Quantum Locally Testable Codes

SIAM J. COMPUT.
Vol. 44, No. 5, pp. 1230–1262

QUANTUM LOCALLY TESTABLE CODES*

DORIT AHARONOV† AND LIOR ELDAR‡

Abstract. We initiate the study of quantum locally testable codes (qLTCs). Classical LTCs are very important in computational complexity. These codes are defined as the linear subspace satisfying a set of local constraints, with the additional requirement that their *soundness*, $R(\delta)$, which is the probability that a randomly chosen constraint is violated, is proportional to the *proximity* δ, where δn is the distance of a word from the code. Excellent LTCs exist in the classical world, and they are tightly related to the celebrated PCP (probabilistically checkable proof) theorem. In quantum complexity, quantum error correcting codes provide central examples in the study of the illusive behavior of multiparticle entanglement, and they have played a crucial role in many computational complexity results. We provide a definition of the quantum analogue of LTCs and motivate it by connecting its central notions in the study of both entanglement and quantum Hamiltonian complexity. A natural question is whether such codes exist, and how good can their soundness be. To the best of our knowledge all quantum codes known today exhibit poor soundness. Moreover, we show that the soundness of CSS codes (which are commonly used quantum codes defined by two classical codes) is governed by the *minimal* soundness of the two classical codes; in the most natural CSS code we examined as a candidate qLTC, namely, the Reed–Muller code, there is a tradeoff between the parameters of the two codes, which prevents the resulting quantum code from being qLTC. These facts seem to suggest a more general phenomenon, by which the soundness of qLTCs is inherently restricted due to multiparticle entanglement. Our main technical contribution consists of two complementary results regarding qLTCs which are stabilizer codes (denoted sLTCs). We first prove a surprising, inherently quantum property of sLTCs. For small constant values of proximity, the better the local expansion of the interaction graph of the constraints, the *less* sound the sLTC becomes. This stands in sharp contrast to the classical setting. The complementary, more intuitive result also holds (and is actually much more involved technically to prove in the quantum case): an upper bound on the soundness when the code is defined on *bad* local expanders. Together we arrive at a quantum upper bound on the soundness of sLTCs set on *any* graph, which does not hold in the classical case. Many open questions are raised regarding what possible parameters are achievable for qLTCs, and their relation to other objects of interest in quantum information theory. In the appendix we also define a quantum analogue of PCPs of proximity (PCPPs) and point out that the result of [E. Ben-Sasson et al., *SIAM J. Comput.*, 36 (2006), pp. 889–974] by which PCPPs imply LTCs with related parameters carries over to the sLTCs. This creates a first link between qLTCs and quantum PCPs [D. Aharonov, I. Arad, and T. Vidick, *ACM SIGACT News Archive*, 44 (2013), pp. 47–79].

Key words. locally testable codes, quantum error correcting codes, quantum PCP, stabilizer codes

AMS subject classifications. 81P40, 81P45, 81P68, 81P70

DOI. 10.1137/140975498

1. Introduction.

Consider the following question: We are given a classical code of n-bit strings, defined by $O(1)$-local constraints (namely, an LDPC code[1]). We are also given a word x which is of Hamming distance $\delta n > 0$ from the code (we say it has *proximity* δ). The underlying question when dealing with locally testable codes

*Received by the editors July 1, 2014; accepted for publication (in revised form) July 23, 2015; published electronically October 15, 2015.

http://www.siam.org/journals/sicomp/44-5/97549.html

†School of Computer Science and Engineering, The Hebrew University, Jerusalem, Israel (doria@cs.huji.ac.il).

‡School of Computer Science and Engineering, The Hebrew University, Jerusalem, Israel, and Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, MA 02139 (leldar@mit.edu).

[1]LDPC stands for low density parity check; this corresponds to the fact that the parity check matrix has only a few coordinates in each row.

is, *what is the probability that a constraint chosen uniformly at random is violated?* We denote by $R(\delta)$ (called the *soundness*) the lower bound on the probability that *any* word of proximity exactly δ will violate a randomly chosen constraint. Locally testable codes (LTCs) are those for which a large distance from the code guarantees with high probability that one could detect that the word is not in the code, namely, those codes with good soundness. The soundness parameter formally captures an intuitive notion of "robustness" of these codes.

LTCs constitute perhaps the most natural and simple instantiation of the following type of problem, defining the topic of *property testing*: We are given a set of strings P, which we refer to as a *property*, and access to coordinates of our choice of a given string x. We would like to decide whether or not x possesses the property, i.e., lies inside P. Clearly, in many cases one needs to have access to all of x's coordinates in order to decide this question. However, when one relaxes the question and only asks whether x is in P or *far* from it (far can be measured in different ways), then it sometimes becomes possible to decide membership in various properties even when reading only a small number of coordinates. In some sense, this is a combinatorial way to study connections between global and local characterization of a certain property.

Indeed, since LTCs which initiated the field of property testing were first defined [33], this has developed into a thriving area of its own. Equally importantly, LTCs play an instrumental role in all proofs of the celebrated PCP (probabilistically checkable proof) theorem [7, 8, 26].[2] The understanding of the limitation and possible constructions of LTCs had developed into a very active field of its own (see, for example, Goldreich's survey [32]).

Here, we initiate the study of quantum LTCs, which, to the best of our knowledge, were not defined or studied before. Our motivating question, like in the classical case, is understanding global versus local behavior of sets, except now we are dealing with quantum states in which multiparticle entanglement plays a crucial role, and in which such global-versus-local tradeoffs are particularly interesting. What new behaviors and limitations emerge when quantum entanglement enters the scene?

1.1. Classical locally testable codes: Background. The behavior of LTCs is usually explored in one of two contexts: as an errorcorrecting code, or in relation to PCPs (see [32]); depending on the context, one is interested in different ranges of distances from the code, namely, different ranges of proximities. In particular, in the context of error correction, the interesting regime of proximities is at most half the distance *of* the code; in this regime, the error can still be corrected. In the context of PCPs, on the other hand, larger distances can be of interest, since a cheating prover may provide witnesses of arbitrary distance from the code.

The following terminology is often used: When the soundness $R(\delta)$ is at least some constant, for any δ larger than some (other) constant, the code will be called *weak*-LTC; but if the requirement is that $R(\delta)$ is bounded from below by a linear function of δ for any δ, we will say that the code is a *strong*-LTC. Note that one cannot expect the soundness to be better than linear in δ if we assume, as is commonly done, that the number of constraints in which each bit is involved is bounded by a constant.

Some well-known classical strong-LTCs include the Reed–Muller code [48], the Hadamard code [6], and Haståd's long code [40], which were used in the PCP proofs of [7, 8, 26]. We mention in passing that these codes are not so satisfying when

[2]The PCP theorem states, roughly, that any NP problem can be cast in a format in which the verifier only needs to read $O(1)$ bits from the proof in order to determine its veracity with some constant (say 2/3) probability.

considering the rate; the Hadamard code and the long code's rates are exponentially and doubly exponentially small, respectively; the best known LTCs [32] which are achieved by combining the results of [26, 15] exhibit constant distance, constant query complexity, and rates which are polylogarithmically small. A major open question in the area of LTCs, called the c^3 problem [34], asks whether good (namely, constant relative rate and distance) weak-LTCs exist. Nevertheless, if one is interested only in the soundness parameter, classical codes perform excellently.

Moreover, it is common wisdom in the PCP community that good soundness is rather easy to achieve for proximities which are below some constant threshold (this is achieved by random codes), and the difficulty in designing such codes arises for proximity values beyond the minimal distance of the code (see Claim 5 in Appendix C for a more precise statement, as well as discussion following the statement of Theorem 1 below).

1.2. Quantum locally testable codes: Definition and motivation. Quantum error correcting codes have played a crucial role in quantum complexity theory (see, e.g., [36, 37, 12, 11]), and their study is a vastly growing field (see, e.g., [35, 24, 55, 46, 47, 20, 30, 17, 42]); they are related to a variety of issues, including, of course, resilience of quantum computations to noise, and their fault tolerance, as well as various protocols in quantum cryptography and quantum communication. Very importantly, quantum error correcting codes provide an excellent probe into understanding the illusive nature of multiparticle entanglement, which is attracting a lot of attention in both the physics and quantum computation communities (see, e.g., [43, 18, 19]).

To the best of our knowledge the quantum analogue of LTCs has not yet been defined. This is of particular interest since the notion of local testability of quantum codes seems to provide a combinatorial handle into a topic of significant interest to the scientific community nowadays, namely, the *robustness* of multiparticle entanglement and its global versus local behavior. We provide here a first definition of quantum locally testable codes (qLTCs) (see Definition 15) and initiate its study by providing some results which indicate an intriguing difference between the quantum and classical behavior with respect to local testability. Later, in the discussion sections, we discuss in much more detail the connection between our definition and various topics of interest such as robustness of entanglement at room temperature, self-correcting codes, and the quantum PCP conjecture.

To define a quantum LTC, we recall that a quantum LDPC code (namely, a code with local constraints; see footnote 1) can be defined as the ground space (namely, the zero eigenspace) of a Hermitian operator acting on n qubits, called a local Hamiltonian $H = \sum_{i=1}^{m} \Pi_i$. The local terms Π_i are projections which act nontrivially only on $O(1)$ qubits. We will refer to the local terms Π_i as *quantum constraints*. As is commonly used in the quantum Hamiltonian complexity literature (see, for example, [2]), the right quantum analogue to the probability for violating a randomly chosen constraint is the average energy per constraint:[3]

$$\frac{1}{m}\langle\psi|H|\psi\rangle = \frac{1}{m}\sum_{i=1}^{m}\langle\psi|\Pi_i|\psi\rangle. \tag{1}$$

[3]Note that this is different from the average energy per qubit, which is perhaps the more commonly used quantity in physics.

To define the notion of soundness for such LDPC codes, we want to require that this quantity be large when the distance from the code is large. We thus need some reasonable notion of a Hamming distance on quantum states. Following standard definitions of distance between orthogonal code states, we say that the distance between two orthogonal states $|\phi_1\rangle, |\phi_2\rangle$ is the maximum weight w, such that if E is any tensor product of one-qubit Pauli matrices, of weight at most w (the weight of a Pauli is the number of coordinates in which it is a nonidentity), then E induces no overlap between them, i.e., $\langle\phi_1|E|\phi_2\rangle = 0$. This definition can of course be extended to a distance between a state and a set of states.

We can now define the soundness of the quantum code. We say that a quantum code C on n qubits is a quantum locally testable code (qLTCs) with soundness $R(\delta)$ if for all $0 < \delta < 1$ we have

$$\min_{|\psi\rangle, dist(|\psi\rangle, C) \geq \delta n} \frac{1}{m}\langle\psi|H|\psi\rangle = R(\delta); \tag{2}$$

i.e., for any state ψ that is within distance exactly δn from the code space, its average energy with respect to the constraints, $\frac{1}{m}\langle\psi|H|\psi\rangle$, is at least $R(\delta)$ (see Definition 15).

Following the classical case, we say that a code is a *weak*-qLTC if its soundness is at least a constant for values of δ larger than some constant, And a *strong*-qLTC would mean linear soundness as a function of the proximity for *all* distances.

1.2.1. Standard quantum codes as qLTCs. Since CSS codes [50] are an important and rather easy to study family of quantum error correcting codes, which are defined using classical codes, and, moreover, since classical LTCs are abundant, it seems natural to start by asking whether we can find qLTCs among known codes in this familiar class.

Recall that CSS codes are defined by a pair of classical codes $\mathcal{L}_x, \mathcal{L}_z \subseteq \mathbf{F}_2^n$, with the added restriction that $\mathcal{L}_x^\perp \subseteq \mathcal{L}_z$ (equivalently, $\mathcal{L}_z^\perp \subseteq \mathcal{L}_x$). This extra restriction allows us to map the defining constraints of such codes into an Abelian stabilizer group with two types of constraints: Pauli X terms, corresponding to $\mathcal{L}_x^\perp$, and Pauli Z terms, corresponding to $\mathcal{L}_z^\perp$. Furthermore, this definition allows for an elegant analysis of the parameters of these codes by looking at the quotients $\mathcal{L}_x/\mathcal{L}_z^\perp$, $\mathcal{L}_z/\mathcal{L}_x^\perp$.

We observe (see Appendix F for the proof) that CSS codes inherit their soundness from the minimal soundness of the two codes that comprise them.

FACT 1. *Let $Q(\mathcal{L}_x, \mathcal{L}_z)$ be a quantum CSS code built from two binary classical codes $\mathcal{L}_x, \mathcal{L}_z \subseteq \{0,1\}^n$. If both $\mathcal{L}_x, \mathcal{L}_z$ have query complexities q_1, q_2 and soundness functions $R_1(\delta), R_2(\delta)$, then Q as a quantum code has query complexity $q = max\{q_1, q_2\}$ and soundness $R(\delta)$, for which*

$$\min\{R_1(\delta), R_2(\delta)\} \geq R(\delta) \geq \min\{R_1(\delta/2), R_2(\delta/2)\}. \tag{3}$$

An important first example to consider is the quantum version of Reed–Muller codes, first defined in [53]. As stated before, classical Reed–Muller codes are strongly locally testable. This is achieved by choosing the degree of the (multivariate) polynomial, whose coefficients encode the data message, as constant. Quantum Reed–Muller codes [53] can be constructed using classical Reed–Muller codes and their dual, in the usual CSS paradigm. By the above fact, the resulting code will inherit its soundness from one of the two classical codes that defines the CSS code—the one with the worse soundness (see Fact 1). However, using a pair of classical RM-codes $C_1 = RM(r_1, m_1), C_2 = RM(r_2, m_2)$ (where r_i is the degree of polynomials over

m_i variables) in their locally testable form, to construct a CSS code, implies that $r_2 = \theta(m - r_1)$, so at least one of r_1, r_2 is $\Omega(m)$. Unfortunately, a result by Alon et al. [4], shows that testing an $RM(r, m)$ code requires $2^{\Omega(r)}$ queries, thus making the classical tests (and the quantum tests, by inheritance of Fact 1) highly nonlocal. This implies that quantum Reed–Muller codes are very far from being locally testable.

Another illuminating example to consider is Kitaev's two-dimensional toric code [43], which turns out to have very poor soundness. The toric code is a CSS codes, and thus its poor soundness can be argued using the poor soundness of its classical codes, but it is easier to explain using the existence of string-like error terms, which is perhaps more familiar. A string operator of length L is a set of one-qubit Pauli operators all of type X, say, applied on all qubits along some "string" in the lattice. Such a string error induces only two violations, one at each end-point of the string. In particular, one could choose a string of length $\theta(\sqrt{n})$, which induces only two violations, thus ruling out strong local testability of the toric code, because local testability for the regime of proximities δ up to the order of $1/\sqrt{n}$ is ruled out.

This upper bound can in fact be made stronger; in fact, the toric code can be shown to be not locally testable even up to $\delta = o(1)$. This is because for any $\delta > 0$ we can devise an error term of weight δn, with a meager $\delta^2 \cdot n$ violated constraints as follows: We divide the two-dimensional grid into boxes of side length $1/\delta$, and in each one deploy a string error of length $\Omega(1/\delta)$ at its center. We make the error sufficiently short so its weight cannot be decreased modulo the centralizer of the stabilizer group of the toric code. The weight of such an error (modulo the code) is δn, since in each square, a fraction δ of the qubits are erred. On the other hand, the number of violations is a fraction $O(\delta^2)$ of the constraints, since this is the ratio in each box. Hence, $R(\delta)$ is $O(\delta^2)$ for subconstant δ's. We observe that this behavior scales with the dimension of the embedding space. For example, for the four-dimensional toric code, using similar arguments, we can upper-bound $R(\delta)$ by $O(\delta^{3/2})$.[4]

One may be tempted to try nonconstant dimensional toric codes, having soundness converging to a linear function of δ as $\delta^{(l-1)/(l-2)}$, where l is the dimension of the embedding space. However, in that case, the query complexity becomes nonconstant as well, thereby undermining its local testability by fiat.

