r \in \tilde{\mathbb{F}_q^*}/\Omega_r$, observe that if $w(\beta) > 0$, then there must be some $x' \in x \Omega_r$ for which all $y' \in F_{x'}$ have $b(y') = \beta_{y' \Omega_r} \cdot y'$. Therefore $\beta_{y' \Omega_r} = b(y')/y'$ for all $y' \in F_{x'}$, which completely determines the value of β . Thus we may simply loop through the r coset elements $x' \in x \Omega_r$, and for each x' compute $w(\beta)$ for β given by $\beta_{y' \Omega_r} = b(y')/y'$ for all $y' \in F_{x'}$; then whichever of these r values of β maximizes $w(\beta)$ is the optimal value $\beta = \beta^{(F_x \Omega_r)}$. This algorithm by definition computes a given $\beta^{(F_x \Omega_r)}$ in time $O(r^2 s \text{ poly log } q)$ time, so it computes all $\beta^{(F_x \Omega_r)}$ for $F_x \Omega_r \in \tilde{\mathbb{F}_q^*}/\Omega_r$ in $O(r \cdot q \text{ poly log } q)$ time. Thus we can compute $\text{dis}(\tilde{b}, \tilde{B}^\perp)$ in $O(r \cdot q \text{ poly log } q)$ time. \square

Proof of Theorem 71. The proof will follow closely the proof of Theorem 71, with the main difference our use of the folded RS list decoding algorithm in Theorem 23 and the folded qTB distance bound in Theorem 63 in place of their unfolded counterparts.

By Proposition 31, it suffices to construct an algorithm $\text{Dec}_{\tilde{C}}$ that takes as input a corrupted codeword $\tilde{a} = \tilde{c} + \tilde{b}$ for some $\tilde{c} \in \tilde{C}$ and some corruption $\tilde{b} \in \widetilde{\mathbb{F}_q^*}$ of Hamming weight $|\tilde{b}| \leq e$, and outputs some $\tilde{c}' \in \tilde{C}$ such that $\tilde{c}' - \tilde{c} \in \tilde{C}^\perp$.

The desired algorithm is given in Algorithm 2. We first show it correctly decodes as described above, and then analyze the running time.

Consider a corrupted codeword $\tilde{a} = \tilde{c} + \tilde{b}$ for some codeword $\tilde{c} = \tilde{\text{ev}}(f) \in \tilde{C}$ and some corruption $\tilde{b} \in \widetilde{\mathbb{F}_q^*}$ of weight $|\tilde{b}| \leq e$. For a given $1 \leq i \leq r - 1$, by definition $a^{(i)}(x) = \omega_r^{-i} a(\omega_r^i x) - a(x)$ equals $c^{(i)}(x) = f^{(i)}(x) = \omega_r^{-i} f(\omega_r^i x) - f(x)$ at every point $x \in \mathbb{F}_q^*$ for which $x, \omega_r^i x \notin \text{supp}(b)$. Meanwhile, within a given coset $x \Omega_r \in \mathbb{F}_q^*/\Omega_r$, the number of ordered pairs of distinct points $y, y' \in x \Omega_r$ such that $y, y' \notin \text{supp}(b)$ is precisely $(r - |b|_{x \Omega_r})(r - |b|_{x \Omega_r} - 1)$. Thus the sum over all $i \in \{1, \dots, r - 1\}$

of the number of points $F_x \in \tilde{\mathbb{F}}_q^*$ where $\tilde{a}^{(i)}(F_x) = \tilde{f}^{(i)}(F_x)$ satisfies

$$\begin{aligned}
& \sum_{i=1}^{r-1} |\{F_x \in \tilde{\mathbb{F}}_q^* : \tilde{a}^{(i)}(F_x) = \tilde{f}^{(i)}(F_x)\}| \\
& \geq \sum_{i=1}^{r-1} |\{F_x \in \tilde{\mathbb{F}}_q^* : F_x, \omega_r^i F_x \notin \text{supp}(\tilde{b})\}| \\
& = \sum_{F_x \Omega_r \in \tilde{\mathbb{F}}_q^*/\Omega_r} |\{(F_y, F_{y'}) \in (F_x \Omega_r)^2 : F_y \neq F_{y'} \text{ and } F_y, F_{y'} \notin \text{supp}(\tilde{b})\}| \\
& = \sum_{F_x \Omega_r \in \tilde{\mathbb{F}}_q^*/\Omega_r} (r - |\tilde{b}|_{F_x \Omega_r})(r - |\tilde{b}|_{F_x \Omega_r} - 1) \\
& \geq \frac{q-1}{rs} \mathbb{E}_{F_x \Omega_r \sim \text{Unif}(\tilde{\mathbb{F}}_q^*/\Omega_r)} [(r - |\tilde{b}|_{F_x \Omega_r})^2] - \left(\frac{q-1}{s} - |\tilde{b}| \right) \\
& \geq \frac{q-1}{rs} \mathbb{E}_{F_x \Omega_r \sim \text{Unif}(\tilde{\mathbb{F}}_q^*/\Omega_r)} [(r - |\tilde{b}|_{F_x \Omega_r})^2] - \left(\frac{q-1}{s} - |\tilde{b}| \right) \\
& = \frac{rs}{q-1} \left(\frac{q-1}{s} - |\tilde{b}| \right)^2 - \left(\frac{q-1}{s} - |\tilde{b}| \right) \\
& \geq \frac{rs}{q-1} \left(\frac{q-1}{s} - e \right)^2 - \left(\frac{q-1}{s} - e \right),
\end{aligned}$$

