UNIVERSITY OF CALIFORNIA SAN DIEGO

Local-to-Global Structure in Quantum Error-Correcting Codes and Entanglement Bootstrap

A dissertation submitted in partial satisfaction of the

requirements for the degree

Doctor of Philosophy

in

Physics

by

Ting-Chun Lin

Committee in charge:

Professor John McGreevy, Chair
Professor Daniel Grier
Professor Tarun Grover
Professor Daniel Rogalski
Professor Yi-Zhuang You

2026

The Dissertation of Ting-Chun Lin is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego
2026

DEDICATION

To 林俊杰 and 陳郁菁.

To space and time.

To 0 and 1.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

First and foremost, I would like to thank my advisor John. A PhD involves a great deal of exploration, and I am grateful that you gave me the space to try, fail, and try again. Your genuine passion for physics has been a source of inspiration for everyone around you, and I will deeply miss our many late-night discussions.

I also thank my mentor Min-Hsiu, with whom I wrote my first paper. The problem you gave me fit me surprisingly well; so much of my later work grew from that seed. Watching you manage many projects simultaneously and advance each with equal care was an education in itself. Without your guidance, I would never have gotten my research off the ground.

I owe a great deal to Bowen and Xiang as well. The shared community we built around entanglement bootstrap — the heated arguments and crazy ideas (yes, instantaneous modular flow and measuring geometries in TQFT) — has been truly precious to me. Though those fond memories will slip through me, I know Xiang will be able to recount every detail for all of us when we meet again.

I want to thank all my friends at UCSD: Rolando, Simon, Owen, Yu-Hsueh, Eugene, and EJ, and my board game companions Chun-Yueh, Dai Xiang, Ting-Yu, and the many wonderful people I met through baseball, UCSD rec, and theatre classes. You all brought so much color and joy to those years.

I have also had the pleasure of working with many wonderful collaborators: Yahya Alavirad, Omar Alrabiah, Shankar Balasubramanian, Madelyn Cain, Alexander Cowtan, Margarita Davydova, Irit Dinur, Iwan Duursma, Quinten Eggerickx, Ryan Gabrys, Louis Golowich, Michael Gullans, Venkatesan Guruswami, Zhiyang He, Yifan Hong, Max Hopkins, Jun-Ting Hsieh, Isaac Kim, Michael Levin, Xingjian Li, Nishad Maskara, Sidhanth Mohanty, Ryan O'Donnell, Daniel Ranard, Samuel Tan, Alexander Vardy, Thomas Vidick, Hsin-Po

Wang, Adam Wills, Dominic Williamson, Andrew Yuan, and Rachel Zhang. Many of our projects started as a shrug — no one knew whether it was even possible. Then, after an uncountable amount of collective confusion, one day we looked up and had already climbed the mountain. Looking back, I still cannot believe how we got there.

Finally, I am grateful to the Simons Center, the UQM Collaboration, and the TASI Summer School for bringing together people with a shared passion for physics and TCS. They broadened my sense of the field, added unexpected pieces to my theoretical world, and sent my thoughts in directions I never would have found alone.

Chapter 2 is based on joint work with Irit Dinur and Thomas Vidick, published as "Expansion of high-dimensional cubical complexes with application to quantum locally testable codes," *Proc. IEEE FOCS*, 2024.

Chapter 3 is based on the sole-authored paper "Transversal non-Clifford gates for quantum LDPC codes on sheaves," *arXiv:2410.14631*, 2024.

Chapter 4 is based on joint work with John McGreevy, published as "Conformal field theory ground states as critical points of an entropy function," *Phys. Rev. Lett.*, 2023.

Chapter 5 is based on joint work with Xiang Li and John McGreevy, "A systematic search for conformal field theories in very small spaces," *arXiv:2509.04596*, 2025.

VITA

2016–2020 B.S. in Physics and Mathematics, Massachusetts Institute of Technology

2021–2026 Ph.D. in Physics, University of California San Diego

PUBLICATIONS

T.-C. Lin, M.-H. Hsieh, "c^3-Locally testable codes from lossless expanders," *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, 2022.

T.-C. Lin, M.-H. Hsieh, "Good quantum LDPC codes with linear time decoder from lossless expanders," *arXiv:2203.03581*, 2022.

M. Hopkins, T.-C. Lin, "Explicit lower bounds against $\omega(n)$-rounds of sum-of-squares," *Proc. IEEE FOCS*, 2022.

I. Duursma, R. Gabrys, V. Guruswami, T.-C. Lin, "Accelerating polarization via alphabet extension," *arXiv:2207.04522*, 2022.

H.-P. Wang, T.-C. Lin, A. Vardy, "Sub-4.7 scaling exponent of polar codes," *IEEE Trans. Inf. Theory*, 2023.

I. Dinur, M.-H. Hsieh, T.-C. Lin, T. Vidick, "Good quantum LDPC codes with linear time decoders," *Proc. ACM STOC*, 2023.

T.-C. Lin, I. H. Kim, M.-H. Hsieh, "A new operator extension of strong subadditivity of quantum entropy," *Lett. Math. Phys.*, 2023.

T.-C. Lin, J. McGreevy, "Conformal field theory ground states as critical points of an entropy function," *Phys. Rev. Lett.*, 2023.

I. H. Kim, M. Levin, T.-C. Lin, D. Ranard, B. Shi, "Universal lower bound on topological entanglement entropy," *Phys. Rev. Lett.*, 2023.

T.-C. Lin, H.-P. Wang, "Optimal self-dual inequalities to order polarized BECs," *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, 2023.

A. Wills, T.-C. Lin, M.-H. Hsieh, "General distance balancing for quantum locally testable codes," *arXiv:2305.00689*, 2023.

T.-C. Lin, A. Wills, M.-H. Hsieh, "Geometrically local quantum and classical codes from subdivision," *arXiv:2309.16104*, 2023.

A. Wills, T.-C. Lin, M.-H. Hsieh, "Tradeoff constructions for quantum locally testable codes," *IEEE Trans. Inf. Theory*, 2024.

I. Dinur, T.-C. Lin, T. Vidick, "Expansion of high-dimensional cubical complexes with application to quantum locally testable codes," *Proc. IEEE FOCS*, 2024.

I. H. Kim, X. Li, T.-C. Lin, J. McGreevy, B. Shi, "Chiral Virasoro algebra from a single wavefunction," *Ann. Phys.*, 2024.

I. H. Kim, T.-C. Lin, D. Ranard, B. Shi, "Strict area law implies commuting parent Hamiltonian," *arXiv:2404.05867*, 2024.

X. Li, T.-C. Lin, M.-H. Hsieh, "Transform arbitrary good quantum LDPC codes into good geometrically local codes in any dimension," *arXiv:2408.01769*, 2024.

T.-C. Lin, "Transversal non-Clifford gates for quantum LDPC codes on sheaves," *arXiv:2410.14631*, 2024.

T.-C. Lin, H.-P. Wang, M.-H. Hsieh, "Proposals for 3D self-correcting quantum memory," *arXiv:2411.03115*, 2024.

I. Kim, X. Li, T.-C. Lin, J. McGreevy, B. Shi, "Conformal geometry from entanglement," *SciPost Phys.*, 2025.

Q. Eggerickx, A. Wills, T.-C. Lin, K. De Greve, "Almost linear decoder for optimal geometrically local quantum codes," *Phys. Rev. Lett.*, 2025.

L. Golowich, T.-C. Lin, "Quantum LDPC codes with transversal non-Clifford gates via products of algebraic codes," *Proc. ACM STOC*, 2025.

J. T. Hsieh, T.-C. Lin, S. Mohanty, R. O'Donnell, "Explicit two-sided vertex expanders beyond the spectral barrier," *Proc. ACM STOC*, 2025.

X. Li, T.-C. Lin, J. McGreevy, B. Shi, "Strict area law entanglement versus chirality," *Phys. Rev. Lett.*, 2025.

X. Li, T.-C. Lin, J. McGreevy, "A systematic search for conformal field theories in very small spaces," *arXiv:2509.04596*, 2025.

S. J. S. Tan, Y. Hong, T.-C. Lin, M. J. Gullans, "Single-shot universality in quantum LDPC codes via code-switching," *arXiv:2510.08552*, 2025.

M.-H. Hsieh, X. Li, T.-C. Lin, "Simplified quantum weight reduction with optimal bounds," *arXiv:2510.09601*, 2025.

X. Li, T.-C. Lin, Y. Alavirad, J. McGreevy, "Chiral gapped states are universally non-topological," *arXiv:2510.23720*, 2025.

A. Wills, T.-C. Lin, R. Y. Zhang, M.-H. Hsieh, "Linear-time encodable and decodable quantum error-correcting codes," *arXiv:2603.04543*, 2026.

A. C. Yuan, A. Cowtan, Z. He, T.-C. Lin, "Parsimonious quantum low-density parity-check code surgery," *arXiv:2603.05082*, 2026.

S. Balasubramanian, M. Davydova, T.-C. Lin, "A passive self-correcting quantum memory in three dimensions," *arXiv:2605.10943*, 2026.

ABSTRACT OF THE DISSERTATION

Local-to-Global Structure in
Quantum Error-Correcting Codes and Entanglement Bootstrap

by

Ting-Chun Lin

Doctor of Philosophy in Physics

University of California San Diego, 2026

Professor John McGreevy, Chair

The local-to-global principle—the idea that global structure is determined by, or can be certified from, local constraints—appears across mathematics, theoretical computer science, and physics under different disguises. This thesis studies two manifestations of this principle: in quantum error-correcting codes, where local parity checks certify the integrity of the global encoded information, and in the entanglement bootstrap, where local reduced density matrices determine the global structure of quantum phases of matter.

Part I constructs quantum low-density parity check (qLDPC) codes. We first build a family of quantum locally testable codes (qLTCs) achieving linear dimension, near-linear distance, and near-constant soundness, by generalizing square-complex qLDPC codes to higher-dimensional cubical chain complexes. We then extend the cubical construction to the sheaf cochain complexes, on which we define Poincaré duality and a cup product. Poincaré duality clarifies the duality between the X- and Z-distances of the corresponding quantum code, while the cup product yields three sheaf codes that jointly support a transversal CCZ gate, providing a general route to non-Clifford gates on qLDPC codes.

Part II develops an entanglement-based characterization of conformal field theories (CFTs). We first show that the ground state of any unitary 1+1D CFT is a critical point of a particular linear combination of subsystem entropies, and we verify numerically that this condition is stringent yet robust across known critical lattice models. We then leverage this characterization to perform a computational search over four-qubit and four-qutrit systems, which recovers known CFTs and identifies roughly 20 new candidate critical points, potentially including new irrational CFTs.

Chapter 1

Introduction

Take any smooth closed surface — a sphere, a torus — and integrate its Gaussian curvature K over every point. The curvature varies from place to place: positive where the surface bends like a bowl, negative where it bends like a saddle, zero where it is flat. Stretch, squeeze, indent the surface however you like. As long as you do not tear or puncture it, the integral does not change:

$$\int_M K \, dA \;=\; 2\pi \, \chi(M). \tag{1.1}$$

The right-hand side, $\chi(M)$, is the *Euler characteristic* — an integer that encodes the topology: $\chi = 2$ for a sphere, $\chi = 0$ for a torus, and more generally, $\chi = 2 - 2g$ for a surface of genus g, where g is the number of holes.

The Gauss–Bonnet theorem (1.1) is striking because the quantity on the left is purely local and differential, while the quantity on the right is purely global and topological. Local curvature, when integrated up, determines the global topology. This is the **local-to-global principle** in one of its most elegant incarnations.

The principle has since spread far beyond differential geometry. It is the organizing idea of modern mathematics, and, as this thesis explores, finds applications in theoretical computer science and quantum field theory.

Local-to-Global in Mathematics

The story of the local-to-global principle in mathematics is, in a sense, the story of 20th-century geometry. Gauss–Bonnet showed that topology can be read off from local curvature. But the deeper realization — that many mathematical objects are *defined* by their local data plus consistency requirements — took time to crystallize.

A *manifold* is defined exactly this way. One does not hand someone a manifold as a single global object; one provides a collection of *local charts* (open sets homeomorphic to $\mathbb{R}^n$) together with transition functions specifying how they agree on overlaps. The manifold is essentially a collection of locally consistent Euclidean patches. Asking whether a collection of such patches assembles into a consistent global space is built into the very definition.

By mid-century, this idea was formalized into *sheaf theory*, initiated by Leray in 1945 and further developed by Cartan, Serre, Grothendieck, and others into a foundational language for cohomology, algebraic geometry, and modern number theory.

A presheaf $\mathcal{F}$ on a topological space X assigns to each open set U a set of *sections* $\mathcal{F}(U)$, together with restriction maps $\mathrm{res}_{UV} \colon \mathcal{F}(U) \to \mathcal{F}(V)$ whenever $V \subseteq U$, satisfying $\mathrm{res}_{UU} = \mathrm{id}_{\mathcal{F}(U)}$ and $\mathrm{res}_{VW} \circ \mathrm{res}_{UV} = \mathrm{res}_{UW}$ whenever $W \subseteq V \subseteq U$. Think of $\mathcal{F}(U)$ as "the data visible from region U."

For a presheaf to be a sheaf, two additional axioms must be satisfied.

Identity axiom. If $\{U_i\}$ covers U, and two sections $s, t \in \mathcal{F}(U)$ agree on every piece, $s|_{U_i} = t|_{U_i}$ for all i, then $s = t$.

Gluability axiom. If $\{U_i\}$ covers U, and we have local sections $s_i \in \mathcal{F}(U_i)$ that agree on every pairwise overlap, $s_i|_{U_i \cap U_j} = s_j|_{U_i \cap U_j}$, then there exists a global section $s \in \mathcal{F}(U)$ restricting to each s_i.

These two axioms say, respectively, that a global section is *uniquely* determined by

its local restrictions, and that it *exists* whenever the local pieces are mutually consistent.

The canonical example is the sheaf of smooth functions on a manifold: a function defined on each chart that agrees on overlaps extends uniquely to a global smooth function. The sheaf condition is, at its core, a concrete expression of when local data determines global structure.

Local-to-Global in Theoretical Computer Science and Quantum Field Theory

This dissertation studies two pairs of results, one in theoretical computer science and one in quantum field theory, each exhibiting local-to-global structure in its own way.

In the theory of **quantum error-correcting codes** (Part I), the question is: Can local parity checks (each involving only a few bits) globally certify that a large encoded message is intact? Specifically, if 99% of the checks pass, is the message necessarily close to a codeword? This property, known as *local testability*, is a robust analogue of the gluability axiom, and it carries significant consequences in complexity theory. Part I explores the quantum analogues of these ideas.

In the study of **quantum many-body ground states** (Part II), the question is: Can the local reduced density matrices of a ground state uniquely determine the ground state? More precisely, can locally consistent quantum data be assembled into a globally consistent quantum state? The *entanglement bootstrap* program conjectures that this is possible under appropriate local entropy conditions, a picture supported by a growing body of results. These conditions were previously established for ground states of topological quantum field theories. Part II investigates their counterparts for ground states of conformal field theories.

Part I Quantum Error-Correcting Codes

We begin with a review of error-correcting codes, classical and quantum, and then reveal their connection to robust versions of the sheaf axioms.

Classical codes

Error-correcting codes address a fundamental question: how can one transmit information reliably in the presence of noise? The answer, going back to Shannon's foundational 1948 work on reliable communication over noisy channels and Hamming's 1950 construction of explicit error-correcting codes, is to introduce redundancy. To transmit k bits of information, one encodes them into a longer string of $n > k$ bits, where the additional $n - k$ bits provide enough redundancy that the original message can still be recovered even after some coordinates are corrupted.

A common realization is through a linear code: the valid encodings, or *codewords*, form a k-dimensional vector subspace $C \subseteq \{0,1\}^n$. Equivalently, C may be described as the kernel of a parity-check matrix H, typically with $n - k$ rows and n columns. Each row of H specifies a parity constraint requiring that a certain subset of coordinates have XOR equal to 0.

The *distance* d of the code is the minimum Hamming distance between distinct codewords. Consequently, any corruption affecting fewer than $d/2$ coordinates can be uniquely corrected by finding the nearest codeword to the received string.

Morally, n is the budget, k is the useful output, and d controls the quality (i.e. the reliability) of the output. Under fixed budget n, one looks for codes that maximize both k and d. In what follows, we focus primarily on the asymptotic regime. A simple benchmark is whether both parameters k and d scale linearly with the block length n. A family of codes is called *good* when this holds.

Low-density parity-check codes and locally testable codes

Low-density parity-check (LDPC) codes, introduced by Gallager in 1960, add a structural constraint: each parity check involves only $O(1)$ bits, and each bit participates in only $O(1)$ checks. In other words, the parity-check matrix H is sparse. This locality was originally motivated by the possibility of efficient decoding algorithms, since sparse constraints permit iterative message-passing processes. It has since become equally important for quantum error correction, where low-weight checks are often essential for physical implementability.

Remarkably, sparsity does not preclude strong asymptotic performance. Gallager showed in 1962, via random constructions, that there exist classical LDPC families with both $k, d = \Omega(n)$. Later, in 1996, Sipser and Spielman gave explicit constructions of such good LDPC codes using expander graphs.

Because each check involves only $O(1)$ bits, each check can be verified by querying very few bits. This efficiency raises a subtler question: if 99% of the checks pass, must the word be 1%-close to a valid codeword? A code with this property is called a *locally testable code* (LTC); its *soundness* s measures the proportionality constant between the fraction of violated checks and the relative distance to the nearest codeword. Constructing good LTC families with $k, d = \Omega(n)$ and $s = \Omega(1)$ remained a long-standing open problem. This was only resolved recently in 2021, independently by Dinur et al. and Panteleev–Kalachev.

Quantum codes

In the quantum setting, an error-correcting code again encodes k logical qubits into n physical qubits. The natural quantum generalization of LDPC codes is a class of *CSS codes*, defined by two sparse parity-check matrices H_X, H_Z satisfying $H_X H_Z^T = 0$. A convenient way to view this condition is through the chain-complex identity $\partial_1 \circ \partial_2 = 0$.

This allows us to identify a quantum CSS code with a chain complex of vector spaces, in which the parity-check matrices serve as the boundary maps:

$$\mathbb{F}_2^{m_x} \xrightarrow{\partial_2} \mathbb{F}_2^{n} \xrightarrow{\partial_1} \mathbb{F}_2^{m_z}.$$

Under this correspondence, many coding-theoretic notions acquire homological meaning: stabilizers correspond to boundaries and coboundaries, logical operators correspond to cycles and cocycles, and the number of encoded qubits is given by the dimension of the associated homology group.

What makes the quantum setting genuinely richer is that higher-dimensional structure becomes indispensable. Classical codewords can be described using only 0-cochains, which are functions. By contrast, quantum logical operators are naturally described using higher-degree cochains. This additional homological structure is what gives the chain-complex framework much of its power in quantum error correction. Developing this perspective is a major theme of Part I.

This homological viewpoint is not merely a reformulation. It has played a central role in some of the major advances in the subject. For decades, it was not known whether there exist *good quantum LDPC codes* — codes with $k, d = \Omega(n)$ and $O(1)$-weight checks. A major breakthrough came in 2020, when Hastings–Haah–O'Donnell introduced the fiber-bundle construction, drawing directly on ideas from algebraic topology. This opened the way for further progress. The question was then fully resolved in 2021 by Panteleev–Kalachev, who proved the existence of good quantum LDPC codes.

Rephrasing the properties as robust versions of identity and gluability axioms

We conclude the introduction by interpreting the main properties of error-correcting codes as robust versions of the sheaf axioms. This interpretation is meaningful primarily for

LDPC codes, which will be our focus below. We use the term local view (or local section) to refer to the restriction of a codeword to the coordinates touched by a check. Each local view is assumed to satisfy the corresponding parity constraint; i.e., $s|_{U_i} \in \mathcal{F}(U_i)$, where $\mathcal{F}(U_i) \subseteq \mathbb{F}_2^{U_i}$ denotes the set of vectors that satisfy the i-th check.

The main dictionary is:

$$\text{Robust identity axiom} \leftrightarrow \text{large distance}$$

$$\text{Robust gluability axiom} \leftrightarrow \text{local testability}$$

The following table gives schematic formulations of the exact and robust axioms in the setting of LDPC codes.

Table 1.1: The identity and gluability axioms, and their robust versions, in the context of error-correcting codes.

	Identity (Uniqueness)	**Gluability (Existence)**
Exact	Two codewords that agree on every local view are identical: $\forall i : s\|_{U_i} = t\|_{U_i} \implies s = t$	A family of local views that are pairwise consistent on all overlaps lifts to a global codeword: $\forall i, j : s_i\|_{U_i \cap U_j} = s_j\|_{U_i \cap U_j} \implies \exists s, s\|_{U_i} = s_i$
Robust	Two codewords that agree on $1-\varepsilon$ fraction of local views are identical: $\mathbb{P}_i\Big[s\|_{U_i} = t\|_{U_i}\Big] > 1-\varepsilon \implies s = t$	A family of local views that are consistent on $1-\varepsilon$ fraction of overlapping pairs lifts to a codeword that agrees with $1-\delta(\varepsilon)$ fraction of the local views: $\mathbb{P}_{\substack{i,j: \\ U_i \cap U_j \neq \emptyset}}\Big[s_i\|_{U_i \cap U_j} = s_j\|_{U_i \cap U_j}\Big] > 1-\varepsilon \implies \exists s,\ \mathbb{P}_i\Big[s\|_{U_i} = s_i\Big] > 1-\delta(\varepsilon)$

We first explain why robust identity is closely related to error correctability. If the code has linear distance $d = \Omega(n)$, then two distinct codewords must differ on a linear number of coordinates. Since each coordinate participates in $O(1)$ local views, and each

local view involves only $O(1)$ coordinates, this discrepancy is visible on a positive fraction of the local views. Consequently, if two codewords agree on all but an ε fraction of local views, with sufficiently small ε, they must in fact be identical. Conversely, any robust uniqueness statement of this form forces distinct codewords to remain far apart, and hence implies large distance.

We next explain why robust gluability is closely related to local testability. Suppose a family of local views is consistent on most overlaps. We can aggregate these local data into a candidate global word by majority vote on each coordinate. Because most overlaps are compatible, the resulting word violates only a small fraction of checks. If the code is locally testable with constant soundness $s = \Omega(1)$, then any word satisfying most checks must be close to a genuine codeword. Hence approximate local consistency produces a nearby global codeword, which is precisely the robust gluability principle. The converse implication may likewise be proven.

We also remark that one may study different quantitative regimes. The formulation above is a high-agreement regime, where ε is small and most local data are mutually consistent. At the opposite extreme, one may ask whether even sparse or noisy patterns of local agreement already force some nontrivial global structure. Such low-agreement phenomena are subtler and often require stronger expansion hypotheses.

Another important detail concerns the choice of sampling distribution on local views and overlaps. The formulation above uses the uniform distribution, in which each local view and each overlapping pair is sampled with equal probability. This is natural when the underlying system is regular. In more general settings, especially in higher dimensions, there are often more refined sampling distributions that better reflect the structure of the cover. Such choices are well studied in the theory of high-dimensional expanders, where one typically samples according to the stationary distribution of a natural random walk on the simplicial complex.

Finally, in quantum and higher-dimensional settings, local views are typically not ordinary functions. Instead, the sheaf values naturally take the form of chain complexes. Compatibility then involves relations across multiple levels rather than a single layer of coordinates. These richer structures are crucial in the study of quantum codes and local testability.

While the sheaf-theoretic framework is somewhat abstract, it provides the intuition behind, and guides the definition of, sheaf codes. In the first part of the thesis, we give concrete definitions and applications of sheaf codes. As we will see, many elegant structures appear naturally, including spectral sequences, homological perturbation theory, Poincaré duality, and cup products.

Part II Conformal Field Theory via Entanglement

We now turn to the review of quantum field theories, their challenges, and how the entanglement bootstrap framework offers a new perspective on their structure. This opens a new avenue in the sheaf-theoretic study of quantum states.

Quantum field theory and conformal field theory

Quantum field theory (QFT) is the language of modern physics. It unifies quantum mechanics with special relativity, underpinning both the Standard Model of particle physics and the theory of phase transitions in condensed matter. Its predictions — from the anomalous magnetic moment of the electron to the critical exponents near the ferromagnetic phase transition in magnets — are among the most precise and far-reaching in all of science.

Yet despite its success, QFT remains mysterious in many ways. The path integral, the central tool for defining a QFT, is not mathematically rigorous in most cases. Strongly coupled theories, where perturbative expansion fails, are difficult to analyze. A complete,

mathematically controlled understanding of Yang–Mills theory in four dimensions remains one of the outstanding challenges of mathematical physics.

Two special classes of QFTs are far better understood, and provide the first concrete proof of concept for the entanglement bootstrap program.

The first is *topological quantum field theory* (TQFT), whose observables depend only on the global topology of spacetime and not on the metric. As a result, TQFTs are often exactly solvable, and their ground states exhibit distinctive entanglement structures.