The apparent difficulty to derive what for classical LTCs is almost standard—namely, strong soundness—even if we allow a vanishing rate raises the question of whether quantum local testability is possible, even in a weak sense. We state our motivating question.

QUESTION 1. *What are possible constructions of qLTCs? In particular, does the additional aspect of multiparticle entanglement in qLTC, compared to their classical counterpart, pose limitations on the possible parameters of LTCs in the quantum setting?*

1.3. Main technical results.

1.3.1. Stabilizer qLTCs. Being probably the richest and most well-studied class of quantum codes, stabilizer codes [35] are compelling and easy to work with. We provide a simpler definition for stabilizer qLTCs (denoted sLTC; see Definition 17) and prove that it coincides with the definition of qLTCs on stabilizer codes in Claim 3. The rest of our technical results concern sLTCs.

[4] As errors we take open two-dimensional manifolds of area $1/\delta^2$, centered in boxes whose volume is $1/\delta^4$. The number of violated constraints in each box scales like the one-dimensional boundary of its error, so it is $O(1/\delta)$. Hence, for an error of fractional weight $(1/\delta^2)/(1/\delta^4) = \delta^2$, the fractional number of violations is $O((1/\delta)/(1/\delta^4)) = O(\delta^3)$, or $R(\delta) = O(\delta^{3/2})$.

1.3.2. Bounds on the soundness of sLTCs. We focus on sLTC's on n qudits, defined by m $k = O(1)$-local check terms, where each qudit participates in $D_L = O(1)$ constraints. For such codes, we provide two upper bounds on the soundness. Both bounds hold for values of proximities which are at most some constant fraction of the minimal distance of the code. This constant is a function of k and D_L. Usually, in the classical setting, it is much easier to derive LTCs whose soundness is good (large) for those *small* proximity values (see, e.g., our Claim 5). Here, we show that in this supposedly *easier* range of parameters, qLTCs are severely limited compared to their classical counterparts.

There is a twofold reason for being interested in linear soundness as a function of the proximity, in the small proximity regime (namely, in strong qLTCs). First, in both the quantum and classical cases, linear soundness as a function of the proximity is the maximal soundness one can hope for (under the commonly used assumption which we make here that each bit or qubit participates in the number of constraints D_L, and that these constraints are local). The fact that classically one can achieve this makes the quantum requirement a reasonable goal to set. There is in fact also a second, purely quantum motivation for studying strong soundness at low ranges of proximities: There are indications that such strong soundness may imply that at subconstant temperatures the equilibrium state of a quantum system is still highly entangled. We refer the reader to the discussion section where we explain this connection between qLTCs and robustness of entanglement.

To make the statement of the results simpler, we observe that the soundness $R(\delta)$ is bounded above by the number of constraints that touch the erred qudits, divided by m: hence it is at most $\delta n D_L/m = k\delta$ (using $D_L n = km$). It turns out that it is more informative to present our results in the following terms.

DEFINITION 1 (relative soundness).

$$r(\delta) = R(\delta)/k\delta. \tag{4}$$

The relative soundness is the soundness normalized by its maximal value (for an exact definition, see Definition 18).

Bound on sLTCs set on local expanders. Our first main result proves that sLTCs exhibit a severe limitation on their relative soundness at small proximities when set on good expanders. More precisely, consider the bipartite graph of the code defined with n bits on the left side, m constraints on the other side, and edges connecting each constraint to all of its bits. This is sometimes referred to as the Tanner graph of the code [54]. Our constraints are all k-local for some $k = O(1)$. We say that the bipartite graph is an ε-local expander if every subset of at most k qubits is examined nearly by as many constraints as it possibly can, namely, by at least $(1-\varepsilon)kD_L$ constraints. Theorem 1 shows that in the quantum setting, when the underlying bipartite graph of the sLTC code is an ε-local expander, the relative soundness is $O(\varepsilon)$. In other words, the better the expansion, the worse the soundness. This holds for all proximities smaller than some fraction of the minimal distance of the code, δ_{min}.

THEOREM 1 (stabilizer qLTCs based on expanding topology are limited). *Let C be a stabilizer code on n d-dimensional qudits, of relative minimal distance δ_{min}, with a k-local generating set $\mathcal{G} \subset \Pi_d^n$, such that each qudit is examined by D_L generators. Put $\delta_0 = \min\left\{\frac{1}{k^3 \cdot D_L}, \frac{1}{2}\delta_{min}\right\}$. Suppose the bipartite interaction graph of $\mathcal{G}$ is ε-local expanding for $\varepsilon < 1/2$. Then, for all $0 < \delta < \delta_0$, we have $r(\delta) \leq 2\varepsilon$.*

See section 3 for exact definitions of stabilizer codes and their generators, and Definition 18 for the exact definition of relative soundness.

Theorem 1 stands in sharp contrast to the classical case. Classically, codes can easily be constructed on good expanders so that for small proximities their soundness is excellent. We provide an explicit example whose relative soundness is arbitrarily close to 1 by plugging the *lossless expanders* constructed in [22] into the expander code construction of Sipser and Spielman [52]. This implies good classical codes with constant query complexity and with almost optimal soundness for any proximity δ smaller than some constant (see Claim 5 in Appendix C).

Bound on the soundness of any sLTC. Our second main result is an upper bound on the relative soundness which holds for sLTCs set on *any* underlying bipartite graph, not necessarily local expanders.

THEOREM 2 (roughly). *Let C be a stabilizer code on n d-dimensional qudits, of k-local terms ($k \geq 4$) with relative distance $\delta_{min} = \Omega(n^{-\varepsilon})$ for some constant $\varepsilon \in [0,1)$, where each qudit interacts with $O(1)$-local terms. Any error of fractional weight $\delta < \frac{1}{2}\delta_{min} \leq 1$ has $r(\delta) \leq \alpha(d)(1-\gamma_{gap})$ for some constant function $\gamma_{gap} = \gamma_{gap}(k,d) > 0$.*

In the above theorem $\alpha(d) = 1 - 1/(d+1)$; this is a technical upper bound on the relative soundness of qLTCs defined on d-dimensional qudits, stemming quite easily from the size of the alphabet d (see subsection 5.1).

The proof of the theorem uses probabilistic bounds in which some exponential-decay behavior "defeats" a linear function. This occurs at $k = 4$, hence the limitation in the statement of the theorem. We point out, though, which the case of $k = 3$, which was left out of the theorem, is in fact irrelevant for qubits/qutrits [3] since any quantum code with such parameters has distance $O(1)$.

Theorem 2 shows that the soundness is further bounded by some seemingly deeper quantum phenomenon. We stress that this upper bound, which is not exhibited in classical codes, is found in the range of parameters of δ (small constants) in which it is supposed to be *easiest* to achieve soundness for LTCs; see, e.g., Claim 5 in Appendix C.

1.3.3. Quantum PCPs of proximity. LTCs are tightly connected [32] to PCPs of proximity (PCPPs), which are proof systems defined very similarly to PCPs (see [14]). For the reader familiar with PCPs, they too consider a verifier who gets access to an untrusted proof; however, PCPPs differ from PCPs in two important aspects: first, they are weaker, in the sense that they are required to reject only inputs that are *far* from the language, whereas in PCPs any input out of the language should be rejected; second, the verifier is charged not only for the number of queries out of the proof, but also for the number of queries out of (part of) the input. For a formal definition, see Appendix H.

Ben Sasson et al. [14] provide a standard construction of an LTC from a PCPP. Given a PCPP for membership in a code, and an error correcting code C, they construct an LTC code C', which inherits its soundness parameter from the soundness parameter of the PCPP and its distance from the code C (see Construction 4.3 and Proposition 4.4 in [14]; see also our Appendix H).

In Appendix H, we suggest a definition of quantum PCPPs and show that a similar result to that of [14] holds in the quantum setting. The meaning of the definition of qPCPP and of the above-described connection, and their relevance and importance to the quantum PCP conjecture, are far from clear (see, for example, [2] for doubts regarding the classical approach to proving the quantum PCP conjecture,

and the direct applicability of quantum error correcting codes in this context). Still we provide these definitions and results in the appendix to make the point that a syntactic connection does carry over also in the quantum regime. Finding a deep meaning to the connection between qLTCs and the quantum local testability of proofs remains an open problem, similarly to the classical case [32].

1.4. Overview of proofs of Theorems 1 and 2.

1.4.1. Bounds on sLTC codes on expanders. To prove Theorem 1, we want to use good local expansion in order to construct an error which will not have a large energy penalty (namely, will not violate too many constraints) but which will be of large weight. More precisely, the error should have a large weight modulo the centralizer of the stabilizer group (see Definition 17), and yet should not violate too many stabilizer generators (recall that an error violates a stabilizer generator, or constraint, if it does not commute with it; see Definition 8).

The key idea is that in a local expander, intersections between stabilizer generators which consist of more than one qudit are rare (see Fact 3). The size of the intersection matters since for two generators that intersect on a single qubit, the restrictions of those operators to that qubit must *commute*, because the two generators commute overall (see Definition 8). We note that it cannot be that *all* generators when restricted to a given qudit commute, because this would mean that this qubit is trivial for the code (see remark at the end of subsection 2.4). An error defined on a qudit in such a way that it commutes with the majority of the generators acting on it will violate only a small fraction of the constraints acting on that qudit.

To extend this to errors of larger weight (up to some small constant fraction of the minimal distance), we apply the above idea to each of the generators in a large "sparse" set of generators, namely, a set in which each two terms are of at least some constant distance apart in the interaction bipartite graph (formally, a 1-independent set of terms; see Definition 20). It is not difficult to see that due to the distance between the generators, the error weight remains large even modulo the centralizer.

1.4.2. Upper bound on soundness for stabilizer sLTCs on any graph. To prove Theorem 2, we want to prove that regardless of the graph they are set on, the relative soundness of sLTCs can never achieve optimal soundness in a well-defined sense. We use the bound of Theorem 1 (the "surprising" side) augmented with a claim that quantum stabilizer codes not only suffer from the quantum effect of Theorem 1 but also cannot avoid the classical effect by which codes with *poor* local expansion have low soundness, namely, that large error patterns are examined by relatively few check terms, so the number of constraints they violate is relatively low. Together, this means that for *any* underlying graph, whether a good or a bad local expander, the relative soundness is nontrivially bounded.

In the classical case, arguing that poor expansion implies poor relative soundness is almost trivial for proximities which are less than half the minimal distance of the code; see below. A similar phenomenon, by which poor local expansion implies poor local testability, holds also in the quantum case, but the proof turns out to be quite nontrivial. Let us clarify what we're trying to show. We want to show that if the expansion is bad, one can construct an error of large weight but which does not have large relative penalty. Let us start with the classical argument.

Suppose we would like to upper-bound the soundness function $r(\delta)$ of a classical code C, for some range of proximity values $(0, \delta_0]$, for $\delta_0 \leq \delta_{min}(C)/2$. Consider a set of bits S whose fractional size is δ_0 and which has expansion error $\varepsilon > 0$, namely,

the number of constraints touching it is ε close to optimal. A priori, the binary word $w \in \{0,1\}^n$ which is the indicator vector on S (i.e., 1 on bits of S, and 0 otherwise) is violated by at most $|S|D_L(1-\varepsilon)$ check terms by the assumption on the expansion. On the other hand, by linearity of the code and the fact that $|S|/n \leq \delta_0$, we have that $dist(w, C) = |S|$. Hence the ratio of violation to distance is at most $D_L(1-\varepsilon)$.

In the quantum setting, however, an analogous Pauli error on a set of qubits S may just seem to be large, whereas it can be actually represented much more succinctly. Formally, for a stabilizer code generated by group $\mathcal{G}$, we need to examine the weight of any Pauli word w *modulo its centralizer*, $C(\mathcal{G})$. If we restrict w to have weight at most $\delta_{min}/2$, this amounts to finding its minimal weight modulo $\mathcal{G}$. We would hence like to devise an error pattern that cannot be downsized significantly modulo $\mathcal{G}$, but would still "sense" the nonexpanding nature of S, and hence have fewer-than-optimal violations. This corresponds to checking that the weight of an error is not reduced modulo the dual code—a requirement which does not appear in the classical setting, and makes the proof much more difficult.

To this end we prove the onion fact (Fact 8), which might be of interest of its own. It states that given an error on at most $k/2$ of the k qudits supporting a generator $g \in \mathcal{G}$, its weight cannot be reduced modulo $\mathcal{G}$ within the k-neighborhood of g (the k-neighborhood is, roughly, the qudits belonging to the set of terms of distance k from g in the interaction graph). The "onion" in the name is due to the fact that the proof (given in subsection 5.3.3) works via some hybrid argument on the onion-like layers $\Gamma^{(i)}(u)$ surrounding the qudits of a generator u.

Building on this fact, our strategy in constructing an error pattern is to concentrate the error on a large set of faraway generators whose k-neighborhoods are nonintersecting (we call those generators "islands"). We now argue as follows. If we draw a random error on the qudits belonging to these "islands", with probability calibrated so that the expected number of errors per "island" is, say, 1 error, the following will occur: On one hand, many islands have more than one error, so they "sense" the suboptimality of expansion. On the other hand, only a meager fraction, exponentially small in k, of the "islands" with at least two errors will have more than $k/2$ errors; only those by the onion fact (Fact 8) can be potentially reduced modulo the centralizer. Hence with high probability the weight of the random error cannot be significantly reduced modulo the centralizer, yet it still has a less-than-optimal number of violations due to the expansion.

On a technical level, in order to actually find a large error that has both a large weight and small energy penalty, we need its size to be diverging in n, so we require that the minimal distance of the code is sufficiently large (diverging in n, though not necessarily linear in n).

1.5. Discussion: Relation to other notions of robustness of entanglement. The definition of local testability in the quantum setting suggests a probe to the study of the robustness of multiparticle entanglement, a topic of much current interest.

Another, perhaps more natural definition of entanglement robustness, which has a more direct physical motivation, is requiring that the Gibbs state at room (namely, constant) temperature of a local Hamiltonian is highly entangled as a mixed state. In physics, the Gibbs state, namely, the state of a system at equilibrium at temperature T, is defined to be the density matrix derived by assigning to an eigenstate of the Hamiltonian with eigenvalue E a probability which is proportional to $e^{-E/T}$. The amount of entanglement in density matrices is hard to define even for two-qubit states,

but we can use the following natural definition: We say that the Gibbs state is highly entangled if, for any distribution D of pure states realizing the density matrix, a state chosen randomly according to D is with all but negligible probability highly entangled.

A stronger formulation is via Hastings' NLTS (no low-energy trivial states) conjecture [41]. This conjecture states roughly that there exist local Hamiltonians such that all their low-energy states (not just all their ground states, as in quantum error correcting codes) are highly entangled. The question of whether NLTS local Hamiltonians exist is a major open question in quantum Hamiltonian complexity [2] as a necessary step towards clarifying the widely open quantum PCP conjecture [1, 2], a quantum analogue of the PCP theorem.[5] It is not hard to see that NLTS systems have a highly entangled Gibbs state for sufficiently small but nonzero temperature. Hence, the NLTS conjecture (and thus also the qPCP conjecture) implies the existence of local Hamiltonians whose Gibbs states are entangled at room temperatures.

We note that NLTS Hamiltonians and qLTCs seem related: while in qLTCs low energies imply closeness to the code, in NLTS Hamiltonians they imply high entanglement, which is well known to be necessary for code states. Indeed, some weak connections between the two notions were already proven.[6] We believe that understanding any one of those notions better would lead to much better understanding of the other.