where the final inequality above holds because the function $\beta \mapsto rs((q-1)/s - \beta)^2/(q-1) - ((q-1)/s - \beta)$ is decreasing for all $\beta \leq e \leq (q-1)/2s$. Then averaging over all $i \in \{1, \dots, r-1\}$, there must be some such i for which

$$|\{F_x \in \tilde{\mathbb{F}}_q^* : \tilde{a}^{(i)}(F_x) = \tilde{f}^{(i)}(F_x)\}| \geq \frac{rs}{(r-1)q-1} \left(\frac{q-1}{s} - e \right)^2 - \frac{1}{r-1} \left(\frac{q-1}{s} - e \right). \quad (26)$$

Recall by Lemma 55, we may decompose $f(X) = g(X) + h(X)$ for $g(X) = \sum_j g_j X^j \in \mathbb{F}_q[X]^{\ell \setminus (\pm 1 + r\mathbb{Z})}$ and $h(X) \in \mathbb{F}_q[X]^{[q-1] \cap (1+r\mathbb{Z})}$. Define $g^{(i)}, h^{(i)}$ analogously to $f^{(i)}$, so that $f^{(i)} = g^{(i)} + h^{(i)}$. Then because the expression in (23) equals 0 for all $j \equiv 1 \pmod{r}$, it follows that $h^{(i)} = 0$ and thus $f^{(i)} = g^{(i)}$. Therefore (26) is equivalent to

$$|\{F_x \in \tilde{\mathbb{F}}_q^* : \tilde{a}^{(i)}(F_x) = \tilde{g}^{(i)}(F_x)\}| \geq \frac{rs}{(r-1)q-1} \left(\frac{q-1}{s} - e \right)^2 - \frac{1}{r-1} \left(\frac{q-1}{s} - e \right). \quad (27)$$

But by (23), the coefficients of $g^{(i)}(X) = \sum_{j \in \ell \setminus (\pm 1 + r\mathbb{Z})} g_j^{(i)} X^j$ are given by $g_j^{(i)} = (\omega_r^{i(j-1) \pmod{r}} - 1)g_j$, and thus g and $g^{(i)}$ have coefficients of the same support, that is, $g_j = 0$ if and only if $g_j^{(i)} = 0$. In particular, it follows that $\deg(g^{(i)}) = \deg(g) < \ell$. In other words, $\text{ev}(g^{(i)}) \in \text{fRS}(q, \ell)$ is a folded Reed-Solomon codeword, so (27) says that $\tilde{a}^{(i)} = \tilde{\text{ev}}(g^{(i)}) + \tilde{b}^{(i)}$ is a corrupted Reed-Solomon

codeword with the corruption $\tilde{b}^{(i)}$ of weight

$$\begin{aligned} |\tilde{b}^{(i)}| &= \frac{q-1}{s} - |\{F_x \in \tilde{\mathbb{F}}_q^* : \tilde{a}^{(i)}(F_x) = \tilde{g}^{(i)}(F_x)\}| \\ &\leq \frac{q-1}{s} - \left(\frac{rs}{(r-1)q-1} \left(\frac{q-1}{s} - e \right)^2 - \frac{1}{r-1} \left(\frac{q-1}{s} - e \right) \right) \\ &= \frac{q-1}{s} \left(1 - \left(\frac{r}{r-1} \left(1 - \frac{es}{q-1} \right)^2 - \frac{1}{r-1} \left(1 - \frac{es}{q-1} \right) \right) \right). \end{aligned}$$

By Theorem 23, the output of running the list decoder $\text{ListDec}_{\text{FRS}(q,\ell,s)}(\tilde{a}^{(i)})$ is a list containing $\tilde{e}\tilde{v}(\tilde{g}^{(i)})$ as long as the right hand side above is at most

$$\frac{q-1}{s} \left(1 - \left(1 + \frac{2}{\sqrt{s}} \right) \left(\frac{\ell}{q-1} \right)^{1-1/\sqrt{s}} \right) - 2,$$

which simplifies to needing that $e \leq e'$ for

$$e' = \frac{q-1}{s} \left(1 - \frac{1}{2r} - \sqrt{\frac{1}{4r^2} + \frac{r-1}{r} \left(\left(1 + \frac{2}{\sqrt{s}} \right) \left(\frac{\ell}{q-1} \right)^{1-1/\sqrt{s}} + \frac{2s}{q-1} \right)} \right). \quad (28)$$