The second, and the focus of Part II of this thesis, is *conformal field theory* (CFT). A CFT is a QFT with an enlarged symmetry. Beyond the Poincaré symmetries of translations, rotations, and boosts, it is also invariant under dilations and special conformal transformations. Together they generate the conformal group. This additional symmetry is so powerful that many CFTs can be studied exactly or through non-perturbative methods such as the conformal bootstrap.

The reason CFTs occupy such a central place in the landscape of QFTs is the *renormalization group* (RG). Imagine gradually zooming out from a physical system. As we do so, short-distance details begin to wash out. Since this information can no longer be resolved, we integrate out its effects, thereby changing the effective description of the system. This process defines a flow on the space of all theories, called the *RG flow*. The fixed points of this flow are precisely the theories that look the same at every length scale: the scale-invariant theories. In many cases, the scale invariance is enhanced to full conformal invariance, so that the fixed points are CFTs.

If one pictures the full space of QFTs, the RG flows are the roads and the CFTs are the hubs. Every road connects two hubs: a generic QFT flows from a UV fixed point toward an IR fixed point, both of which are CFTs. Deforming a CFT by a relevant operator leads down a new road toward another hub in the infrared. The vast and seemingly intractable landscape of quantum field theories is thus organized by the smaller, better-understood

class of CFTs. Understanding CFTs is not merely a curiosity; they provide the scaffold on which the map of QFTs can be drawn.

Before developing the theory of entanglement bootstraps for TQFTs and CFTs, we present a different motivation through a puzzle arising from the perspective of sheaf theory.

Sheaf theory for quantum states

As we have seen, functions on a manifold are a canonical example of a sheaf. But what about quantum states? A natural way to construct a presheaf is from reduced density matrices. Let X be a set of vertices, each with a qudit degree of freedom. For every subset $U \subseteq X$, let $\mathcal{F}(U)$ be the set of density matrices on the subsystem U. For $V \subseteq U$, the restriction map $\mathcal{F}(U) \to \mathcal{F}(V)$ is given by tracing out the degrees of freedom in $U \setminus V$:

$$\mathrm{res}_{UV} : \rho_U \mapsto \rho_V = \mathrm{tr}_{U \setminus V} \rho_U. \tag{1.2}$$

However, does this form a sheaf? In particular, the identity axiom asks whether two states with the same reduced density matrices on every local region must be identical. The gluability axiom asks whether a family of local density matrices that is consistent on every pairwise overlap can be realized as the marginals of a global state.

For functions these are trivial. For quantum states, however, both fail. The identity axiom fails because quantum states can be globally distinct yet locally identical: the four ground states of the toric code on a torus all have the same reduced density matrices on any contractible region. The gluability axiom fails because of quantum monogamy: it is impossible to find a global state ρ_{ABC} whose marginals ρ_{AB} and ρ_{BC} are both EPR pairs, even though their shared marginal ρ_B is consistent.

These failures might seem fatal to the hope of applying sheaf theory to quantum states. But the right response is not to abandon the sheaf axioms. Rather, it is to identify the additional structure in physical ground states that restores them. This motivates the

search for local entropy conditions that recover the identity and gluability axioms.

Before turning to these conditions, we make one final digression. Developing a sheaf theory for quantum states is not merely a mathematical exercise. It also has implications for algorithms and the Yang–Mills mass gap problem.

Efficient representation and fast algorithms for physics

Given a Hamiltonian, we would hope to know its properties. Is it gapped or gapless? What is the ground state energy density? Which phase of matter does it belong to? The common theoretical perspective is that such tasks are intractable in general: one cannot expect efficient algorithms that work for all Hamiltonians. However, the Hamiltonians that arise in nature are not arbitrary. Even if the worst-case problem is hard, one may still hope that typical physically relevant cases are tractable. Indeed, if the ground state of a local Hamiltonian has local density matrices with the right structure, so that the sheaf axioms are restored, then the state can be efficiently represented and characterized.

We present a concrete example of this phenomenon. A quantum state on n qubits is, in general, a vector in a Hilbert space of dimension 2^n, which is exponentially large. Describing such a state in full is intractable for large n. But if the state is captured by its local density matrices, then it admits a cover by $\Theta(n)$ local regions, each involving only $O(1)$ qubits. Thus the state is specified by only $\Theta(n)$ parameters. Moreover, if the state is translationally invariant, then these local density matrices are related by translations, so the state is specified by only $O(1)$ parameters. In this case, an exponentially large wavefunction is replaced by a constant amount of local information. Remarkably, existing evidence suggests that this constant-size local description still retains enough information to determine the relevant physics of the state, including its phase of matter and whether that phase is gapped.

This perspective also provides a strategy for approaching the Yang–Mills mass gap

conjecture. The conjecture has two parts: first, that the QFT is well-defined in the UV, and second, that the resulting IR theory has a mass gap. One may hope to address the second part by showing that the ground state of the theory has local density matrices satisfying the approximate A0 and A1 conditions, as discussed below.

Recovering the sheaf axioms: the A0 and A1 conditions

Given the practical and theoretical motivations above, we now review a concrete set of local entropy conditions that restore the sheaf axioms for quantum states. These conditions were first proposed by Kim [1, 2] and later developed by Shi, Kato, and Kim [3] into a classification framework for topological order. They were further applied by Kim and Ranard [4] to classify gapped ground states with gapped boundaries in 2+1D. The two local conditions are called **A0** (local purifiability) and **A1** (extendability).

Although the framework applies in arbitrary dimension, we focus below on the case of a 2+1D system, where the state lives on $\mathbb{R}^2$. Technically, the state is defined on a lattice, but we will ignore this detail and treat the state as living in a continuum. A precise lattice formulation can be found in [5].

The conditions are defined in terms of the von Neumann entropy of reduced density matrices on bounded-size regions.

A0. For any bipartite region BC that is a topological disk of diameter at most $L = 5$, with C strictly contained in B, and with the annulus B having thickness at least $\ell = 1$ (as in Figure 1.1, left),

$$S(\rho_{BC}) + S(\rho_C) = S(\rho_B). \tag{1.3}$$

A1. For any tripartite region BCD that is a topological disk of diameter at most $L = 5$, with C strictly contained in BD, with the annulus B having thickness at least

$\ell = 1$, and with both BC and CD topological disks (as in Figure 1.1, right),

$$S(\rho_{BC}) + S(\rho_{CD}) = S(\rho_B) + S(\rho_D). \tag{1.4}$$

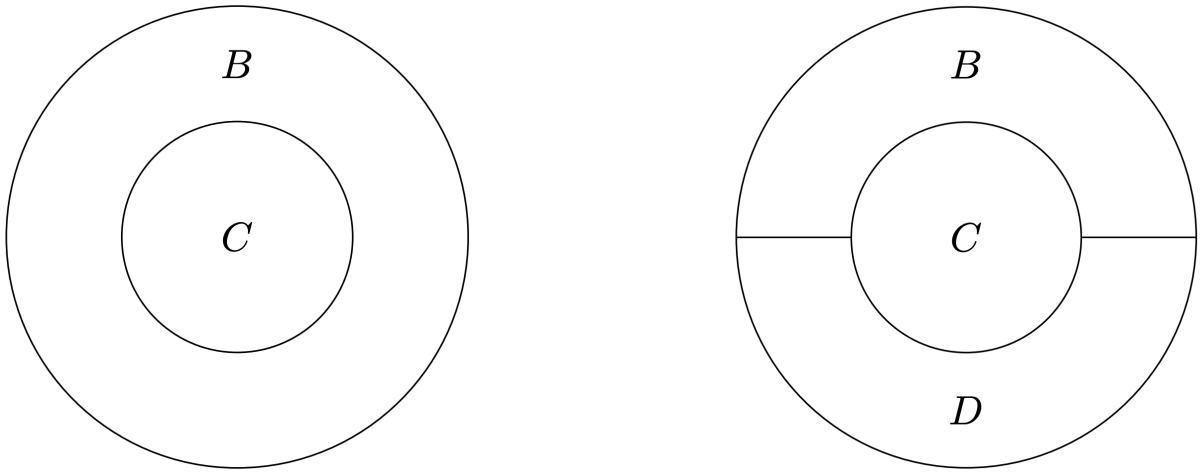

$S(\rho_{BC}) + S(\rho_C) = S(\rho_B)$ $\qquad$ $S(\rho_{BC}) + S(\rho_{CD}) = S(\rho_B) + S(\rho_D)$

Figure 1.1: Left: the region configuration for **A0**. The inner disk is C, the outer annulus is B, and the full disk is BC. Right: the region configuration for **A1**. The inner disk is C, while the outer annulus is partitioned into two connected regions B and D. Regions BC and CD each form a topological disk.

The conditions A0 and A1 strongly constrain the local entropy structure of the state. In particular, they correspond to saturation of two well-known inequalities in quantum information: subadditivity,

$$S(\rho_E) + S(\rho_C) \geqslant S(\rho_{EC}) \iff S(\rho_{BC}) + S(\rho_C) \geqslant S(\rho_B), \tag{1.5}$$

and strong subadditivity,

$$S(\rho_{BC}) + S(\rho_{EB}) \geqslant S(\rho_B) + S(\rho_{EBC}) \iff S(\rho_{BC}) + S(\rho_{CD}) \geqslant S(\rho_B) + S(\rho_D). \tag{1.6}$$

Here E denotes the environment in the purification of the local reduced density matrix, and we use the identity $S(\rho_X) = S(\rho_{\overline{X}})$ for a pure state.

The saturation of inequalities imposes significant constraints on the state. For A0, saturation forces ρ_{BC} to decompose as

$$\rho_{BC} \;=\; W_B(\rho_{B_1} \otimes |\psi_{B_2C}\rangle\langle\psi_{B_2C}|)\, W_B^\dagger,$$

where $W_B \colon \mathcal{H}_{B_1} \otimes \mathcal{H}_{B_2} \to \mathcal{H}_B$ is an isometry. The uncertainty in C is completely purified within B, hence the name *local purifiability*.

For A1, let A be the upper half-plane disjoint from BCD. Saturation implies that the conditional mutual information $I(A : C|B) = 0$, which implies that ρ_{ABC} forms a *quantum Markov chain*, and can be reconstructed from its marginals as

$$\rho_{ABC} = \rho_{AB}\, \rho_B^{-1}\, \rho_{BC}. \tag{1.7}$$

This allows ρ_{AB} to be extended to ρ_{ABC}, hence the name *extendability*.

With A0 and A1 in hand, modified versions of the identity and gluability axioms hold. Let $\mathcal{F}_{q,0}(U_i)$ and $\mathcal{F}_{q,1}(U_i)$ denote the set of density matrices on U_i that satisfy A0 and A1 respectively. And let $\mathcal{F}_q(U_i) = \mathcal{F}_{q,0}(U_i) \cap \mathcal{F}_{q,1}(U_i)$. More explicitly, $\rho_i \in \mathcal{F}_{q,0}(U_i)$ means that ρ_i satisfies A0 for every bipartite region $BC = U_i$, and $\rho_i \in \mathcal{F}_{q,1}(U_i)$ means that ρ_i satisfies A1 for every tripartite region $BCD = U_i$. This requirement is stronger than necessary: it suffices to require A0 and A1 for finitely many choices of partitions. This is easier to state in the discrete setting; see [5] for more details.

Given a region U, let

$$U^- := \{x \in U : \mathrm{dist}(x, \partial U) \geqslant \ell\} \tag{1.8}$$

denote its ℓ-interior.

Theorem 1.1 (Modified identity axiom)**.** *Let U be contractible, equivalently topologically a disk. Let $\{U_i\}$ be a cover of U by topological disks of diameter at most L, such that $\{U_i^-\}$ also forms a cover of U^-. If two states have local density matrices in $\mathcal{F}_q(U_i)$ and agree on every U_i, then they coincide on U^-.*

Remark 1.2. Replacing $\mathcal{F}_q(U_i)$ with $\mathcal{F}_{q,0}(U_i)$ is not sufficient for the identity axiom to hold. For example, consider a pair of nonabelian anyons separated by distance at least $2L$ on the vacuum background, so that any disk of diameter at most L can contain at most one of the two anyons. One can show that the reduced density matrices on all such disks satisfy A0, but it does not satisfy A1 when one of the anyons is placed in region C. Now, the pair of nonabelian anyons can carry a nontrivial fusion space, and different states in this fusion space are indistinguishable on every disk of diameter at most L. Thus there can be two different global states that agree on all regions in the cover. This shows why A0 alone does not imply the identity axiom, and why A1 cannot be dropped.

Remark 1.3. On the other hand, A1 alone is also not sufficient. The obstruction comes from superselection sectors. A simple example is provided by the cat states $|0\rangle^{\otimes n} \pm |1\rangle^{\otimes n}$. Their reduced density matrices satisfy A1 but not A0, and they are locally indistinguishable but globally distinct.

Note, however, that each superselection sector itself satisfies both A0 and A1. In fact, one can show that if the local density matrices satisfy A1, then the global state decomposes into superselection sectors, and each sector satisfies both A0 and A1.

Theorem 1.4 (Modified gluability axiom)**.** *Let U be any region and let $\{U_i\}$ be a cover of U by topological disks of diameter at most L. Any collection of local density matrices in $\mathcal{F}_{q,1}(U_i)$ that agree on all overlapping pairs lifts to a global density matrix on U.*

Thus, by restricting the local density matrices to lie in $\mathcal{F}_q(U_i)$, i.e. those that satisfy A0 and A1, we recover modified versions of the sheaf axioms for quantum states, as summarized in Table 1.2.

In many applications, one would also like the global state to arise as the ground state of a local Hamiltonian. Under the A0 and A1 conditions, such a realization is provided by the Hamiltonian reconstruction result of [6].

Theorem 1.5 (Hamiltonian reconstruction)**.** *Let U be any region and let $\{U_i\}$ be a cover of U by topological disks of diameter at most L, such that $\{U_i^-\}$ also forms a cover of U^-. Let ρ be a state on U such that $\rho|_{U_i} \in \mathcal{F}_q(U_i)$ for all i. Then, there exists a local Hamiltonian on U whose ground state σ satisfies $\sigma|_{U_i^-} = \rho|_{U_i^-}$ for all i.*

Remark 1.6. We conjecture that Hamiltonian reconstruction still holds under the weaker assumption that the local density matrices lie in $\mathcal{F}_{q,0}(U_i)$, i.e. that they satisfy A0 but not necessarily A1. We do not know how to prove this yet.

Table 1.2: The identity and gluability axioms for classical functions, quantum states, and their robust version. The robust versions remain to be proven and should be read schematically: for example, the L^1-norm may be replaced by other appropriate norms depending on the setting.

	Identity (Uniqueness)	**Gluability (Existence)**
Classical	$\forall i : s\|_{U_i} = t\|_{U_i} \in \mathcal{F}_c(U_i)$ $\implies s = t.$	$\forall i,j : s_i\|_{U_i\cap U_j} = s_j\|_{U_i\cap U_j}$, $s_i \in \mathcal{F}_c(U_i)$ $\implies \exists s, s\|_{U_i} = s_i.$
Quantum	$\forall i : \rho\|_{U_i} = \sigma\|_{U_i} \in \mathcal{F}_q(U_i)$ $\implies \rho\|_{U^-} = \sigma\|_{U^-}$ for contractible U^-.	$\forall i,j : \rho_i\|_{U_i\cap U_j} = \rho_j\|_{U_i\cap U_j}$, $\rho_i \in \mathcal{F}_q(U_i)$ $\implies \exists \rho : \rho\|_{U_i} = \rho_i.$
Robust Quantum 1	Given $\sigma\|_{U_i} \in \mathcal{F}_q(U_i)$, $\{\rho\|_{U^-} : \|\rho\|_{U_i} - \sigma\|_{U_i}\|_1 < \varepsilon$, $\rho\|_{U_i} \in \mathcal{F}_q(U_i)\}$ is contractible for contractible U^-.	$\forall i,j : \|\rho_i\|_{U_i\cap U_j} - \rho_j\|_{U_i\cap U_j}\|_1 < \varepsilon$, $\rho_i \in \mathcal{F}_q(U_i)$ $\implies \exists \rho\ \forall i : \|\rho\|_{U_i} - \rho_i\|_1 < \delta(\varepsilon)$, $\rho\|_{U_i} \in \mathcal{F}_q(U_i).$
Robust Quantum 2	Given $\sigma\|_{U_i} \in \mathcal{F}_{q,\varepsilon}(U_i)$, $\{\rho\|_{U^-} : \|\rho\|_{U_i} - \sigma\|_{U_i}\|_1 < \varepsilon$, $\rho\|_{U_i} \in \mathcal{F}_{q,\varepsilon}(U_i)\}$ is contractible for contractible U^-.	$\forall i,j : \|\rho_i\|_{U_i\cap U_j} - \rho_j\|_{U_i\cap U_j}\|_1 < \varepsilon$, $\rho_i \in \mathcal{F}_{q,\varepsilon}(U_i)$ $\implies \exists \rho\ \forall i : \|\rho\|_{U_i} - \rho_i\|_1 < \delta(\varepsilon).$

The entanglement bootstrap definition for TQFTs and the robust axioms

Given A0 and A1, it is natural to ask which physical systems satisfy these conditions. They are expected to describe a class of gapped ground states known as string-net states, including the toric code and quantum double models. These models are topologically ordered and are described by 2+1D TQFTs.

However, not all systems regarded as topologically ordered satisfy A0 and A1. For example, quantum Hall states and chiral spin liquids exhibit topological phenomena, but they fail to satisfy A0 and A1 exactly [7]. Instead, they satisfy robust versions of these axioms, where the entropy equalities are relaxed to inequalities with small error ε. Many questions remain open in this approximate regime, as summarized in Table 1.2.

More broadly, the entanglement bootstrap framework suggests a new route toward defining TQFTs. Since TQFTs describe many gapped phases, one might have hoped that all gapped phases admit a TQFT description. However, the discovery of fracton phases shows that TQFTs form a more restrictive class than general gapped phases of matter. In this framework, A1 provides the key additional condition that distinguishes TQFTs from more general gapped phases.

On the mathematical side, the standard definition of a TQFT is functorial, formulated in terms of cobordisms. However, connecting this definition directly to physical systems remains challenging. For example, in the cobordism-based formulation, the chiral central charge is defined only modulo 24, rather than as a real number.

Thus A0 and A1, together with their robust versions, provide a new, physically motivated definition of TQFTs from local entanglement structure, while also pointing toward a formulation that could be made mathematically rigorous.

The natural next frontier is to extend this philosophy beyond gapped phases to

gapless ones, which is the subject of the second part of the thesis. We obtain local entropy conditions for 1+1D CFTs that are analogues of A0 and A1, and use them to systematically explore the space of CFTs.

Organization of this Dissertation

Chapter 2 presents the t-dimensional cubical complex construction and proves systolic and co-systolic expansion at every level, yielding the first quantum locally testable code family with linear dimension, near-linear distance, and near-constant soundness.

Chapter 3 develops the cup product on sheaf cochain complexes, proves the triorthogonality criterion for transversal CCZ gates, establishes Poincaré duality, and connects the framework to magic-state distillation via the companion construction of Golowich–Lin.

Chapter 4 derives the vector fixed-point equation for CFT ground states, extracts all six global conformal generators as entanglement-Hamiltonian combinations, and validates the characterization numerically on critical lattice models.

Chapter 5 describes the systematic CFT search algorithm, applies it to four-qubit and four-qutrit systems, recovers known critical points, and presents roughly 20 new candidate critical points in the qutrit sector.

Bibliography

[1] I. H. Kim, On the informational completeness of local observables, *arXiv:1405.0394* (2014).

[2] I. H. Kim, Conservation Laws from Entanglement, https://www.physics.usyd.edu.au/quantum/Coogee2015/Presentations/Kim.pdf (2015).

[3] B. Shi, K. Kato, and I. H. Kim, Fusion rules from entanglement, *Ann. Phys.* **418**, 168164 (2020).

[4] I. H. Kim and D. Ranard, Classifying 2D topological phases: mapping ground states to string-nets, *arXiv:2405.17379* (2024).

[5] B. Shi, Anyon theory in gapped many-body systems from entanglement, Ph.D. thesis, The Ohio State University (2020).

[6] I. H. Kim, T.-C. Lin, D. Ranard, and B. Shi, Strict area law implies commuting parent Hamiltonian, *arXiv:2404.05867* (2024).

[7] X. Li, T.-C. Lin, J. McGreevy, and B. Shi, Strict area law entanglement versus chirality, *Phys. Rev. Lett.* **134**, 180402 (2025).

[8] I. H. Kim, B. Shi, K. Kato, and V. V. Albert, Chiral central charge from a single bulk wave function, *Phys. Rev. Lett.* **128**, 176402 (2022).

[9] I. H. Kim, B. Shi, K. Kato, and V. V. Albert, Modular commutator in gapped quantum many-body systems, *Phys. Rev. B* **106**, 075147 (2022).

Chapter 2

Quantum Locally Testable Codes via Higher-Dimensional Cubical Complexes

Chapter 2 is based on joint work with Irit Dinur and Thomas Vidick (arXiv:2402.07476).

2.1 Introduction

Expander graphs are bounded-degree graphs with strong connectivity, or more precisely *expansion*, properties. Explicit constructions of expander graphs are non-trivial, but, by now, abound; their use is ubiquitous in algorithms design, complexity theory, combinatorics, and other areas. High-dimensional expanders generalize the expansion requirements of expander graphs to higher-dimensional structures, composed of vertices and edges as well as higher-dimensional faces. Constructions of high-dimensional expanders (HDX) are difficult and comparatively few techniques are known. In recent years HDX have found impactful applications to randomized algorithms, complexity theory, and the

design of error-correcting codes, among others. In addition to their intrinsic interest as combinatorial/geometric objects, the increasing range of applications further motivates the study of HDX.

In this paper we introduce a natural generalization of a family of two-dimensional (graphs are one-dimensional) expanders introduced recently and independently in [1, 2] and used to simultaneously construct the first good classical locally testable codes (LTC) [1, 2] and the first qood quantum low-density parity check codes (qLDPC) [2]. We extend this construction to arbitrary dimensions and prove lower bounds on its *cycle* and *co-cycle* expansion, which we define below. The bounds that we obtain depend on underlying parameters of the complex, namely the (spectral) expansion of an associated family of Cayley graphs and a certain robustness-like property of an associated family of constant-sized *local codes*. The main application of our construction, which has been our primary motivation, is towards the construction of quantum locally testable codes (qLTC). By instantiating our complex using a suitable abelian lift of an expander, as described in [3], and leveraging recent results on product expansion of tuples of random codes over large fields [4], we obtain the following result.

Theorem 2.1 (Informal, see Corollary 2.18)**.** *There exists an explicit family of* $[n, k, d]$ *quantum LDPC codes where* $n \to \infty$ *and* $k = \Omega(n), d = \Omega(n/(\log n)^3)$*, such that the codes are CSS codes with parity checks of weight* $O(1)$ *and are qLTC with soundness* $\rho = \Omega(1/(\log n)^3)$*.*

We refer to the main text for the definition of a quantum CSS code and the soundness of a qLTC. In the theorem, by "explicit" we mean that the parity check matrix for the code of dimension n in our family can be computed in time polynomial in n. We note that the soundness parameter in the theorem can be improved to a constant by allowing a larger check weight $O((\log n)^3)$.[1] Consequently, one can obtain an explicit CSS code with

[1]We omit the proof of this fact, which follows from standard arguments; briefly, one uses (explicit)

$k = \Omega(n)$, $d = \Omega(n/(\log n)^3)$, parity checks of weight $O((\log n)^3)$, and soundness $\rho = \Omega(1)$. Additionally, the distance parameter can also be improved to the optimal through distance amplification [5]. In particular, we obtain a CSS code with $k, d = \Omega(n)$, parity checks of weight $O(\text{poly} \log n)$, and soundness $\rho = \Omega(1)$.

In contrast, previous best constructions of qLTC had $k = 1$, $d = \Theta(\sqrt{n})$, $\rho = \Omega(1/(\log n)^2)$ for [6] and $\rho = \Omega(1/\log n)$ for [7]. Additional tradeoffs are possible (see Section 2.1.4 for further discussion); however, prior to our work no qLTC codes with even polynomial (let alone linear) dimension were known while keeping the locality and soundness at most polylogarithmic and at least inverse-polylogarithmic respectively.

2.1.1 Construction

For any integer $t \geqslant 1$ we construct a t-dimensional cubical chain complex that generalizes the Sipser-Spielman construction of expander codes [8] (case $t = 1$) and the lifted product codes from [9] (case $t = 2$) to higher dimensions $t \geqslant 3$.