There are strong indications that qLTCs may imply the existence of multiparticle entanglement in the Gibbs distribution at temperatures proportional to δ_{min} for a somewhat different, more combinatorial notion of multiparticle entanglement than the one usually used in the literature. This will be explained in a follow-up work by one of us [29]. Very roughly, the idea in [29] is to characterize the existence of multiparticle entanglement in a state by providing a lower bound on the ability of classical low-depth circuits to generate the *correlations* exhibited by measuring that state. It is then shown that if the soundness is linear in δ for a wide enough range of values of δ starting from some sufficiently small fraction of δ_{min}, then any quantum state satisfying at least $1 - O(\delta_{min})$ of the local tests of the code is highly entangled under the above definition. This indicates that possibly *all* quantum states with energy at most $O(\delta_{min})$, as in the Gibbs state, may be highly entangled.

qLTC's can be related also to the well-known physical notion of self-correcting memories [21, 25, 39, 23, 38, 56]; in fact, they can be viewed as a strengthening of that notion. A self-correcting memory is a medium in which a quantum state is maintained almost in tact for a long time without active error correction, even at constant temperatures, because any transition between two orthogonal code states will encounter a high energy barrier. In qLTCs, the requirement is that not only should such transitions encounter a high energy barrier at some point, but also that at any point along the transition the energy barrier scale with the distance from the code space. We note that a strong qLTC of linear minimal distance would imply a self correcting memory, far better (in terms of the energy barrier) than the state-of-the-art self-correcting memories known today [38, 56].

1.6. Related work. Theorem 1 is related to a recent result of the current authors [3] which investigated low energy states of local Hamiltonian set of expanders.

[5] The qPCP question states, roughly, that it is quantum-NP hard to approximate the ground energies of local Hamiltonians even to within a constant fraction. There has been much recent work attempting to make progress on the qPCP question (see the recent survey [2] and references therein).

[6] One can show that qLTCs do not have tensor-product states with small (constant) mean energy.

It was shown in [3] that when a quantum local Hamiltonian, whose terms mutually commute, is set on a good local expander, then already at low energies we can find states which are almost nonentangled. As the expansion improves, this happens at lower and lower energies. In Theorem 1 we show that as the expansion improves, the soundness deteriorates, which means that already states with low energy can be very far from the code.

Another result of similar spirit was derived by Brandão and Harrow [16] for noncommuting 2-local Hamiltonians on standard expanders. In both results good expansion poses a limitation on the expressiveness of quantum constraint systems.

We note that both the proof of our Theorem 1 and the result of [3] start with Facts 2 and 3 regarding the percentage of unique neighbors in good local expanders; however, the proofs proceed from that point onward in very different directions.

It is interesting to compare this behavior to the results of Dinur and Kaufman [28], who showed that classical LTC codes *must* be set on a good local expander. More precisely, given a code with soundness $R(\delta) = \rho \cdot \delta$ for all $\delta > \delta_0$ for some constant δ_0, the edge expansion of the underlying graph is at least $c\rho$ for some constant c. This might seem to provide another classical contrast to our Theorem 1, in addition to our Claim 5. However, [28] does not use bipartite graph expansion but rather the graph in which an edge connects any two nodes that participate in a common constraint; the two notions of expansion are very different, and hence direct comparison to the [28] result is not possible.

1.7. Conclusions and open questions. Our results raise the following fundamental question: could it be that the notion of quantum local testability, and more generally, the notion of testing global quantum properties using local proves, is inherently limited?

Our upper-bound results are unfortunately restricted to stabilizer codes; however, those are very general structures which are known to be capable of exhibiting very intricate quantum behaviors (including the existence of good error correcting quantum codes). It seems thus natural to ask how general those results are, and whether quantum local testability is somehow inherently limited, at least when the properties to be tested are represented by linear subspaces.

Many open questions arise regarding qLTCs. Can we find qLTCs with much better soundness than those mentioned in this article? Do qLTCs exist with parameters which are as good as those that are described in [32], even disregarding the rate, namely, constant relative distance, constant query complexity, and constant soundness for all proximities larger than some constant $\delta_0 > 0$? If not, can we prove appropriate upper bounds on qLTCs? Can we rule out the c^3 conjecture in the quantum setting?

The upper bounds we provided here point to an inherently quantum phenomenon, which constitutes an obstacle against local testability for qLTCs in the low-proximity range of parameters. Both of our main theorems reflect, in fact, a deeper phenomenon called *monogamy of entanglement*, which was identified also in [3] for commuting local Hamiltonian and in [16] for 2-local general Hamiltonians. Essentially, this phenomenon limits the amount of entanglement that a single qudit with $O(1)$ quantum levels can "contain." Whether Theorem 2 hints at a more profound limitation on quantum local testability of codes that holds also for larger values of δ calls for further research.

Finally, we remark that though in the classical setting LTCs have been instrumental in PCP theory [7, 8, 26], it is yet unclear whether qLTCs can be applied directly for qPCPs. Still, their study is likely to shed light on the important and tightly

related open problem of robustness of multiparticle entanglement, tightly related to the qPCP question.

Organization of paper. In section 2 we provide the necessary background on quantum error correcting codes and on local expanders. Section 3 provides definitions of quantum locally testable codes (qLTCs) and stabilizer qLTCs, as well as basic results. Section 4 provides bounds on the soundness of quantum LTCs on local expanders, and section 5 provides an absolute bound on soundness of stabilizer LTCs regardless of the expansion of their underlying graph. Finally, in the appendices we provide several proofs which are on the more technical side. In Appendix H we provide our definition of quantum PCPPs and the construction and proof of the induced qLTC.

2. Preliminaries.

2.1. Notation.

NOTATION 1. *A quantum system or code is defined on a finite-dimensional Hilbert space $\mathcal{H}$, comprised of n d-dimensional qudits, i.e., $\mathcal{H} = \mathbf{C}^{d^{\otimes n}}$.*

NOTATION 2. *A bipartite graph is denoted by $G = (L, R; E)$: L is the set of left vertices of size $|L| = n$ (corresponding to qudits), R is the set of right vertices $|R| = m$ (corresponding to constraints), and E is the set of edges between L and R. D_L will denote the left degree of a bipartite graph. k will denote the locality of the constraints, namely, the right degree of the graph.*

NOTATION 3. *Given $S \subseteq R$ (or L) in a bipartite graph, $\Gamma(S)$ denotes the neighbor set of S in L (or R). $\mathcal{N}(q)$ will denote the qudit-neighborhood of a qudit q in L, namely, all the qudits participating in all the constraints acting on q (so, $\mathcal{N}_q = \Gamma^{(2)}(q)$). We will use ε to denote the expansion error for bipartite graphs (as in Definition* 13*).*

NOTATION 4. *We will use $\delta = \delta(w, C)$ to denote the relative distance of a word w from a code C, sometimes referred to as proximity. δ_{min} denotes the relative minimal distance of the code.*

2.2. The Pauli groups.

DEFINITION 2 (Pauli group). *The group Π^n is the n-fold tensor product of Pauli operators $A_1 \otimes A_2 \otimes \cdots \otimes A_n$, where $A_i \in \{I, X, Y, Z\}$, along with multiplicative factors $\pm 1, \pm i$ with matrix multiplication as group operation.*

The Pauli group can be generalized to particles of any dimensionality d.

DEFINITION 3 (Pauli group generalized to F_d). *Let $X_d^k : |i\rangle \mapsto |(i+k) \pmod d\rangle$, $P_d^\ell |j\rangle \mapsto w_d^{j\ell} |j\rangle$ be the generalized bit and phase flip operators on the d-dimensional Hilbert space, where $w_d = e^{2\pi i/d}$ is the primitive dth root of unity. Let Π_d be the group generated by these operators and all roots of unity of order d. The group Π_d^n is the n-fold tensor product of Pauli operators $A_1 \otimes A_2 \otimes \cdots \otimes A_n$, where $A_i \in \{X_d^k P_d^\ell\}$ along with the multiplicative factors w_d^{jl}.*

The weight of a Pauli operator is defined to be the number of locations where it is a nonidentity.

2.3. General quantum error correction.

DEFINITION 4 (quantum code). *A quantum code on n qudits is given by a set of (m) projections Π_i. The code is defined to be the simultaneous 0 eigenstates of all those projections.*

DEFINITION 5 (quantum error detection 1 [45]). *Let $C \subseteq \mathcal{H}$ be a quantum code on n qudits. Let Π_C be the orthogonal projection onto C. We say that the set of errors $\mathcal{E}$ is detectable by C if, for any $E \in \mathcal{E}$, we have*

$$\Pi_C E \Pi_C = \gamma_E \Pi_C, \tag{5}$$

where γ_E is some constant which may depend on E.

DEFINITION 6 (quantum error detection 2 [45]). *A set $\mathcal{E}$ is detectable by C if, for any $|\psi\rangle, |\phi\rangle \in C$ with $\langle\psi|\phi\rangle = 0$, and any $E \in \mathcal{E}$, $\langle\psi|E|\phi\rangle = 0$.*

CLAIM 1 (see [45]). *Definitions 6 and 5 are equivalent.*

The proof can be found in the appendix. Definition 6 gives rise to the following natural definition.

DEFINITION 7 (minimal distance of a code [45]). *Let C be a quantum code detecting error set $\mathcal{E} \subset \Pi_d^n$. C has minimal distance $\delta_{min}(C)$ if, for any two orthogonal code states $|\phi\rangle, |\psi\rangle$ and any $E \in \mathcal{E}$ of weight at most $\delta_{min}(C) - 1$, we have $\langle\phi|E|\psi\rangle = 0$.*

2.4. Stabilizer quantum error correcting codes.

DEFINITION 8 (stabilizer code). *A stabilizer code C is defined by an Abelian subgroup $A = A(\mathcal{G}) \subset \Pi_d^n$, generated by a set $\mathcal{G} \subset \Pi_d^n$. The code space is defined as the mutual 1-eigenspace of all elements in $\mathcal{G}$ (we require that $-I \notin \mathcal{G}$ so that this code space is not empty). An element $E \in \Pi_d^n$ is said to be an error if it does not commute with at least one element of $\mathcal{G}$, i.e., $E \notin \mathbf{Z}(\mathcal{G})$, where $\mathbf{Z}(\mathcal{G})$ is the centralizer of $\mathcal{G}$. An element $E \in \Pi_d^n$ is said to be a logical operation if it commutes with all of $\mathcal{G}$, but is not generated by $\mathcal{G}$, i.e., $E \in \mathbf{Z}(\mathcal{G}) - A$. A stabilizer code is said to be k-local if each term $g \in \mathcal{G}$ is an element of Π_d^n with weight exactly k.*[7]

To fit with the terminology of Definition 4, consider for each generator g the projection Π_g which projects on the orthogonal subspace to the 1-eigenspace of g.

DEFINITION 9 (succinct representation). *A k-local set of generators $\mathcal{G}$ is said to be succinct if there does not exist a different generating set $\mathcal{G}'$, such that $A(\mathcal{G}) = A(\mathcal{G}')$ and $wt(g) < k$ for some $g \in \mathcal{G}'$.*

The following is a well-known fact [35] which will be useful later on, and which we prove in Appendix D.

LEMMA 1 (stabilizer decomposition). *Let C be a stabilizer code on n qudits, and consider the sets $EC = \{E|\phi\rangle, |\phi\rangle \in C\}$ with $E \in \Pi_d^n$. Then two sets EC, $E'C$ are either orthogonal or equal to each other, and $\{EC\}_{E \in \Pi_d^n}$ span the entire Hilbert space. Moreover, consider the partition of the entire Hilbert space to sets of states which are common eigenvectors of all generators of C with exactly the same set of eigenvalues for each generator. Then this partition is exactly the partition derived by the EC's, and two orthogonal EC's have two lists of eigenvalues which differ on at least one generator. In particular, any n qudit state $|\psi\rangle$ may be written as a sum of orthogonal vectors*

$$|\psi\rangle = \sum_i E_i |\eta_i\rangle,$$

where $E_i \in \Pi_d^n$ and $|\eta_i\rangle \in C$.

[7] In the literature, a k-local term usually implies that the support of the term is *at most* k qudits. Here for simplicity of analysis we specify that it is exactly k qudits.

DEFINITION 10 (weight of an error in stabilizer codes). *Let C be a stabilizer code on n d-dimensional qudits, with generating set $\mathcal{G} \subset \Pi_d^n$. For $E \in \Pi_d^n$, we denote*

1. *the number of locations in which E is a nonidentity by $wt(E)$;*
2. *the weight of E modulo the group $A(\mathcal{G})$ by $wt_{\mathcal{G}}(E)$:*

$$wt_{\mathcal{G}}(E) = \min_{f \in A(\mathcal{G})} \{wt(fE)\}.$$

3. *the weight of E modulo the centralizer $\mathbf{Z}(\mathcal{G})$ by $wt_{\mathbf{Z}(\mathcal{G})}(E)$:*

$$wt_{\mathbf{Z}(\mathcal{G})}(E) = \min_{z \in \mathbf{Z}(\mathcal{G})} \{wt(zE)\}.$$

The above claims give rise to the following definition of distance in a stabilizer code.

DEFINITION 11 (minimal distance of a stabilizer code). *Let C be a k-local stabilizer code on n d-dimensional qudits, with generating set $\mathcal{G} \subset \Pi_d^n$. The minimal distance of C is defined as the minimal weight of any logical operation on C:*

$$\delta_{min}(C) = \min_{E \in \mathbf{Z}(\mathcal{G}) - A(\mathcal{G})} wt(E).$$

CLAIM 2 (equivalence of distance definitions). *A stabilizer code C has $\delta_{min}(C) \geq \rho$ by Definition 11 iff it has distance $\geq \rho$ by Definition 7.*

The proof is given in Appendix E. A code C on n qudits is said to have a constant relative distance $\delta > 0$ if its distance is at least δn. We will make use of the following assumption, which we isolate so that we can refer to it later on.

Remark. If there is a qudit q such that all states in the code look like $|\alpha\rangle$ tensor with some state on the remaining qudits, for some fixed one-qudit state $|\alpha\rangle$ of that qudit q, we say that q is *trivial* for the code. We will assume in the remainder of the paper that for all codes we handle, no qudits are trivial for the code, since such qudits can be simply discarded.

2.5. Interaction graphs and their expansion. We assume in the rest of the paper that each qudit participates in exactly D_L constraints. We define bipartite expanders similar to [52, 22], who used them to construct locally testable classical codes. Note that we require expansion to hold only for sets of constant size k.

DEFINITION 12 (bipartite interaction graph). *Let C be a quantum code on n d-dimensional qudits, whose check terms $\{\Pi_i\}_i$ are k-local. We define the bipartite interaction graph of C $G = G(C) = (L, R; E)$ as follows: the nodes L correspond to the qudits, the nodes R correspond to the check terms, and the set of edges connect each constraint $\Pi_i \in R$ to all the qudits in L on which it acts nontrivially. We note that G is left D_L-regular and right k-regular.*

DEFINITION 13 (bipartite expansion). *Let $G = (L, R; E)$ be a bipartite graph that is left D_L-regular, right k-regular. A subset of qudits $S \subseteq L$ is said to be ε-expanding if $|\Gamma(S)| \geq |S| D_L(1 - \varepsilon)$, where $\Gamma(S)$ is the set of neighbors of S in this graph. ε is called the expansion error for this set. G is said to be ε-local expanding if every subset $S \subseteq L$, $|S| \leq k$ has expansion error at most ε.*

We state two technical facts on good bipartite expanders that will be useful later on. The proofs are in Appendix B.

FACT 2. *Consider $S \subseteq L$ in a bipartite graph $G(L, R : E)$, and let S be ε-expanding for $\varepsilon < \frac{1}{2}$. Then a fraction at most 2ε of all vertices of $\Gamma(S)$ have degree strictly larger than 1 in S.*

FACT 3. *Let $S \subseteq L$ in a bipartite graph $G = (L, R; E)$, such that S is ε-expanding for $\varepsilon < \frac{1}{2}$. Then there exists a vertex $q \in S$, such that the fraction of neighbors of q with at least two neighbors in S is at most 2ε.*

3. Locally testable quantum codes. In this section we define locally testable quantum codes, both in the general case and in the specific case of stabilizer codes. We then show that our definitions coincide for stabilizer codes.

3.1. Local testability of general quantum codes. We first generalize Definition 7 from a definition of distance *of a code* to a definition of distance *from a code*.