But the above holds by the definition of e , so the list \mathcal{L}_i in Algorithm 2 will contain $g^{(i)}$, and thus after the i th iteration of the for loop, the list \mathcal{L} will contain $\tilde{e}\tilde{v}(g(X)) = \tilde{e}\tilde{v}(\sum_j (\omega_r^{(j-1)i} - 1)^{-1} g_j^{(i)} X^j)$. Thus if Algorithm 2 outputs $\tilde{e}\tilde{v}(g') \in \mathcal{L}$, then $g' \in \mathbb{F}_q[X]^{\ell \setminus (\pm 1 + r\mathbb{Z})}$ and

$$\text{dis}(\tilde{e}\tilde{v}(g') - \tilde{a}, \tilde{B}^\perp) \leq \text{dis}(\tilde{e}\tilde{v}(g) - \tilde{a}, \tilde{B}^\perp) \leq |\tilde{e}\tilde{v}(g+h) - \tilde{a}| = |\tilde{e}\tilde{v}(f) - \tilde{a}| = |\tilde{b}| \leq e.$$

But as $\tilde{B}^\perp \subseteq \tilde{C}^\perp$, it follows that $\text{dis}(\tilde{e}\tilde{v}(g) - \tilde{a}, \tilde{C}^\perp) \leq e$ and $\text{dis}(\tilde{e}\tilde{v}(g') - \tilde{a}, \tilde{C}^\perp) \leq e$, so $\text{dis}(\tilde{e}\tilde{v}(g) - \tilde{e}\tilde{v}(g'), \tilde{C}^\perp) \leq 2e$. But by the definition of e along with Theorem 63, \mathcal{C} has distance $\min_{c \in \mathcal{C} \setminus \mathcal{C}^\perp} |c| > 2e$, so it follows that $\tilde{e}\tilde{v}(g) - \tilde{e}\tilde{v}(g') \in \tilde{C}^\perp$. Thus $\text{Dec}_{\tilde{C}}(\tilde{a})$ outputs some $\tilde{e}\tilde{v}(g') \in \tilde{e}\tilde{v}(g) + \tilde{C}^\perp$, as desired.

It remains to show that Algorithm 2 runs in time $q^{O(\sqrt{s})}$. Theorem 23 implies that the calls to $\text{ListDec}_{\text{FRS}(q,\ell,s)}$ run in $q^{O(\sqrt{s})}$ time, while Lemma 73 implies that the arg min computation in the final line of the algorithm runs in $q^{O(1)}$ time. The rest of the algorithm by definition also runs in $q^{O(1)}$ time, so the result follows. \square

8 Impossibility of Quantum Locally Correctable Codes

In this section, we show that quantum codes are inherently unable to perform local recovery from a large number of erasures. Thus the local decoding capabilities of qLRCs is in some sense close to optimal for quantum codes. This result is in contrast to the classical setting, where there exist locally correctable codes (LCCs), which can recover every given code symbol from a constant number r of other code symbols, even after a linear number of code symbols have been erased.

Our impossibility result is stated below. Informally, it states that any qudit in a quantum code that can be recovered from two disjoint sets of other qudits must be useless for error correction; that is, such a qudit is entirely unentangled from the remainder of the code state, and contains no information about the encoded message.

Theorem 74. *Let \mathcal{C} be a quantum code of block length n and dimension $k > 0$. Assume that for some $i \in [n]$, there exist two subsets $I_i^1, I_i^2 \subseteq [n]$ satisfying $I_i^1 \cap I_i^2 = \{i\}$ such that for each $b = 1, 2$, there is an associated recovery channel $\text{Rec}_i^b : \mathcal{M}(I_i^b \setminus \{i\}) \rightarrow \mathcal{M}(I_i^b)$ with the guarantee that for every code state $\psi \in \mathcal{C}$,*

$$\text{Rec}_i^b \otimes I_{[n] \setminus I_i^b}(\psi_{[n] \setminus \{i\}}) = \psi.$$

Then there exists a 1-qudit density matrix $\alpha \in \mathcal{M}(1)$ such that every $\psi \in \mathcal{C}$ can be decomposed as $\psi = \alpha_i \otimes \psi_{[n] \setminus \{i\}}$.

In comparison, qLRCs require each qudit $i \in [n]$ to be recoverable from just a single set $I_i \setminus \{i\}$ of other qudits. Thus Theorem 74 shows that it is impossible to extend a qLRC to even just have a second disjoint recovery set for each qudit.