The construction is based on two ingredients. Firstly, there is the cubical complex X, which is a purely geometric structure (technically, a graded incidence poset with certain good expansion properties). Secondly, there is a system of local coefficients (a sheaf $\mathcal{F}$, in the terminology of [10]) that is constructed from a family of classical codes of constant dimension and good "robustness" properties. The chain complex is a sequence of coboundary maps from i-cochains to $i + 1$-cochains,

$$C^0(X, \mathcal{F}) \to C^1(X, \mathcal{F}) \to \cdots \to C^t(X, \mathcal{F}).$$

We expand on these ingredients below.

parity samplers to combine parity checks to form higher-weight checks. This process does not change k and d.

Cubical complex The (higher-dimensional) cubical complex X can be constructed from any set G of size $N = |G|$, and finite subsets of permutations $A_1, \ldots, A_t$ of G of size $n = |A_i|$.[2] The sets A_i should be closed under inverse and such that permutations taken from different sets commute.

For example, in the case of $t = 2$ we can take G to be any finite group and let A_1 act by multiplication on the left, and A_2 by multiplication on the right. This choice, with constant-size sets A_1 and A_2 of expanding generators, underlies the square complex from [1, 2]. For larger $t \geqslant 3$, it is not clear how to generalize the fact that multiplication on the right commutes, as an action, with multiplication on the left, $(ag)b = a(gb)$. A natural approach is to take G an abelian group, such as $G = \mathbb{Z}_2^{\log N}$, and let A_i be subsets of generators of G acting on it by multiplication. This choice has the disadvantage that each A_i must have size at least logarithmic in $|G|$ (otherwise the graphs won't expand enough). A better cubical complex with *constant* degree can be obtained through an abelian lift [3]. We describe this approach, which gives us the best parameters overall, in Section 2.3.5. Other, more structured choices may be possible.

Irrespective of the specific choice of G and $\{A_i\}$, the resulting complex $X(G; \{A_i\}_{i=1}^t)$ is 2^t-partite, with vertex set $G \times \{0,1\}^t$ and k-dimensional faces associated with the $2^{|S|}$ vertices $\{g \cdot \prod_{j \in T} a_j : T \subseteq S\}$, where $S \subseteq \{1, \ldots, t\}$ has size $|S| = k$ and $a_j \in A_j$. Edges coming from permutations in A_i are said to have direction i, and more generally, a face has type S if S is the set of directions of edges it contains. See Section 2.3.1 for a complete description, and Figure 2.1 for an illustration.

The only requirement on the permutation sets A_i, other than pairwise commutation, is that the associated graphs $\mathrm{Cay}(G, A_i)$, with vertex set G and edge set $\{g, \pi(g)\}$ for $\pi \in A_i$, have good spectral expansion[3]. As we will show, this suffices to imply expansion of

[2] In general it is possible to make the integer n also depend on $i \in \{1, \ldots, t\}$; for simplicity we assume that there is no such dependence.

[3] Traditionally, $\mathrm{Cay}(G, A_i)$ is a connected graph. In our construction, what we call "$\mathrm{Cay}(G, A_i)$" may

a natural set of high-dimensional random walks on the cubical complex that are used in the proof of (co)-cycle expansion.

Local coefficients The cubical complex X is turned into a chain complex $C(G; \{A_i\}_{i=1}^t)$ by endowing it with a system of local coefficients. In the terminology of [10], we put a sheaf structure over the geometrical complex X. To each face f of the complex, which can be a vertex, an edge, a square, a cube, etc., we associate a small coefficient space that is in some sense dual to it. For example, suppose $t = 3$ and our complex X is a cubical complex. Fix a 1-dimensional face, namely an edge. Its space of coefficients is a space of 2-dimensional matrices. If the edge is labeled by direction 1, then the matrix dimensions are labeled using the other directions of the complex, namely $2, 3$. These spaces are described in detail in Section 2.3.2. We introduce a system of linear maps $\mathcal{F}_{f \to f'}$ from the coefficient space of a k-face f to that of a $(k+1)$-face f'. These maps will be constructed from a collection of parity check matrices $h_1, \ldots, h_t$, where each $h_i \in \mathbb{F}_q^{m_i \times n}$ for some $m_i \leqslant n$ and some finite field $\mathbb{F}_q$. These matrices will be used as follows. Suppose for example that, as before, $t = 3$, f is an edge and f' is a face that contains the edge f, as well as another (pair of) edges in direction 2 (and a fourth edge, parallel to f). Then the map $\mathcal{F}_{f \to f'}$ is obtained by taking the coefficient $x(f) \in \mathbb{F}_q^{m_2 \times m_3}$ associated with f, expanding it in the second direction by applying h_2^T to obtain $h_2^T x(f) \in \mathbb{F}_q^{n \times m_3}$, and then restricting to the row indexed by the face f' (which is uniquely identified by its edge in direction 2, labeled by an element of A_2 which has size n) and returning the resulting row vector, an element of the coefficient space $\mathbb{F}_q^{m_3}$ associated with the face f'. These maps are described in much more detail in Section 2.3.3 (where they are denoted co-res$_{f,f'}$).

be disconnected. The specific notion we require is introduced in Definition 2.47 — informally, Cay(G, A_i) should be a union of not too many disjoint copies of spectral expanders.

2.1.2 Two-way robustness and product expansion

Our proofs of cycle and co-cycle expansion of the complex follow from the expansion of the graphs $Cay(G, A_i)$ and from a property of the local parity check matrices $(h_1, \ldots, h_t)$ which we call *two-way robustness*. A collection of linear transformations $h_i : \mathbb{F}_q^n \to \mathbb{F}_q^{k_i}$, $1 \leqslant i \leqslant t$, is said to be *two-way robust* if the t-fold homological product of $h_1^T, \ldots, h_t^T$ is *coboundary expanding* in all dimensions $i < t$, and the same holds also for $(h_1^\perp)^T, \ldots, (h_t^\perp)^T$, where for each i, $h_i^\perp : \mathbb{F}_q^n \to \mathbb{F}_q^{n-k_i}$ is (some) matrix that satisfies $h_i^\perp h_i^T = 0$, i.e. the code with parity check matrix $h_i^\perp$ is dual to the code with parity check matrix h_i. We refer to Section 2.5.2 for the precise definition. The reason for considering the t-fold homological product is that this is precisely the local view of the complex from a vertex. Namely, the view when restricting to cells that contain a fixed vertex. So our requirements for co-cycle expansion boils down to a local condition (two-way robustness) on the restriction of the code to cells around a vertex and a global condition (spectral expansion of each $\mathrm{Cay}(G, A_i)$). This is in perfect analogy to previous works on co-cycle expansion of simplicial complexes [11, 12] and of cubical complexes [1, 2]. In both types of complexes, like here, global co-cycle expansion followed from some form of spectral expansion together with coboundary expansion at the links.

In case $t = 2$, the notion of two-way robustness coincides with product expansion as in [2] (and with agreement testability, see [1]). However, for $t \geqslant 3$ these notions diverge, as studied in [13]. In case $t = 2$, Panteleev and Kalachev proved that a pair of random maps h_1, h_2 satisfy, with high probability, a version of two-way robustness. The robustness parameters were later improved in [14, 15], but again, only for a pair of random maps, namely $t = 2$. For three or more maps it was conjectured in [14] that random maps are product expanding. Recently, Kalachev and Panteleev [16] showed that any (constant) number of random maps are product expanding, provided that the maps are over a large enough field — of size $2^{\mathrm{poly}(n)}$, see Theorem 2.36 for a precise statement. Because for us n

can be taken a constant, this large field size is not an obstacle.

Our proof of global cycle and co-cycle expansion, however, requires *two-way robustness*, which is a property that is a priori stronger than product expansion. While product expansion addresses coboundary expansion in level $i = t - 1$, two-way robustness requires coboundary expansion for all levels $i < t$. In Section 2.5.3 we prove a reduction from the latter to the former, thereby completing all requirements on the sheaf structure. An interesting component in this reduction is a proof that expansion with respect to coefficients in a vector space V implies the same with respect to V^M, with no dependence on the expansion parameter on M.

2.1.3 Main result on expansion

To state our main result we need to define the notions of distance and expansion that we consider. We will use the term systolic (resp. co-systolic) distance to denote the lowest weight of a k-cycle (co-cycle) that is not a k-boundary (co-boundary). Namely, if $C(X)$ is a chain complex with boundary map ∂ and coboundary map δ, then

$$\mu_{\text{syst}}(k) \;=\; \min\Big\{|x| \;:\; x \in \ker \partial_k - \text{im}\, \partial_{k+1}\Big\}, \qquad \mu_{\text{cosyst}}(k) \;=\; \min\Big\{|x| \;:\; x \in \ker \delta_k - \text{im}\, \delta_{k-1}\Big\}\,.$$

Here we employ the usual notation $\partial_k : C_k(X) \to C_{k-1}(X)$ and $\delta_k : C^k(X) \to C^{k+1}(X)$, with $C_k(X)$ and $C^k(X)$ the (isomorphic) spaces of chains and co-chains respectively.

We establish lower bounds on the *systolic* and *co-systolic* distance of the complex, as well as lower bounds on the cycle and co-cycle expansion, which we define next. The lower bounds are attained via bounding a locally-minimal version of distance (see Definition 2.4), using the strategy initiated in [11, 12], and used also in [2, 1] as well as several other works. The *cycle expansion* of the complex is the smallest size of the boundary of a k-chain, relative to the distance of the k-chain to the set of k-cycles. More formally,

$$\varepsilon_{\text{cyc}}(k) \;=\; \min\left\{\frac{|\partial_k(x)|}{dist(x, \ker \partial_k)} \;:\; x \in C_k(X) - \ker \partial_k\right\},$$

and similarly, co-cycle expansion is defined by

$$\varepsilon_{\mathrm{cocyc}}(k) \;=\; \min\left\{ \frac{|\delta_k(x)|}{dist(x, \ker \delta_k)} \;:\; x \in C_k(X) - \ker \delta_k \right\}.$$

In the definitions above, $dist(\cdot,\cdot)$ and $|\cdot|$ refers to the Hamming weight, or, in the case of a sheaf complex, the *block-Hamming weight*. For precise definitions of all these notions, we refer to Section 2.2.1. We remark that co-cycle expansion was first studied in [11] in the simplicial setting (earlier works [17, 18] introduced the notion of coboundary expansion, which coincides with co-cycle expansion for exact chains). The term co-systolic expansion was introduced in [12], and combined both co-systolic distance and co-cycle expansion. In retrospect it makes sense to separate into two definitions as above.

Our main result is a general result on (co-)systolic expansion of dimension-t chain complexes of the form $C(G; \{A_i\})$ described in Section 2.1.1 above. We bound the expansion parameters of the complex as a function of the spectral expansion of the graphs $\mathrm{Cay}(G, A_i)$ and the two-way robustness of the local codes $\{h_i\}$. We note that for $t = 1$ and n a constant, this result recovers the special case that Sipser-Spielman codes based on good enough expander graphs and good local codes have a linear distance. For $t = 2$ and n a constant, we recover the expansion properties of square complexes that underlie the constructions of locally testable and quantum low-density parity check codes from [1, 2]. We refer to Theorem 2.15 for a statement of this general result.

Here we instantiate the general result with a specific choice of $(G, \{A_i\})$ that are obtained from an abelian lift as in [3] (this construction is described in Section 2.3.5), and local codes with two-way robustness that can be obtained by exhaustive search (since in this construction the degree n is constant; existence is guaranteed as explained in Section 2.1.2 above). This gives an explicit construction of a (co-)systolic expander of dimension t, as stated in the following corollary.

Theorem 2.2 (Main construction, restated as Corollary 2.17)**.** *Let $t \geqslant 2$ be an integer.*

There is an explicit construction of a chain complex $C(G;\{A_i\})$ and local codes $\{h_i\}$ such that the k-chain space C_k has dimension $\Theta(N)$ for every $0 \leqslant k \leqslant t$ and furthermore

1. *The co-chain complex C^* has co-systolic distance $\mu_{\mathrm{cosyst}}(k) = \Theta(|X(k)|/(\log N)^{t-1})$ for every $0 \leqslant k \leqslant t-1$ and co-cycle expansion $\varepsilon_{\mathrm{cocyc}}(k) = \Theta(1/(\log N)^{t-1})$ for every $0 \leqslant k \leqslant t-2$;*

2. *The chain complex C_* has systolic distance $\mu_{\mathrm{syst}}(k) = \Theta(|X(k)|/(\log N)^{t-1})$ for every $1 \leqslant k \leqslant t$ and cycle expansion $\varepsilon_{\mathrm{cyc}}(k) = \Theta(1/(\log N)^{t-1})$ for every $2 \leqslant k \leqslant t$.*

In the theorem, the constants implicit in the $\Theta(\cdot)$ notation depend exponentially on t, so our bounds get exponentially worse as t increases. The main application we have in mind, to quantum locally testable codes, requires only $t = 4$, and we did not attempt to optimize the dependence on t.

2.1.4 Towards good quantum locally testable codes

The existence of good quantum error-correcting codes is one of the pillars that underlie the promise of quantum computing to deliver impactful applications in the long term. Good codes with *low-density parity checks* are particularly sought after for their potential application to fault-tolerance [19] and quantum complexity theory [20]. In the past few years a flurry of works obtained better understanding and better parameters [21, 22, 23, 9] culminating in a construction of good quantum LDPC codes [2]. This was followed by a couple of related variants [24, 25, 15]. These constructions rely on a well-known connection between the design of quantum codes and chain complexes. In particular, all existing constructions of good qLDPC rely on virtually the same length-3 chain complex, the "square complex" from [1] (which is also a special case of "balanced product codes" from [23] and of "lifted product codes" from [9]).

Quantum *locally testable* codes (qLTC) were introduced in [26] as a natural quantization of the notion of local testability for classical codes. Although the formal definition, first given in [27], is somewhat technical, quantum LTCs have the same intuitive basis as their classical counterparts — informally, whenever a word (quantum state) is at a certain relative distance δ from the codespace, this word must violate a fraction $\rho\,\delta$, for some constant $\rho > 0$ referred to as the *soundness* of the LTC, of the parity checks. The first constructions of qLTCs with soundness that scales better than inverse polynomial were given in [6, 7]. These achieve soundness that scales as $1/\log^2(n)$ and $1/\log(n)$ respectively; however, the distance of the code is small ($\Theta(\sqrt{n})$), the codes have constant dimension, and the weight of the parity checks is logarithmic.

Prior to our work quantum LTCs with good locality and soundness and with either reasonable dimension (say, above poly-logarithmic) or super-quadratic ($n^{\frac{1}{2}+\varepsilon}$) distance were not known to exist. Such codes are expected to have applications in quantum complexity theory. Until recently the most prominent was the famous "NLTS conjecture" [27]; which was known to follow from qLTC with good distance and soundness (without regard to dimension). However, the NLTS theorem was recently proven without the use of qLTC [20], relying only on the recent qLDPC constructions from [2, 25].

While there is no formal connection with the long-standing quantum PCP conjecture [28], it is natural to hope that progress on the former may eventually lead to progress on the latter. One reason is that LTCs are strongly tied to PCPs in the classical world, which, perhaps more philosophically, is due to the inherent local-to-global nature of both questions, where a global property (being a codeword in the case of LTC and being a minimum energy state in the case of qPCP) is to be tested by random local tests. The local to global aspect manifests in high dimensional expanders, shedding some insight as to the connection between these objects.

Recall our main result on quantum locally testable codes, which is stated as Theo-

rem 2.1 above. This result is obtained from Theorem 2.2, for the case $t = 4$, using standard arguments in the construction of quantum (CSS) codes. This result falls short of constructing good quantum LTCs since the parameters differ from ideal parameters by an inverse poly-logarithmic factor. Still, these parameters mark a significant improvement compared to previously known qLTCs [6, 7], which achieve $k = 1$, $d = \Theta(\sqrt{n})$, and $\rho = \Omega(1/\log n)$. (Additional tradeoffs are possible between the parameters, for example increasing the dimension k at the cost of a decrease of the relative distance; we refer to [29, 5] and in particular Table 3 in [5] for details.

The codes from Theorem 2.1 are explicit CSS codes, and in particular have an efficient encoder. Moreover, although we do not show this explicitly, we expect these codes to have linear-time decoders and noisy syndrome decoders (often referred to as single-shot decoders), along the same lines as [15, 30]. Although there is no formal connection known, we observe that qLTC may be of additional relevance to the problem of noisy syndrome decoding. This is because a *simpler* problem than decoding from a noisy syndrome is deciding, given an approximate bound on the size of the syndrome, if the state is close to the codespace or not—which is precisely the problem solved by qLTCs. This observation suggests that, beyond the basic LDPC requirement, the LTC condition may have applications to e.g. fault-tolerance that have not yet been fully explored.

2.1.5 Techniques

Our construction of a higher-dimensional cubical chain complex, when instantiated at $t = 2$, is not identical to the chain complex of [2], but rather can be seen as a "90° rotation" of it as in [15]. The vector spaces participating in the chain complex are ordered by geometric dimension, a feature that facilitates generalization to $t > 2$. The price is that our construction is not self-dual, in the sense that the geometric properties of the complex are different from those of the co-complex. Therefore, the argument for showing co-cycle

expansion is different from the argument for showing cycle expansion. Nevertheless, the latter is based on a reduction to the former for a *different* co-complex that is, in some sense, dual to the original one.

Cocycle expansion Our proof for co-cycle expansion follows the general template of earlier works on cocycle expansion in HDX [31, 12], and the analogous statements for qLDPC codes and for LTCs [2, 1].

Let us outline this argument for showing cocycle expansion at the highest level of the complex, i.e. for co-chains in dimension $k = t - 1$, which is the most interesting case. Consider an element $x \in C^{t-1}(X)$ such that $\delta(x) = 0$, and assume it is locally minimal, that is, assume its weight is minimal with respect to adding coboundaries of local elements. We let $\mathcal{A}$ denote the set of *active faces*, which is the subset of geometric faces $f \in X(t-1)$ such that $x(f) \neq 0$. The core of the proof consists in showing that the condition $\delta(x) = 0$ implies that the set $\mathcal{A}$ does not expand according to a natural random walk W on $X(t-1)$. However, the random walk is such that it *is* expanding (this is where the assumption that $\mathrm{Cay}(G, A_i)$ expands is used); therefore either $\mathcal{A} = \emptyset$ or $|\mathcal{A}|$ is large, which is the desired conclusion. To show that the set $\mathcal{A}$ does not expand, we show that most neighbors of an $f \in \mathcal{A}$ according to W lie in $\mathcal{A}$. This step uses the definition of δ to argue that the condition $\delta(x) = 0$ implies that many neighbors of f under W must be in $\mathcal{A}$. The robustness condition is used crucially at this step.

Indeed, one of the contributions of this work is in coming up with the appropriate definition for the random walk W, together with the definition for two-way robustness so that this argument goes through.

We note that the walk W is defined as a mixture of $(t-1)$ different walks, each of which goes "down" $j \in [t-1]$ steps in the complex to a $(t-1-j)$-dimensional sub-face of f, takes one parallel step, and then "up" again j steps to a "neighbor" of f satisfying

certain conditions. As such, this walk is similar to iterations of the "down-up" random walk that is familiar in the combinatorial analysis of high-dimensional expanders. It is often used to circumvent certain non-expanding link-like obstacles that would appear in a more naive choice of random walk W. In this sense the extension of the arguments from [15, 2, 1] from $t = 2$ to higher dimension $t > 2$ is analogous to the move from [31] to [12]. We refer to Section 2.7 for details on this part of the argument.

Cycle expansion Our argument for cycle expansion is completely different from the argument for co-cycle expansion described above. In fact, the argument is an indirect reduction to the case of co-cycle expansion, generalizing a reduction that appeared in [15], for the case $t = 2$. Moving to higher dimensions $t > 2$ is more involved, consisting of a nice arrow-chasing argument on a double chain complex which we describe next.

The idea is to define a new "dual" chain, whose coefficients are local views from geometric objects of the main chain. In the dual chain, the coefficient at a vertex tells us everything about the k-faces touching that vertex. We define two kinds of boundary maps for this dual chain, and show that they commute nicely with each other, giving rise to a double chain complex (shown in Figure 2.3). The argument proceeds as follows. Given a low weight k-chain x such that $\partial x = 0$, we wish to find a preimage z such that $\partial z = x$. We first look at the dual chain $\{x_v\}$ obtained by local views of x. We take a preimage z_v separately for each x_v. This is possible because the local maps are exact and $\partial x = 0$ implies $\partial x_v = 0$ for all v. The preimages z_v are also local views, but they don't necessarily agree with each other, and this is where things become interesting. We move to look at the differences between local views of adjacent vertices, assigning these differences as local views to the *edges*. These new local views have zero boundary because of the commutativity of the arrow diagram. As a result we are in a similar situation as before, except that we moved one dimension up, from vertices to edges. We proceed inductively

(again using the exactness of the (local) boundary map to find local preimages, take their differences, and so on). This inductive arrow-chasing ends in one of two cases. Either we reach a stage where all local views agree, or we push all the way to the end of the chain without reaching agreement. In the former case, we have obtained a global "corrections" chain which can then be propagated all the way back down, leading to a global solution z such that $\partial z = x$. The interesting case however is when one reaches the end of the chain, without agreement. In this case the key observation is that local views, which lie on the higher-dimensional faces of the complex (i.e. t-cubes) are tensor codewords (because their boundary is 0) and can thus be decoded. The "decoded" object can be interpreted as a co-cycle at the $(t-1)$-st level of a complex with dual local codes $h_1^\perp, \ldots, h_t^\perp$. Applying the co-distance argument, we find a pre-image co-cycle. When re-encoded, this pre-image co-cycle provides corrections which can once again be propagated down the chain to define z.

2.1.6 Discussion

In this work we construct quantum LTCs with parameters that are "almost" good, up to polylogarithmic factors. One approach for removing the poly-logarithmic factors is to venture away from cubical complexes and study chain complexes based on simplicial high-dimensional expanders, as done in [21]. The advantage is that constant degree expanding complexes are known in any dimension t. In contrast, our complexes are also constant-degree, but unfortunately their expansion guarantees hold only for sets of density up to $1/\operatorname{poly}\log n$. In [21], the distance of the quantum code follows from the systolic and co-systolic distance of the complex. Unfortunately, the systolic distance is inherently sublinear (in their construction it is polylogarithmic). This obstacle theoretically should also appear for cubical complexes, but it is circumvented through the addition of local codes. Why not add local codes to high-dimensional expanders? This seems like a reasonable

direction; the main difficulty is that, unlike in the cubical case, the local structure around a vertex does not have a natural product structure. Two-way robustness seems quite tricky when the structure is not product, and is currently not known.

Generally speaking the connection between HDX and quantum codes has been extremely fruitful, leading to insights such as the distance balancing method from [21]. This connection crucially requires one to consider both the complex and the co-complex, to reflect the symmetry between X and Z parity checks in the definition of a quantum CSS code. Most constructions of HDX, including ours, do not exhibit a perfect symmetry between complex and co-complex. For our complexes, we are nevertheless able to reduce cycle expansion to co-cycle expansion of a *related* complex. This proof technique could be of independent interest, and lead to new constructions of HDX exhibiting expansion in both directions.

Quantum LDPC codes are intrinsically linked with the study of high-dimensional surfaces, and indeed their recent construction leads to new topological objects [32] (and vice-versa [33, 34]). It will be interesting to determine if our higher-dimensional construction has similar consequences. Conversely, one may hope that a construction of a high-dimensional cubical complex with constant degree could be obtained by leveraging connections with group theory, a central source of constructions of simplicial complexes and HDX [35]. Cubical complexes with constant degree and strong expansion properties are known to exist [36], but these don't seem to support a system of local coefficients that would lead to a useful chain complex.

Of course one of the main focus points of the area, which we leave entirely open, is the possible application of qLTC constructions to quantum complexity theory. By providing an explicit family of candidate good qLTC we uncover a concrete object that may form the basis for later explorations of the connections with complexity.

Organization We start with some preliminaries regarding notation, chain complexes, quantum codes and local testability in Section 2.2. In Section 2.3.4 we give an overview of our results. In Section 2.3 we describe our high-dimensional cubical complex. In Section 2.5 we introduce an auxiliary "local" chain complex, that will be used in the analysis. In this section we prove that a random tuple of codes, for large enough field size, is two-way robust (relying on the recent result on product expansion of such codes). In Section 2.6 we describe and analyze some random walks on the cubical complex, that will also be used in the analysis. Finally, in Section 2.7 and Section 2.8 we prove our main results, lower bounds on the locally co-minimal distance and distance respectively at each level of our chain complex.

Acknowledgments We thank Louis Golowich for pointing us to [3] which leads to a construction with constant (rather than logarithmic) degree. We thank Hayata Yamasaki and Quynh T. Nguyen for pointing out an error in the previous rate analysis. ID is supported by ERC grant 772839, and ISF grant 2073/21. TV is supported by a research grant from the Center for New Scientists at the Weizmann Institute of Science, AFOSR Grant No. FA9550-22-1-0391, and ERC Consolidator Grant VerNisQDevS (101086733). TCL is supported in part by funds provided by the U.S. Department of Energy (D.O.E.) under the cooperative research agreement DE-SC0009919 and by the Simons Collaboration on Ultra-Quantum Matter, which is a grant from the Simons Foundation (652264 JM).