DEFINITION 14 (distance from a quantum code). *Let C be a quantum code on a Hilbert space $\mathcal{H}$ of n qubits. For any two orthogonal states $|\phi\rangle, |\psi\rangle \in \mathcal{H}$, we define the Hamming distance between them $\delta(|\phi\rangle, |\psi\rangle)$ as the maximal integer ρ, such that for any $E \in \Pi_d^n$, with $wt(E) \leq \rho - 1$, we have $\langle\psi|E|\phi\rangle = 0$. Similarly, given a state $|\phi\rangle$ orthogonal to C, we say that the distance of $|\phi\rangle$ from C denoted by $\delta(|\phi\rangle, C)$ is the minimum over all $|\psi\rangle \in C$ of $\delta(|\phi\rangle, |\psi\rangle)$.*

We note here that the distance of a state *from the code* in the above can be larger than the distance *of the code*. This is akin to the classical case, where locally testable codes are required to identify words far from the code, even if they cannot be (uniquely) decoded, so that these codes can be used as proof systems.

DEFINITION 15 (quantum locally testable codes (qLTCs)). *Let $R = R(\delta)$ be some function $R(\delta) : [0,1] \to [0,1]$; this is called the soundness function. Let C be a quantum code on n d-dimensional qudits, defined as the ground space of $H = \sum_{i=1}^{m} \Pi_C^i$, where Π_C^i are m k-local projections for some constant k. We say that C is quantum locally testable with soundness $R(\delta)$ if*

$$\forall \delta_0 > 0, |\Psi\rangle : \quad \delta(|\Psi\rangle, C) \geq \delta_0 n \Rightarrow \frac{1}{m}\langle\Psi|H|\Psi\rangle \geq R(\delta_0). \tag{6}$$

The query complexity of the code is defined to be k.

DEFINITION 16 (strong qLTC). *We call a qLTC code C strongly qLTC if there exists a constant $R > 0$, such that*

$$\min_{|\psi\rangle \in C^\perp} \frac{\frac{1}{m}\langle\psi|H|\psi\rangle}{\delta(|\psi\rangle, C)} \geq R. \tag{7}$$

One could also be interested in local testability in different regimes of the parameters δ. We say that a code is locally testable in a certain regime of proximities (δ_0, δ_1) if there exists a constant $R > 0$, such that for any $0 < \delta_0 < \delta_1 < 1$,

$$\min_{\delta_0 n \leq \delta(|\psi\rangle, C) \leq \delta_1 n} \frac{\frac{1}{m}\langle\psi|H|\psi\rangle}{\delta(|\psi\rangle, C)} \geq R. \tag{8}$$

3.2. Local testability of quantum stabilizer codes. We now show that local testability defined above (Definition 15) has a natural interpretation in the context of stabilizer codes.

DEFINITION 17 (local testability for stabilizer codes (sLTCs)). *Let $R(\delta)$ be some function $R(\delta) : [0,1] \to [0,1]$. We say that a stabilizer code C on n d-dimensional qudits is an sLTC with query complexity k and soundness $R(\delta)$ if there exists a generating set $\mathcal{G}$ for C, where each element has support k, such that the following holds: for any $E \in \Pi_d^n$ with $wt_{\mathbf{Z}(\mathcal{G})}(E) \geq \delta n$, a uniformly random generator $g \in \mathcal{G}$ does not commute with E with probability at least $R(\delta)$.*

3.2.1. Equivalence of definitions of locally testable codes. We now show that the definition of sLTCs (Definition 17) is in fact a special case of the general qLTCs (Definition 15).

CLAIM 3.

1. *If C is a stabilizer code with generating set $\mathcal{G}$, which is an sLTC with query complexity k and soundness $R(\delta)$, then the set of projections $\{P_g\}_{g\in\mathcal{G}}$, where $I - P_g$ is the projection on the 1-eigenspace of g defines a qLTC with query complexity k and soundness $R(\delta)$.*
2. *If C is a qLTC with query complexity k and soundness $R(\delta)$, defined by a set of projections $\{P_g\}_{g\in\mathcal{G}}$, such that the set $\{I - 2P_g\}_{g\in\mathcal{G}}$ generates an Abelian subgroup of Π_d^n, then C is also an sLTC with query complexity k and soundness $R(\delta)$.*

Proof. **sLTC $\Rightarrow$ qLTC**. By definition of a stabilizer code, for any $|\phi\rangle \in C$, we have $g|\phi\rangle = |\phi\rangle$ for all $g \in \mathcal{G}$, so $P_g|\phi\rangle = 0$ for all $g \in \mathcal{G}$. Next, consider a state $|\phi\rangle$ orthogonal to C, such that $\delta(|\phi\rangle, C) \geq \delta n$. We would now like to show that a projection chosen randomly from $\{P_g\}_{g\in G}$ is violated by $|\phi\rangle$ with probability at least $R(\delta)$. Consider the following orthogonal decomposition of ϕ as implied by Lemma 1:

$$|\phi\rangle = \sum_i \alpha_i|\alpha_i\rangle = \sum_i \alpha_i E_i|\eta_i\rangle, \quad \alpha_i \neq 0, \tag{9}$$

where $E_i \in \Pi_d^n$, $|\eta_i\rangle \in C$, and $E_i|\eta_i\rangle$ are orthogonal. We claim that for each i, $wt_{\mathbf{Z}(\mathcal{G})}(E_i) \geq \delta n$; otherwise, it is easy to see that there exists some $E' \in \Pi_d^n$, $wt(E') < \delta n$, such that for at least one i, we have $E'E_i \in \mathbf{Z}(\mathcal{G})$. Since for any $J \in \mathbf{Z}(\mathcal{G})$, $JC = C$, we have that, alternatively, $E'|\alpha_i\rangle \in C$. Since E' is unitary, and the $|\alpha_i\rangle$'s are orthogonal, then the $E'|\alpha_i\rangle$'s are orthogonal, and thus $E'|\phi\rangle$ has a nonzero projection on C—contrary to the assumption that $\delta(|\phi\rangle, C) \geq \delta n$.

If E_i and $g \in \mathcal{G}$ do not commute, $E_i g = \omega g E_i$ for some $\omega \neq 1$. In particular, $E_i|\eta_i\rangle$ is an ω-eigenstate of g. This means it is orthogonal to the 1-eigenspace of g, and therefore

$$\langle\alpha_i|P_g|\alpha_i\rangle = 1. \tag{10}$$

Yet, by the sLTC property of C, for each i, E_i does not commute with a fraction at least $R(\delta)$ of the generators of $\mathcal{G}$. Thus, a randomly chosen check term is violated by $|\alpha_i\rangle$ with probability at least $R(\delta)$, so

$$\frac{1}{|\mathcal{G}|}\sum_{g\in\mathcal{G}}\langle\alpha_i|P_g|\alpha_i\rangle \geq R(\delta). \tag{11}$$

Since by Lemma 1 the decomposition above coincides with the simultaneous eigenbasis of $\mathcal{G}$, we have

$$\frac{1}{|\mathcal{G}|}\langle\phi|\sum_{g\in\mathcal{G}} P_g|\phi\rangle = \frac{1}{|\mathcal{G}|}\sum_i\sum_{g\in\mathcal{G}}|\alpha_i|^2\langle\alpha_i|P_g|\alpha_i\rangle \geq R(\delta). \tag{12}$$

qLTC $\Rightarrow$ sLTC. First, by definition, the set of states that are in the mutual ground space of the P_g's are stabilized (i.e., eigenvalue 1) with respect to (w.r.t.) the terms $\mathcal{G}$, and vice versa. Now, let $E \in \Pi_d^n$, whose weight modulo $\mathbf{Z}(\mathcal{G})$ is at least δn. Let $|\phi\rangle \in C$ be any code state, and denote $|\psi\rangle = E|\phi\rangle$. We claim that $\delta(|\psi\rangle, C) \geq \delta n$. Otherwise there exists $E' \in \Pi^n$, $wt(E') < \delta n$, such that $E'|\psi\rangle$ has a nonzero projection on C, and hence $E'E|\phi\rangle$ has a nonzero projection on C, so by Lemma 1, we have that $E'EC = C$. Therefore, $E'E$ commutes with all $\mathcal{G}$, and hence

$E'E \in \mathbf{Z}(\mathcal{G})$, which implies that $wt_{\mathbf{Z}(\mathcal{G})}(E) < \delta n$, a contradiction. By the qLTC property of C, we have

$$\langle\psi| \sum_{g\in\mathcal{G}} P_g|\psi\rangle \geq |\mathcal{G}| \cdot R(\delta). \tag{13}$$

Since $|\psi\rangle = E|\phi\rangle$, then for any generator g, $g|\psi\rangle = gE|\phi\rangle = \omega Eg|\phi\rangle = \omega E|\phi\rangle$ for some $\omega \in \mathbf{C}$. So for any $g \in \mathcal{G}$, $|\psi\rangle$ is some eigenstate of g. Hence $|\psi\rangle$ is either in the 1-eigenspace of P_g or in its 0-eigenspace, so by (13) it violates a fraction at least $R(\delta)$ of all generators $\mathcal{G}$. ∎

4. Bound on the soundness of stabilizer LTCs on local expanders. In this section we prove Theorem 1. We define the *relative soundness* formally.

DEFINITION 18 (relative soundness). *Define* $r : [0,1] \mapsto [0,1]$ *as follows:*

$$r(\delta) = R(\delta)/\Theta(\delta),\ \Theta(\delta) \equiv \min\{\delta k, 1\}.$$

Here, we will be interested in $\delta < 1/k$, for which $r(\delta) = R(\delta)/k\delta$.

4.1. A useful fact about restrictions of stabilizers.

DEFINITION 19 (restriction of stabilizers). *For an* $E \in \Pi_d^n$, *let* $E|_q$ *denote the* q*th component of the tensor product* E, *and let* $E|_{-q}$ *denote the tensor product of all terms except the* q*th. Similarly, for a generating set* $\mathcal{G}$, *we denote by* $\mathcal{G}|_q$ *the set* $\{g|_q,\ g \in \mathcal{G}\}$, *and similarly for* $\mathcal{G}|_{-q}$.

We now prove a useful fact: the restrictions to a given qudit q of all the generators of a stabilizer code with absolute distance strictly larger than 1 cannot all commute.

FACT 4. *Let* C *be a stabilizer code with absolute minimal distance strictly larger than* 1. *Then for any qudit* q *and any generator* g *acting on* q, *there exists another generator* $h(q)$ *acting on* q *such that* $[g|_q, h|_q] \neq 0$.

Proof. Assume on the contrary that there is a qudit q and a generator g such that for all other generators h, we have $[g|_q, h|_q] = 0$. Let $Q = g|_q$. We have that $Q' = Q \otimes I_{-q}$, namely, the tensor product with identity on the other qubits, commutes with all $g \in \mathcal{G}$, and thus $Q' \in \mathbf{Z}(\mathcal{G})$. However, Q' cannot be inside $A(\mathcal{G})$, since otherwise q is in some constant state (the 1-eigenvector of Q) $|\alpha\rangle$ for all code states, and thus q is trivial for the code (see remark at the end of subsection 2.4). Hence, $Q' \in \mathbf{Z}(\mathcal{G}) - A(\mathcal{G})$, so the minimal distance of the code by Definition 11 is 1, in contradiction to our assumption. ∎

4.2. Proof of Theorem 1. In the proof we will make use of "sparse" sets of constraints, defined as follows (we later generalize this definition to t-independent sets; see Definition 22).

DEFINITION 20 (1-independent set of constraints). *For a given constraint* u, *consider* $\Gamma^3(u)$, *the set of qudits acted upon by constraints which act on qudits in* u. *A set of constraints* U *is said to be* 1*-independent if, for any two constraints* $u, w \in U$, $\Gamma^3(u) \cap \Gamma^3(w) = \emptyset$.

Proof of Theorem 1.

Generating the error. We want to construct an error $E \in \Pi_d^n$, $wt_{\mathbf{Z}(\mathcal{G})}(E) \geq \delta n$, that will not violate too many constraints in $\mathcal{G}$. Let C be a stabilizer code with a k-local generating set $\mathcal{G}$, such that the bipartite interaction graph of C is an ε-local bipartite expander. Let U be a 1-independent set of constraints of size δn. For values $\delta \leq \frac{1}{k^3 D_L}$ a 1-independent must exist, by a simple greedy algorithm that iteratively discards all constraints intersecting a given constraint. For a given constraint $u \in U$

and $i \in [k]$, let $\alpha_i(u)$ denote the fraction of generators $g \in \mathcal{G}$ that act on a qudit i in u and intersect u in at least one other qudit, out of all generators that act on i. Then for each $u \in U$ we define $q(u)$ to be a qudit of minimal $\alpha_i(u)$ over all $i \in [k]$. Let $T = \{q(u) | u \in U\}$. Let us define an error pattern:

$$E = \bigotimes_{u \in U} u|_{q(u)}. \tag{14}$$

We first note that $E \notin \mathbf{Z}(\mathcal{G})$. This is true by Fact 4: for each qudit q in the support of E, $E|_q$ does not commute with $h|_q$ for some $h \in \mathcal{G}$. But since T is induced by a 1-independent set, h does not touch any other qudit in the support of E except q, so this implies $[h, E] = [h|_q, E|_q] \neq 0$. We will now show that E has large weight modulo $\mathbf{Z}(\mathcal{G})$, but is penalized by a relatively small fraction of $\mathcal{G}$.

Weight analysis. By definition, we have that $wt(E) = |T| = |U| = \delta n$. We claim that

$$wt_{\mathbf{Z}(\mathcal{G})}(E) = |T|. \tag{15}$$

Since δ was chosen to be smaller than half of the minimal distance of the code C, $wt_{\mathbf{Z}(\mathcal{G})}(E) = wt_{\mathcal{G}}(E)$, and so it suffices to lower-bound $wt_{\mathcal{G}}(E)$.

Suppose on the contrary that $wt_{\mathcal{G}}(E) < |T|$. Then there exists $\Delta \in A(\mathcal{G})$, such that $E' = \Delta E$ has $wt(E') < |T|$. Since the weight of E' is strictly smaller than that of E, there must be one qudit q_0 in T, such that on the neighborhood $\mathcal{N}(q_0)$ the weight of E' is strictly smaller than that of E, which is 1; namely, E' must be equal to the identity on all the qudits in the qudit neighborhood of q_0. Here, we have used the fact that the qudit neighborhoods of different qudits in T are nonintersecting. This is true by the fact that the qudits were chosen by picking one qudit from each constraint out of a 1-independent set of constraints (Definition 20). This means that Δ must be equal to the inverse of E on this neighborhood. But this inverse is exactly the following: It is equal to $E|_{q_0}^{-1}$ on q_0, and to the identity on all other qudits in the neighborhood. By construction, $E|_{q_0}$ on q_0 (and therefore also $E^{-1}|_{q_0} = \Delta_{q_0}$) does not commute with $h|_{q_0}$ for some $h \in \mathcal{G}$. Since Δ is an identity on all qudits of h other than q_0, this implies that Δ does not commute with h, in contradiction to the fact that $\Delta \in A(\mathcal{G})$.

Soundness analysis. We upper-bound the number of generators that do not commute with E. For each $u \in U$, the number of generators $g \in \mathcal{G}$ that do not commute with $E|_{q(u)}$ is at most the number of generators that share at least two qudits with u. By Fact 3 there exists a qudit $q \in \Gamma(u)$ such that the fraction of its check terms with at least two qudits in $\Gamma(u)$ is at most 2ε; since we chose $q(u)$ to be the qudit that minimizes that fraction over all qudits on which u acts, we have that for $q(u)$, the fraction of terms acting on it that intersect u with at least 2 qudits is at most 2ε. Thus, the absolute number of generators acting on $q(u)$ that intersect u in at least two qudits is at most $2\varepsilon D_L$. Hence the overall number of generators violated by E is at most $2\varepsilon |T| D_L$. By (15) this is equal to $2\varepsilon D_L wt_{\mathbf{Z}(\mathcal{G})}(E)$. Using $D_L n = mk$, we have $R(\delta) \leq 2\varepsilon k \delta$ and so $r(\delta) \leq 2\varepsilon$. ∎

We now show that a slightly stronger version of the above theorem holds. This version will be used for showing Theorem 2.