This result is in contrast to the classical case, where code components can be recovered from many different disjoint subsets of components. Indeed, classically there exist locally correctable code (LCCs), which have the property that even after erasing any constant fraction of the code components, each erased component can be recovered from some constant number of unerased components.

Theorem 74 shows that no such code LCC can exist quantumly. Specifically, assume a quantum code \mathcal{C} is such that each code qudit i can be recovered from qudits $I_i \setminus \{i\}$. Then Theorem 74 shows that for every $i \in [n]$, if the qudits in I_i are erased, then it is impossible to locally recover qudit i without performing a more global decoding operation that also recovers other qudits in I_i .

Proof of Theorem 74. Define two additional qudits $i_1 = n$ and $i_2 = n + 1$. For $b = 1, 2$, define $J^b = I_i^b \cup \{i_1\} \setminus \{i\}$. Now for any given pure code state $\psi = |\psi\rangle \langle \psi| \in \mathcal{C}$, let $\rho \in \mathcal{M}([n+2])$ be the state obtained by applying Rec_i^1 but placing the recovered copy of qudit i in register i_1 , and also applying Rec_i^2 but placing the recovered copy of qudit i in register i_2 . That is,

$$\rho = (\text{Rec}_i^1)_{J^1} (\text{Rec}_i^2)_{J^2} (\psi).$$

Note that by definition J^1, J^2 , and $\{i\}$ are all disjoint, so in particular $(\text{Rec}_i^1)_{J^1}$ and $(\text{Rec}_i^2)_{J^2}$ act on disjoint sets of qudits and therefore commute. Letting $\psi_{[n] \setminus \{i\} \cup \{i_b\}}$ denote the state ψ with qudit i moved to position i_b , it follows that

$$\begin{aligned} \text{Tr}_{i_1}(\text{Tr}_i \rho) &= (\text{Rec}_i^2)_{J^2} \text{Tr}_{i_1}(\text{Rec}_i^1)_{J^1}(\psi_{[n] \setminus \{i\}}) \\ &= (\text{Rec}_i^2)_{J^2} \text{Tr}_{i_1}(\psi_{[n] \setminus \{i\} \cup \{i_b\}}) \\ &= (\text{Rec}_i^2)_{J^2}(\psi_{[n] \setminus \{i\}}) \\ &= \psi_{[n] \setminus \{i\} \cup \{i_2\}}. \end{aligned} \tag{29}$$

Therefore we have shown that tracing over i_1 in $\text{Tr}_i \rho$ yields the pure state $\psi_{[n] \setminus \{i\} \cup \{i_2\}} \cong \psi$, and thus there is a tensor decomposition

$$\text{Tr}_i \rho = \alpha_{i_1} \otimes \psi_{[n] \setminus \{i\} \cup \{i_2\}}$$

for some 1-qudit density matrix $\alpha \in \mathcal{M}(1)$.

But the same reasoning used to show (29) also implies that $\text{Tr}_{i_2}(\text{Tr}_i \rho) = \psi_{[n] \setminus \{i\} \cup \{i_1\}}$, so

$$\begin{aligned} \psi_{[n] \setminus \{i\} \cup \{i_1\}} &= \text{Tr}_{i_2}(\text{Tr}_i \rho) \\ &= \text{Tr}_{i_2}(\alpha_{i_1} \otimes \psi_{[n] \setminus \{i\} \cup \{i_2\}}) \\ &= \alpha_{i_1} \otimes \psi_{[n] \setminus \{i\}}, \end{aligned}$$

or equivalently,

$$\psi = \alpha_i \otimes \psi_{[n]\setminus\{i\}}.$$

It remains to be shown that α_i is the same for all $\psi \in \mathcal{C}$. This conclusion follows Lemma 26, which states that if a quantum code can recover from erasures on a given set of qudits (here Rec_i^b recovers from erasures on qudit i), then the reduced density matrix of a code state on those qudits contains no information about the encoded message state. \square

9 Acknowledgments

We thank Thiago Bergamaschi for helpful discussions.