2.2 Preliminaries

2.2.1 Chain Complexes

Chain complexes are algebraic constructs that have proven helpful to connect the study of quantum codes with notions of high-dimensional expansion. Although there are

more general definitions, for us a chain complex C will always be specified by a sequence of finite-dimensional vector spaces $\{C_i\}_{i\in\mathbb{Z}}$ over a finite field $\mathbb{F}_q$, termed *chain spaces*, together with linear maps $\partial_i : C_i \to C_{i-1}$, termed *boundary maps*, that satisfy the condition $\partial_{i-1}\partial_i = 0$ for all i. (It will always be the case that $C_i = \{0\}$ for all $|i|$ large enough.) (We will focus on the case where $\mathbb{F}_q$ has characteristic 2, which allows us to view a linear map over $\mathbb{F}_q$ as a linear map over $\mathbb{F}_2$.) Since C_i is a finite vector space over $\mathbb{F}_q$, it always takes the form $C_i = \mathbb{F}_q^{D_i}$ for some integer D_i. Elements of C_i are called *i-chains*. Taking the standard inner product on $\mathbb{F}_q^{D_i}$, we can define the dual space $C^i := (C_i)^* \simeq \mathbb{F}_q^{D_i} = C_i$. Elements of C^i are called *i-cochains*, and the map $\partial_i^* : C^{i-1} \to C^i$ is denoted δ_{i-1} and called the *co-boundary* map.

Given a chain complex C, we can associate to it subsets of its i-chains called i-cycles (elements of $\ker \partial_i$) and i-boundaries (elements of $\operatorname{im} \partial_{i-1}$). We can also define further algebraic objects such as homology groups and various notions of distance and expansion. We avoid surveying all such quantities here, referring the interested reader to e.g. [2]. Instead, we focus on the definitions that are essential for this paper.

Sheaf complexes We consider chain complexes that are obtained by attaching *local coefficient spaces* $\mathcal{F}$ (or, in the terminology of [10], a *sheaf*) to a *graded poset* X. While ultimately these are "usual" chain complexes as above, their definition from two distinct objects provides a convenient way to separate the "global (geometric) structure", provided by X, and the "local (algebraic) structure", provided by $\mathcal{F}$.

We give the definitions.[4] Recall that a *poset* is a set X equipped with a partial order $\preceq$. A *graded poset* is a poset X together with a map $\rho : X \to \mathbb{Z}$ called a *rank function* such that for any $f, g \in X$, if $f \preceq g$ then $\rho(f) \leqslant \rho(g)$, and furthermore if $f \prec g$ and

[4]We give definitions that are sufficiently general to capture our constructions, but are more restrictive than the most general setting considered in the literature, such as in [2, 10]. In particular, we emphasize that we only consider complexes defined over $\mathbb{F}_q$ with characteristic 2.

there is no $f' \in X$ such that $f \prec f'$ and $f' \prec g$ (a condition which we write as $f \prec\!\!\prec g$), then $\rho(g) = \rho(f) + 1$. Given a graded poset X we define $X(i) = \{f \in X : r(f) = i\}$. Finally, we say that a graded poset X is an *incidence poset* if for every $f \preceq f''$ such that $r(f'') = r(f) + 2$, there is an *even* number of $f' \in X$ such that $f \prec\!\!\prec f' \prec\!\!\prec f''$.

Given a graded poset X with rank function r, a *system of local coefficients*, or *sheaf* $\mathcal{F}$ for X is given by two collections of objects. Firstly, for each $f \in X$ we specify a *local coefficient space* V_f. Although in general V_f may have a general group structure, here we only consider the case where each V_f is a finite dimensional $\mathbb{F}_q$-vector space $V_f \simeq \mathbb{F}_q^{m(f)}$ for some "local dimension" parameter $m(f)$. Secondly, for every $f \succeq f'$ there is a homomorphism (of vector spaces) $\mathcal{F}_{f \to f'} : V_f \to V_{f'}$ such that whenever $f \succeq f' \succeq f''$, we have that $\mathcal{F}_{f' \to f''} \circ \mathcal{F}_{f \to f'} = \mathcal{F}_{f \to f''}$.

Given a graded incidence poset X and a local system of coefficients $\mathcal{F}$ on X we define a *chain complex* $C_*(X, \mathcal{F})$ as follows. For each $i \in \mathbb{Z}$, define $C_i(X, \mathcal{F}) = \oplus_{f \in X(i)} V_f$ and $C_*(X) = \oplus_i C_i(X, \mathcal{F})$. Next, define linear maps $\partial_i : C_i(X, \mathcal{F}) \to C_{i-1}(X, \mathcal{F})$ by letting, for $f \in X(i)$ and $u \in V_f$, $\partial_i(u) = \sum_{f' \prec\!\!\prec f} \mathcal{F}_{f \to f'}(u)$. It is easily verified that these maps satisfy the chain complex condition $\partial_{i-1}\partial_i = 0$ as a result of the compatibility condition for the sheaves, and the incidence condition on the poset. We write $C^*(X, \mathcal{F}) = \oplus_i C^i(X, \mathcal{F})$ for the co-complex. Sometimes, when the sheaf $\mathcal{F}$ is clear from context we write $C_i(X)$ and $C^i(X)$ for $C_i(X, \mathcal{F})$ and $C^i(X, \mathcal{F})$ respectively.

co-cycle and cycle expansion The definition of co-cycle expansion is based on definitions from the domain of simplicial complexes. Coboundary expansion was introduced in [17, 18], and the extension to cycle and co-cycle expansion is direct. First, we introduce the notation

$$|x| = \sum_{f \in X} 1_{x(f) \neq 0} \,, \tag{2.1}$$

where for $x \in C^*(X, \mathcal{F})$ and $f \in X$ we write $x(f) \in V_f$ for the coefficient of x associated

with the face f. Eq. (2.1) is an analogue of the Hamming weight which counts the number of faces on which x has a nonzero coefficient. $|x|$ is called the *block-Hamming weight*, or *weight* for short, of x. Note that this quantity is in general *smaller* than the Hamming weight, which would sum the Hamming weights of each vector $x(f)$, $f \in X$.

Definition 2.3 (Co-cycle expansion)**.** Let $C^*(X, \mathcal{F})$ be a co-chain sheaf complex of dimension t. For integer $0 \leqslant i < t$, let $\varepsilon_{\text{cocyc}}(i)$ be the co-cycle expansion

$$\varepsilon_{\text{cocyc}}(i) \;=\; \min\left\{ \frac{|\delta(x)|}{\min_{y \in \ker \delta_i} |x - y|} \;:\; x \in C^i(X) - \ker \delta_i \right\} .$$

An analogous definition of *cycle expansion* ε_{cyc} is given in the obvious way, with the co-boundary map δ replaced by the boundary map ∂.

Co-systolic and systolic distance We define the the co-systolic distance of $C^i(X)$ to be the smallest weight of a co-cycle that is not a coboundary,

$$\mu_{\text{cosyst}}(i) \;=\; \min\left\{ |x| \;:\; x \in \ker \delta_i - \operatorname{im} \delta_{i-1} \right\} ,$$

The systolic distance μ_{syst} is defined similarly by replacing the coboundary map by the boundary map ∂.

Local minimality An important definition, that will be used throughout, is that of *co-local minimality* of a co-chain $u \in C^i(X)$, and of the *locally co-minimal distance* of a space $C^i(X, \mathcal{F})$.

The following definition is essentially from [31].

Definition 2.4. For $i > 0$, an element $x \in C^i(X, \mathcal{F})$ is *locally co-minimal* if $|x + \delta(y)| \geqslant |x|$ for all $y \in C^{i-1}(X, \mathcal{F})$ supported on $X_{\geqslant v}(i-1) = \{f \in X(i-1) : f \succeq v\}$ for some $v \in X(0)$.

This notion is referred to as *local* co-minimality because $X_{\geqslant v}(i-1) = \{f \in X(i-1) : f \succeq v\}$ has a finite size and is considered 'local' compared to the 'global' $X(i-1)$. Notice that the weight x cannot be lowered through any 'local' change $\delta(y)$.

Next we define a notion of locally co-minimal distance in the natural way:

$$\mathsf{d}_{\mathrm{coloc}}(i) \;=\; \min\left\{|x| \;:\; x \in \ker \delta_i - \{0\},\, x \text{ is locally co-minimal}\right\}.$$

A lower bound on the locally co-minimal distance immediately implies bounds on the co-systolic distance and co-cycle expansion parameters of the complex.

Lemma 2.5. *Let $C^*(X, \mathcal{F})$ be a co-chain sheaf complex of dimension t. Then for any $0 \leqslant i \leqslant t-1$,*

$$\mu_{\mathrm{cosyst}}(i) \;\geqslant\; \mathsf{d}_{\mathrm{coloc}}(i)\,,$$

and for any $0 \leqslant i \leqslant t-2$

$$\varepsilon_{\mathrm{cocyc}}(i) \;\geqslant\; \min\left\{\frac{1}{\max_{v\in X(0)} |X_{\geqslant v}(i)|}\,,\; \frac{\mathsf{d}_{\mathrm{coloc}}(i+1)}{|X(i)|}\right\}.$$

Proof. By definition, $\mu_{\mathrm{cosyst}}(i)$ is the smallest norm of an element $x \in C^i(X)$ such that $\delta(x) = 0$ and $x \notin \operatorname{im}\delta$. We claim that such an x must be locally minimal. This is because if it is not locally minimal, then there is $y \in C^{i-1}(X)$ such that if $x' = x + \delta(y)$ then $|x'| < x$. But x' still satisfies $\delta(x') = 0$ and $x' \notin \operatorname{im}\delta$, which contradicts the minimality of x. Therefore $\mu_{\mathrm{cosyst}}(i) \geqslant \mathsf{d}_{\mathrm{coloc}}(i)$.

We now consider the co-cycle expansion. Let $x \in C^i(X)$ such that $x \neq 0$. If $|\delta(x)| \geqslant \mathsf{d}_{\mathrm{coloc}}(i+1)$ then because $|x| \leqslant |X(i)|$ there is nothing to show. So assume $|\delta(x)| < \mathsf{d}_{\mathrm{coloc}}(i+1)$. If $\delta(x) = 0$ then again there is nothing to show. If $\delta(x) \neq 0$ then it cannot be locally co-minimal. So let $v \in X(0)$ and $y \in X_{\geqslant v}(i)$ be such that $|\delta(x) + \delta(y)| < |\delta(x)|$. Let $x^{(1)} := x + y$ and note that $|x - x^{(1)}| \leqslant |X_{\geqslant v}(i)|$. Repeating this process at most $|\delta(x)|$ times eventually gives some $x^{(\ell)}$ such that $|\delta(x^{(\ell)})| \leqslant |\delta(x)|$ and $\delta(x^{(\ell)})$ is locally co-minimal.

Therefore $\delta(x^{(\ell)}) = 0$ and $x^{(\ell)} \in \ker \delta$. Finally, $|x - x^{(\ell)}| \leqslant \max_{v \in X(0)} |X_{\geqslant v}(i)||\delta(x)|$, as desired. □

2.2.2 Quantum LDPC Codes

Since our construction falls within the framework of CSS codes [37, 38], we restrict our attention to such codes. A *quantum CSS code* is uniquely specified by a length-3 chain complex

$$X: \ \mathbb{F}_2^{X(0)} \xrightarrow{H_X^T} \mathbb{F}_2^{X(1)} \xrightarrow{H_Z} \mathbb{F}_2^{X(2)} \ , \tag{2.2}$$

where $X(1)$ is identified with the qubits, $|X(1)| = n$, and each element of $X(0)$ (resp. $X(2)$) is identified with an X-parity check (resp. Z-parity check). We represent the linear maps H_X, H_Z through their matrices $H_X \in \mathbb{F}_2^{X(0)\times X(1)}$ and $H_Z \in \mathbb{F}_2^{X(2)\times X(1)}$, and note that the transpose in (2.2) is taken according to the canonical basis of each space, the same basis that we fixed to identify $(\mathbb{F}_2^m)^* \simeq \mathbb{F}_2^m$.

Let $\sigma_X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ and $\sigma_Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ denote the usual Pauli matrices. The code associated with the chain complex X, which we denote $\mathcal{C} = CSS(H_X, H_Z)$, is a subspace of the space $(\mathcal{C}^2)^{\otimes n}$ of n qubits, with stabilizer generators $\sigma_X(a) = \sigma_X^{a_1} \otimes \cdots \otimes \sigma_X^{a_n}$ for a ranging over the $n_X = |X(0)|$ rows of H_X and $\sigma_Z(b) = \sigma_Z^{b_1} \otimes \cdots \otimes \sigma_Z^{b_n}$ for b ranging over the $n_Z = |X(2)|$ rows of H_Z. The code has dimension

$$k = \dim \ker H_Z - \dim \ \mathrm{im}\, H_X^T \tag{2.3}$$

and distance $d = \min(d_X, d_Z)$ where

$$d_X = \min_{c^1 \in \ker H_Z - \mathrm{im}\, H_X^T} |c^1|_H \ , \quad d_Z = \min_{c_1 \in \ker H_X - \mathrm{im}\, H_Z^T} |c_1|_H \ , \tag{2.4}$$

where here $|c|_H$ denotes the Hamming weight. We summarize these parameters by writing that $\mathcal{C} = CSS(H_X, H_Z)$ is an $[n, k, d]$ quantum (CSS) code. Finally, the code is called

low-density parity check (LDPC) if each row vector of the matrices H_X, H_Z have constant Hamming weight.[5]

2.2.3 Local testability

We start with the definition of local testability for a classical code $\mathcal{C} = \ker H$, where $H \in \mathbb{F}_2^{m\times n}$ is a parity check matrix. Recall that such a code is termed an $[n, k, d]$ code, where $k = \dim \ker H$ and $d = \min\{|c| : c \in \ker H - \{0\}\}$. We give a (somewhat restricted) definition of local testability that depends on the choice of H, which implicitly specifies a natural "tester" for the code. Intuitively, a code is locally testable if, the further a word x is from the code, the higher the probability of the "natural tester" — which selects a row of H uniformly at random and checks the corresponding parity of x — to reject.

Definition 2.6. The (classical) code $\mathcal{C} = \ker H$ is called locally testable *with soundness* ρ if for all $x \in \mathbb{F}_2^n$ it holds that

$$\frac{1}{m}|Hx|_H \geqslant \rho\frac{d_H(x, \mathcal{C})}{n} , \tag{2.5}$$

where we use d_H to denote the distance function associated with the Hamming weight (to distinguish it from d, which is the distance function associated with the block Hamming weight $|\cdot|$).

This definition can be restated in terms of expansion of the natural length-2 chain complex associated with H,

$$X\colon\ \mathbb{F}_2^{X(0)} \xrightarrow{\delta=H} \mathbb{F}_2^{X(1)} ,$$

where $X(0) = \{1, \ldots, n\}$ and $X(1) = \{1, \ldots, m\}$. Then the definition says that for every element $x \in \mathbb{F}_2^{X(0)}$, the (relative) weight $\frac{1}{m}|\delta(x)|_H$ is at least a constant times the (relative)

[5]Of course this notion only makes sense when considering a family of codes with asymptotically growing size n; in which case we mean that the row weights of H_X and H_Z should remain bounded by a universal constant as $n \to \infty$.

distance of x from $\ker \delta$. In the study of chain complexes, the parameter ρ is referred to as the *co-cycle expansion* of the complex, and denoted $h^1(X, \mathbb{Z}_2)$. (We will not use this notation.)

Now consider a quantum code $\mathcal{C}$. A general definition of local testability, and of the associated soundness parameter ρ, exists and is given in [27], which formalizes the same intuition, between distance to the code space and probability of rejection by a natural tester, as for the classical setting. This definition is somewhat cumbersome to work with and we will not need it, so we omit it. Instead, we collect the following two facts from the literature.

Firstly, it is known that local testability of a quantum CSS code is directly linked to local testability of the two associated check matrices H_X, H_Z. Formally, we have the following.

Lemma 2.7 (Adapted from Claim 3 in [26])**.** *If the quantum CSS code $\mathcal{C} = CSS(H_X, H_Z)$ is locally testable with soundness ρ (in short, ρ-qLTC) then the classical codes with parity-check matrices H_X and H_Z are each locally testable with soundness at least $\rho/2$. Conversely, if the classical codes with parity-check matrices H_X and H_Z are locally testable with soundness ρ then the quantum CSS code $\mathcal{C} = CSS(H_X, H_Z)$ is locally testable with soundness at least ρ.*

Because the lemma is tight up to a factor 2, and we will not be concerned with constant factors, we can take the local testability of H_X and H_Z as our definition of local testability for the quantum code $CSS(H_X, H_Z)$.

Definition 2.8 (Definition of local testability for quantum CSS codes)**.** The (quantum) CSS code $\mathcal{C} = CSS(H_X, H_Z)$ is called locally testable *with soundness* ρ if both $\mathcal{C}_X = \ker H_X$ and $\mathcal{C}_Z = \ker H_Z$ are (classical) codes with soundness at least ρ.

We emphasize that this definition is local to our paper, and is equivalent to the general accepted definition of a quantum LTC up to a multiplicative factor 2 in the soundness.

To translate between the notion of (co-)cycle expansion that we formulate our main analytic results in, and quantum local testability properties of the quantum codes that can be built from the underlying complex, we use the following lemma.

Lemma 2.9. *Let $C(X, \mathcal{F})$ be a sheaf complex, and $i \in \mathbb{Z}$. Let $\mathcal{C} = CSS(\partial_i, \delta_i)$ be the quantum CSS code associated with the linear maps ∂_i and δ_i over $\mathbb{F}_q$.*[6] *(When constructing codes, ∂_i and δ_i are treated as $\mathbb{F}_2$-linear maps.) Let $D_i = \dim C_i(X, \mathcal{F})$ and let $M_i = \max_{f \in X(i)} \dim V_f$. Suppose further that $C(X, \mathcal{F})$ has systolic and co-systolic distance in dimension i at least μ_{syst} and μ_{cosyst} respectively, and cycle and co-cycle expansion in dimension i at least ε_{cyc} and $\varepsilon_{\text{cocyc}}$ respectively.*

Then $\mathcal{C}$ is an $[n, k, d]$ quantum code, where

$$k \;=\; \dim_{\mathbb{F}_2} C_i(X) - \dim_{\mathbb{F}_2} \operatorname{im} \delta_i^T - \dim_{\mathbb{F}_2} \operatorname{im} \partial_i^T$$

and

$$d \;\geqslant\; \min\{\mu_{\text{syst}},\, \mu_{\text{cosyst}}\}\,.$$

Moreover, the code $\mathcal{C}$ is ρ-qLTC where

$$\rho \;\geqslant\; \left(\frac{1}{\log_2 q}\right) \min\left\{\frac{D_i}{D_{i-1}} \frac{\varepsilon_{\text{cyc}}}{M_i},\; \frac{D_i}{D_{i+1}} \frac{\varepsilon_{\text{cocyc}}}{M_i}\right\}. \tag{2.6}$$

We note that the bound in (2.6) may be somewhat loose, due to the interplay between the use of the Hamming weight and the block Hamming weight, which leads to the factor $\frac{1}{M_i}$ on the right-hand side. In our constructions, M_i is thought of as polylogarithmic, or even constant, in the length of the code.

[6] Here, we slightly abuse notation and use $C_i \simeq C^i$ to view both maps ∂_i, δ_i acting on the same space.

Proof. The formula for the dimension k is clear from (2.3) $k = \dim\ker\delta_i - \dim\ \mathrm{im}\,\partial_i^T$ and the fact that for the linear operator $\delta_i : \mathbb{F}_2^n \to \mathbb{F}_2^m$, $\dim\ker\delta_i = n - \dim\ \mathrm{im}\,\delta_i = n - \dim\ \mathrm{im}\,\delta_i^T$. Intuitively, this is saying that the number of logical qubits is equal to the number of physical qubits minus the number of independent stabilizer generators.

For the distance, the bound follows from the definition (2.4) of the distance of a quantum CSS code, the definition of the parameters μ_{syst} and μ_{cosyst} and the fact that the block weight $|c|$ is a lower bound on the Hamming weight of c.

To show the bound on the soundness parameter ρ, consider first the soundness of the classical code $\mathcal{C}_X = \ker H_X = \ker\partial_i$. Recall that $\partial_i : C_i(X,\mathcal{F}) \to C_{i-1}(X,\mathcal{F})$. Let $c \in C_i(X,\mathcal{F})$. Then

$$\begin{aligned}\frac{1}{D_{i-1}}|H_X c|_H &\geqslant \frac{1}{D_{i-1}}|\partial_i c| \\ &\geqslant \frac{1}{D_{i-1}}\,\varepsilon_{\mathrm{cyc}}\,d(c,\ker\partial_i) \\ &\geqslant \frac{D_i}{D_{i-1}}\frac{\varepsilon_{\mathrm{cyc}}}{M_i}\frac{d_H(c,\mathcal{C}_X)}{D_i}\ .\end{aligned}$$

Here the first inequality follows since $|\cdot|_H \geqslant |\cdot|$, the second inequality uses the definition of $\tilde{\varepsilon}_i$, and the third inequality uses that $|x| \geqslant \frac{1}{M_i}|x|_H$ for $x \in C_i(X)$. Let $|c|_{\mathbb{F}_2}$ be the Hamming weight of c when viewed as a vector over $\mathbb{F}_2$. Because $|c|_{\mathbb{F}_2} \leqslant |c|_H \leqslant (\log_2 q)|c|_{\mathbb{F}_2}$, we obtain the lower bound on the soundness of the code $\mathcal{C}_X$.

The argument showing a lower bound on the soundness of the code $\mathcal{C}_Z = \ker H_Z$ is analogous, and we omit it.

□

2.3 The complex

Let $t \geqslant 1$ be an integer. Our construction is based on a graded incidence poset that is obtained from a t-dimensional cubical complex

$$X = X(0) \cup \cdots \cup X(t) \ .$$

The cubical complex can be generated from a set G of size $N = |G|$, and finite subsets of permutations $A_1, \ldots, A_t$ of G of size $n = |A_i|$. The sets A_i should be closed under inverse and such that permutations taken from different sets commute. We denote this complex $X = X(G; \{A_i\})$. The requirement that the sets A_i pairwise commute can be understood as a requirement that allows us to iterated the *balanced product* construction from [23], which a priori was only introduced for the product of two graphs (or more generally chain complexes) with naturally commuting left and right group actions.

A concrete example that one may keep in mind is the case where G is a finite abelian group and $A_1 \ldots, A_t \subseteq G$ are generating subsets. Another example, when $t = 2$, is where G is any finite group and $A_1, A_2 \subset G$ act by multiplication on the left and right respectively. The best construction that we identified is based on a third example, derived from an expander construction of [3]. This construction was pointed out to us by Louis Golowich. We describe it in more detail in Section 2.3.5.

2.3.1 Geometry

We describe our general construction of a complex X, in the process introducing notation that will be used throughout. See also Figure 2.1.

Faces For $0 \leqslant k \leqslant t$ the k-dimensional faces of X, i.e. the elements of $X(k)$, are partitioned according to their *type*, which can be any subset $S \subseteq \{1, \ldots, t\}$ of size $|S| = k$. Let $\overline{S}$ denote the complement of S in $\{1, \ldots, t\}$. For every such S, a face f of type S is

uniquely specified as

$$f \;=\; [\, g;\; (a_j)_{j\in S},\; (b_j)_{j\in \overline{S}}\,]\,, \tag{2.7}$$

where $g \in G$, for each $j \in S$, $a_j \in A_j$, and for each $j \in \overline{S}$, $b_j \in \{0,1\}$. We let $X(S)$ denote the set of faces of type S. Thus $|X(S)| = |G| n^k 2^{t-k}$. Geometrically, the k-face $f = [g; a, b]$ contains the 2^k 0-faces (vertices)

$$\Big\{ \Big(g \cdot \prod_{j : b'_j = 1} a_j \,;\, b' \| b \Big) \;\Big|\; b' \in \{0,1\}^S \Big\}\,, \tag{2.8}$$

where we use the notation $b' \| b$ to denote string concatenation with re-ordering of indices in the obvious way. Here and in the following we use the notation $g \cdot a$ to denote the element $a(g)$, for $a \in A_j$ a permutation of G. Note that the element $g \cdot \prod_j a_j$ is well-defined because we assumed that permutations a_j taken from different sets A_j pairwise commute.

For a face $f = [g; a, b]$ we write $S = \mathtt{type}(f)$ for its type, and (slightly overloading notation compared to (2.7)) write

$$f_0 = g \qquad \text{and} \qquad f_j = \begin{cases} a_j & j \in S \\ b_j & j \notin S \end{cases}$$

for $j = 1, \ldots, t$.