Claim 4. *Let C be a stabilizer code, with a k-local succinct generating set, where each qubit is examined by D_L constraints. If there exists a 1-independent set of constraints $U \subseteq R$, such that $|U| = \delta n$ for some $0 < \delta < 1/k$ and $\Gamma(U)$, the set of qudits*

that the constraints in U act on, satisfies $|\Gamma(\Gamma(U))| \geq |\Gamma(U)|D_L(1-\varepsilon)$, then for any $\delta' \leq \delta$ we have that $r(\delta') \leq 2\varepsilon$.

Proof. For a set $S \subseteq L$, let $\Gamma_1(S)$ denote the number of neighbors of S having a single neighbor in S, and let $\Gamma_{\geq 2}(S) \equiv \Gamma(S) - \Gamma_1(S)$. Put $S = \Gamma(U)$, and let $S = \bigsqcup_{i=1}^{k} S_i$ denote a partition of S into k disjoint sets, where each S_i takes a single (arbitrary) qubit from each $\Gamma(u)$, $u \in U$. By assumption, $|\Gamma(S)| \geq |S|D_L(1-\varepsilon)$, whereas the total degree of S is $|S|D_L$. By an argument similar to the pigeonhole principle, $|\Gamma_{\geq 2}(S)| \leq |S|D_L\varepsilon$, so $|\Gamma_1(S)| \geq |S|D_L(1-2\varepsilon)$. Since each unique neighbor of S examines exactly one partition S_j, there exists a partition S_0 examined by at least $|S_0|D_L(1-2\varepsilon) = \delta n D_L(1-2\varepsilon)$ constraints from $\Gamma_1(S)$.

Now, given any $\delta' \leq \delta$, let S_0' be a subset of S_0 of size $\delta' n$, maximizing the ratio $\Gamma_1(S')/|S'|$ over all sets $S' \subseteq S_0$ of this size. Since each element of $\Gamma_1(S)$ examines just one element of S, such a set exists, with ratio at least $D_L(1-2\varepsilon)$. A tensor-product error $\mathcal{E}$ defined by taking, for each $u \in U$, the restriction to its qubit in S_0', we have by the same arguments leading to (15) that $wt_{\mathbf{Z}(\mathcal{G})}(E) = \delta' n$, whereas the maximal penalty is at most $2\varepsilon D_L \delta' n$, because the penalty arises only from nonunique neighbors. Therefore $R(\delta') \leq \frac{1}{m} 2\varepsilon D_L \delta' n = 2\varepsilon k \delta'$, and since $\delta' \leq \delta < 1/k$ it follows that $r(\delta') \leq 2\varepsilon$. ∎

5. An upper bound on soundness. We now show an absolute constant strictly less than 1, upper-bounding the relative soundness of any good quantum stabilizer code generated by k-local generators, whose qudits are acted upon by D_L stabilizers each. We start with an easy alphabet-based upper bound.

5.1. Alphabet-based bound on soundness. In attempting to understand soundness of good stabilizer codes, one must first account for limitations on the soundness that seem almost trivial and occur even when there is just a single error.

DEFINITION 21 (single-error soundness). *Let $t(d) = 1/(d+1)$. For prime integer d, the single error relative soundness in dimension d is defined to be $\alpha(d) = 1 - t(d)$.*

The motivation for the above definition is as follows. For any qudit q, there always exists $Q \in \Pi_d$, $Q \neq I$, such that a fraction at least $t(d)$ of the generators touching q commute to Q when restricted to q. If we consider a single-qudit error on q to be equal to Q, then it would commute with $t(d)$ of the generators acting on q; thus they can violate at most $\alpha(d)$ of the constraints acting on q. Hence, one can expect that it is possible to construct an error of linear weight whose relative soundness $r(\delta)$ is bounded by the single-error relative soundness using qudits whose neighboring constraints are far from each other. Indeed, we show the following fact.

FACT 5 (alphabet bound on soundness). *Let d be some prime number. For any stabilizer code C on n d-dimensional qudits with a k-local succinct generating set $\mathcal{G}$, whose left-degree is D_L, and relative minimal distance δ_{min}, we have $r(\delta) \leq \alpha(d)$ for any $\delta \leq \min\left\{1/(k^3 D_L), \frac{1}{2}\delta_{min}\right\}$.*

Proof. We first proof the theorem for $d = 2$. Similarly to Theorem 1, given the parameters assumed in the statement of the fact, there exists a 1-independent set of constraints U of size δn. For each constraint $u \in U$ we select arbitrarily some qubit $q = q(u) \in \Gamma(u)$ and examine the restrictions to q of all stabilizers acting nontrivially on q. Let $P(q)$ denote the set of all such restrictions. Let $MAJ(q)$ denote the element of Π_d that appears a maximal number of times in $P(q)$. We then set $E = \bigotimes_{u \in U} MAJ(q(u))$. We first realize that E is an error: we want to show that there exists a generator g such that E and g do not commute. Otherwise, E commutes with all generators. Since by construction, each generator intersects E with at most

one qudit, this means that the restrictions to q also commute: $[E|_q, g|_q] = 0$ for all $q(u)$ acted upon by E. This is a contradiction by Fact 4; hence, there must be a generator which does not commute with E, so E is indeed an error. Similarly to the proof of (15) in the proof of Theorem 1, we also have $wt_{\mathbf{Z}(\mathcal{G})}(E) = \delta n$. Furthermore, for each qudit q, the fraction of generators on q whose restriction to q does not commute with $E|_q$ is at most $\alpha(d)$, since the number of appearances of $E|_q = MAJ(q)$ in $P(q)$ is at least $t(d) = 1 - \alpha(d)$. Hence the number of violated constraints is at most $\alpha(d) \cdot |U| \cdot D_L = \alpha(d)\delta n D_L$. Since $\delta < 1/k$ it follows that $r(\delta) \leq \alpha(d)$.

Consider the case of $d > 2$ for prime d. In this case, for each Pauli $\mathcal{E} = X_d^a \cdot P_d^b \in \Pi_d^n$, where at least one of a, b is nonzero, we have that the $d-1$ unique powers of $\mathcal{E}$ of the form $\mathcal{E}^t$, $t \in [1, \ldots, d-1]$, commute with $\mathcal{E}$. Hence, the $d^2 - 1$ nonidentity elements of Π_d^n can be partitioned into $d+1$ equivalence classes $S_1, \ldots, S_{d+1}$, where in each class all terms commute. In this case, the error $\mathcal{E}$ defined on the set U is computed by taking, for each $u \in U$, some representative Pauli term $\mathcal{E}_i \in S_i$ if S_i is the equivalence class that appears the largest number of times in the restriction of all generators to u, out of all equivalence classes $S_1, \ldots S_{d+1}$. In such a case, it follows that $r(\delta) \leq 1 - \frac{1}{d+1} = \alpha(d)$. ∎

We will assume from this point onward that d is prime.[8]

5.2. Separation from alphabet-based soundness. In this section we show that the alphabet-based bound on the relative soundness in fact cannot be achieved, and the relative soundness is further bounded by a constant factor strictly less than 1, which is due to what seems to be an inherently quantum phenomenon. We will use the geometry of the underlying interaction graph to achieve this separation by treating expanding instances and nonexpanding instances differently. Before stating the main theorem of this section, we require a generalization of Definition 20 and a simple fact.

DEFINITION 22 (t-independent set of constraints). *Let C be a quantum code with a set of k-local constraints, whose underlying bipartite graph is $G(C) = (L, R; E)$. A set of constraints $U \subseteq R$ is said to be t-independent if for any $a, b \in U$ we have $\Gamma^{(2t+1)}(a) \cap \Gamma^{(2t+1)}(b) = \emptyset$.*

The following fact can be easily derived by a greedy algorithm.

FACT 6. *Let $\eta = \eta(k, D_L) = k^{-(2k+1)} D_L^{-(2k-1)}$. For any quantum code C whose bipartite graph $G(C)$ is left D_L-regular and right k-regular, there exists a k-independent set of size at least ηn.*

Proof. Pick a constraint u, remove all constraints in $\Gamma^{(4k)}(u)$, and repeat. The number of constraints we have removed for each constraint is $(kD_L)^{2k}$. Hence, we can proceed for $m/(kD_L)^{2k}$ steps. We get that the fraction of constraints is at least $k^{-(2k)} D_L^{-(2k)}$, and since $mk = nD_L$, we get the desired result. ∎

THEOREM 2. *Let C be a stabilizer code on n d-dimensional qudits, with $\delta_{min} = \Omega(n^{-\varepsilon})$ for some constant $\varepsilon \in [0, 1)$, and a k-local ($k \geq 4$) succinct generating set $\mathcal{G} \subset \Pi_d^n$, where the right degree of the interaction graph of $\mathcal{G}$ is D_L. Then there exists a function $\gamma_{gap} = \gamma_{gap}(k) > \min\left\{10^{-3}, 0.01/k\right\}$ such that for any $\delta \leq \min\{\delta_{min}/2, \eta/10\}$ (for η as defined in Fact 6) we have $r(\delta') \leq \alpha(d)\,(1 - \gamma_{gap})$ for some $\delta' \in (0.99\delta, 1.01\delta)$.*

[8] For composite d, the proof of Fact 5 does not hold. Using Bézout's identity, one can prove a version of Fact 5 with a weaker bound, $\alpha(d) = 1 - \theta(1/d^2)$. The results in the remainder of the paper, and in particular Theorem 2, also need to be adapted for composite d. The analysis is somewhat cumbersome, and we omit it for simplicity.

The proof of the theorem will use, on one hand, Claim 4, which upper-bounds the soundness of expanding instances, and on the other hand a lemma on nonexpanding instances which tries to "mimic" the behavior of the classical setting, in which nonexpanding topologies suffer from poor soundness. We now state this lemma.

LEMMA 2. *Let C be a stabilizer code on n qudits of dimension d, with $\delta_{min} = \Omega(n^{-\varepsilon})$ for some constant $\varepsilon \in [0,1)$, and a k-local ($k \geq 4$) succinct generating set $\mathcal{G}$, where the left degree of the interaction graph of $\mathcal{G}$ is D_L. Let $\gamma_{gap} = \gamma_{gap}(k) = \min\{10^{-3}, 0.01/k\}$. If there exists a k-independent set U of size $|U| = \delta n$, with $\delta < \frac{1}{2}\delta_{min}$, such that the bipartite expansion error of $\Gamma(U)$ is at least $\varepsilon = 0.32$, i.e., $|\Gamma(\Gamma(U))| = |\Gamma(U)|D_L(1-\varepsilon')$ for some $\varepsilon' \geq 0.32$, then*

$$r(\delta') \leq \alpha(d) \cdot (1 - \gamma_{gap}) \tag{16}$$

for some $\delta' \in (0.099\delta, 0.101\delta)$.

The proof of this lemma is technically nontrivial, and we defer it to a separate section. From this lemma, it is easy to show Theorem 2.

Proof of Theorem 2. The parameters of the theorem allow us to apply directly Fact 6; hence there exists a k-independent set S of size at least ηn for η as defined in Fact 6. Since $\delta \leq \eta/10$ there exists a k-independent set S of size 10δ. Now, one of the following holds:

1. S has expansion error at least 0.32. By Lemma 2, we have
$$r(\delta_0) < \alpha(d)(1 - \gamma_{gap}) \tag{17}$$
for some $\delta_0 \in (0.099 \cdot (10\delta), 0.101 \cdot (10\delta)) = (0.99\delta, 1.01\delta)$, and $\gamma_{gap}(k)$ from Lemma 2, which is at least $\min\{10^{-3}, 0.01/k\}$.
2. The set S is ε-expanding for $\varepsilon < 0.32$, in which case, since S is in particular R-independent, then by Claim 4, the soundness function is upper-bounded by $r(\delta') \leq 2\varepsilon < 2/3 - 0.01 \leq \alpha(d) - 0.01$ for all $\delta' \leq |S|/n$. In particular $r(\delta_0) < \alpha(d)(1 - 0.01/k)$.

Taking the higher of these two bounds, we get the desired upper bound for $r(\delta_0)$. □

5.3. Proof of Lemma 2. In the following we first define the error. We provide the proof that the expected penalty of this error is small in Fact (7), then state and prove the onion fact in section 5.3.3 and use it to prove Fact 9, in which we show that the error has large weight modulo the group. Finally we combine all the above to finish the proof of the lemma.

5.3.1. Constructing the error. Let $U \subseteq R$ be a k-independent set as promised by the conditions of the lemma. Then $|U| = \delta n$, and denoting $S = \Gamma(U)$, we have that $|S| = \delta nk$. Therefore, $|\Gamma(S)| = |S|D_L(1-\varepsilon')$ for some $\varepsilon' \geq 0.32$. Let $\mathcal{E}$ be the following random error process: for each qudit of S independently, we apply I with probability $1-p$ for $p = 1/(10k)$, and one of the other elements of Π_d with equal probability $p \cdot t(d)$, where t is defined in Definition 21.

$$\mathcal{E} = \bigotimes_{i \in S} \mathcal{E}_i, \text{ where } \mathcal{E}_i = \begin{cases} I_i & \text{with probability } 1 - 1/(10k), \\ X_d^k P_d^l & \text{with probability } t/(10k). \end{cases} \tag{18}$$

We note here that the choice of p is such that, on average, each k-tuple has only a small number of errors; the expectation of the number of errors is an absolute constant $1/10$ (not a fraction of k). This will help, later on, to lower-bound the weight of the error modulo the group.

5.3.2. Analyzing penalty. We first claim that on average $\mathcal{E}$ has a relatively small penalty w.r.t. $\mathcal{G}$, using the fact that the expansion error is at least 0.32 as in the condition of Lemma 2. For any $\mathcal{E}$, let penalty($\mathcal{E}$) denote the number of generators of $\mathcal{G}$ that do not commute with $\mathcal{E}$.

FACT 7.

$$\mathbf{E}_{\mathcal{E}}\left[\text{penalty}(\mathcal{E})\right] \leq p\alpha|S|D_L\left(1-0.02/k\right). \tag{19}$$

Proof. Let $G=(L,R;E)$ denote the bipartite graph corresponding to $\mathcal{G}$, with R being the generators of $\mathcal{G}$ and L the qudits. Let $S=\Gamma(U)$ be as before. Let the error process $\mathcal{E}$ be the one defined above. For any constraint $c\in\Gamma(S)$ which is violated when applied to this error, observe that there must be a qudit $i\in\text{supp(c)}$ such that $[c|_i,\mathcal{E}_i]\neq 0$. We now would like to bound the number of constraints violated by $\mathcal{E}$ using this observation and linearity of expectation.

For an edge $e\in E$ connecting a qudit i in S and a constraint c in $\Gamma(S)$, let $x(e)$ denote the binary variable which is 1 iff the error term $\mathcal{E}_i$ does not commute with $c|_i$. In other words, an edge marked by 1 is an edge whose qudit causes its constraint to be violated. By construction, for each $e\in E$ which connects the qudit i and the constraint c we have

$$\mathbf{E}_{\mathcal{E}}[x(e)] = p(1-t). \tag{20}$$

This is true since a constraint c restricted to the qudit i, $c|_i$, does not commute with the error restricted to the same qudit i, $\mathcal{E}_i$, iff $\mathcal{E}_i$ is both a nonidentity (which happens with probability p) and not equal to $c|_i$.

If we had just now added $x(e)$ over all edges going out of S (whose number is $|S|D_L$), then by linearity of expectation this would have given an upper bound on the expected number of violated constraints equal to

$$\sum_e p(1-t) = p|S|D_L\alpha(d). \tag{21}$$

Unfortunately this upper bound does not suffice; to strengthen it we would now like to take advantage of the fact that many of those edges go to the same constraint, due to the fact that the expansion is bad; thus, instead of simply summing these expectation values, we take advantage of the fact that two qudits touching the same constraint cannot contribute twice to its violation. Observe that it may even be the case that some edges may cause constraints to become "unviolated," so the actual bound may be even lower.

Let $E_{inj}\subseteq E$ be an arbitrary subset of the edges between S and $\Gamma(S)$ chosen by picking a single edge for each constraint in $\Gamma(S)$. For an edge $e\in E$ let $c(e)$ denote the constraint incident on e, and let $e_{inj}(c(e))$ denote the edge in E_{inj} that is connected to $c(e)$.