References

- [AEL95] Noga Alon, Jeff Edmonds, and Michael Luby. Linear time erasure codes with nearly optimal recovery. In *Proceedings of IEEE 36th Annual Foundations of Computer Science*, pages 512–519, October 1995. ISSN: 0272-5428.
- [AN22] Anurag Anshu and Chinmay Nirkhe. Circuit Lower Bounds for Low-Energy States of Quantum Code Hamiltonians. In Mark Braverman, editor, *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*, volume 215 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 6:1–6:22, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. ISSN: 1868-8969.
- [BGG22] Thiago Bergamaschi, Louis Golowich, and Sam Gunn. Approaching the Quantum Singleton Bound with Approximate Error Correction, December 2022. arXiv:2212.09935 [quant-ph].
- [DEL⁺22] Irit Dinur, Shai Evra, Ron Livne, Alexander Lubotzky, and Shahar Mozes. Locally testable codes with constant rate, distance, and locality. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2022, pages 357–374, New York, NY, USA, June 2022. Association for Computing Machinery.
- [DHLV23] Irit Dinur, Min-Hsiu Hsieh, Ting-Chun Lin, and Thomas Vidick. Good Quantum LDPC Codes with Linear Time Decoders. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, STOC 2023, pages 905–918, New York, NY, USA, June 2023. Association for Computing Machinery.
- [EI76] R. J. Evans and I. M. Isaacs. Generalized Vandermonde Determinants and Roots of Unity of Prime Order. *Proceedings of the American Mathematical Society*, 58(1):51–54, 1976. Publisher: American Mathematical Society.
- [Fre04] P. E. Frenkel. Simple proof of Chebotarev’s theorem on roots of unity, July 2004. arXiv:math/0312398 version: 3.
- [GGI05] Daniel Goldstein, Robert M. Guralnick, and I. M. Isaacs. Inequalities for Finite Group Permutation Modules. *Transactions of the American Mathematical Society*, 357(10):4017–4042, 2005. Publisher: American Mathematical Society.

- [GHSY12] Parikshit Gopalan, Cheng Huang, Huseyin Simitci, and Sergey Yekhanin. On the Locality of Codeword Symbols. *IEEE Transactions on Information Theory*, 58(11):6925–6934, November 2012.
- [GI01] V. Guruswami and P. Indyk. Expander-based constructions of efficiently decodable codes. In *Proceedings 42nd IEEE Symposium on Foundations of Computer Science*, pages 658–667, October 2001. ISSN: 1552-5244.
- [GI02] Venkatesan Guruswami and Piotr Indyk. Near-optimal linear-time codes for unique decoding and new list-decodable codes over smaller alphabets. In *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*, STOC '02, pages 812–821, New York, NY, USA, May 2002. Association for Computing Machinery.
- [GI03] Venkatesan Guruswami and Piotr Indyk. Linear time encodable and list decodable codes. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, STOC '03, pages 126–135, New York, NY, USA, June 2003. Association for Computing Machinery.
- [GK16] Venkatesan Guruswami and Swastik Kopparty. Explicit subspace designs. *Combinatorica*, 36(2):161–185, April 2016.
- [GKO⁺18] Sivakanth Gopi, Swastik Kopparty, Rafael Oliveira, Noga Ron-Zewi, and Shubhangi Saraf. Locally Testable and Locally Correctable Codes approaching the Gilbert-Varshamov Bound. *IEEE Transactions on Information Theory*, 64(8):5813–5831, August 2018.
- [GR08] Venkatesan Guruswami and Atri Rudra. Explicit Codes Achieving List Decoding Capacity: Error-Correction With Optimal Redundancy. *IEEE Transactions on Information Theory*, 54(1):135–150, January 2008.
- [GRS22] Venkatesan Guruswami, Atri Rudra, and Madhu Sudan. Essential coding theory. *Draft available at <http://www.cse.buffalo.edu/atri/courses/coding-theory/book>*, 2022.
- [GS98] V. Guruswami and M. Sudan. Improved decoding of Reed-Solomon and algebraic-geometric codes. In *Proceedings 39th Annual Symposium on Foundations of Computer Science (Cat. No.98CB36280)*, pages 28–37, November 1998. ISSN: 0272-5428.
- [GXY19] Venkatesan Guruswami, Chaoping Xing, and Chen Yuan. How Long Can Optimal Locally Repairable Codes Be? *IEEE Transactions on Information Theory*, 65(6):3662–3670, June 2019.
- [HRZW20] Brett Hemenway, Noga Ron-Zewi, and Mary Wootters. Local List Recovery of High-Rate Tensor Codes and Applications. *SIAM Journal on Computing*, 49(4):FOCS17–157, January 2020. Num Pages: FOCS17-195 Publisher: Society for Industrial and Applied Mathematics.
- [HSX⁺12] Cheng Huang, Huseyin Simitci, Yikang Xu, Aaron Ogus, Brad Calder, Parikshit Gopalan, Jin Li, and Sergey Yekhanin. Erasure Coding in Windows Azure Storage. In *Proceedings of the 2012 USENIX conference on Annual Technical Conference*, USENIX ATC'12, page 2, USA, June 2012. USENIX Association.