Partial order We introduce the following partial order on faces of X. This partial order is the one that is induced by set inclusion, when each face is seen as a set of vertices as in (2.8). Formally, given two faces f, f' and $j \in \{1, \ldots, t\}$ we write $f' \prec_j f$ if all the following hold:

- $f'_j \in \{0,1\}$ while $f_j \in A_j$,
- $f_i = f'_i$ for all $i \notin \{0, j\}$,
- $f'_0 = \begin{cases} f_0 & f'_j = 0 \\ f_0 \cdot f_j & f'_j = 1 \end{cases}$.

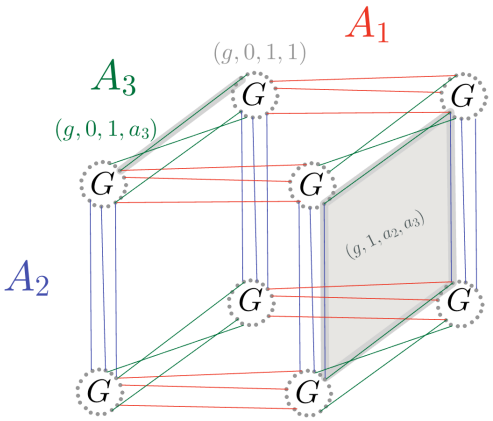

Figure 2.1: A cubical complex with $t = 3$ and three sets of permutations A_1, A_2, A_3. The underlying graph can be seen to be $2^t = 8$-partite. We label the faces as (f_0, f_1, f_2, f_3).

We sometimes suppress j from the notation and write $f' \prec f$ to indicate that there is some j for which $f' \prec_j f$. Note that if $f' \prec f$ then necessarily f is one dimension higher than f'. We also write $f' \prec f$ if there is a sequence $f' \prec \cdots \prec f$, and $f' \preceq f$ if $f' \prec f$ or $f' = f$.

Lemma 2.10. *The relation $\prec$ defines a transitive graded partial order on faces. Moreover, with this relation X is an incidence poset.*

Proof. The relation $\preceq$ is reflexive and transitive. It is antisymmetric because if $f' \preceq f$ and $f \neq f'$ then f must be of dimension at least one more than f'. To show that it is graded, we verify that the covering relation associated with $\preceq$ is precisely $\prec$.

For the moreover part, note that if $f \prec f' \prec f''$, then there are exactly two distinct positive $i \neq j$ such that $f''_j \neq f_j$ and $f''_i \neq f_i$, and this determines exactly two possible values for f'. □

Definition 2.11. Given the set G and the sets of permutations $A_1, \ldots, A_t$, we let $X(G; \{A_i\})$ denote the graded incidence poset $X = \cup_{0 \leqslant k \leqslant t} X(k)$. Here the rank function is given by $\rho(f) = |\texttt{type}(f)|$, for $f \in X$.

We end the section by introducing convenient notation.

Incidence maps The up incidence map $\mathfrak{u}$ takes a k-face f to all $(k+1)$-faces that cover it,

$$\mathfrak{u}(f) = \{f' \mid \ f' \succ f\} \ ,$$

and is naturally extended to sets of faces by taking unions. Similarly,

$$\mathfrak{d}(f) = \{f' \mid \ f' \prec f\} \ .$$

Links The upwards link of a face $v \in X$ is a sub-complex consisting of all faces above v,

$$X_{\geqslant v} = \{f \in X \mid \ f \succeq v\} .$$

Note that since the structure of X is not simplicial, we do not remove v from a face $f \succeq v$ when defining the link. Similarly, the downwards link of a face $v \in X$ is a sub-complex consisting of faces below v,

$$X_{\leqslant v} = \{f \in X \mid \ f \preceq v\} .$$

For a subset $S \subseteq \{1, \ldots, t\}$ we denote $X_{\geqslant v}(S) = X_{\geqslant v} \cap X(S)$ and $X_{\leqslant v}(S) = X_{\leqslant v} \cap X(S)$. Similarly, $X_{\geqslant v}(k) = X_{\geqslant v} \cap X(k)$ and $X_{\leqslant v}(k) = X_{\leqslant v} \cap X(k)$.

2.3.2 Local coefficients

To each face we associate a local coefficient space. For $i \in \{1, \ldots, k\}$ let $m_i \geqslant 1$ be an integer and $\hat{A}_i = \{1, \ldots, m_i\}$. The local coefficient space associated with a face of type S is the vector space over a finite field $\mathbb{F}_q$ of characteristic 2[7]

$$V_S = \mathbb{F}_q^{\prod_{j \in \overline{S}} \hat{A}_j} . \tag{2.9}$$

[7]Our construction can be applied more generally to arbitrary fields. However, we focus on finite fields with characteristic 2, so that the corresponding code can be viewed as a code over $\mathbb{F}_2$. Moreover, the construction of the (co)boundary maps becomes more straightforward in this context, and we will comment further in Section 2.3.3.

So the local coefficient space of a t-dimensional face is $V_{\{1,\ldots,t\}} = \mathbb{F}_q$, and the local coefficient space of a vertex is $(\prod_{j=1}^{t} m_j)$-dimensional. We write $V_f = V_{\mathtt{type}(f)}$ for the local coefficient space of the face f, and $\mathcal{F} = \{V_f\}$ for the collection of local coefficient spaces.

For $i \in \{0, \ldots, t\}$ we denote by $C_i(X, \mathcal{F})$, or $C_i(X)$ for short, the vector space of i-chains, that assign to every i-face f a "coefficient" from V_f:

$$C_i(X, \mathcal{F}) \;=\; \bigoplus_{f \in X(i)} V_f \;.$$

Similarly, we use $C^i(X)$ to denote the i-cochains

$$C^i(X, \mathcal{F}) \;=\; \bigoplus_{f \in X(i)} V_f^* \;,$$

where $V_f^* = \{g : V_f \to \mathbb{F}\}$ is the dual vector space. Throughout we fix an identification of each V_S^* with V_S through the use of a self-dual basis which we label $\{a : \; a \in \prod_{j \in \overline{S}} \hat{A}_j\}$, so $V_f^* \simeq V_f$ canonically.

Definition 2.12. Given a set G and subsets of permutations $A_1, \ldots, A_t$ we let $C_*(G; \{A_i\}; \{h_i\})$ be the length-$(t+1)$ chain complex

$$C_*(G; \{A_i\}; \{h_i\}) \;=\; \bigoplus_i C_i(X, \mathcal{F}) \;.$$

We often write $C_*(X, \mathcal{F})$ or even $C_*(X)$ for this complex, when the parameters of the complex are clear from context. We use $C^*(X) = \oplus_i C^i(X)$ to denote the dual complex.

The (co-)boundary maps associated with C_* are defined in the next sub-section. Before proceeding, we can already compute the dimension of the space of i-chains $C_i(X, \mathcal{F})$.

Lemma 2.13. *Let $C_*(X) = C_*(G; \{A_i\}; \{h_i\})$, $0 \leqslant i \leqslant t$ and $D_i = \dim\, C_i(X, \mathcal{F})$. Then*

$$D_i \;=\; N\, n^i 2^{t-i} \sum_{\substack{T \subseteq \{1,\ldots,t\} \\ |T| = t-i}} \prod_{j \in T} m_j \;.$$

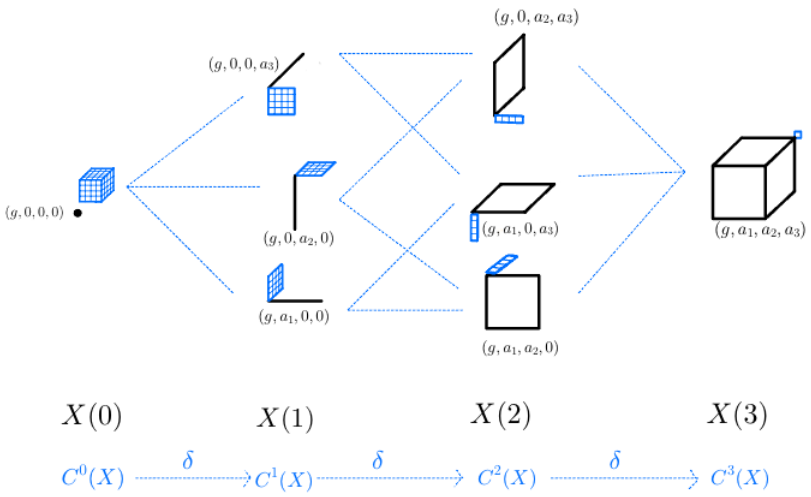

Figure 2.2: The co-chain complex, localized at a vertex $f = (g, 0, 0, 0)$; with $t = 3$. The geometric elements of $X_{\geqslant f}$ are drawn in black. The associated coefficient spaces are drawn in blue.

Proof. We have that $D_i = \sum_{f \in X(i)} \dim V_f$. For each $f \in X(i)$, $|\mathtt{type}(f)| = i$ and hence by (2.9), $\dim V_f = \prod_{j \notin \mathtt{type}(f)} m_j$. Moreover, for $|S| = i$ we have $|X(S)| = Nn^i 2^{t-i}$, where N is to choose $g \in G$, n^i is to choose a_j for $j \in S$, and 2^{t-i} is to choose b_j for $j \notin S$. □

2.3.3 (Co)boundary maps

For $i \in \{1, \ldots, t\}$ let $h_i \in \mathbb{F}_q^{m_i \times n}$. Assume that $m_i \leqslant n_i$ and that h_i has full row rank, i.e. the rows of h_i are linearly independent. We index the rows of h_i using the set $\hat{A}_i = \{1, \ldots, m_i\}$, and the columns of h_i using the set A_i, and view h_i as a linear map $h_i : \mathbb{F}_q^{A_i} \to \mathbb{F}_q^{\hat{A}_i}$.

Restriction and co-restriction maps Given $i \notin \mathtt{type}(f')$ and $f' \prec_i f$, define the *co-restriction map*

$$\text{co-res}_{f',f} : \quad V_{f'}^* \to V_f^*$$

as follows. For any $z \in V_{f'}^*$, co-res$_{f',f}(z)$ is obtained by applying $I_{-i} \otimes h_i^T$, resulting in a vector in $\mathbb{F}_q^{(\prod_{j \notin \mathtt{type}(f')} \hat{A}_j) \times A_i}$, and then setting the A_i-coordinate to f_i. Formally,

$$\text{co-res}_{f',f}(z) = \Big((I_{-i} \otimes h_i^T)z\Big)[\cdots, f_i, \cdots] \,. \tag{2.10}$$

Similarly, for $i \in \mathtt{type}(f')$ and $f' \succ_i f$ define the *restriction map*

$$\text{res}_{f',f} : V_{f'} \to V_f$$

for any $z \in V_{f'}$ by

$$\text{res}_{f',f}(z) = z \otimes (h_i f_i') \,, \tag{2.11}$$

which is an element of $\mathbb{F}_q^{\prod_{j \notin \mathtt{type}(f')} \hat{A}_j} \otimes \mathbb{F}_q^{\hat{A}_i} \cong \mathbb{F}_q^{\prod_{j \notin \mathtt{type}(f)} \hat{A}_j}$. We abuse the notation and write f_i' as the canonical basis vector of $\mathbb{F}_q^{A_i}$ corresponding to $f_i' \in A_i$. We verify below that the maps co-res and res are adjoint to each other (see (2.13)). We extend the notation by defining, for any sequence $f_i \prec f_{i+1} \cdots \prec f_k$ where $f_j \in X(j)$,

$$\text{res}_{f_k,f_i}(z) \;=\; \text{res}_{f_{i+1},f_i}(\cdots \text{res}_{f_k,f_{k-1}}((z))) \,.$$

We note that this leads to a well-defined quantity, i.e. res$_{f_k,f_i}(z)$ is independent of the path $f_i \prec f_{i+1} \cdots \prec f_k$, as can easily be verified from the definition (2.11). This is because application of a linear map along different coordinates, or restriction on different coordinates, are operations that commute. A similar extension is performed to define co-restriction maps co-res$_{f',f}$ for general $f' \prec f$.

The coboundary map The coboundary map $\delta : C^*(X) \to C^*(X)$ is defined as

$$\delta(z)(f) = \sum_{f' \prec f} \text{co-res}_{f',f}(z(f')) \,, \tag{2.12}$$

where co-res$_{f',f}$ is as in (2.10). In words, at a face $f \in X$, we go over all faces f' below f and map $z(f')$ to an element in V_f^* via the restriction map res$_{f',f}$, and then sum all the

contributions thus obtained. Because $\mathbb{F}_q$ has characteristic 2, $\delta \circ \delta = 0$, as required by the definition of the coboundary map.[8]

The boundary map The boundary map $\partial : C_*(X) \to C_*(X)$ is defined as

$$\partial(z)(f) = \sum_{f' \succ f} \mathrm{res}_{f',f}(z(f')).$$

The next lemma shows that the map ∂ is the adjoint of the map δ with respect to the natural inner product on $C(X)$.

Lemma 2.14. *For any $z \in C_*(X)$ and $z' \in C^*(X)$ it holds that*

$$\langle z, \delta(z') \rangle \;=\; \langle \partial(z), z' \rangle \; .$$

Proof. It is enough to verify the equality for $z \in V_f$ and $z' \in V_{f'}^*$. By definition $\delta : V_{f'}^* \to \oplus_{f'' \succ f'} V_{f''}^*$ and $\partial : V_f : \oplus_{f'' \prec f} V_{f''}$. Thus both sides of the equality are non-zero only if $f' \prec f$. Suppose that this is the case, and let i be such that $\texttt{type}(f) = \texttt{type}(f') \cup \{i\}$. It remains to verify that

$$\langle z, \text{co-res}_{f',f}(z') \rangle \;=\; \langle \mathrm{res}_{f,f'}(z), z' \rangle \; . \tag{2.13}$$

According to (2.10)

$$\text{co-res}_{f',f}(z') = (I_{-i} \otimes h_i^T) z'[\cdots, f_i, \cdots] \; ,$$

and according to (2.11)

$$\mathrm{res}_{f,f'}(z) = z \otimes (h_i f_i) \; .$$

Thus the left-hand side of (2.13) evaluates to

$$\langle z, (I_{-i} \otimes h_i^T) z'[\cdots, f_i, \cdots] \rangle = \langle z \otimes f_i, (I_{-i} \otimes h_i^T) z' \rangle$$

[8] For general fields $\mathbb{F}_q$ with characteristics not equal to 2, the maps requires additional signs so that $\delta \circ \delta = 0$. Specifically, $\delta(z)(f) = \sum_{f' \prec f} (-)^{s(f,f')} \text{co-res}_{f',f}(z(f'))$ for some s.

$$= \langle (I_{-i} \otimes h_i)(z \otimes f_i), z' \rangle$$
$$= \langle z \otimes (h_i f_i), z' \rangle \ ,$$

which matches the right-hand side. □

2.3.4 Statement of results

Having described our construction, we can state our main result about it.

Theorem 2.15. *Let $(G; \{A_i\})$ be a set and collections of pairwise commuting permutations on it such that the graphs $\mathrm{Cay}(G, A_i)$ are each λ-expanding up to size $r|G|$ (see Definition 2.47). Let $N = |G|$ and $n = |A_1| = \cdots = |A_t|$. Let $(h_1, \ldots, h_t)$ be a collection of matrices $h_j \in \mathbb{F}_q^{m_j \times n}$ such that the family $(h_1, \ldots, h_t)$ is two-way κ-robust (see Definition 2.34). Let $C = C_*(G; \{A_j\}; \{h_j\})$ be the length-$(t+1)$ chain complex over $\mathbb{F}_q$ from Definition 2.12. Then there is a $c = \kappa^t \exp(-O(t^2))$ such that C has the following properties.*

1. *The space C_k of k-chains has dimension $D_k = N\ n^k 2^{t-k} \sum_{\substack{T \subseteq \{1,\ldots,t\} \\ |T| = t-k}} \prod_{j \in T} m_j$.*
2. *The co-chain complex C^* has co-systolic distance $\mu_{\mathrm{cosyst}}(k)$ for every $0 \leqslant k \leqslant t-1$ and co-cycle expansion $\varepsilon_{\mathrm{cocyc}}(k)$ for every $0 \leqslant k \leqslant t-2$, where*

$$\mu_{\mathrm{cosyst}}(k) \geqslant \frac{c-\lambda}{2^{2t} n^t} r|X(k)| \qquad \text{and} \qquad \varepsilon_{\mathrm{cocyc}}(k) \geqslant \frac{c-\lambda}{(2nt)^{2t+1}} r \ . \tag{2.14}$$

3. *The chain complex C_* has systolic distance $\mu_{\mathrm{syst}}(k)$ for every $1 \leqslant k \leqslant t$ and cycle expansion $\varepsilon_{\mathrm{cyc}}(k)$ for every $2 \leqslant k \leqslant t$, where*

$$\begin{aligned} \mu_{\mathrm{syst}}(k) &\geqslant \frac{1}{(2nt)^t} \mu_{\mathrm{cosyst}}(t-k) \ , \\ \varepsilon_{\mathrm{cyc}}(k) &\geqslant \min\left\{ \frac{\varepsilon_{\mathrm{cocyc}}(t-k)}{2(t^2 2^{2t} n^{t+1})^t} \ , \ \frac{1}{(2nt)^t} \frac{\mu_{\mathrm{cosyst}}(t-k+1)}{|X(k)|} \right\} . \end{aligned} \tag{2.15}$$

We note that we did not attempt to overly optimize the bounds in Theorem 2.15; in particular the dependence on t may be loose. To understand the bounds, one can consider that the main asymptotic parameter is $N = |G|$, which grows to infinity. We think of t as a constant, such as $t = 4$. The parameter n is also ideally a constant, but the parameter r is not, $r = \Omega((\log N)^{-(t-1)})$. This is chosen such that λ is sufficiently small compared to κ so that the bounds in (2.14) are positive. This is the origin of the loss of the polylogarithmic factor in distance and soundness.

Proof of Theorem 2.15. The first item is shown in Lemma 2.13. For the second item, we first note the bound

$$\mathsf{d}_{\mathrm{coloc}}(k) \geqslant \frac{1 - \lambda C_1}{C_2} r |X(k)| \,,$$

valid for $0 \leqslant k < t$, that follows from Proposition 2.56. The prefactor on the right-hand side can be re-written as

$$\begin{aligned} \frac{1 - \lambda C_1}{C_2} &= \big(1/C_1 - \lambda\big)\frac{C_1}{C_2} \\ &\geqslant \frac{1}{2^{2t} n^t}(\kappa^t 2^{-4t^2} - \lambda) \,, \end{aligned}$$

where the second line uses $t^2 2^{2t} \leqslant 2^{3t^2}$ for $t \geqslant 1$. We then use Lemma 2.5 to conclude the claimed bounds. For $\varepsilon(k)$, we also use $|X(k+1)|/|X(k)| \leqslant 2nt$ and $|X_{\geqslant v}(k)| \leqslant (2n)^t$, which follow from Lemma 2.44. Finally, the bounds on $\tilde{\varepsilon}(k)$ and $\tilde{\mu}(k)$ follow immediately from Proposition 2.63. □

We can apply Theorem 2.15 to obtain quantum codes with parameters described in the following corollary. Since a quantum CSS code can be defined from any length-3 complex, we have the choice to locate the qubits on the space of k-chains for any $i \in \{1, \ldots, t-1\}$. However, to obtain a bound on the local testability soundness of the quantum code, we need a complex of length at least 5. To simplify the discussion of parameters we assume

that, of the codes $\mathcal{C}_i = \ker h_i$, exactly i have rate close to 1, and the remaining $(t-i)$ have rate close to 0. This ensures that most of the dimension of the space C_i is concentrated on a single type of i-face, and guarantees that the resulting code has positive dimension. Of course, other settings of parameters are possible.

Corollary 2.16. *With the same setup as in Theorem 2.15, assume $t \geqslant 4$ and fix $2 \leqslant i \leqslant t-2$. Assume that the parameters $m_1, \ldots, m_t$ and n are chosen such that $m_j = \Theta(n)$ for all j, and furthermore $m_j/m_{j'}$ is at least a sufficiently large constant (depending on t) when $j \in \{1, \ldots, i\}$ and $j' \in \{i+1, \ldots, t\}$. Assume that, for these parameters there is a family $(h_1, \ldots, h_t)$ of two-way κ-robust codes, where κ is small enough, as a function of λ and t, so that the lower bounds in (2.14) are positive.*[9] *Let $\mathcal{C} = CSS(\delta_i, \partial_i)$ be the CSS code associated with the space C_i and the boundary and co-boundary operators on it, where C_i is viewed as a $\mathbb{F}_2$-vector space and δ_i, ∂_i are viewed as $\mathbb{F}_2$-linear maps. Then $\mathcal{C}$ is an $[\ell_i, k_i, d_i]$ CSS code with parity checks of weight $O((\log_2 q)tn)$ and such that*

1. $\ell_i = (\log_2 q) D_i = (\log_2 q) N \cdot \Theta(n)^t$,
2. $k_i = \Omega((\log_2 q)\ell_i)$,
3. $d_i = rN \cdot (nt)^{-O(t)} \exp(-O(t^2))$.

Moreover, $\mathcal{C}$ is a qLTC with soundness $\rho \geqslant r(\log_2 q)^{-1}(nt)^{-O(t)} \exp(-O(t^2))$.

Proof. The first item follows by definition of $\mathcal{C}$. For the second item, we will prove the lower bound on k_i in Section 2.4 (see end of Section 2.4.1). As shown in Lemma 2.9 the distance is $d_i \geqslant \min\{\mu_{\text{cosyst}}(i), \mu_{\text{syst}}(i)\}$, giving the stated bound by Theorem 2.15. Similarly, the soundness is bounded using Lemma 2.9. □

[9] We emphasize that one can generally expect to achieve λ that is an arbitrarily small constant, or even an inverse polynomial in the degree n. For constant t, the requirement on κ is simply that it should be a small enough constant, or even just that $\kappa \geqslant 1/n^{\varepsilon}$ for small enough $\varepsilon > 0$.

As discussed earlier, if n, t, q, r are all constant then we obtain a family of quantum LTC with constant weight parity checks, constant rate, relative distance and soundness. While the parameters n, t, q can be chosen to be constants, we do not know how to construct the required 4-dimensional cubical complex, with the right expansion properties, for constant r. However, as discussed in Section 2.3.5 we can construct it for $r = \Omega((\log |G|)^{-(t-1)})$.

2.3.5 Instantiating G and the $\{A_i\}$

We provide two examples of G and pairwise commuting $A_1, ..., A_t$. The first example is easier to understand; the second example has better parameters, and in particular leads to our main application, stated as Theorem 2.1 in the introduction.

For both constructions, given target rates $(n - m_j)/n$ for the local codes $\mathcal{C}_j$ we fix a family $(h_1, \ldots, h_t)$ of parity check matrices for codes $\mathcal{C}_1, \ldots, \mathcal{C}_j \subseteq \mathbb{F}_q^n$ such that the $\mathcal{C}_j$ have the desired rate and moreover the family $(h_1, \ldots, h_t)$ is two-way κ-robust according to Definition 2.34. It follows from Theorem 2.36 and Lemma 2.38 that by choosing q as a large enough power of 2, a random t-tuple of codes satisfies this property for some $\kappa > 0$ independent of n.

First construction

The first construction is obtained from Cayley graphs over $(\mathbb{Z}/2\mathbb{Z})^{\log N}$.

Let $G = (\mathbb{Z}/2\mathbb{Z})^{\log N}$. Let $n = \lceil c \log N \rceil$, for a sufficiently large constant c such that there exist sets $A_1, \ldots, A_t$ of size $n = |A_j|$ such that the graphs $\mathrm{Cay}(G, A_j)$ have spectral expansion at most λ, where λ is sufficiently small compared to the robustness parameter κ. Since κ is a constant independent of n then we can always make λ small enough by increasing c, independently of n. This is because it is well-known that for any target λ (possibly depending on κ) there is a constant $c(\lambda) = O(1/\lambda^2)$ and subsets A_j of G of size $|A_j| = c(\lambda) \log |G|$ such that $\mathrm{Cay}(G, A_j)$ is λ-expanding. (This can be shown by a

probabilistic argument [39], and also by explicit constructions [40].)

Note that this construction is not "fully explicit," because we do not have a way to check in time poly(N) that a given choice of codes $(h_1, \ldots, h_t)$ is indeed two-way κ-robust. This is because according to Theorem 2.36 we need to choose $q = 2^{(n+3)^t}$, which is super-polynomial in N.

Second example

The second example is beyond Cayley graphs, but can be described as abelian lifts and is tightly related to Cayley graphs over $\mathbb{Z}/\ell\mathbb{Z}$.

Let $G_0 = (V_0, E_0)$ be an n-regular Cayley expander graph on n' vertices. Based on known explicit constructions of expanders, n can be chosen any large enough constant, and n' arbitrarily large. In [3, Theorem 1.3] the authors describe an H-lift of this graph for an abelian group H, $|H| = \exp(\Theta(n'))$, which is given by associating with each edge $e \in E_0$ an element $s \in H$. This graph, which we call R, has $N = |V_0||H|$ vertices and has second largest eigenvalue $O(\sqrt{n}\log n)$, making it nearly-Ramanujan. Furthermore, the graph R can be constructed efficiently in time polynomial in N.