We now bound the expectation by subtracting $x(e)$ from the sum if the Boolean variable $x(e_{inj}(c(e)))$ is 1; this avoids counting the violation of the same constraint twice due to the two edges. We have

$$\mathbf{E}_{\mathcal{E}}\left[\text{penalty}\right] \leq \mathbf{E}_{\mathcal{E}}\left[\sum_{e\in E_{inj}} x(e) + \sum_{e\notin E_{inj}} \left(1-x(e_{inj}(c(e)))\right)\cdot x(e)\right]. \tag{22}$$

Expanding the above by linearity of expectation gives us

$$\mathbf{E}\left[\text{penalty}\right] \leq \sum_{e\in E_{inj}} \mathbf{E}_{\mathcal{E}}\left[x(e)\right] + \sum_{e\notin E_{inj}} \mathbf{E}_{\mathcal{E}}\left[x(e)\right] - \sum_{e\notin E_{inj}} \mathbf{E}_{\mathcal{E}}\left[x(e_{inj}(c(e)))\cdot x(e)\right] \tag{23}$$

$$= \sum_{e\in E} \mathbf{E}_{\mathcal{E}}\left[x(e)\right] - \sum_{e\notin E_{inj}} \mathbf{E}_{\mathcal{E}}\left[x(e_{inj}(c(e)))\cdot x(e)\right]. \tag{24}$$

The first summand in (24) was already computed in (21). We now lower-bound the correction given by the second term, using the fact that for any $e \notin E_{inj}$

$$\mathbf{E}_{\mathcal{E}}\left[x(e_{inj}(c(e))x(e)\right] = \mathbf{E}_{\mathcal{E}}[x(e_{inj}(c(e)))]\mathbf{E}_{\mathcal{E}}[x(e)], \tag{25}$$

since $\mathcal{E}$ is independent between different qudits. We can thus substitute (20) in (24) and get

$$\mathbf{E}_{\mathcal{E}}\left[\text{penalty}\right] \leq p\alpha|S|D_L - |S|D_L\varepsilon(p\alpha)^2, \tag{26}$$

where we have used the fact that $|E\backslash E_{inj}| = |S|D_L\varepsilon$. This is equal to

$$p\alpha|S|D_L(1-p\alpha\varepsilon). \tag{27}$$

Using $p = 1/(10k), \varepsilon \geq 0.32$, $\alpha(d) \geq 2/3$, we get the desired bound. ∎

5.3.3. The onion fact.

FACT 8 (onion fact). *Let C be a stabilizer code on n qudits with a succinct generating set $\mathcal{G}$ of locality k, such that $\delta_{min}(C) \geq k$. Let $E \in \Pi_d^n$ such that* $\text{supp}(E) \subseteq \Gamma(u)$ *for some generator $u \in \mathcal{G}$. Finally let $\Delta \in A(\mathcal{G})$, and let $E_{\mathcal{G}} = \Delta \cdot E$. Then, for any $i \in [k]$, if $wt(E|_{\Gamma(u)}) = i$, then $wt(E_{\mathcal{G}}|_{\Gamma^{(2k+1)}(u)}) \geq \min\{i, k-i\}$.*

Proof. If $\Delta|_{\Gamma(u)} = I$, then

$$wt\left(E_{\mathcal{G}}|_{\Gamma^{(2k+1)}(u)}\right) \geq wt\left(E_{\mathcal{G}}|_{\Gamma(u)}\right) = wt\left(E|_{\Gamma(u)}\right) = i, \tag{28}$$

so in this case we are done.

Otherwise, $\Delta|_{\Gamma(u)}$ is a nonidentity and so has at least one nonidentity coordinate. Since Δ is a nonidentity, by the assumption on the succinctness of $\mathcal{G}$ we have $wt(\Delta) \geq k$.

Moreover, we claim that $wt\left(\Delta|_{\Gamma^{(2k+1)}(u)}\right) \geq k$. Otherwise, consider the following process. Start with the generator u, and consider the qudits in $\Gamma(u)$. Now add the qudits in $\Gamma^{(3)}(u)$ (namely, the qudits that are acted upon by generators intersecting u). Then add the next level, and so on for k levels, at which point we have added all qudits belonging to $\Gamma^{(2k+1)}(u)$. By the pigeonhole principle, if $wt\left(\Delta|_{\Gamma^{(2k+1)}(u)}\right) < k$, then there must exist a level t, $1 \leq t \leq k$, such that Δ has zero support on qudits added in this level.

We now claim that $\tilde{\Delta} = \Delta|_{\Gamma^{(2(t-1)+1)}(u)}$ is in the centralizer $\mathbf{Z}(\mathcal{G})$ but its weight is less than k. This, together with the fact that $\tilde{\Delta} \notin A(\mathcal{G})$, shown in the next paragraph, contradicts the assumption that $\delta(C) \geq k$. To see that $\tilde{\Delta}$ is in the centralizer, we observe first that Δ commutes with all elements of $\mathcal{G}$ that act only on qudits in $\Gamma^{(2t-1)}(u)$, and since $\tilde{\Delta}$ agrees with Δ on $\Gamma^{(2t-1)}(u)$, $\tilde{\Delta}$ also commutes with them. We also observe that $\tilde{\Delta}$ trivially commutes with all elements in $\mathcal{G}$ whose support does not intersect $\Gamma^{(2t-1)}(u)$. Hence we only need to worry about those terms that act on

at least one qudit in $\Gamma^{(2t+1)}(u) - \Gamma^{(2t-1)}(u)$ and at least one qudit in $\Gamma^{(2t-1)}(u)$. Let v be some such term. Note that v does not act on any qudit outside $\Gamma^{(2t+1)}(u)$ by definition. We know that Δ commutes with v. But by the choice of t, we know that Δ is trivial on those qudits added at the tth level, and hence Δ restricted to $\Gamma^{(2t+1)}(u)$ (which contains the qudits of v) is the same as Δ restricted to $\Gamma^{(2t-1)}(u)$. And so Δ restricted to $\Gamma^{(2t-1)}(u)$ commutes with v.

We showed that $\tilde{\Delta}$ is in $\mathbf{Z}(\mathcal{G})$. If it also belongs to $A(\mathcal{G})$, this contradicts succinctness of $\mathcal{G}$; otherwise it is in $\mathbf{Z}(\mathcal{G}) - A(\mathcal{G})$, implying the distance of C is at most $k-1$, contrary to assumption. This means that $wt\left(\Delta|_{\Gamma^{(2k+1)}(u)}\right) \geq k$. Therefore, we now know by the triangle inequality on the Hamming distance that

$$\begin{aligned} wt\left(E_{\mathcal{G}}|_{\Gamma^{(2k+1)}(u)}\right) &\geq wt\left(\Delta|_{\Gamma^{(2k+1)}(u)}\right) - wt\left(E|_{\Gamma^{(2k+1)}(u)}\right) \\ &= wt\left(\Delta|_{\Gamma^{(2k+1)}(u)}\right) - wt\left(E|_{\Gamma(u)}\right) \geq k - i. \end{aligned} \tag{29}$$

Taking the minimal of the bounds from (28), (29) completes the proof. ∎

5.3.4. Analyzing error weight. First, we consider the case that $\delta_{min} = \Omega(1)$, and hence $|S| = \Omega(n)$. We note that the expected weight of $\mathcal{E}$ is $p|S|$, and since $|S|$ is linear in n, by Chernoff the probability that the weight of $\mathcal{E}$ is smaller by more than a constant fraction than this expectation is $2^{-\Omega(n)}$. We need to show a similar bound on the weight modulo the centralizer group; given that $\delta < \delta_{min}/2$ we only need to bound the weight modulo $A(\mathcal{G})$. Let $\Delta \in A$ be some element in the stabilizer group and let $\mathcal{E}_{\mathcal{G}} = \Delta \cdot \mathcal{E}$. We now need to lower-bound $wt(\mathcal{E}_{\mathcal{G}})$.

FACT 9. *For integer k, let $\hat{k} = \lfloor k/2 \rfloor + 1$. Let $y(k) : [4, \infty] \to \mathbf{R}$ be the function*

$$y(k) = \begin{cases} 1 - 2^{(-\hat{k}+1)log(k)+k-2.3\hat{k}+4.54}, & k \geq 12, \\ 0.9999, & 6 \leq k \leq 11, \\ 0.9992, & k = 5, \\ 0.9985, & k = 4. \end{cases} \tag{30}$$

We claim

$$\mathrm{Prob}_{\mathcal{E}}\left(wt(\mathcal{E}_{\mathcal{G}}) < |S|py(k)\right) = 2^{-\Omega(n)}. \tag{31}$$

Proof (sketch; the detailed proof can be found in Appendix G). The proof builds on the onion fact (Fact 8) as follows: The onion fact shows that "islands" with fewer than $k/2$ errors cannot "lose" error weight modulo the centralizer of $\mathcal{G}$. The proof uses standard probabilistic arguments to argue that the random error pattern we chose is such that the vast majority of islands have fewer than this threshold error weight, and so the overall error weight is virtually unharmed.

Then we claim that the proof of Fact 9 can be easily extended to the case where $\delta_{min} = \Omega(n^{-\varepsilon})$ for some $\varepsilon \in [0, 1)$. This is because the Chernoff bound used in that proof is sufficiently strong to retain the same asymptotic estimates. ∎

5.3.5. Concluding the proof of Lemma 2.

Proof. By Fact 7 the average penalty of $\mathcal{E}$ is small, i.e.,

$$\mathbf{E}\left[Penalty(\mathcal{E})\right] \leq |S|D_L p\alpha(1 - 0.02/k) \triangleq P. \tag{32}$$

Yet, by Fact 9 with probability exponentially close to 1, we have

$$wt(\mathcal{E}_{\mathcal{G}}) \geq |S|py(k) \triangleq W_{low} \geq |S|p \cdot 0.99. \tag{33}$$

Similarly, by the Hoeffding bound with probability exponentially close to 1, we have

$$wt(\mathcal{E}_{\mathcal{G}}) < |S|p(1+0.01) \triangleq W_{high}. \tag{34}$$

Since all penalties are nonnegative, we conclude that *conditioned* on $|wt(\mathcal{E}_{\mathcal{G}})/(|S|p) - 1| < 0.01$, we have $\mathbf{E}\left[Penalty(\mathcal{E})\right] \leq P + 2^{-\Omega(n)}$. Therefore, there must exist an error $\mathcal{E}$ whose weight modulo $\mathcal{G}$ deviates by a fraction at most 0.01 from $|S|p$, and whose penalty is at most $P + 2^{-\Omega(n)}$.

We would like to bound the soundness of this error, which is the ratio of the penalty to its relative weight times D_L. We get that its soundness is at most

$$r = \frac{P + 2^{-\Omega(n)}}{D_L W_{low}} \leq \frac{1}{D_L} \cdot \frac{|S| D_L p \alpha (1 - 0.019/k)}{|S| p y(k)} = \alpha \left(\frac{1 - 0.019/k}{y(k)} \right). \tag{35}$$

We now note that in the last expression, for all $k \geq 12$, the ratio $\frac{1-0.019/k}{y(k)}$ is at most $1 - 0.01/k$. For all values of $4 \leq k < 12$ we substitute the appropriate value of $y(k)$ and get similarly that the ratio $\frac{1-0.019/k}{y(k)}$ is at most $1 - 10^{-3}$. Hence, the soundness of the error r is at most $\alpha(d)(1 - \gamma_{gap})$, where γ_{gap} is as defined in the statement of Lemma 2. ∎

Appendix A. Proof of Claim 1. We prove that Definitions 6 and 5 are equivalent.

Proof. If Definition 5 holds, then for any $E \in \mathcal{E}$ and any two orthogonal states of the code $|\phi\rangle, |\psi\rangle$, we have

$$\langle \phi | E | \psi \rangle = \langle \phi | \Pi_C E \Pi_C | \psi \rangle = \gamma_E \langle \phi | \Pi_C | \psi \rangle = \gamma_E \langle \phi | \psi \rangle = \gamma_E 0 = 0. \tag{36}$$

On the other hand, suppose that for any two orthogonal states $|\phi\rangle, |\psi\rangle$ in the code, and any $E \in \mathcal{E}$, we have $\langle \phi | E | \psi \rangle = 0$. Choose some orthogonal basis of the code C $\{|b_i\rangle\}_{i=1}^m$. Then for each of these basis vectors, we have $\langle b_i | E | b_j \rangle = 0$ for $i \neq j$. Hence, in particular, the operator $E|_C$, i.e., E restricted to C, is a diagonal matrix $\mathrm{diag}(\lambda_1, \ldots, \lambda_m)$. We claim further that $E|_C = \gamma_E I$ for some constant γ_E, and hence $\Pi_C E \Pi_C = \gamma_E \Pi_C$. Suppose, on the contrary that there exist two eigenvalues of $E|_C$ that are different, say $\lambda_1 \neq \lambda_2$. Consider the orthogonal states $|\phi\rangle = \frac{1}{\sqrt{2}}(|b_1\rangle + |b_2\rangle)$, $|\psi\rangle = \frac{1}{\sqrt{2}}(|b_1\rangle - |b_2\rangle)$. Then $|\phi\rangle, |\psi\rangle$ are in the code by linear closure, and are orthogonal, and yet

$$\langle \phi | E | \psi \rangle = \frac{1}{2}\langle b_1 | E | b_1 \rangle - \frac{1}{2}\langle b_1 | E | b_2 \rangle + \frac{1}{2}\langle b_2 | E | b_1 \rangle - \frac{1}{2}\langle b_2 | E | b_2 \rangle = \frac{1}{2}(\lambda_1 - \lambda_2) \neq 0, \tag{37}$$

contrary to our assumption on E. ∎

Appendix B. Proofs of geometrical facts on local expanders.

B.1. Proof of Fact 2. For $S \subseteq R$ let $\Gamma_1(S) \subseteq \Gamma(S)$ denote the subset of the neighbors of S with exactly one neighbor in S. Similarly, let $\Gamma_{\geq 2}(S)$ denote the subset of neighbors with at least two neighbors in S.

Proof. The average degree of a vertex in $\Gamma(S)$ w.r.t. $|S|$ is at most $\frac{D_L S}{D_L S(1-\varepsilon)} = \frac{1}{1-\varepsilon}$. Let α_1 denote the fraction $|\Gamma_1(S)|/|\Gamma(S)|$, where $\Gamma_1(S)$ is the set of neighbors of S with degree exactly 1 with respect to S. Then

$$\frac{1}{1-\varepsilon} \geq \alpha_1 1 + (1 - \alpha_1)m, \tag{38}$$

where m is the average degree of a vertex with at least two neighbors in S. Then by simple algebra

$$\alpha_1(m) \geq 1 - \frac{1}{m-1} \cdot \frac{\varepsilon}{1-\varepsilon}, \tag{39}$$

so $\alpha_1(m)$ is a monotonously increasing function of m, and since $m \geq 2$, then α_1 is minimized for $m = 2$. Hence,

$$\alpha_1 \geq 1 - \frac{\varepsilon}{1-\varepsilon}, \tag{40}$$

and since $\varepsilon < 1/2$, we have

$$\alpha_1 \geq 1 - \varepsilon(1+2\varepsilon) \geq 1 - 2\varepsilon. \tag{41}$$ ∎

B.2. Proof of Fact 3. By definition, we have $|\Gamma(S)| \geq |S|D_L(1-\varepsilon)$. Let $E_{inj} \subseteq E(S)$ be a subset of the edges incident on S such that each $u \in \Gamma(S)$ has a single neighbor in S connected by an edge of E_{inj}. Then E_{inj} is of size $|\Gamma(S)|$ which is at least $|S|D_L(1-\varepsilon)$. Also $|E(S)| = |S|D_L$, thus $|E(S) - E_{inj}| \leq |S|D_L\varepsilon$. Therefore $|\Gamma_{\geq 2}(S)| \leq |S|D_L\varepsilon$. Hence, $\Gamma_1(S) = \Gamma(S) - \Gamma_{\geq 2}(S)$ is of size at least $|S|D_L(1-\varepsilon) - |S|D_L\varepsilon = |S|D_L(1-2\varepsilon)$. Therefore, when $\varepsilon < 1/2$ there exists a vertex $v \in S$ with at least $D_L(1-2\varepsilon)$ neighbors in $\Gamma_1(S)$. Since v has D_L neighbors in $\Gamma(S)$, then the fraction of neighbors of v with at least two neighbors in S is at most 2ε when $\varepsilon < \frac{1}{2}$. ∎

Appendix C. Existence of arbitrarily sound classical LTCs on local expanders.