- [HW18] Brett Hemenway and Mary Wootters. Linear-time list recovery of high-rate expander codes. *Information and Computation*, 261:202–218, August 2018.
- [KMRZS16] Swastik Kopparty, Or Meir, Noga Ron-Zewi, and Shubhangi Saraf. High-rate locally-correctable and locally-testable codes with sub-polynomial query complexity. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, STOC '16, pages 202–215, New York, NY, USA, June 2016. Association for Computing Machinery.
- [LZ22] Anthony Leverrier and Gilles Zémor. Quantum Tanner codes. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 872–883. IEEE Computer Society, October 2022.
- [MLR⁺14] Subramanian Muralidhar, Wyatt Lloyd, Sabyasachi Roy, Cory Hill, Ernest Lin, Weiqi Liu, Satadru Pan, Shiva Shankar, Viswanath Sivakumar, Linpeng Tang, and Sanjeev Kumar. f4: Facebook’s Warm BLOB Storage System. In *Proceedings of the 11th USENIX conference on Operating Systems Design and Implementation*, OSDI'14, pages 383–398, USA, October 2014. USENIX Association.
- [MOP21] Sidhanth Mohanty, Ryan O’Donnell, and Pedro Paredes. Explicit Near-Ramanujan Graphs of Every Degree. *SIAM Journal on Computing*, pages STOC20–1, February 2021. Num Pages: STOC20-23 Publisher: Society for Industrial and Applied Mathematics.
- [NC10] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, December 2010.
- [PK22] Pavel Panteleev and Gleb Kalachev. Asymptotically good Quantum and locally testable classical LDPC codes. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2022, pages 375–388, New York, NY, USA, June 2022. Association for Computing Machinery.
- [SS96] M. Sipser and D.A. Spielman. Expander codes. *IEEE Transactions on Information Theory*, 42(6):1710–1722, November 1996.
- [Tao04] Terence Tao. An uncertainty principle for cyclic groups of prime order, July 2004. arXiv:math/0308286 version: 6.
- [TB14] Itzhak Tamo and Alexander Barg. A Family of Optimal Locally Recoverable Codes. *IEEE Transactions on Information Theory*, 60(8):4661–4676, August 2014.
- [Vad12] Salil P. Vadhan. Pseudorandomness. *Foundations and Trends[®] in Theoretical Computer Science*, 7(1–3):1–336, December 2012.
- [WLH23] Adam Wills, Ting-Chun Lin, and Min-Hsiu Hsieh. Tradeoff Constructions for Quantum Locally Testable Codes, September 2023. arXiv:2309.05541 [quant-ph] version: 1.
- [Zha19] Guanghui Zhang. On the Chebotarëv theorem over finite fields. *Finite Fields and Their Applications*, 56:97–108, March 2019.

A Technical Lemmas

Below we prove the technical lemmas used in the Section 6.2 for bounding the distance of the fqTB codes.

The following proof of Lemma 65 is well known, and is included for completeness.

Proof of Lemma 65. To show that \det has a root of multiplicity $m - t$ at $x = (x_{ij})_{i,j \in [m]}$ of rank t , it suffices to show that all d th derivatives of \det vanish at x for all $d \in [m - t]$. But by definition every d th derivative of $\det(X_{ij})_{i,j \in [m]}$ either vanishes, or is equal (up to a sign) of the derivative of a $(m - d) \times (m - d)$ submatrix of $(X_{ij})_{i,j \in [m]}$. For $X_{ij} = x_{ij}$ and $d \in [m - t]$, every such submatrix has rank at most $\text{rank}(x_{ij})_{i,j \in [m]} = t < m - d$, so every such submatrix is not full rank and therefore has determinant 0. Thus every d th derivative of $\det(X_{ij})_{i,j \in [m]}$ for $d \in [m - t]$ vanishes at $X_{ij} = x_{ij}$, and thus \det has a root of multiplicity $m - t$ at x . \square

The following lemma bounds the term ϵ defined by (18) in Theorem 63.

Lemma 75. *Define ϵ as in (18). If we set $s = cr^2$ for $c \geq 2$ and let $\lambda = 1 - (\ell - 1)/(q - 1)$, then*

$$\epsilon \leq \left(1 + \frac{1}{c}\right) \sqrt{\frac{\lambda}{r}}. \quad (30)$$

Proof. By definition

$$\epsilon \leq \max_{1 \leq m \leq r} \lambda \cdot \frac{m - 1}{r} < \lambda.$$

Thus if $\lambda \leq 1/r$, then

$$\epsilon \leq \lambda = \sqrt{\lambda^2} \leq \sqrt{\frac{\lambda}{r}},$$

from which (30) follows.