Let $G_1 = (V_1, E_1)$ be the double cover of G_0, with vertex set $V_1 = V_0 \times \{0,1\}$. Furthermore, let us assume a fixed numbering of the n edges out of each vertex, so that for each $i \in [n]$ we have an action on $H \times V_0 \times \{0,1\}$ obtained by taking an element (h, v, b) to $(s \cdot h, v', 1-b)$ where $(v', 1-b)$ is the i-th neighbor of (v, b) in G_1 and $s \in H$ is associated with the i-th edge out of v. This set of n actions gives a n-regular graph that is the H-lift of G_1 given by the association of H-elements to edges.

We now move to the t-fold product, letting $G = H \times V_0^t$, and letting $X(0) = G \times \{0,1\}^t$. We have t sets $A_1, \ldots, A_t$ of permutations on $X(0)$, where $A_j = \left\{s_j^1, \ldots, s_j^d\right\}$ and

$$(h, v_1, \ldots, v_t, b_1 \ldots, b_t) \xrightarrow{s_j^i} (s \cdot h, v_1, \ldots, v'_j, \ldots, v_t, b_1, \ldots, 1-b_j, \ldots, b_t)$$

where $(v'_j, 1-b_j)$ is the i-th neighbor of (v_j, b_j) in G_1, and s is associated with the i-th neighbor of v_j.

It is easy to check that every pair of permutations $s \in A_j$ and $s' \in A_{j'}$ pairwise commute, and therefore X_0 together with $A_1, \ldots, A_t$ give rise to a t-dimensional complex X as described in Section 2.3.1.

We make the following two observations about this construction. Firstly, for any fixed $j \in [t]$ the graph obtained by connecting every $v \in X(0)$ to the n vertices obtained by the n permutations from A_j is a graph consisting of V_1^{t-1} copies of the expander graph R. Since $|V_1| = O(\log N)$, we obtain that the graphs $\mathrm{Cay}(G, A_i)$ are λ-expanding up to size $r|G|$, where λ is the normalized second eigenvalue of R and $r = \Omega(1/(\log N)^{t-1})$. Secondly, n can be chosen as any constant large enough such that if λ is the normalized second eigenvalue of R then λ is small enough compared to κ for the lower bounds in (2.15) to be positive. This is always possible because the universal constant κ obtained from Theorem 2.35 does not depend on n (only on t), whereas λ goes to zero as n grows. Because n is a constant, we can exhaustively search for a tuple $(h_1, \ldots, h_t)$ that is two-way κ-robust.

Combining these observations lead to the following main instantiation of Theorem 2.15.

Corollary 2.17. *Let $t \geqslant 2$ be an integer. There is a polynomial-time algorithm that on input an integer N (in unary) returns a description of $(G; \{A_i\})$ and $\{h_i\}$ as above such that the following hold.*

1. *The space C_k of k-chains has dimension $D_k = \Theta(N)$;*
2. *The co-chain complex C^* has co-systolic distance $\mu_{\mathrm{cosyst}}(k) = \Theta(|X(k)|/(\log N)^{t-1})$ for every $0 \leqslant k \leqslant t-1$ and co-cycle expansion $\varepsilon_{\mathrm{cocyc}}(k) = \Theta(1/(\log N)^{t-1})$ for every $0 \leqslant k \leqslant t-2$;*

3. *The chain complex C_* has systolic distance $\mu_{\mathrm{syst}}(k) = \Theta(|X(k)|/(\log N)^{t-1})$ for every $1 \leqslant k \leqslant t$ and cycle expansion $\varepsilon_{\mathrm{cyc}}(k) = \Theta(1/(\log N)^{t-1})$ for every $2 \leqslant k \leqslant t$.*

Here, the constants implicit in the $\Theta(\cdot)$ notation can depend on t.

We also obtain a corresponding instantiation of Corollary 2.16, which leads to the previously stated Theorem 2.1.

Corollary 2.18. *There is a polynomial-time algorithm that on input an integer N (in unary) returns explicit descriptions of parity check matrices H_X and H_Z on $\Theta(N)$ qubits such that the following hold:*

1. *The weight of each row of H_X and H_Z is $\Theta(1)$;*
2. *The code $\mathcal{C} = CSS(H_X, H_Z)$ has dimension $\Theta(N)$ and distance $\Theta(N/(\log N)^3)$;*
3. *$\mathcal{C}$ is a qLTC with soundness $\rho = \Omega(1/(\log N)^3)$.*

2.4 A lower bound on the code rate

A natural approach to bounding the code dimension is to count the number of qubits and the number of X and Z checks. The number of qubits, X and Z checks are $\dim_{\mathbb{F}_2} C_i(X, \mathcal{F})$, $\dim_{\mathbb{F}_2} C_{i-1}(X, \mathcal{F})$, and $\dim_{\mathbb{F}_2} C_{i+1}(X, \mathcal{F})$, respectively. This leads to a lower bound on the code dimension: $k \geqslant \dim_{\mathbb{F}_2} C_i(X, \mathcal{F}) - \dim_{\mathbb{F}_2} C_{i-1}(X, \mathcal{F}) - \dim_{\mathbb{F}_2} C_{i+1}(X, \mathcal{F})$. However, when applied to our construction, this bound yields a negative value, making it trivial.

To obtain a nontrivial lower bound on the code dimension, we need to take into account the linear dependencies among the X and Z checks. Rather than simply counting the total number of checks, we refine our estimate by counting only the number of linearly independent checks. We formalize this refinement by establishing a correspondence between

two chain complexes: our original complex, where the i-cochain consists of i-cells, and an alternative complex that combines cells of different dimensions. This relation is analogous to the relation between the constructions in [15] and [2]. Conceptually, the new chain complex reorganizes the qubits and checks to remove redundant checks that are linearly dependent on others. Within this modified framework, the counting of qubits and checks leads to a nontrivial lower bound on the code dimension.

2.4.1 Shifting the complex

We now describe the construction of the new chain complex. Recall that our main construction defines $C^i(X, \mathcal{F}) = \bigoplus_{f \in X(i)} V_f$ where we take all cells of dimension i to form C^i. Notice that the dimension of a cell f can be expressed as $|\texttt{type}(f)|$ and the space of all possible types $\texttt{type}(f) \subseteq [t]$ can be viewed as a t-dimensional hypercube. Therefore, the process of forming a chain complex from a cubical complex can be rephrased as choosing a diagonal direction of the t-dimensional hypercube. In the case of $C^i(X, \mathcal{F})$, the diagonal direction goes from $(0, 0, ..., 0)$ to $(1, 1, ..., 1)$.

More generally, given a subset $T \subseteq [t]$, it specifies a direction from 1_T to $1_{[t]-T}$, where 1_T is the vector whose i-th coordinate is 1 if $i \in T$ and 0 otherwise. This induces the following chain complex:

$$C^i(X, \mathcal{F}, T) = \bigoplus_{f \in X, |\texttt{type}(f) \Delta T| = i} V_f \tag{2.16}$$

where Δ is the symmetric difference of the two sets. In particular, $C^i(X, \mathcal{F}) = C^i(X, \mathcal{F}, \emptyset)$.

The coboundary map of this new chain complex, $\delta : C^i(X, \mathcal{F}, T) \to C^{i+1}(X, \mathcal{F}, T)$, is again a map between faces with incidence relations, just like the original chain complex. In particular, the coboundary maps can be decomposed into maps consist of $\delta_{f,f'} : V_f \to V_{f'}$, for $f, f' \in X$, $|\texttt{type}(f)\Delta T| = i$, and $|\texttt{type}(f')\Delta T| = i + 1$. If $f \succ f'$, we set the map to be co-res$_{f,f'}$, as defined in (2.10). If $f \prec f'$, we set the map to be res$_{f,f'}$, as defined in (2.11).

If there is no incidence relation between f and f', we set the map to be 0.

The coboundary maps satisfy $\delta \circ \delta = 0$ because, in some sense, the corestriction and restriction maps commute with themselves and with each other. To check $\delta \circ \delta = 0$, it is equivalent to check $\sum_{f' \in X} \delta_{f',f''}\delta_{f,f'} = 0$ for every $f, f'' \in X$. The LHS is nontrivial only when f and f'' are both incident to a common face and

- $\texttt{type}(f'') = \texttt{type}(f) \cup \{j,k\}$,
- $\texttt{type}(f'') = \texttt{type}(f) \cup \{j\} - \{k\}$,
- $\texttt{type}(f'') = \texttt{type}(f) - \{j,k\}$,

for $j, k \in [t]$ and $j \neq k$. In each case, there are exactly two f', say f'_1, f'_2, with nontrivial $\delta_{f',f''}$ and $\delta_{f,f'}$. What remains is to verify is the identity $\delta_{f'_1,f''}\delta_{f,f'_1} = \delta_{f'_2,f''}\delta_{f,f'_2}$, which follows from the definitions of the restriction and corestriction maps given in Section 2.3.3.

This new family of chain complexes is, in a certain sense, equivalent to the original chain complex from our main construction. In particular, for any subset $T \subseteq [t]$, there exists a sheaf $\mathcal{F}_T$ (which we will specify below) such that the chain complex $C^*(X, \mathcal{F})$ is chain homotopy equivalent to $C^*(X, \mathcal{F}_T, T)$. Furthermore, this chain homotopy equivalence is "sparse", which implies that the code parameters induced from the two chain complexes are the same up to constant factors. That means if we only care about the asymptotic scaling of the code parameters it is sufficient to study $C^*(X, \mathcal{F})$. See [22, App A] and [41, Sec 2.1] for more detail on the notion of chain homotopy equivalence.

We now formally define $\mathcal{F}_T$. We write $\mathcal{F}(\{h_i\}_{i=1}^t)$ to denote the sheaf constructed from $h_i : \mathbb{F}_q^{A_i} \to \mathbb{F}_q^{\hat{A}_i}$. Recall that $h_i^\perp : \mathbb{F}_q^{A_i} \to \mathbb{F}_q^{A_i - \hat{A}_i}$ is a linear map that satisfies $h_i^\perp h_i^T = 0$, i.e. the code with parity check matrix $h_i^\perp$ is dual to the code with parity check matrix h_i. Let $\mathcal{F}_T$ be the sheaf $\mathcal{F}(\{h'_i\}_{i=1}^t)$, where $h'_i = h_i$ if $i \notin T$ and $h'_i = h_i^\perp$ if $i \in T$. In some sense, the structures are flipped when $i \in T$.

Rather than establishing a full chain homotopy equivalence, Instead, for the purpose of deriving a lower bound on the dimension it suffices to prove a weaker statement: that the homologies of the two complexes have the same dimension. In the next subsection we show the following.

Theorem 2.19. *For any* $0 \leqslant i \leqslant t$ *and* $T \subseteq [t]$

$$\dim\ H^i(X, \mathcal{F}) = \dim\ H^i(X, \mathcal{F}_T, T). \tag{2.17}$$

Using the theorem, we can show the lower bound on the rate k_i claimed in Corollary 2.16. For this we require the following simple estimate.

Lemma 2.20. *Let* $D'_i = \dim\ C^i(X, \mathcal{F}(h'), T)$. *Then*

$$D'_i = N \sum_{S \subseteq [t], |S \Delta T| = i} n^{|S|} 2^{t-|S|} \prod_{j \notin S} m'_j \tag{2.18}$$

where m'_j *is the dimension of the codomain of* h'_j.

Proof. We have that $D'_i = \sum_{f \in X, |\mathtt{type}(f) \Delta T| = i} \dim\ V_f$. For each f, $\dim\ V_f = \prod_{j \notin \mathtt{type}(f)} m'_j$. Moreover, we have $|X(\mathtt{type}(f))| = N n^{|\mathtt{type}(f)|} 2^{t-|\mathtt{type}(f)|}$, where N is to choose $g \in G$, $n^{|\mathtt{type}(f)|}$ is to choose a_j for $j \in \mathtt{type}(f)$, and $2^{t-|\mathtt{type}(f)|}$ is to choose 0 or 1 for $j \notin \mathtt{type}(f)$. □

The results above lead to the bound on the code dimension claimed in Corollary 2.16.

Proof of the lower bound on k_i *in Corollary 2.16.* Choose $m_j = \nu n$ for $1 \leqslant j \leqslant i$ and $m_j = (1-\nu)n$ for $i+1 \leqslant j \leqslant t$. Let $T = \{i+1, ..., t\}$. We have the following bound on k_i

$$k_i = \dim\ H^i(X, \mathcal{F}) = \dim\ H^i(X, \mathcal{F}_T, T) \geqslant D'_i - D'_{i-1} - D'_{i+1} \tag{2.19}$$

Notice that $m'_j = \nu n$ for $1 \leqslant j \leqslant t$. Thus, $D'_i \geqslant N n^t$ from the term $S = \{1, ..., t\}$ alone.

Notice that $\sum_{j=0}^{t} D'_j = N\prod_{j=1}^{t}(n+2m'_j) = Nn^t(1+2\nu)^t$. Thus, $D'_{i-1} + D'_{i+1} \leqslant \sum_{j=0}^{t} D'_j - D'_i \leqslant Nn^t(1+2\nu)^t - Nn^t$ which implies $k_i \geqslant D'_i - D'_{i-1} - D'_{i+1} \geqslant 2Nn^t - Nn^t(1+2\nu)^t$. Therefore, as long as $\nu < (2^{1/t}-1)/2$, we have $k_i = \Theta(\ell_i)$. When evaluating the dimension over $\mathbb{F}_2$, this leads to $\Theta((\log_2 q)\ell_i)$. □

2.4.2 Proof of Theorem 2.19

We prove Theorem 2.19 using the following claims.

Claim 2.21. *For any $0 \leqslant i \leqslant t$ and $0 < j \leqslant t$,*

$$\dim\, H^i(X, \mathcal{F}_{[j-1]}, [j-1]) \leqslant \dim\, H^i(X, \mathcal{F}_{[j]}, [j]). \tag{2.20}$$

Claim 2.22. *For any $0 \leqslant i \leqslant t$ and $0 \leqslant j \leqslant t$,*

$$\dim\, H^i(X, \mathcal{F}_{[j]}, [j]) = \dim\, H^{t-i}(X, \mathcal{F}^{\perp}_{[t]-[j]}, [t]-[j]), \tag{2.21}$$

where $\mathcal{F}^{\perp}$ is the sheaf where every h_k is replaced with $h_k^{\perp}$.

Proof of Theorem 2.19. We first show the following corollary.

Corollary 2.23. *For any $0 \leqslant i \leqslant t$ and $0 < j \leqslant t$,*

$$\dim\, H^i(X, \mathcal{F}_{[j-1]}, [j-1]) = \dim\, H^i(X, \mathcal{F}_{[j]}, [j]).$$

Proof of Corollary 2.23. By Claim 2.21, we have $\dim H^i(X, \mathcal{F}_{[j-1]}, [j-1]) \leqslant \dim H^i(X, \mathcal{F}_{[j]}, [j])$. Additionally, we have

$$\dim H^{t-i}(X, \mathcal{F}^{\perp}_{[t]-[j]}, [t]-[j]) \leqslant \dim H^{t-i}(X, \mathcal{F}^{\perp}_{[t]-[j-1]}, [t]-[j-1]),$$

where we reverse the ordering of the t directions of the cubical complex, which maps $[t-j]$ to $[t]-[j]$ and $[t-j+1]$ to $[t]-[j-1]$. Together with Claim 2.22, we have

$$\dim H^i(X, \mathcal{F}_{[j]}, [j]) \leqslant \dim H^i(X, \mathcal{F}_{[j-1]}, [j-1]).$$

Thus, $\dim H^i(X, \mathcal{F}_{[j-1]}, [j-1]) = \dim H^i(X, \mathcal{F}_{[j]}, [j])$ as desired. □

By applying the corollary inductively, for any $0 \leqslant j \leqslant t$, we have

$$\dim H^i(X, \mathcal{F}) = \dim H^i(X, \mathcal{F}_\emptyset, \emptyset) = \dim H^i(X, \mathcal{F}_{[j]}, [j]).$$

Rearranging the order of the t directions of the cubical complex maps $[j]$ to T, which gives the desired result, $\dim H^i(X, \mathcal{F}) = \dim H^i(X, \mathcal{F}_T, T)$. □

We now prove the two claims. We first prove Claim 2.22, which is more straightforward.

Proof of Claim 2.22. The main observation is that the two chain complexes $C^*(X, \mathcal{F}_{[j]}, [j])$ and $C_{t-*}(X, \mathcal{F}^\perp_{[t]-[j]}, [t]-[j])$ are exactly the same. In particular,

$$\begin{aligned}
& C^i(X, \mathcal{F}_{[j]}, [j]) \\
= & \bigoplus_{f \in X, |\mathtt{type}(f) \Delta [j]| = i} V_f \\
= & \bigoplus_{f \in X, |\mathtt{type}(f) \Delta [j]| = i} \mathbb{F}_q^{\prod_{k \in \mathtt{type}(f) \setminus [j]} \hat{A}_k} \otimes \mathbb{F}_q^{\prod_{k \in \mathtt{type}(f) \cap [j]} A_k - \hat{A}_k}.
\end{aligned}$$

On the other hand,

$$\begin{aligned}
& C_{t-i}(X, \mathcal{F}^\perp_{[t]-[j]}, [t]-[j]) \\
= & \bigoplus_{f \in X, |\mathtt{type}(f) \Delta ([t]-[j])| = t-i} V_f^\perp \\
= & \bigoplus_{f \in X, |\mathtt{type}(f) \Delta [j]| = i} V_f^\perp \\
= & \bigoplus_{f \in X, |\mathtt{type}(f) \Delta [j]| = i} \mathbb{F}_q^{\prod_{k \in \mathtt{type}(f) \setminus ([t]-[j])} A_k - \hat{A}_k} \otimes \mathbb{F}_q^{\prod_{k \in \mathtt{type}(f) \cap ([t]-[j])} \hat{A}_k}
\end{aligned}$$

$$= \bigoplus_{f\in X, |\mathtt{type}(f)\Delta[j]|=i} \mathbb{F}_q^{\prod_{k\in\mathtt{type}(f)\cap[j]} A_k - \hat{A}_k} \otimes \mathbb{F}_q^{\prod_{k\in\mathtt{type}(f)\setminus[j]} \hat{A}_k}.$$

Thus, $C^i(X, \mathcal{F}_{[j]}, [j]) \cong C_{t-i}(X, \mathcal{F}^{\perp}_{[t]-[j]}, [t]-[j])$. Similarly, one can verify that the coboundary maps and the boundary maps agree. Therefore,

$$\dim H^i(X, \mathcal{F}_{[j]}, [j]) = \dim H_{t-i}(X, \mathcal{F}^{\perp}_{[t]-[j]}, [t]-[j]).$$

Because the dimension of homology equals to the dimension of cohomology, we have the desired result, $\dim H^i(X, \mathcal{F}_{[j]}, [j]) = \dim H^{t-i}(X, \mathcal{F}^{\perp}_{[t]-[j]}, [t]-[j])$. □

We now prove Claim 2.21.

Proof of Claim 2.21. To establish the result, we first construct a family of maps $\alpha^i : C^i(X, \mathcal{F}_{[j-1]}, [j-1]) \to C^i(X, \mathcal{F}_{[j]}, [j])$, and then show that these maps form a chain map, i.e. $\delta\alpha^i = \alpha^{i+1}\delta$. Finally, we prove that the induced map on cohomology is injective, which implies the desired result $\dim H^i(X, \mathcal{F}_{[j-1]}, [j-1]) \leqslant \dim H^i(X, \mathcal{F}_{[j]}, [j])$.

We start by constructing the maps $\alpha^i : C^i(X, \mathcal{F}_{[j-1]}, [j-1]) \to C^i(X, \mathcal{F}_{[j]}, [j])$, i.e. $\bigoplus_{f\in X, |\mathtt{type}(f)\Delta[j-1]|=i} V_{[j-1],f} \to \bigoplus_{f'\in X, |\mathtt{type}(f')\Delta[j]|=i} V_{[j],f'}$, where $V_{[j-1],f}$ is the local coefficients induced from $\mathcal{F}_{[j-1]}$. The map can be decomposed into maps consist of $V_{[j-1],f} \to V_{[j],f'}$ between $|\mathtt{type}(f)\Delta[j-1]| = i$ and $|\mathtt{type}(f')\Delta[j]| = i$. Most of the maps are the zero maps except for a few pairs of f and f' with a certain geometric relation.

To describe such a geometric relation, we establish a natural coordinate system in $\mathbb{R}^t$, where we assign each cell a point p_f, which is its center and takes value in $\{0, \frac{1}{2}, 1\}^t$. In particular, for the cell $f = [g; (a_k)_{k\in S}, (b_k)_{k\in\overline{S}}]$, the k-th coordinate of its center is

- 0 if $k \notin S, b_k = 0$;
- $\frac{1}{2}$ if $k \in S$;
- 1 if $k \notin S, b_k = 1$.

The map $\alpha : V_{[j-1],f} \to V_{[j],f'}$ is nonzero only when the faces f, f' are incident to each other and $p_{f'} - p_f = \frac{1}{2} \cdot e_j$, where e_j is the standard vector with value 0 at every coordinate except the j-th coordinate which has value 1. More explicitly, suppose $f = [g; (a_k)_{k\in S}, (b_k)_{k\in\overline{S}}]$. If $j \in S$, then $f' = [g \cdot a_j; (a_k)_{k\in S-\{j\}}, (b_k)_{k\in\overline{S}}, b_j = 1]$. Otherwise, if $j \notin S$, then $f' = [g; (a_k)_{k\in S}, a'_j, (b_k)_{k\in\overline{S}-\{j\}}]$ for some $a'_j \in A_j$. In the first case $f \succ f'$ and $p_{f,j} = \frac{1}{2}$, $p_{f',j} = 1$; In the second case $f \prec f'$ and $p_{f,j} = 0$, $p_{f',j} = \frac{1}{2}$.

For each pair f, f' that satisfies the above condition, we now define the map $\alpha : V_{[j-1],f} \to V_{[j],f'}$ i.e. $\mathbb{F}_q^{\prod_{k\in\overline{S}} \hat{A}_{[j-1],k}} \to \mathbb{F}_q^{\prod_{k\in\overline{S\Delta\{j\}}} \hat{A}_{[j],k}}$, where $\hat{A}_{[j],k}$ is $\hat{A}_k$ if $k \notin [j]$ and $A_k - \hat{A}_k$ if $k \in [j]$. Note that $\hat{A}_{[j-1],k} = \hat{A}_{[j],k}$ when $k \neq j$ and $\hat{A}_{[j-1],j} = \hat{A}_j$, $\hat{A}_{[j],j} = A_j - \hat{A}_j$. We again split into two cases. If $j \in S$, the domain and codomain simplify into $\mathbb{F}_q^{\prod_{k\in\overline{S}} \hat{A}_{[j],k}} \to \mathbb{F}_q^{\prod_{k\in\overline{S}} \hat{A}_{[j],k}} \otimes \mathbb{F}_q^{A_j - \hat{A}_j}$, and the map is obtained by extending the A_j-coordinate from a_j to $\mathbb{F}_q^{A_j}$ and applying $h_j^{\perp} : \mathbb{F}_q^{A_j} \to \mathbb{F}_q^{A_j - \hat{A}_j}$. Otherwise, if $j \notin S$, the domain and codomain simplify into $\mathbb{F}_q^{\prod_{k\in\overline{S}} \hat{A}_{[j],k}} \otimes \mathbb{F}_q^{\hat{A}_j} \to \mathbb{F}_q^{\prod_{k\in\overline{S}} \hat{A}_{[j],k}}$, and the map is obtained by applying $h_j^T : \mathbb{F}_q^{\hat{A}_i} \to \mathbb{F}_q^{A_i}$ and restricting the A_j-coordinate to a'_j. The first case is like applying the restriction map, except that it uses $h^{\perp}$ instead of h, while the second case is like applying the co-restriction map. This completes the construction of the maps α.

Next, we show that these maps form a chain map. Specifically, we need to verify that

$$\delta\alpha = \alpha\delta.$$

The proof is similar to the proof for $\delta \circ \delta = 0$. What we need to check is

$$\sum_{f'\in X} \delta_{f',f''}\alpha_{f,f'} = \sum_{f'\in X} \alpha_{f',f''}\delta_{f,f'}$$

for every $f, f'' \in X$. Since the maps are nonzero only when the faces are incident to each other, there are three possible relative positions between f and f''

- $p_{f''} - p_f = \frac{1}{2} \cdot e_j \pm \frac{1}{2} \cdot e_k$ for some $k \neq j$,
- $p_{f''} - p_f = e_j$,

- $p_{f''} - p_f = 0$.