CLAIM 5. *For any $\varepsilon \in (0, 1/2)$, $\rho \in (0,1)$ there exists $\delta = \delta(\rho,\varepsilon)$, and an explicit infinite family of codes $\{C_\varepsilon(n)\}_{n\in\mathbf{N}}$ of n bits, of constant fractional rate ρ, and constant fractional distance $\delta_{min} = \delta_{min}(\rho,\varepsilon)$, whose check terms are $O(1)$-local on average, and all errors of weight less than δn have soundness $r(\delta) \geq 1 - 3\varepsilon$. Moreover, the underlying graph of these codes is an ε-local expander.*

Proof. The construction of [22] generates explicitly for any ε, ρ a left-D_L-regular bipartite graph $G = (L, R; E)$ such that $|R|/|L| = 1 - \rho$, and for any subset $S \subseteq L$, $|S| \leq |L|\delta$ the neighbor set of S is of size at least $|S|D_L(1-\varepsilon)$, where D_L is the left degree of G. Note that since the left degree is D_L, the average right degree is $D_L|L|/|R| = D_L\frac{1}{1-\rho}$, which is a constant, given that D_L is a constant.

The code is defined by assigning to each right node a parity check over its incident vertices. Let us lower-bound the fractional rate of this code: it is at least $\rho = (|L| - |R|)/|L|$, since each constraint in R at most halves the dimension of the code space. The fractional minimal distance of the code is at least δ, since any nonzero word of fractional weight at most δ is rejected, since there exists at least one check term that "sees" just a single bit at state 1, by Fact 2. Hence, these are so-called "good" codes.

Furthermore, their soundness is at least $1 - 3\varepsilon$ since an error on a set of bits S of size $|S| \leq \delta n$ is examined by at least $|S|D_L(1-\varepsilon)$ constraints. By Fact 2 at least $1-2\varepsilon$ of those constraints examine S in exactly one location; all constraints that touch a given error set S in exactly one location will be violated. Hence the total number of constraints that will be violated is at least $|S|D_L(1-\varepsilon)(1-2\varepsilon) \geq |S|D_L(1-3\varepsilon)$. Therefore, the soundness function $R(\delta')$ is at least $(1-3\varepsilon)\delta' k$ for all $\delta' \in [0, \delta]$. ∎

Appendix D. Proof of Lemma 1: Decomposition to cosets of a stabilizer code. For any $E \in \Pi_d^n$ and any $g \in \mathcal{G}$, we have $Eg = \omega gE$, where $\omega \in \mathbf{C}$. Therefore, for any $|\eta\rangle$ in C, we have that $E|\eta\rangle$ is an ω-eigenstate of g. Then for any $E \in \Pi_d^n$, we have that EC is some simultaneous eigenspace of $\mathcal{G}$. But, since Π_d^n spans over $\mathbf{C}$ all unitaries on n qudits, it must be that every simultaneous eigenspace of $\mathcal{G}$ is equal to EC for some $E \in \Pi_d^n$. In particular, any state $|\phi\rangle$ may be written as a sum,

$$|\phi\rangle = \sum_i E_i|\eta_i\rangle, \tag{42}$$

where $E_i \in \Pi_d^n$ and $|\eta_i\rangle \in C$. ∎

Appendix E. Proof of Claim 2: Equivalence of definitions of minimal distance of a code. We prove that a stabilizer code C has $\delta(C) \geq \rho$ by Definition 11 iff it has distance $\geq \rho$ by Definition 7.

If the minimal weight of a Pauli in $\mathbf{Z}(\mathcal{G}) - A(\mathcal{G})$ has weight at least ρ, then all terms $E \in \Pi_d^n$ of weight strictly less than ρ (namely, at most $\rho-1$) are either generated by $\mathcal{G}$ or outside $\mathbf{Z}(\mathcal{G})$. Take any two orthogonal code states $|\phi\rangle, |\psi\rangle$. If $E \in A(\mathcal{G})$, then all code states are stabilized by E, so we have $\langle\phi|E|\psi\rangle = 1 \cdot \langle\phi|\psi\rangle = 0$. If $E \notin \mathbf{Z}(\mathcal{G})$, E does not commute with some generator, so, in particular, E does not preserve the simultaneous 1-eigenspace of all generators, namely, the code. By Lemma 1, this implies that EC is orthogonal to C. Thus we have in this case as well $\langle\phi|E|\psi\rangle = 0$. Hence the minimal distance of the code, according to Definition 7, is at least d.

Proving the converse, assume that $\delta(C) < \rho$, i.e., $\min_{E \in \mathbf{Z}(\mathcal{G}) - A(\mathcal{G})} wt(E) < \rho$. Then there exists $E \in \Pi_d^n$, of weight less than ρ, that commutes with all generators of $\mathcal{G}$ but not generated by them, so there exists some state $|\phi\rangle \in C$, such that $E|\phi\rangle \neq |\phi\rangle$, yet $E|\phi\rangle \in C$ (see [35, p. 27]). Thus, there exists a nonzero projection of $E|\phi\rangle$ on some other code state $|\psi\rangle$ orthogonal to $|\phi\rangle$. Therefore, $\langle\psi|E|\phi\rangle \neq 0$, contrary to Definition 6. ∎

Appendix F. Quantum LTCs built from CSS codes inherit classical parameters: Proof of Fact 1. We define the check matrices H_x, H_z of the quantum CSS code by taking exactly the locally testable check matrices of the corresponding classical codes $\mathcal{L}_x, \mathcal{L}_z$. So the generating group $\mathcal{G}$ of the stabilizer code Q corresponds either to rows of the matrix H_x (translated to Pauli X) or to rows of the matrix H_z (translated to Pauli Z). Hence by definition, the query complexity is at most $\max\{q_1, q_2\}$. We now consider the soundness function:

1. *Lower bound.* Consider an operator $w \in \Pi^n$, with $wt_{C(\mathcal{G})}(w) \geq \delta n$. This implies that the Pauli operator of minimal weight Δ—so that $w \cdot \Delta$ commutes with all of $\mathcal{G}$—has $wt(\Delta) \geq \delta n$. Decomposing such Δ to a product of two tensor products of X and Z operators, Δ_x, Δ_z, implies that of the binary strings that correspond to these operators, at least one of them has distance at least $\frac{1}{2}\delta n$ from $\mathcal{L}_x, \mathcal{L}_z$. By local testability of these codes, this string violates a fraction at least $\min\{R_1(\delta/2), R_2(\delta/2)\}$ of all check terms (H_x, H_z). Hence Δ violates, as a quantum operator, the same fraction of check terms. Therefore $R(\delta) \geq \min\{R_1(\delta/2), R_2(\delta/2)\}$.
2. *Upper bound.* Let us examine, say, $\mathcal{L}_x$. Let w be any word in $\mathcal{F}_2^n$ at distance at least δn from $\mathcal{L}_x$. Let us take the word in $\mathcal{F}_2^{2n}$ comprising $[w, 0]$ and examine its corresponding Pauli word, denoted by w'. (By corresponding here we mean by the usual isomorphism between the n-fold tensor product Pauli group and the $2n$-fold additive group modulo 2.) Then w' is composed

entirely of Pauli X operators and has $wt_{C(\mathcal{G})}(w') \geq \delta n$, i.e., the minimal-weight Pauli operator Δ, such that $\Delta \cdot w' \in C(\mathcal{G})$ is composed entirely of Pauli X operators and has weight exactly $\delta(w, \mathcal{L}_x)$. By quantum local testability of Q we have that w' violates a fraction $R(\delta)$ of the constraints. But w' can only violate constraints of H_x (i.e., Pauli Z constraints), because it is composed entirely of Pauli X operators, thereby trivially commuting with all H_z (Pauli X constraints). This implies that w must violate a fraction at least $R(\delta)$ of the constraints of H_x. Hence $R_1(\delta) \geq R(\delta)$. The same holds by symmetry for $R_2(\delta)$. ∎

Appendix G. Lower bound on weight: Proof of Fact 9. Let $x \sim B(k, p = 1/(10k))$ denote a random variable which is the sum of k independent and identically distributed Boolean variables, each equal to 1 with probability p; in other words, x is a binomial process; $B(i) = \text{Prob}(x = i)$. Let U be a k-independent set of size $\Omega(n)$, and let $\mathcal{E}$ be the error process defined in subsection 5.3.1. Let $U_i = \left\{u \in U | wt(\mathcal{E}|_{\Gamma(u)}) = i\right\}$ be the set of generators which have exactly i erroneous qudits. Using the Hoeffding bound, for a given $i \in [k]$ and a given any constant $\chi > 0$, we have

$$\text{Prob}_{\mathcal{E}}\left(\left|\frac{|U_i|}{|U|} - B(i)\right| \geq \chi\right) = 2^{-\Omega(n)}. \tag{43}$$

By the union bound, we have that for any constant $\chi > 0$

$$\text{Prob}_{\mathcal{E}}\left(\exists i, \textit{such that} \left|\frac{|U_i|}{|U|} - B(i)\right| \geq \chi\right) = 2^{-\Omega(n)}. \tag{44}$$

Since the set U is a k-independent set, then the sets $\left\{\Gamma^{(k)}(u)\right\}_{u \in U}$ are nonintersecting, so

$$wt(\mathcal{E}_{\mathcal{G}}) \geq \sum_{u \in U} wt\left(\mathcal{E}_{\mathcal{G}}|_{\Gamma^{(k)}(u)}\right). \tag{45}$$

By the onion fact (Fact 8), for each $u \in U_i$ we have $wt\left(\mathcal{E}_{\mathcal{G}}|_{\Gamma^{(k)}(u)}\right) \geq \min\{i, k-i\}$, hence

$$wt(\mathcal{E}_{\mathcal{G}}) \geq \sum_{i \in [k]} |U_i| \min\{i, k-i\} = \frac{|S|}{k} \sum_{i \in [k]} \frac{|U_i|}{|U|} \min\{i, k-i\} \tag{46}$$

using $k|U| = |S|$. Using (44) with probability close to 1 we have

(47)

$$wt(\mathcal{E}_{\mathcal{G}}) \geq \frac{|S|}{k} \sum_{i \in [k]} (B(i) - \chi) \min\{i, k-i\} \geq \frac{|S|}{k} \left(\sum_{i \in [k]} B(i) \min\{i, k-i\} - 2^{-k^2}\right)$$

for $\chi = 2^{-k^2}/k^2$.

We separate the rest of the proof into two cases: $k \geq 12$ and $4 \leq k < 12$. We start with the case $k \geq 12$. Recall $\hat{k} = \lfloor k/2 \rfloor + 1$. Let

$$A_{loss} = \sum_{i \geq \hat{k}} B(i)(2i - k). \tag{48}$$

Then by (47) we have that with probability exponentially close to 1

$$wt(\mathcal{E}_{\mathcal{G}}) \geq \frac{|S|}{k}\left(\sum_{i\in[k]} B(i)i - A_{loss} - 2^{-k^2}\right) = \frac{|S|}{k}\left(pk - 2^{-k^2} - A_{loss}\right). \tag{49}$$

In the rest of the proof for $k \geq 12$ we upper-bound A_{loss} and substitute in the above equation to derive the desired result. Using an upper bound of the binomial, we have

(50)
$$B(\hat{k}) = \binom{k}{\hat{k}} p^{\hat{k}}(1-p)^{k-\hat{k}} \leq 2^k \cdot (10k)^{-\hat{k}}(1-p)^{k-\hat{k}} \leq k^{-\hat{k}}10^{-\hat{k}}2^k \leq 2^{-\hat{k}log(k)+k-3.3\hat{k}}.$$

For any $i \geq \hat{k}$ and $p < 1/2$ we have

$$B(i+1) = B(i)\left(\frac{k-i}{i+1}\right)\left(\frac{p}{1-p}\right) < B(i)\frac{p}{1-p} < 2pB(i). \tag{51}$$

Substituting (51) and (50) into the expression for A_{loss} we have

$$A_{loss} = \sum_{i\geq\hat{k}}^{k} B(i)(2i-k) \leq 2^{-\hat{k}log(k)+k-3.3\hat{k}} \sum_{i\geq\hat{k}}^{k} (2p)^{(i-\hat{k})}(2i-k) \tag{52}$$

$$\leq 2^{-\hat{k}log(k)+k-3.3\hat{k}+1+\hat{k}} \sum_{i\geq\hat{k}}^{k} p^{(i-\hat{k})}(i - \lfloor k/2 \rfloor). \tag{53}$$

Changing summation $i - \lfloor k/2 \rfloor \mapsto j$ we have that the above is at most

$$2^{-\hat{k}log(k)+k-2.3\hat{k}+1} \sum_{j\geq 1}^{\lceil k/2 \rceil} p^{j-1} j \leq 2^{-\hat{k}log(k)+k-2.3\hat{k}+1} \sum_{j\geq 1}^{\lceil k/2 \rceil} p^{j-1} k \tag{54}$$

(55)
$$\leq 2^{-\hat{k}log(k)+k-2.3\hat{k}+1} k \sum_{j\geq 1}^{\lceil k/2 \rceil} p^{j-1} \leq 2^{-\hat{k}log(k)+k-2.3\hat{k}+1} k \cdot 1.1 \leq 2^{(-\hat{k}+1)log(k)+k-2.3\hat{k}+1.2},$$

where in the last inequality we bound the sum by $\sum_{i\geq 0} p^i$ and set $p = 1/(10k) \leq 1/100$, using $k \geq 12$. Substituting this value into (49) we have that with probability $2^{-\Omega(n)}$ close to 1

$$wt(\mathcal{E}_{\mathcal{G}}) \geq \frac{|S|}{k}\left(pk - 2^{-k^2} - 2^{(-\hat{k}+1)log(k)+k-2.3\hat{k}+1.2}\right) \tag{56}$$

$$\geq \frac{|S|}{k}\left(pk - 2^{(-\hat{k}+1)log(k)+k-2.3\hat{k}+1.21}\right), \tag{57}$$

where in the last inequality we used again $k \geq 12$. Continuing, using $p = \frac{1}{10k}$, the above bound is equal to

$$= |S|p\left(1 - 2^{(-\hat{k}+1)log(k)+k-2.3\hat{k}+1.21+log_2(10)}\right) \geq |S|py(k) \tag{58}$$

for all $k \geq 12$. For values of $4 \leq k < 12$ we directly substitute k into (47), evaluate, and show it is at least $|S|py(k)$. ∎

Appendix H. Quantum PCP of proximity.

H.1. Classical PCPs of proximity. We begin by presenting the definitions following [14]. A pair language L is a subset of $\{0,1\}^n \times \{0,1\}^\ell$ for $\ell = poly(n)$. For a pair language L, let $L(x) = \{y | (x,y) \in L\}$.

DEFINITION 23 (PCP of proximity (PCPP)). *For functions $s, \delta : Z^+ \to [0,1]$, a verifier $V = V(x)$ is a probabilistically checkable proof of proximity (PCPP) system for a pair language L with proximity parameter δ and soundness error s if the following two conditions hold for every pair of strings $(x,y) \in \{0,1\}^n \times \{0,1\}^\ell$:*

1. *Completeness: If $(x,y) \in L$, there exists π such that $V(x)$ accepts oracle $y \circ \pi$ with probability 1.*
2. *Soundness: If y is $\delta(|x|)$-far from $L(x)$, then for every π, the verifier $V(x)$ accepts oracle $y \circ \pi$ with probability at most $s(|x|)$.*

If s and δ are not specified, then both can be assumed to be constants in $(0,1)$. The query complexity of the verifier V, $q(V)$, is defined to be the number of coordinates that V queries out of y and π.