Therefore assume that $\lambda \geq 1/r$. By definition

$$\epsilon \leq \max_{1 \leq m \leq r} \min \left\{ \frac{\lambda}{r} \cdot m, \frac{1}{m} + \frac{1}{s} \cdot m \right\}.$$

Because $s = cr^2 \geq r^2$ by assumption, the expression $1/m + m/s$ is decreasing in m for $1 \leq m \leq r$, while the expression $\lambda m/r$ is increasing in m . Thus because $m^* = 1/\sqrt{\lambda/r - 1/s}$ is the unique positive real number for which $\lambda m^*/r = 1/m^* + m^*/s$, it follows that

$$\begin{aligned} \epsilon &\leq \frac{\lambda}{r} \cdot m^* = \frac{1}{m^*} + \frac{1}{s} \cdot m^* \\ &= \sqrt{\frac{\lambda}{r}} \cdot \frac{1}{\sqrt{1 - \frac{r}{\lambda s}}} \\ &\leq \sqrt{\frac{\lambda}{r}} \cdot \frac{1}{\sqrt{1 - \frac{1}{c}}} \\ &\leq \left(1 + \frac{1}{c}\right) \sqrt{\frac{\lambda}{r}}, \end{aligned}$$

where the second inequality above holds because $s = cr^2$ and $\lambda \geq 1/r$, and the final inequality holds because $1/\sqrt{1-x} \leq 1+x$ for all $x \leq 1/2$. Thus (30) holds. \square

The following bound is used in the proof of Theorem 69.

Lemma 76. *For all real numbers $0 \leq x \leq 1/6$ and $1/2 \leq y \leq 1$, it holds that*

$$\frac{1}{2} \left(1 - x - \sqrt{x^2 + (1-2x)y} \right) \leq 1 - x - \sqrt{x^2 + (1-2x)\sqrt{y}}.$$

Proof. Let $f(x, y)$ denote the LHS minus the RHS of the inequality in the lemma statement. Letting $D = [0, 1/6] \times [1/2, 1]$, then our goal is to show that $f(x, y) \leq 0$ for all $(x, y) \in D$. By definition $f(x, 1) = 0$ for all $0 \leq x \leq 1/6$, so it is sufficient to show that $\frac{\partial f}{\partial y} \geq 0$ for all $(x, y) \in D$. Now by definition

$$\frac{\partial f}{\partial y} = -\frac{(1-2x)/2}{2\sqrt{x^2 + (1-2x)y}} + \frac{(1-2x)/(2\sqrt{y})}{2\sqrt{x^2 + (1-2x)\sqrt{y}}},$$

so setting the above ≥ 0 and rearranging gives that we must show

$$(1-2x-x^2)y - (1-2x)\sqrt{y} + x^2 \leq 0$$

But solving the quadratic equation on the LHS above gives roots at $\sqrt{y} = 1$ and $\sqrt{y} = x^2/(1-2x-x^2) \leq 1/23$ for $0 \leq x \leq 1/6$, so indeed the above inequality holds for all $(x, y) \in D$, and thus all $(x, y) \in D$ have $\frac{\partial f}{\partial y} \geq 0$ and therefore $f(x, y) \leq 0$, as desired. \square

The following bound is used in the proof of Theorem 71.

Lemma 77. *For all real numbers $0 \leq x \leq 1/6$ and $0 \leq y \leq 1/2$, it holds that*

$$\frac{1}{2}(y - \sqrt{2xy}) \leq 1 - x - \sqrt{x^2 + (1-2x)(1-y)}.$$

Proof. Let $f(x, y)$ denote the LHS minus the RHS of the inequality in the lemma statement. Letting $D = [0, 1/6] \times [0, 1/2]$, then our goal is to show that $f(x, y) \leq 0$ for all $(x, y) \in D$. Computing the partial derivatives

$$\begin{aligned} \frac{\partial f}{\partial x} &= -\frac{1}{2\sqrt{2}} \cdot \sqrt{\frac{y}{x}} + 1 - \frac{1-x-y}{\sqrt{x^2 + (1-2x)(1-y)}} \\ \frac{\partial f}{\partial y} &= \frac{1}{2} \left(1 - \sqrt{\frac{x}{2y}} \right) - \frac{1-2x}{2\sqrt{x^2 + (1-2x)(1-y)}} \\ \frac{\partial^2 f}{\partial x \partial y} &= -\frac{1}{4\sqrt{2}} \cdot \frac{1}{\sqrt{xy}} - \frac{(1-x-y)(1-2x)}{2(x^2 + (1-2x)(1-y))^{3/2}} + \frac{1}{\sqrt{x^2 + (1-2x)(1-y)}}, \end{aligned}$$

it follows that the statements below hold for all $(x, y) \in D$:

1. $f(x, 0) = 0$
2. $f(0, y) = y/2 + \sqrt{1-y} - 1 \leq 0$
3. $\frac{\partial f}{\partial x}(x, 0) = 0$

4. $\frac{\partial}{\partial y} \frac{\partial f}{\partial x}(x, y) = \frac{\partial^2 f}{\partial x \partial y}(x, y) \leq 0$ if $x \leq 1/50$ or $y \leq 1/20$
5. $\frac{\partial f}{\partial x}(x, y) \leq 1$
6. $\frac{\partial f}{\partial y}(x, y) \leq 1/2$

Combining item 3 and item 4 above implies that $\frac{\partial f}{\partial x}(x, y) \leq 0$ for all $(x, y) \in D$ such that either $x \leq 1/50$ or $y \leq 1/20$. Combining this fact with item 2 above implies that $f(x, y) \leq 0$ for all $(x, y) \in D$ such that either $x \leq 1/50$ or $y \leq 1/20$.