The first case was analyzed in the proof of $\delta \circ \delta = 0$. For the second case, there is only one possible f' on both sides. So we need to check $\delta_{f',f''}\alpha_{f,f'} = \alpha_{f',f''}\delta_{f,f'}$ for that f'. Indeed, both sides are written as $\text{res}^{\perp}_{f',f''}\text{co-res}_{f,f'}$, where $\text{res}^{\perp}_{f',f''}$ is the restriction map using $h^{\perp}$.

For the third case, we will show that both sides are 0. Let us first focus on the left hand side $\sum_{f' \in X} \delta_{f',f''}\alpha_{f,f'}$. Suppose that for some f', $\delta_{f',f''}\alpha_{f,f'}$ is nonzero, then $p_{f'} - p_f = \frac{1}{2} \cdot e_j$ by the construction of α. Since $p_{f''} - p_f = 0$, this implies $p_{f''} - p_{f'} = -\frac{1}{2} \cdot e_j$. If we review the definition of δ for $C^*(X, \mathcal{F}_{[j]}, [j])$, either $p_{f''} - p_{f'} = \pm\frac{1}{2} \cdot e_k$ for some $k \neq j$ or $p_{f',j} = \frac{1}{2}$ to $p_{f'',j} = 0$ or 1. Thus, the map could be nonzero only when $p_{f,j} = p_{f'',j} = 0, p_{f',j} = \frac{1}{2}$. In this case, α acts as a co-restriction map using h_j and δ acts as a restriction map using $h_j^{\perp}$. Overall, we have $\sum_{f' \in X} \delta_{f',f''}\alpha_{f,f'} = (h_j^{\perp}h_j^T) \otimes I_{\dots} = 0$, where $I_{\dots}$ is the identity operator on the other components.

The analysis is similar for the right hand side $\sum_{f' \in X} \alpha_{f',f''}\delta_{f,f'}$. Suppose that for some f', $\alpha_{f',f''}\delta_{f,f'}$ is nonzero, then $p_{f''} - p_{f'} = \frac{1}{2} \cdot e_j$ by the construction of α. Since $p_{f''} - p_f = 0$, this implies $p_{f'} - p_f = -\frac{1}{2} \cdot e_j$. If we review the definition of δ for $C^*(X, \mathcal{F}_{[j-1]}, [j])$, either $p_{f'} - p_f = \pm\frac{1}{2} \cdot e_k$ for some $k \neq j$ or $p_{f',j} = \frac{1}{2}$ to $p_{f,j} = 0$ or 1. Thus, the map could be nonzero only when $p_{f,j} = p_{f'',j} = 1, p_{f',j} = \frac{1}{2}$. In this case, δ acts as a co-restriction map using h_j and α acts as a restriction map using $h_j^{\perp}$. Overall, we have $\sum_{f' \in X} \alpha_{f',f''}\delta_{f,f'} = (h_j^{\perp}h_j^T) \otimes I_{\dots} = 0$.

Finally, we show that the induced map $\alpha^i : H^i(X, \mathcal{F}_{[j-1]}, [j-1]) \to H^i(X, \mathcal{F}_{[j]}, [j])$ is injective. In particular, we need to show that for any $x \in Z^i(X, \mathcal{F}_{[j-1]}, [j-1])$, if $\alpha(x) \in B^i(X, \mathcal{F}_{[j]}, [j])$, then $x \in B^i(X, \mathcal{F}_{[j-1]}, [j-1])$. The remainder of the proof relies on two claims.

Claim 2.24. *For $0 \leqslant i \leqslant t$. If $\alpha(x) \in B^i(X, \mathcal{F}_{[j]}, [j])$, then there exists $z \in C^{i-1}(X, \mathcal{F}_{[j-1]}, [j-1])$, such that $\alpha(x) = \delta(\alpha(z))$.*

Claim 2.25. *For $0 \leqslant i \leqslant t$. If $w \in Z^i(X, \mathcal{F}_{[j-1]}, [j-1])$ satisfies $\alpha(w) = 0$, then $w \in B^i(X, \mathcal{F}_{[j-1]}, [j-1])$.*

(Note that $C^i(...)$ is defined to be $\{0\}$ for $i < 0$.)

We first prove the theorem assuming the two claims. By Claim 2.24, there exists z such that $\alpha(x) = \delta(\alpha(z)) = \alpha(\delta(z))$, where we use the fact that α is a chain map. This implies, $\alpha(w) = 0$ for $w = x - \delta(z) \in C^i(X, \mathcal{F}_{[j-1]}, [j-1])$. Because $x \in Z^i(X, \mathcal{F}_{[j-1]}, [j-1])$ and $\delta(z) \in B^i(X, \mathcal{F}_{[j-1]}, [j-1])$, we have $w \in Z^i(X, \mathcal{F}_{[j-1]}, [j-1])$. By Claim 2.25, this implies $w \in B^i(X, \mathcal{F}_{[j-1]}, [j-1])$. Since $x = w + \delta(z)$, we have the desired result $x \in B^i(X, \mathcal{F}_{[j-1]}, [j-1])$.

We now prove the two claims in order.

Proof of Claim 2.24. For $i = 0$, $B^i(X, \mathcal{F}_{[j]}, [j]) = \{0\}$, which means $\alpha(x) = 0$. Thus, we set $z = 0$.

For $i > 0$, since $\alpha(x) \in B^i(X, \mathcal{F}_{[j]}, [j])$, there exists $y \in C^{i-1}(X, \mathcal{F}_{[j]}, [j])$ such that $\alpha(x) = \delta(y)$. We claim that y can be expressed as $y' + \delta(s)$, where $s \in C^{i-2}(X, \mathcal{F}_{[j]}, [j])$ and y' is supported only on faces f such that $p_{f,j} \neq 0$. (When $i = 1$, y already has the desired support. So we can set $y' = y$.) This follows from the fact that part of the δ map that relates faces with the j-th coordinate equal to $\frac{1}{2}$ to those faces with the j-th coordinate equal to 0 is surjective. In particular, for every f with $p_{f,j} = 0$, consider the map $\bigotimes_{f' \succ_j f} V_{[j],f'} \to V_{[j],f}$ restricted from δ. By the definition of δ on $C^*(X, \mathcal{F}_{[j]}, [j])$, the map can be written as $h_j^{\perp} \otimes I_{...}$. Such a map is surjective, because $h_j^{\perp}$ is surjective. Therefore, we can find $s \in C^{i-2}(X, \mathcal{F}_{[j]}, [j])$, that cancels out the contribution of y on faces f where $p_{f,j} = 0$.

Since $\alpha(x) = \delta(y) = \delta(y')$, it suffices to find $z \in C^{i-1}(X, \mathcal{F}_{[j-1]}, [j-1])$ such that $\alpha(z) = y'$. In our construction, z is chosen to be supported on the faces where the j-th coordinate is either 0 or $\frac{1}{2}$. We first describe the faces with $p_{f,j} = \frac{1}{2}$. To define the value of

z on those faces, we utilize the restriction of α to $\bigotimes_{f \succ_j f'} V_{[j-1],f} \to V_{[j],f'}$ where $p_{f',j} = 1$. By the definition of α, this map can be written as $h_j^{\perp} \otimes I_{...}$, which is surjective. Therefore, we can demand $z|_{f \succ_j f'}$ to be the preimage of $y'|_{f'}$.

We now describe the faces with $p_{f,j} = 0$. To define the value of z on those faces, we utilize the restriction of α to $V_{[j-1],f} \to \bigotimes_{f' \succ_j f} V_{[j],f'}$ where $p_{f',j} = \frac{1}{2}$. By the definition of α, this map can be written as $h_j^T \otimes I_{...}$, which is not surjective. Nevertheless, we claim that $y'|_{f' \succ_j f}$ is in the image of $h_j^T \otimes I_{...}$, which allows us to define $z|_f$ as its preimage.

To show that $y'|_{f' \succ_j f}$ is in the image of $h_j^T \otimes I_{...}$, we use the property that $\alpha(x) = \delta(y')$ and the fact about their support. By construction $\alpha(x)$ and y' are not supported on the faces f where $p_{f,j} = 0$. That means $\alpha(x)|_f = (h_j^{\perp} \otimes I_{...})(y'|_{f' \succ_j f})$ which is equal to 0. Therefore, $y'|_{f' \succ_j f}$ is in the kernel of $h_j^{\perp} \otimes I_{...}$, which means it is in the image of $h_j^T \otimes I_{...}$, as claimed. □

Proof of Claim 2.25. The goal is to construct $v \in C^{i-1}(X, \mathcal{F}_{[j-1]}, [j-1])$ such that $w = \delta(v)$. For $i = 0$, we set $v = 0$. For $i > 0$, v is chosen to be supported on the faces where the j-th coordinate is 1. For every f with $p_{f,j} = 1$, consider the map $V_{[j-1],f} \to \bigotimes_{f' \succ_j f} V_{[j-1],f'}$ restricted from δ. By the definition of δ on $C^*(X, \mathcal{F}_{[j-1]}, [j-1])$, the map can be written as $h_j^T \otimes I_{...}$, which is not surjective. Nevertheless, we claim that $w|_{f' \succ_j f}$ is in the image of $h_j^T \otimes I_{...}$, which allows us to define $v|_f$ as its preimage.

To show that $w|_{f' \succ_j f}$ is in the image of $h_j^T \otimes I_{...}$, we use the property that $\alpha(w) = 0$. The restriction of the domain of α to $\{f' \succ_j f\}$ gives the map $\bigotimes_{f' \succ_j f} V_{[j-1],f'} \to V_{[j],f}$. By the definition of α, the map can be written as $h_j^{\perp} \otimes I_{...}$. Therefore, $w|_{f' \succ_j f}$ is in the kernel of $h_j^{\perp} \otimes I_{...}$, which means it is in the image of $h_j^T \otimes I_{...}$, as claimed.

We now show that $\delta(v) = w$. We check that the equality at faces with j-th coordinate equal to $0, \frac{1}{2}, 1$ in order. For faces with $p_{f,j} = 0$, we claim that both sides have a value of 0. We have $\delta(v)|_f = 0$ because v is supported on faces where the j-th coordinate is 1. This implies that $\delta(v)$ is supported on faces where the j-th coordinate is either 1 or $\frac{1}{2}$. To

check $w|_f = 0$, we use the fact that $\alpha(w) = 0$. Consider the map $V_{[j-1],f} \to \bigotimes_{f' \succ_j f} V_{[j],f'}$ restricted from α. By the definition of α, the map can be written as $h_j^T \otimes I_{\dots}$, which is injective. Since $\alpha(w)|_{f' \succ_j f} = 0$, we obtain $w|_f = 0$.

We now study the faces with $p_{f,j} = \frac{1}{2}$. For every f' with $p_{f',j} = 1$, by the construction of v, $w|_{f \succ_j f'} = (h_j^T \otimes I_{\dots})(v|_{f'})$. By the definition of δ, $(h_j^T \otimes I_{\dots})(v|_{f'}) = \delta(v)|_{f \succ_j f'}$. Thus, $\delta(v)$ and w agrees on the faces with $p_{f,j} = \frac{1}{2}$.

Finally, we study faces with $p_{f,j} = 1$. Let $w' = w - \delta(v)$. To show $w'|_f = 0$, we utilize the fact that $\delta(w') = 0$, and that w' is supported within the faces with $p_{f,j} = 1$. $\delta(w') = 0$ follows from $w \in Z^i(X, \mathcal{F}_{[j-1]}, [j-1])$, and the second statement follows from the previous parts of the proof. Consider the map $V_{[j-1],f} \to \bigotimes_{f' \succ_j f} V_{[j-1],f'}$ restricted from δ. By the definition of δ, the map can be written as $h_j^T \otimes I_{\dots}$, which is injective. Since $\delta(w')|_{f' \succ_j f} = 0$, we obtain $w'|_f = 0$, which implies $\delta(v)$ and w agrees on faces with $p_{f,j} = 1$. This completes the proof. □

□

2.5 The local chain

In this section we define a "local" chain complex $C(L_S)$, for each subset $S \subseteq \{1, \dots, t\}$. This chain complex will be used as a tool in the analysis of the complex $C_*(X, \mathcal{F})$ introduced in the previous section. In particular, as we will show in Lemma 2.27, the chain complex $C(L_S)$ is isomorphic to the "local view" of $C_*(X, \mathcal{F})$ from any face f such that $\texttt{type}(f) = \overline{S}$, i.e. the subcomplex $C_*(X_{\succcurlyeq f}; \mathcal{F})$. In Section 2.5.2 we introduce an important property of the local chain complex, *robustness*.

2.5.1 The complex

We start by defining a 1-dimensional complex $L_i = L_i(0) \cup L_i(1)$, where $i \in \{1, \ldots, t\}$, as follows. There is a single 0-dimensional face $L_i(0) = \{\emptyset\}$ and there are n_i 1-dimensional faces which we label using elements of A_i, $L_i(1) = A_i$. The incidence structure is the obvious one, $\emptyset \prec a$ for any $a \in A_i$.

Next we introduce coefficient spaces $V_\emptyset = \mathbb{F}_q^{\hat{A}_i}$ and $V_a = \mathbb{F}_q$ for each $a \in A_i$. The coboundary map is defined, for $x \in \mathbb{F}_q^{\hat{A}_i}$ and (with a slight abuse of notation, which we will repeat) a_i the a_i-th canonical basis vector of $\mathbb{F}_q^{A_i}$, as

$$\delta_{\{i\}}(x)(a_i) \;=\; (h_i^T(x))(a_i) \;.$$

The boundary map $\partial_{\{i\}} = \delta_{\{i\}}^T$ is then

$$\partial_{\{i\}}(y)(\emptyset) \;=\; h_i \, y \;,$$

where $y \in \oplus_a V_a \simeq \mathbb{F}_q^{A_i}$.

Above we used a subscript $\{i\}$ to distinguish them from the (co)boundary maps associated with the complex $C(X)$. We denote the resulting complex as $C(L_{\{i\}})$.

More generally, for $S \subseteq \{1, \ldots, t\}$ we define an $|S|$-dimensional complex $C(L_S)$ as follows. For $k \in \{0, \ldots, |S|\}$ the k-faces are $L_S(k) = \cup_{T \subseteq S, |T|=k} L_S(T)$, where the faces $L_S(T)$ *of type* T are in bijection with elements $\prod_{i \in T} A_i$. The incidence structure is given by $f' \succ f$ iff f' is of type T' such that $|T'| = |T| + 1$ and $T \subseteq T'$. The local coefficient space associated to a face f of type T is $V_f = \mathbb{F}_q^{\prod_{i \in S-T} \hat{A}_i}$. For $x \in V_f$, the coboundary map is defined as

$$\delta_S(x) \;=\; \sum_{j \in S-T} (I_{-j} \otimes h_j^T)(x) \;\in\; \bigoplus_{f' \succ f} V_{f'} \;,$$

and the boundary map is

$$\partial_S(x) \;=\; \sum_{j \in T} (I_{-j} \otimes h_j)(x) \;\in\; \bigoplus_{f' \prec f} V_{f'} \;.$$

Remark 2.26. One can verify that the complex $C(L_S)$ is the homological product [42] of the 1-dimensional complexes $C(L_{\{i\}})$, for $i \in S$. In particular, we see that $C_0(L_S) \simeq \otimes_{i\in S} C_0(L_{\{i\}})$, $C_1(L_S) \simeq \oplus_{i\in S}(C_1(L_{\{i\}}) \otimes_{j\in S-\{i\}} C_0(L_{\{j\}}))$, etc.

Lemma 2.27 (Local and global)**.** *For any face $f = [g; a, b]$ of type S there is a natural isomorphism between $C(X_{\geqslant f})$ and $C(L_{\overline{S}})$. In particular, for every $T \subseteq \overline{S}$, $C(X_{\geqslant f}(S\cup T)) \simeq C(L_{\overline{S}}(T))$.*

Proof. A face $g \in X_{\geqslant f}(S \cup T)$ is given by $a_T \in A_T$ and can be written as $g = (g_i)$ where $g_i = \begin{cases} f_i & i \notin T \\ a_i & i \in T \end{cases}$. In shorthand, $g = (f_{-T}||a_T)$. By definition, for a face g of type $S \cup T$, $V_g = \mathbb{F}_q^{\prod_{j\notin(S\cup T)} \hat{A}_j} = V_{S\cup T}$ so

$$C(X_{\geqslant f}(S \cup T)) = \bigoplus_{g\in X_{\geqslant f}(S\cup T)} V_g = \bigoplus_{a\in A_T} V_{(f_{-T}||a_T)} = \bigoplus_{a\in A_T} V_{S\cup T} \ .$$

An element of $C(X_{\geqslant f}(S \cup T))$ is thus a tuple $x = (x_{\bar{a}})_{\bar{a}\in A_T}$ where $x_{\bar{a}} \in \mathbb{F}_q^{\prod_{j\notin(S\cup T)} \hat{A}_j}$. Such a tuple is naturally identified with $x' \in \mathbb{F}_q^{\prod_{j\in T} A_j \prod_{j\in\overline{S}\setminus T} \hat{A}_j} = C(L_{\overline{S}}(T))$ by setting $x'_{\bar{a},\bar{a}'} = (x_{\bar{a}})_{\bar{a}'}$ for $\bar{a} \in \prod_{j\in T} A_j$ and $\bar{a}' \in \prod_{j\in\overline{S}\setminus T} \hat{A}_j$.

It is easy to check that under this isomorphism the boundary and coboundary maps of $C(L_{\overline{S}})$ and of $C(X_{\geqslant f})$ coincide. □

We will use the following.

Lemma 2.28. *For any $S \subseteq \{1, \ldots, t\}$ the chain complex $C(L_S)$ is exact at $i = 0, \ldots, |S|-1$, i.e. for any $x \in C_i(L_S)$ such that $\partial_S(x) = 0$ there is a $y \in C_{i+1}(L_S)$ such that $\partial_S(y) = x$.*

Proof. This follows from the Künneth formula, using that $C(L_S)$ is the homological product of 1-dimensional complexes $C(L_{\{i\}})$ which satisfy $\dim H_0(L_{\{i\}}) = 0$. □

Lemma 2.29. *For any $S \subseteq \{1, \ldots, t\}$ and for any $x \in C_{|S|}(L_S)$ such that $\partial_S(x) = 0$, x must be in the tensor code $\bigotimes_{j\in S} \ker(h_j)$.*

Proof. Observe that $C_{|S|}(L_S) \cong \mathbb{F}_q^{\prod_{j\in S} A_j}$, so an element $x \in C_{|S|}(L_S)$ is an $|S|$-dimensional tensor. The condition $\partial_S x = 0$ for $x \in C_{|S|}(L_S)$ implies that for every $j \in S$ and every $(a_j)_{j\in S} \in \prod_{j\in S} A_j$, the column vector $x(a_{-j}, \cdot)$ where we fix all coordinates except the j-th belongs to $\ker(h_j)$. This is exactly the definition of the space $\bigotimes_{j\in S} \ker(h_j)$. □

2.5.2 Robustness

Let $S \subseteq \{1, \ldots, t\}$ and $C(L_S)$ be the dimension-$|S|$ complex described in the previous section. For $i \in S$ let

$$\mathcal{C}_i^{\perp} = \mathrm{im}\,(h_i^T) \subseteq \mathbb{F}_q^{A_i} \ .$$

Let $d_i^{\perp}$ be the minimum distance of $\mathcal{C}_i^{\perp}$, i.e. $d_i^{\perp} = \min\{|x|_H : x \in \mathcal{C}_i^{\perp}\}$, where recall that $|\cdot|_H$ denotes the Hamming weight. For a k-face $f \in L_S(k)$ of type T and $x \in V_f = \mathbb{F}_q^{\prod_{i\in S\setminus T} \hat{A}_i}$ we let $|x| = 1_{x\neq 0}$. For $x \in C^k(L_S)$ we let $|x| = \sum_{f\in L_S(k)} |x(f)|$, i.e. the number of faces f such that $x(f)$ is nonzero. We refer to $|\cdot|$ as the *block-(Hamming)-weight*.

Definition 2.30. Let $k \in \{1, \ldots, |S|\}$. An element $x \in C^k(L_S)$ is called *minimal* if for any $y \in C^{k-1}(L_S)$, $|x + \delta_S(y)| \geqslant |x|$.

Definition 2.31 (Robustness)**.** Let $S \subseteq \{1, \ldots, t\}$ and $\ell = |S|$. For $0 \leqslant k \leqslant \ell - 1$ let $\kappa_{\ell,k}$ be a positive real. We say that $C(L_S)$ is $\kappa_{\ell,k}$*-robust* (implicitly, *at level* k) if for any $x \in C^k(L_S)$ such that x is minimal it holds that

$$|\delta_S(x)| \ \geqslant \ \kappa_{\ell,k}\, n\, |x| \ .$$

Remark 2.32. The condition of κ-robustness is stronger than the notion of product-expansion from [14] because it applies to all levels $k < |S|$ as opposed to only $k = |S| - 1$. Nevertheless, in Lemma 2.38 below we show a reduction from the later to the former. Furthermore,

both notions are equivalent to *co-boundary* expansion of the complex $C(L_S)$. For $k \in \{1, \ldots, |S|-1\}$ the co-boundary expansion coefficient $h^k(L_S)$ is defined as

$$h^k(L_S) = \min_{x \in C^k(L_S) - \mathrm{im}\,(\delta_S)} \frac{|\delta_S(x)|}{\min_{y \in \mathrm{im}\,(\delta_S)} |x+y|} .$$

We have that $h^k(L_S) \leqslant \kappa_{\ell,k} n$ because if x is minimal then $\min_{y \in \mathrm{im}\,(\delta_S)} |x+y| = |x|$. Conversely, if x achieves the minimum in the definition of $h^k(L_S)$ then $x' = x + y$ for any $y \in \mathrm{im}\,(\delta_S)$ that minimizes $|x+y|$ will be such that $|\delta_S(x')| = h^k(L_S)|x'|$.

Remark 2.33. For the case $|S| = 1$, the condition of $x \in C^0(L_S)$ being locally minimal is vacuous, because there are no (-1)-faces. Moreover, in that case $S = \{i\}$ for some i, $\delta_S(x) = h_i^T(x)$, $|x| = 1_{x \neq 0}$ and $|\delta_S(x)| = |\delta_S(x)|_H$, the Hamming weight. Thus in this case the condition of robustness is equivalent to the condition of distance, i.e. for each $i \in \{1, \ldots, t\}$ we have that $L_{\{i\}}$ is $(d_i^\perp/n)$-robust, where $d_i^\perp$ is the distance of the dual code $\mathcal{C}_i^\perp$.

Definition 2.34. Let $\kappa > 0$. We say that the family $\{h_i\}$ is *two-way κ-robust* if the following conditions hold.

- When based on the matrices $\{h_1, \ldots, h_t\}$, the complex $C(L_S)$ is κ-robust for each $S \subseteq \{1, \ldots, t\}$ and at each level $k = 0, \ldots, |S|-1$.
- When based on the matrices $\{h_1^\perp, \ldots, h_t^\perp\}$, the complex $C(L_S)$ is κ-robust for each $S \subseteq \{1, \ldots, t\}$ and at each level $k = 0, \ldots, |S|-1$.

We show the following:

Theorem 2.35. *For each collection of intervals $I_1, \ldots, I_\ell \subseteq (0,1)$, there exists $\kappa > 0$ such that for all $n \in \mathbb{N}$ there exist parity check matrices $h_i \in \mathbb{F}_q^{(1-\rho_i)n \times n}$, where $q = 2^{(n+3)^\ell}$, such that $\rho_i \in I_i$ and $\{h_1, \ldots, h_t\}$ is two-way κ-robust.*

We note that the theorem was previously known for $\ell = 1$ and $\ell = 2$ over fixed q (for example $\mathbb{F}_2$). The case for $\ell = 1$ is known as the Gilbert–Varshamov bound and the case for $\ell = 2$ was shown recently in [14, 15] (and earlier for a sub-constant robustness parameter in [2]). Our proof crucially relies on a recent result by Panteleev and Kalachev on the existence of code tuples with product expansion, provided that the field size q is large enough.

Theorem 2.36 (Panteleev and Kalachev [4])**.** *For each collection of intervals $I_1, \ldots, I_\ell \subseteq (0,1)$, there exists $\rho > 0$ such that for all $n \in \mathbb{N}$ there exist codes $\mathcal{C}_1 \ldots, C_\ell \subset \mathbb{F}_q^n$, where $q = 2^{(n+3)^\ell}$, such that $\frac{1}{n}\dim(\mathcal{C}_i) \in I_i$ and both collections $\mathcal{C}_1, \ldots, \mathcal{C}_\ell$ and $\mathcal{C}_1^\perp, \ldots, C_\ell^\perp$ are ρ-product expanding.*

In Lemma 2.38 in the following section we show that any ρ-product expanding family is also ρ'-robust, for positive ρ'. This proves Theorem 2.35.

2.5.3 Product expansion implies robustness

In this section we show that product expansion implies robustness. Recall the definition.