In the above definition, V is not charged for reading x but is charged for reading y even though it is part of the input. This is a more stringent restriction than in the case of a PCP proof; however, the requirements on the proof system are weaker: V is supposed to reject only words which are *far* from words in the language. As in the case of PCP, we would like the query complexity to be as low as possible and, optimally, $q(V) = O(1)$.

A good pair language to keep in mind is CIRCUIT-VAL, i.e., the pairs (x,y), where x is a circuit on n bits of polynomial size, and y is a string on n bits, and $(x,y) \in L$ if $x(y) = 1$, i.e., the circuit x evaluates to 1 on input y. Though this problem lies in P, a simple argument (Proposition 2.4 in [14]) shows that a PCPP for CIRCUIT-VAL implies a PCP for the NP complete decision language CIRCUIT-SAT, the set of all x, for which there exists y, such that $x(y) = 1$.

H.2. From PCPPs to LTCs. Given a PCPP, [14] provides a standard construction of an LTC with related parameters as follows. Given is a PCPP for membership in a code, namely, for the pair language of (C, w), a code and a member in that code. Suppose the proximity parameter of the PCPP is δ, the soundness s, and the query complexity k. Suppose also that we are given a code C with distance D. Then one can construct an error correcting code C' which is an LTC with k-local constraints, whose weighted distance is D, and whose soundness is proportional to the soundness s.

C' is defined as the strings $w \circ \pi$ for all w in C, where π is the proximity proof of w. If one defines the distance by weighting only the coordinates in the first register, then C' trivially has the same distance as C.[9] The local test for the code C' as an LTC are the k-local tests performed by the verifier of the PCPP. Consider now a word $w' \circ \pi'$ which is δ-far from any $w \circ \pi$ in the code C', where the distance is measured again by taking into account only the coordinates of the left register. This means that w' is δ-far from a word in the code C; then the tests will reject the word $w' \circ \pi'$ with probability s, which will thus be the soundness of the code for proximity δ.

H.3. Quantum PCPPs. We now define the quantum analogue of PCPs of proximity. We consider quantum pair languages $L \subseteq \{0,1\}^n \otimes \mathcal{H}_{prf}$, where $\mathcal{H}_{prf}$ is a

[9] In [14] this choice of definition of distance is referred to as equivalent to the one used in [14], in which many repetitions of the string w are taken, so that the weight of the error on the second, proof, register becomes negligible.

Hilbert space of ℓ d-dimensional qudits for $\ell = poly(n)$. For a quantum pair language L let $L(x) = \text{Span}\{|\psi\rangle \in \mathcal{H}_{prf}, (x, |\psi\rangle) \in L\}$.

DEFINITION 24 (quantum PCP of proximity). *Fix functions $s, \delta : Z^+ \to [0,1]$. Let $V = V(x)$ be a function from n bit strings x to sets of m k-local projections $\{\Pi_i\}_{i=1}^m$, each acting on $\mathcal{H}_{prf} \otimes \mathcal{H}_{pxmty}$. V is a quantum probabilistically checkable proof of proximity (qPCPP) system, for a quantum pair language L, with proximity parameter δ and soundness error s, if the following two conditions hold for every pair $(x, |\psi\rangle)$:*

1. *Completeness: If $(x, |\psi\rangle) \in L$, there exists $|w\rangle \in \mathcal{H}_{pxmty}$, such that for all check terms $\Pi_i \in V(x)$*

$$\Pi_i \left(|\psi\rangle \otimes |w\rangle\right) = 0. \tag{59}$$

2. *Soundness: If $|\phi\rangle$ is a quantum state in $\mathcal{H}_{prf} \otimes \mathcal{H}_{pxmty}$ whose reduced density matrix to $\mathcal{H}_{prf}$ is supported on states, each of distance at least $\delta(|x|)$ from $L(x)$, then*

$$\frac{1}{m}\sum_i \langle\phi|\Pi_i|\phi\rangle \geq s(|x|) : . \tag{60}$$

The query complexity of the verifier V is k.

H.4. From qPCPPs to qLTCs. Given is a qPCPP for membership in a quantum code on ℓ qudits, namely, for the pair language L composed of pairs $(C, |\psi\rangle)$: a code (described by n bits) and an ℓ-qudit state in the code. Suppose L has a qPCP of proximity with parameters δ, s for some functions $s, \delta : Z^+ \to [0,1]$, with projections Π_i. Let C' be the code space $\subseteq \mathcal{H}_{prf} \otimes \mathcal{H}_{pxmty}$, defined as

$$\text{Span}\left\{|\phi\rangle \otimes |\Pi(\phi)\rangle \text{ such that } |\phi\rangle \in C\right\}, \tag{61}$$

where $|\Pi(\phi)\rangle$ is some proof of proximity for $|\phi\rangle$ from the qPCPP. Let $dist_{prf}$ denote the distance from the code space as in Definition 14, except it only counts nonidentity Paulis acting on the register $\mathcal{H}_{prf}$.

CLAIM 6. *C' is a qLTC with query complexity k and soundness $R(\delta) = s$ (where the proximity δ is defined with respect to the distance $dist_{prf}$).*

Proof. Set $\{\Pi_i\}_{i=1}^m$ as the check terms for $L(C)$. These are k-local terms, so C' has query complexity k. By definition of the qPCPP, for any state $|\phi\rangle$ in the code space of C', we have $\Pi_i|\phi\rangle = 0$ for any $\Pi_i \in V(C')$. Let us assume now that $dist_{prf}(|\phi\rangle, C') \geq \delta(|C|) \cdot \ell$. Then, by Definition 14, for any Pauli operator E acting on $\mathcal{H}_{prf} \otimes \mathcal{H}_{pxmty}$, whose support on $\mathcal{H}_{prf}$ is at most $\delta(|C|) \cdot \ell - 1$, we have that $E|\phi\rangle$ is still orthogonal to C'. In particular, for any E whose support is contained in $\mathcal{H}_{prf}$, and whose weight is at most $\delta(|C|) \cdot \ell - 1$, we have that $E|\phi\rangle$ is still orthogonal to C'. It is easy to see that the reduced state of $|\phi\rangle$ to $\mathcal{H}_{prf}$ is a mixture of orthogonal states $\{|\eta_i\rangle\}_i$, each of which is at least $\delta(|C|) \cdot \ell$-far from C. By virtue of the soundness of the qPCPP, $|\phi\rangle$ will be rejected by the check terms Π_i with probability at least $s(|C|)$. ∎

Acknowledgments. The authors would like to thank Eli Ben-Sasson, Irit Dinur, and Tali Kaufman for insightful discussions.

REFERENCES

[1] D. Aharonov, I. Arad, Z. Landau, and U. Vazirani, *The Detectability Lemma and Quantum Gap Amplification*, preprint, arXiv:0811.3412.

[2] D. Aharonov, I. Arad, and T. Vidick, *The quantum PCP conjecture*, ACM SIGACT News Archive, 44 (2013), pp. 47–79.

[3] D. Aharonov and L. Eldar, *The commuting local Hamiltonian problem on local expanders is in NP*, preliminary version of the results appeared in the preprint quant-ph arXiv:1301.3407.

[4] N. Alon, T. Kaufman, M. Krivelevich, S. Litsyn, and D. Ron, *Testing Reed–Muller codes*, IEEE Trans. Inform. Theory, 51 (2005), pp. 4032–4039.

[5] S. Arora, *Probabilistic Checking of Proofs and the Hardness of Approximation Problems*, Ph.D. thesis, UC Berkeley, 1994.

[6] S. Arora and B. Barak, *Computational Complexity: A Modern Approach*, Cambridge University Press, Cambridge, UK, 2009.

[7] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy, *Proof verification and intractability of Approximation problems*, J. ACM, 45 (1998), pp. 501–555.

[8] S. Arora and S. Safra, *Probabilistic checking of proofs: A new characterization of NP*, J. ACM, 45 (1998), pp. 70–122.

[9] L. Babai, L. Fortnow, L. Levin, and M. Szegedy, *Checking computation in polylogarithmic time*, in Proceedings of the 23rd ACM Symposium on Theory of Computation, New Orleans, 1991.

[10] D. Bacon, *Operator quantum error correcting subsystems for self-correcting quantum memories*, Phys. Rev. A, 73 (2005), 012340.

[11] H. Barnum, C. Crepeau, D. Gottesman, A. Smith, and A. Tapp, *Authentication of quantum messages*, in Proceedings of the 43rd Annual IEEE Symposium on the Foundations of Computer Science (FOCS '02), IEEE Press, 2002, pp. 449–458.

[12] H. Buhrman, R. Cleve, J. Watrous, and R. de Wolf, *Quantum fingerprinting*, Phys. Rev. Lett., 87 (2001), 167902.

[13] E. Ben-Sasson, A. Chiesa, D. Genkin, E. Tromer, and M. Virza, *Verifying program executions succinctly and in zero knowledge*, in Proceedings of the 33rd International Cryptology Conference (CRYPTO 2013).

[14] E. Ben-Sasson, O. Goldreich, P. Harsha, M. Sudan, and S. P. Vadhan, *Sound PCPs of proximity, shorter PCPs, and applications to coding*, SIAM J. Comput., 36 (2006), pp. 889–974.

[15] E. Ben-Sasson and M. Sudan, *Short PCPs with polylog query complexity*, SIAM J. Comput., 38 (2008), pp. 551–607.

[16] F. G. S. L. Brandão and A. W. Harrow, *Product-state approximations to quantum ground states*, in Proceedings of the 45th ACM Symposium on Theory of Computing (STOC 2013), pp. 871–880.

[17] S. Bravyi and M. B. Hastings, *Homological Product Codes*, preprint, arXiv:1311.0885, 2013.

[18] S. Bravyi and M. B. Hastings, *A short proof of stability of topological order under local perturbations*, Commun. Math. Phys., 307 (2011), pp. 609–627.

[19] S. Bravyi, M. B. Hastings, and S. Michalakis, *Topological quantum order: Stability under local perturbations*, J. Math. Phys., 51 (2010), 093512.

[20] S. Bravyi, D. Poulin, and B. M. Terhal, *Tradeoffs for reliable quantum information storage in 2D systems*, Phys. Rev. Lett., 104 (2010), 050503.

[21] S. Bravyi and B. Terhal, *A no-go theorem for a two-dimensional self-correcting quantum memory based on stabilizer codes*, New J. Phys., 11 (2009).

[22] M. R. Capalbo, O. Reingold, S. Vadhan, and A. Wigderson, *Randomness conductors and constant-degree lossless expanders*, in Proceedings of the ACM Symposium on Theory of Computing (STOC), 2002, pp. 659–668.

[23] S. Chesi, D. Loss, S. Bravyi, and B. M. Terhal, *Thermodynamic stability criteria for a quantum memory based on stabilizer and subsystem codes*, New J. Phys., 12 (2010), 025013.

[24] A. Couvreur, N. Delfosse, and G. Zémor, *A construction of quantum LDPC codes from Cayley graphs*, in Proceedings of the IEEE International Symposium on Information Theory Proceedings (ISIT), 2011, pp. 643–647.

[25] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, *Topological quantum memory*, J. Math. Phys., 43 (2002), pp. 4452–4505.

[26] I. Dinur, *The PCP theorem by gap amplification*, J. ACM, 54 (2007), 12.
[27] I. Dinur and T. Kaufman, *On the Structure of NP-hard 3-SAT Instances, and a Similar Question for LTCs*, talk at the Fourth Israel CS Theory Day Thursday, March 24, 2011.
[28] I. Dinur and T. Kaufman, *Locally Testable Codes and Expanders*, preprint.
[29] L. Eldar, *Quantum Systems with Approximation-Robust Entanglement*, preprint, 2015.
[30] E. Fetaya, *Bounding the distance of quantum surface codes*, J. Math. Phys., 53 (2012), 062202.
[31] K. Friedl and M. Sudan, *Some improvements to total degree tests*, in Proceedings of the 3rd Israel Symposium on Theoretical and Computing Systems (Tel Aviv, Israel), 1998.
[32] O. Goldreich, *Short Locally Testable Codes and Proofs: A Survey in Two Parts*, in Property Testing: Current Research and Surveys, Lecture Notes in Comput. Sci. 6390, Springer, New York, 2010, pp. 65–104.
[33] O. Goldreich, S. Goldwasser, and D. Ron, *Property testing and its connection to learning and approximation*, J. ACM, 45 (1998), pp. 653–750.
[34] O. Goldreich and M. Sudan, *Locally testable codes and PCPs of almost-linear length*, J. ACM, 53 (2006), pp. 558–655.
[35] D. Gottesman, *Stabilizer Codes and Quantum Error Correction*, Ph.D. thesis, Caltech, Pasadena, CA, 2004.
[36] D. Gottesman, *A theory of fault-tolerant quantum computation*, Phys. Rev. A, 57 (1998), pp. 127–137.
[37] D. Gottesman, *On the theory of quantum secret sharing*, Phys. Rev. A, 61 (2000), 042311.
[38] J. Haah, *Local stabilizer codes in three dimensions without string logical operators*, Phys. Rev. A, 83 (2011), 042330.
[39] J. Haah and J. Preskill, *Logical-operator tradeoff for local quantum codes*, Phys. Rev. A, 86 (2012), 032308.
[40] J. Håstad, *Some optimal inapproximability results*, J. ACM, 48 (2001), pp. 798–859.
[41] M. B. Hastings, *Trivial Low Energy States for Commuting Hamiltonians, and the Quantum PCP Conjecture*, preprint, arXiv:1201.3387v1, 2012.
[42] M. B. Hastings, *Decoding in Hyperbolic Spaces: LDPC Codes with Linear Rate and Efficient Error Correction*, preprint, arXiv:1312.2546, 2013.
[43] A. Yu. Kitaev, *Fault-tolerant quantum computation by anyons*, Ann. Physics, 303 (2003), pp. 2–30.
[44] A. Yu. Kitaev, A. H. Shen, and M. N. Vyalyi, *Classical and Quantum Computation*, Grad. Stud. Math. 47, AMS, Providence, RI, 2002.
[45] E. Knill, R. Laflamme, A. Ashikhmin, H. Barnum, L. Viola, and W. H. Zurek, *Introduction to Quantum Error Correction*, preprint, http://arxiv.org/abs/quant-ph/0207170, 2002.
[46] A. A. Kovalev and L. P. Pryadko, *Fault Tolerance of "Bad" Quantum Low Density Parity Check Codes*, preprint, arXiv:1208.2317, 2012.
[47] A. A. Kovalev and L. P. Pryadko, *Quantum "Hyperbicycle" Low Density Parity Check Codes with Finite Rate*, preprint, arXiv:1212.6703, 2012.
[48] F. J. MacWilliams and N. J. A. Sloane, North-Holland, Amsterdam, 1977.
[49] K. Michnicki, *3-d Quantum Stabilizer Codes with a Power Law Energy Barrier*, preprint, arXiv:1208.3496, 2012.
[50] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (10th anniversary ed.), Cambridge University Press, Cambridge, UK, 2010.
[51] R. Rubinfeld and M. Sudan, *Robust characterizations of polynomials with applications to program testing*, SIAM J. Comput., 25 (1996), pp. 252–271.
[52] M. Sipser and D. Spielman, *Expander codes*, IEEE Trans. Inform. Theory, 42 (1996), pp. 1710–1722.
[53] A. Steane, *Quantum Reed–Muller Codes*, preprint, arXiv:quant-ph/9608026, 1996.
[54] R. M. Tanner, *A recursive approach to low-complexity codes*, IEEE Trans. Inform. Theory, 27 (1981), pp. 533–547.
[55] J. P. Tillich and G. Zémor, *Quantum LDPC codes with positive rate and minimum distance proportional to $n^{1/2}$*, in Proceedings of the IEEE International Symposium on Information Theory (ISIT) 2009, pp. 799–803.
[56] B. Yoshida, *Feasibility of self-correcting quantum memory and thermal stability of topological order*, Ann. Phys., 326 (2011), pp. 2566–2633.