Therefore if we find some sufficiently large integer $n \geq 100$ such that $f(x, y) \leq -2/n$ for all $(x, y) \in ([1/100, 1/6] \cap \mathbb{Z}/n) \times ([1/40, 1/2] \cap \mathbb{Z}/n)$, then item 5 and item 6 above imply that $f(x, y) \leq 0$ for all $(x, y) \in [1/50, 1/6] \times [1/20, 1/2]$, which combined with our conclusion above that $f(x, y) \leq 0$ when $x \leq 1/50$ or $y \leq 1/20$, implies the desired inequality $f(x, y) \leq 0$ for all $(x, y) \in D$. But we may numerically verify that $n = 1000$ satisfies the desired property, which completes the proof. \square

B Omitted Proofs

This section provides proofs that were omitted in the main text.

Proof of Proposition 46. The desired decoding algorithm for \mathcal{C} simply applies π_G^{-1} to unpermute the symbols, then applies the decoder for \mathcal{C}_{in} to each of the n_{out} inner code blocks, and finally applies the decoder for \mathcal{C}_{out} to the resulting state. By definition this decoder runs in $\text{poly}(n \log q)$ if the outer and inner codes' decoding algorithms both run in $\text{poly}(n \log q)$ time. Thus we just need to verify the correctness of this decoding algorithm.

Assume that the corruptions occur on some set $T \subseteq [n]$ of $|T| \leq \alpha n$ components of \mathcal{C} . Also assume for a contradiction that our decoder for \mathcal{C} fails to recover the original message, which can only occur if after applying π_G^{-1} , at least $\alpha_{\text{out}} n_{\text{out}}$ of the inner code blocks have $\geq \alpha_{\text{in}} n_{\text{in}}$ corruptions, so that all of these inner code blocks are “overloaded” and their inner decodings fail to recover the correct value of the outer code component. Let $S \subseteq [n] \cong [n_{\text{out}}] \times [n_{\text{in}}/\Delta]$ denote the set of all the folded components in these overloaded blocks. Then because $\geq \alpha_{\text{out}} n_{\text{out}}$ inner code blocks are overloaded, we have

$$|S| \geq \alpha_{\text{out}} n_{\text{out}} \cdot \frac{n_{\text{in}}}{\Delta} = \alpha_{\text{out}} n \quad (31)$$

while because each of the $|S|\Delta/n_{\text{in}}$ overloaded inner code block has $\geq \alpha_{\text{in}} n_{\text{in}}$ of its n_{in} symbols corrupted, we have

$$|E(S, T)| \geq \frac{|S|\Delta}{n_{\text{in}}} \cdot \alpha_{\text{in}} n_{\text{in}} = \alpha_{\text{in}} \Delta |S|,$$

as each corrupted inner code symbol in an overloaded inner code block corresponds to an edge from S to T . But the expander mixing lemma implies that

$$\begin{aligned} |E(S, T)| &\leq \frac{\Delta |S| |T|}{n} + \lambda \Delta \sqrt{|S| |T|} \\ &\leq \alpha \Delta |S| + \lambda \Delta \sqrt{|S| \cdot \alpha n} \\ &< \left(\alpha_{\text{in}} - \lambda \sqrt{\frac{\alpha_{\text{in}}}{\alpha_{\text{out}}}} \right) \Delta |S| + \lambda \Delta \sqrt{|S| \cdot \alpha_{\text{in}} n}, \end{aligned}$$

where we have applied the definition of $\alpha = \alpha_{\text{in}} - \lambda\sqrt{\alpha_{\text{in}}/\alpha_{\text{out}}}$. Combining the two above inequalities gives that

$$\alpha_{\text{in}}\Delta|S| < \left(\alpha_{\text{in}} - \lambda\sqrt{\frac{\alpha_{\text{in}}}{\alpha_{\text{out}}}}\right)\Delta|S| + \lambda\Delta\sqrt{|S| \cdot \alpha_{\text{in}}n},$$

which simplifies to

$$|S| < \sqrt{\alpha_{\text{out}}n},$$

contradicting (31). Thus the assumption that the decoder fails was false, so the decoder must succeed in correcting errors on any set $|T|$ of $|T| \leq \alpha n$ components of \mathcal{C} . \square