Definition 2.37 (Product expansion, [14])**.** A family of codes $\{\mathcal{C}_i\}_{i\in[\ell]}$ is said to have ρ-product expansion iff the co-cycle expansion of $L_{[\ell]}$ at level $\ell - 1$ is at least ρ, namely, if $\varepsilon_{\text{cocyc}}(\ell - 1) \geqslant n \cdot \rho$.

To be precise, the definition in [14] is given in slightly different terms, but it is exactly equivalent to the definition above, as shown in [14, Appendix B].

For our results, we need a bit more: as stated in Definition 2.34, we need not only coboundary expansion at level $\ell - 1$, but also for *all* levels $k \leqslant \ell - 1$. In the next lemma we show that the latter follows from the former. For convenience, assume for the remainder of this section that the family $\{\mathcal{C}_i\}_i$ is fixed, and we omit it from our notation.

Lemma 2.38. *Let $\ell \geqslant 1$ and S such that $|S| = \ell$. For every $\rho > 0$ there is some $\rho' > 0$ (depending on ρ and ℓ) such that the following holds. Suppose that for every $S' \subsetneq S$ and $\ell' = |S'| \geqslant 1$ the complex $L_{S'}$ satisfies $\kappa_{\ell',\ell'-1} \geqslant \rho$. Then for all $0 \leqslant k \leqslant \ell - 1$, $\kappa_{\ell,k} > \rho'$.*

The proof of the lemma is given in the next two subsections.

Proof of Lemma 2.38

We prove Lemma 2.38 by induction on ℓ. For $\ell = 1$ it is true because the only relevant case, $k = 0$, is covered by product expansion. We assume for convenience that $S = \{1, \ldots, \ell\} = [\ell]$. Assume that we have a lower bound on $\kappa_{\ell,k}$ for all $0 \leqslant k < \ell$. We show a lower bound on $\kappa_{\ell+1,k}$, for any $k < \ell$ (the case $k = \ell$ is covered by product expansion).

Recall the sheaf $\mathcal{F}_\ell$, which is such that for any $s \in L_{[\ell]}$, $\mathcal{F}_\ell(s) = \mathbb{F}_q^{\prod_{i \in [\ell] \setminus \mathtt{type}(s)} \hat{A}_i}$. Let $C(L_{[\ell+1]}, \mathcal{F}_{\ell+1})$ denote the chain complex. We partition $L_{[\ell+1]}$ into two parts, L_0 and L_1, as follows.

1. L_0 consists of all faces whose type is contained in $[\ell] = \{1, \ldots, \ell\}$. Clearly $L_0 \cong L_{[\ell]}$. Moreover, for each face $s \in L_0$,

$$(\mathcal{F}_\ell(s))^{\hat{A}_{\ell+1}} = (\mathbb{F}_q^{\prod_{i \in [\ell] \setminus \mathtt{type}(s)} \hat{A}_i})^{\hat{A}_{\ell+1}} \cong \mathbb{F}_q^{\prod_{i \in [\ell+1] \setminus \mathtt{type}(s)} \hat{A}_i} = \mathcal{F}_{\ell+1}(s).$$

   So

$$C^k(L_0, \mathcal{F}_{\ell+1}) \cong C^k(L_{[\ell]}, \mathcal{F}_\ell^M) , \tag{2.22}$$

   for $M = |\hat{A}_{\ell+1}|$. Here, $\mathcal{F}_\ell^M$ is the sheaf where each space $\mathcal{F}_\ell(s)$ is replaced by the cartesian product $\mathcal{F}_\ell(s)^M$. We will rely on Proposition 2.39 below that upgrades coboundary expansion with respect to coefficients from $\mathcal{F}$ to the same with respect to $\mathcal{F}^M$ for any $M \in \mathbb{N}$.

2. L_1 consists of faces whose type contains $\ell+1$. These faces can be written as $(s; j)$ for a face $s \in L_{[\ell]}$ and an element $j \in A_{\ell+1}$. Thus $L_1(k)$ is partitioned into $n = |A_{\ell+1}|$ copies

of $L_{[\ell]}(k)$. Put differently, for each $s \in L_{[\ell]}$ we have an n-tuple $(\mathcal{F}_{\ell+1}((s;j)))_{j\in A_{\ell+1}}$ of coefficients that naturally lives in $\mathcal{F}_\ell(s)^n$. So

$$C^k(L_1, \mathcal{F}_{\ell+1}) \cong C^k(L_{[\ell]}, \mathcal{F}_\ell^n).$$

Here we are using the fact that there is a face $(s;j) \in L_1$ for each $j \in A_{\ell+1}$, and the coefficient space of this face satisfies

$$\mathcal{F}_{\ell+1}((s;j)) = \mathbb{F}_q^{\prod_{i\in[\ell+1]\setminus \mathtt{type}((s;j))} \hat{A}_i} = \mathbb{F}_q^{\prod_{i\in[\ell]\setminus \mathtt{type}(s)} \hat{A}_i} = \mathcal{F}_\ell(s).$$

For $x \in C^k(L_1, \mathcal{F}_{\ell+1})$, we can decompose $x = \sum_{s\in L_1} x(s)$ according to the value $j \in A_{\ell+1}$ that s takes on the $(\ell+1)$-st coordinate, as

$$x = \sum_{j\in A_{\ell+1}} y_j \circ j \ , \qquad y_j \in C^k(L_{[\ell]}, \mathcal{F}_\ell) \tag{2.23}$$

where $y_j \circ j \in C^k(L_{[\ell+1]}, \mathcal{F}_{\ell+1})$ is defined by

$$\forall s \in L_{[\ell]}(k) \ , \quad \forall i \in A_{\ell+1} \ , \qquad (y_j \circ j)(s;i) = \begin{cases} y_j(s) & i = j \\ 0 & i \neq j \end{cases}.$$

Let $x = \delta y \in C^{k+1}(L_{[\ell+1]}, \mathcal{F}_{\ell+1})$ be such that $|x| = \varepsilon|X(k+1)|$. Write $x_0 = x|_{L_0} \in C^k(L_0, \mathcal{F}_{\ell+1})$ and observe that $x_0 = (\delta y)|_{L_0} = \delta_\ell(y|_{L_0})$, where δ_ℓ denotes the coboundary map in $C(L_{[\ell]}, \mathcal{F}_{\ell+1}) \cong C(L_{[\ell]}, \mathcal{F}_\ell^M)$. (This is because δ_ℓ only adds coordinates in one of the first ℓ directions.) To complete the argument we need a bound on $\bar{\kappa}_{\ell,k}$, the coboundary expansion constant of the chain complex $C^k(L_{[\ell]}, \mathcal{F}_\ell^M)$. Let $\bar{\kappa}_{\ell,k}$ be the infimum of all coboundary expansion constants of $C^k(L_{[\ell]}, \mathcal{F}_\ell^M)$, taken over all natural integers M. We call $\bar{\kappa}_{\ell,k}$ the *collective robustness* because of its similarity to the notion of *collective cosystolic expansion* studied in [43, Definition 1.2].

In Section 2.5.3 below we show the following.

Proposition 2.39. *Let $\kappa = \min_{i<\ell'\leqslant\ell} \kappa_{\ell',i}$. There is a constant $A > 0$ depending only on ℓ such that $\bar{\kappa}_{\ell,k} \geqslant A \cdot \kappa^{2\ell}$.*

In words, this proposition roughly says that if the chain complex $C^*(L_{[\ell]}, \mathcal{F}_\ell)$ has constant robustness $\kappa_{\cdot,\cdot}$, then the chain complex has constant collective robustness $\bar{\kappa}_{\cdot,\cdot}$.

By the inductive hypothesis we have a lower bound for $\kappa_{\ell',i}$ for all $i < \ell' \leqslant \ell$, so by the proposition, there is $\tilde{y}_0 \in C^k(L_{[\ell]}, \mathcal{F}_\ell^M)$ such that $\delta_\ell \tilde{y}_0 = x_0$ and

$$\bar{\kappa}_{\ell,k} n \cdot |\tilde{y}_0| \leqslant |x_0| \leqslant |x|. \tag{2.24}$$

Let $x' = x - \delta\tilde{y}_0 \in C^k(L_1, \mathcal{F}_{\ell+1})$. By construction $x' \in \operatorname{im}\delta$ since $x = \delta(y - \tilde{y}_0)$, and also, using $n = |A_{\ell+1}|$,

$$|x'| \leqslant |x| + |\delta\tilde{y}_0| \leqslant |x| + n\ell \cdot |\tilde{y}_0| \leqslant |x| + \frac{n\ell}{\bar{\kappa}_{\ell,k} n}|x| \leqslant (1 + \frac{\ell}{\bar{\kappa}_{\ell,k}})|x| \ . \tag{2.25}$$

Claim 2.40. *There are chains $z_j \in C^{k-1}(L_0, \mathcal{F}_\ell)$ for each $j \in A_{\ell+1}$, such that*

$$x' = \sum_{j \in A_{\ell+1}} (\delta_\ell z_j) \circ j \ . \tag{2.26}$$

Proof. We have $x' = \delta y'$ for $y' = y - \tilde{y}_0$. Since $x'|_{L_0} = 0$, we deduce that $y'|_{L_0} \in \ker \delta_\ell$. By exactness of δ_ℓ there is some z' such that $y'|_{L_0} = \delta_\ell z'$. Letting $y'' = y' - \delta z'$, we still have $\delta y'' = x'$ but now $y''|_{L_0} = y'|_{L_0} - \delta z'|_{L_0} = \delta_\ell z' - \delta z'|_{L_0} = 0$. In conclusion, $y'' \in C^k(L_1)$, so by (2.23) it can be written as $y'' = \sum_{j\in A_{\ell+1}} z_j \circ j$, and then we can write $x' = \delta y'' = \sum_{j\in A_{\ell+1}} \delta_\ell z_j \circ j$. □

Let $y'' = \sum_{j\in A_{\ell+1}} \tilde{z}_j \circ j$, where $\tilde{z}_j$ are chosen to have smallest weight while satisfying $\delta_\ell \tilde{z}_j = \delta_\ell z_j$. By induction $\kappa_{\ell,k-1} n \cdot |\tilde{z}_j| \leqslant |\delta_\ell z_j|$, and

$$|y''| = \sum_j |\tilde{z}_j| \leqslant \sum_j \frac{|\delta_\ell z_j|}{n\kappa_{\ell,k-1}}$$

$$= \frac{|x'|}{\kappa_{\ell,k-1} n}$$
$$\leqslant \frac{1}{\kappa_{\ell,k-1} n}(1 + \frac{\ell}{\bar{\kappa}_{\ell,k}})|x| \,, \tag{2.27}$$

where the second line is by (2.26) and the third is by (2.25). This allow us to define $\tilde{y} = \tilde{y}_0 + y''$, which satisfies $\delta\tilde{y} = x$ and

$$|\tilde{y}| \leqslant |\tilde{y}_0| + |y''| \leqslant \left(\frac{1}{\bar{\kappa}_{\ell,k} n} + \frac{1}{\kappa_{\ell,k-1} n}(1 + \frac{\ell}{\bar{\kappa}_{\ell,k}}) \right) |x|$$

which follows from (2.24) and (2.27). Thus,

$$\kappa_{\ell+1,k} \geqslant \frac{1}{\frac{1}{\bar{\kappa}_{\ell,k}} + \frac{1}{\kappa_{\ell,k-1}}(1 + \frac{\ell}{\bar{\kappa}_{\ell,k}})}$$

which concludes the induction step. □

Proof of Proposition 2.39

In this subsection we prove Proposition 2.39, which lower bounds the coboundary expansion of a direct product sheaf based on the coboundary expansion of the original sheaf. To lighten the notation, we fix ℓ and let $X = L_\ell$, and we fix $M \in \mathbb{N}$ and denote by $\bar{\mathcal{F}}$ the sheaf that to each face f of L_ℓ associates the coefficient space $\bar{\mathcal{F}}(f) = (\mathcal{F}(f))^M$. The coboundary operators in this complex simply operate component-wise. Moreover, the weight of a chain is *still* the number of faces with non-zero values. Let $\bar{\kappa}_{\ell,k}$ be the robustness of $C(X, \bar{\mathcal{F}})$ at level k, as in Definition 2.31.

Definition 2.41. Let $0 < p \leqslant 1$ and $0 \leqslant r \leqslant i$. Given an i-chain z we say that an r-face e is *p-heavy* for z if

$$|z(X_{\geqslant e})| \geqslant p \cdot |X_{\geqslant e}(i)| \,.$$

If e is not p-heavy then we say that it is *p-light*.

Proof of Proposition 2.39. Let $\mathbf{y}_1 \in C^k(L_{[\ell]}, \bar{\mathcal{F}})$ and let $x = \delta\mathbf{y}_1$. We will find $\mathbf{y} \in C^k(L_{[\ell]}, \bar{\mathcal{F}})$ such that $\delta\mathbf{y} = x$ and $|\mathbf{y}| \leqslant \frac{\varepsilon}{\bar{\kappa}_{\ell,k}} \cdot |X(k)|$, where $|x| = \varepsilon \cdot |X(k+1)|$. We

may assume wlog that $\varepsilon \leqslant \frac{1}{3}(\frac{\kappa}{4})^{2\ell}$. This is because if not, then as long as $A \leqslant \frac{1}{3(16)^\ell}$ the proposition follows trivially as

$$\bar{\kappa}_{\ell,k} n \cdot |\mathbf{y}_1| \leqslant \bar{\kappa}_{\ell,k} |X(k+1)| \leqslant \varepsilon |X(k+1)| = |x| \ .$$

For each $\alpha \in \mathbb{F}_q^M$, we denote $\alpha \cdot \mathbf{y}_1$ the chain in $C^k(X, \mathcal{F})$ that satisfies $(\alpha \cdot \mathbf{y}_1)(s) = \alpha \cdot (\mathbf{y}_1(s))$ for every $s \in X(k)$. By linearity, $\delta(\alpha \cdot \mathbf{y}_1) = \alpha \cdot (\delta \mathbf{y}_1) = \alpha \cdot x$. Observe that for any face s, $x(s) = 0$ implies $\alpha \cdot x(s) = 0$ so $|\alpha \cdot x| \leqslant |x|$. Let $\hat{y}_\alpha \in C^k(L_{[\ell]}, \mathcal{F})$ be the minimum-weight chain such that $\delta \hat{y}_\alpha = \alpha \cdot x$. By coboundary expansion of $C(L_{[\ell]}, \mathcal{F})$,

$$\forall \alpha, \quad \kappa_{\ell,k} n \cdot |\hat{y}_\alpha| \leqslant |\delta \hat{y}_\alpha| = |\alpha \cdot x| \leqslant |x|. \tag{2.28}$$

Let $\mathsf{e}_1, \ldots, \mathsf{e}_M \in \mathbb{F}_q^M$ be the standard basis and define

$$\forall s \in X(k), \quad \mathbf{y}(s) = (\hat{y}_{\mathsf{e}_1}(s), \ldots, \hat{y}_{\mathsf{e}_M}(s)). \tag{2.29}$$

Clearly $\delta \mathbf{y} = (\mathsf{e}_1 \cdot x, \ldots, \mathsf{e}_M \cdot x) = x$. To conclude the proof we will show the following claim.

Claim 2.42. *Let $\rho = \frac{1}{3}(\frac{\kappa}{4})^{2\ell}$. Let $H \subseteq X(k)$ be the set of k-faces that are above some ρ-heavy face of x:*

$$H \ = \ \{ s \in X(k) \,|\, \exists e \prec s, \ e \text{ is } \rho\text{-heavy for } x \} \ . \tag{2.30}$$

Then

$$\forall s \notin H, \forall \alpha, \beta \in \mathbb{F}_q^M, \qquad \hat{y}_\alpha(s) + \hat{y}_\beta(s) = \hat{y}_{\alpha+\beta}(s) \ . \tag{2.31}$$

Moreover,

$$|H| \ \leqslant \ 2^k \ \frac{\varepsilon}{\rho} \cdot |X(k)| \ . \tag{2.32}$$

We first show how the claim can be used to conclude the proof. For any face $s \notin H$ and for any $\alpha \in \mathbb{F}_q^M$, (2.31) implies that $\hat{y}_\alpha(s) = \sum_i \alpha_i \hat{y}_{\mathsf{e}_i}(s)$. If $\mathbf{y}(s) \neq 0$ then there is some

i such that $\hat{y}_{\mathsf{e}_i}(s) \neq 0$ which means that $\mathbb{P}_\alpha[\hat{y}_\alpha(s) \neq 0] \geqslant 1/2$, where the probability is taken over a uniformly random choice of α. We conclude that for every $s \in \operatorname{supp}(\mathbf{y}) \setminus H$, at least half of the α have $\hat{y}_\alpha(s) \neq 0$. Reversing the order of the quantifiers, there is some α such that $\hat{y}_\alpha(s) \neq 0$ for at least half of $s \in \operatorname{supp}(\mathbf{y}) \setminus H$. Namely $\frac{1}{2}|\operatorname{supp}(\mathbf{y}) \setminus H| \leqslant |\hat{y}_\alpha|$, or

$$|\mathbf{y}| \leqslant 2|\hat{y}_\alpha| + |H|. \tag{2.33}$$

Using the bound (2.32) on the size of H,

$$\begin{aligned} n|\mathbf{y}| \leqslant 2n|\hat{y}_\alpha| + n|H| &\leqslant 2(\kappa_{\ell,k})^{-1}|x| + 2^k n\frac{\varepsilon}{\rho}|X(k)| \\ &\leqslant \Big(\frac{2}{\kappa_{\ell,k}} + \frac{2^k t}{\rho}\Big)|x| \,, \end{aligned}$$

where the last line uses $(n/t)|X(k)| \leqslant |X(k+1)|$ and the definition of ε. Since by definition $\kappa_{\ell,k} \geqslant \kappa$, this shows that $\bar{\kappa}_{\ell,k} = \Omega_k(1)\kappa^{2\ell}$, as desired.

It remains to prove Claim 2.42. We start by showing the upper bound (2.32) on the size of H. Towards this, consider the random process of choosing a face in $X(k+1)$ by first choosing uniformly at random an r-face e and then a random $(k+1)$-face containing it, $s \succeq e$. If e is ρ-heavy for x, then the probability that $x(s) \neq 0$ is at least ρ. The fraction of r-faces that are heavy for x is thus at most ε/ρ. A k-face has exactly $\binom{k}{r}$ r-faces below it, so it has this many chances to contain a heavy r-face. By a union bound, the probability of this event is at most $\binom{k}{r} \cdot \varepsilon/\rho$. Summing over all $r \leqslant k$,

$$|H| \leqslant \sum_{r=1}^{k} \binom{k}{r} \frac{\varepsilon}{\rho} \cdot |X(k)| = 2^k \ \frac{\varepsilon}{\rho} \cdot |X(k)| \,.$$

It remains to prove (2.31). First observe that for all α, β,

$$\delta\hat{y}_\alpha + \delta\hat{y}_\beta - \delta\hat{y}_{\alpha+\beta} = \alpha \cdot x + \beta \cdot x - (\alpha+\beta) \cdot x = 0.$$

So $\hat{y}_\alpha + \hat{y}_\beta - \hat{y}_{\alpha+\beta} \in \ker \delta_k = \operatorname{im} \delta_{k-1}$ and therefore we can write

$$\hat{y}_\alpha + \hat{y}_\beta - \hat{y}_{\alpha+\beta} = \delta z \tag{2.34}$$

for some $z \in C^{k-1}(X, \mathcal{F})$. Let z be such a chain with minimal weight. We assume that $\delta z(s_0) \neq 0$ for some $s_0 \in X(k)$, and show that $s_0 \in H$. This will show (2.31).

We begin with the following claim that will help find a face e below s_0 that is heavy for δz.

Claim 2.43. *Let $1 \leqslant r \leqslant i < \ell$ and let $z \in C^i(X)$. Define $f_{i,r}, g_{i,r} : [0,1] \to [0,1]$ by*

$$f_{i,r}(p) = \frac{\kappa_{\ell-r,i-r}}{2} \cdot \frac{i-r+1}{r(\ell-r+1)} \cdot p \,, \quad \text{and} \quad g_{i,r}(p) = \frac{\kappa_{\ell-r,i-r}}{2} \cdot \frac{i-r+1}{\ell-i} \cdot p \,.$$

Then for any $0 < p \leqslant 1$, if $e \in X(r)$ is p-heavy for z yet each $e' \prec e$ is $f_{i,r}(p)$-light for z, then e is $g_{i,r}(p)$-heavy for δz.

We defer the proof of the claim to the end of this section. Since $\delta z(s_0) \neq 0$, there must be some $e_0 \prec s_0$ such that $z(e_0) \neq 0$. Clearly $e_0 \in X(k-1)$ itself is $p_0 = 1$-heavy for $z \in C^{k-1}(X)$. If all $e \prec e_0$ are $p_1 = f_{k-1,k-2}(p_0)$-light for z, we let $s_1 = e_0$. By Claim 2.43, s_1 is $q_1 = g_{k-1,k-2}(1)$-heavy for δz. Otherwise, choose $e_1 \in X(k-2)$, $e_1 \prec e_0$ a p_1-heavy face for z. If all $e \prec e_1$ are $p_2 = f_{k-1,k-3}(p_1)$-light, we let $s_1 = e_1$. By Claim 2.43, s_1 is $q_2 = g_{k-1,k-3}(p_1)$-heavy for δz. Otherwise we continue taking $e_j \prec \cdots e_1 \prec e_0$ such that $e_j \in X(k-1-j)$ is p_j-heavy for z, where

$$\begin{aligned} p_j &= f_{k-1,k-j-1} \circ \cdots \circ f_{k-1,k-2}(1) \\ &= \kappa_{\ell-k+2,1} \cdots \kappa_{\ell-k+j+1,j} \cdot \frac{1}{2^j} \cdot \frac{2 \cdots (j+1)}{\ell \cdots j(\ell-j+1)} \\ &= \frac{1}{2^j} \cdot \kappa_{\ell-k+2,1} \cdots \kappa_{\ell-k+j+1,j} \cdot \frac{j+1}{(\ell-j+1) \cdots \ell} \,. \end{aligned}$$

We can continue this process until either $j = k$ or for some $j < k$ all $e \prec e_j$ are p_{j+1}-light for z. In the latter case Claim 2.43 applies and we find s_1 that is $q_{j+1} = g_{k-1,k-j-2}(p_j)$-heavy for δz. Note that if $j = k$ then there is a single face $e_{j+1} \in X(0)$ such that $e_{j+1} \preceq e_0$, and being p_{k-1}-light is the same as the fractional weight of x being smaller than p_{k-1}, which holds by assumption (given the assumption made on ε at the start of the proof).

To summarize the argument, letting q be the minimum of all q_j, for $j \in \{k-1, \ldots, 1\}$, we have found $s_1 \preceq e_0$ such that s_1 is at least q-heavy for $\delta(z)$. Moreover, $q = \Omega(\kappa^\ell)$, with the implicit constant depending on ℓ only.

By (2.34), and relying on the triangle inequality, one of $\hat{y}_\alpha, \hat{y}_\beta$, or $\hat{y}_{\alpha+\beta}$ have at least $1/3$ of the weight of δz above s_1. Assume without loss of generality that it is $\hat{y}_\alpha$, so s_1 is $q/3$-heavy for $\hat{y}_\alpha$.

We now repeat the same argument as before. This yields a sequence $s_2 = e_{j'} \prec e_{j'-1} \prec \cdots \prec s_1$ such that s_2 is $\Omega(\kappa^\ell q) = \Omega(\kappa^{2\ell})$-heavy for $\delta\hat{y}_\alpha = \alpha \cdot x$. Thus s_2 is also heavy for x. Since $s_0 \succeq s_1 \succeq s_2$ this implies $s_0 \in H$,as desired. □

Proof of Claim 2.43. By definition,

$$|z(X_{\geqslant e}(i))| \geqslant p \cdot |X_{\geqslant e}(i)| = p \cdot \binom{\ell - r}{i - r} \cdot n^{i-r}. \tag{2.35}$$

Let δ_e denote the coboundary operator of $X_{\geqslant e}$ and notice that $X_{\geqslant e}(i) \cong L_{[\ell-r]}(i-r)$. (Note that we allow $r = i$, in which case this complex has a single element ϕ in level $i - r = 0$.) By coboundary expansion of $L_{[\ell-r]}(i-r)$,

$$\begin{aligned}
|\delta_e z(X_{\geqslant e}(i))| &\geqslant n \cdot \kappa_{\ell-r,i-r} \cdot |z(X_{\geqslant e}(i))| \\
&\geqslant n \cdot \kappa_{\ell-r,i-r} \cdot p \cdot |X_{\geqslant e}(i)| \\
&= \kappa_{\ell-r,i-r} \cdot \frac{\binom{\ell-r}{i-r}}{\binom{\ell-r}{i-r+1}} \cdot p \cdot |X_{\geqslant e}(i+1)| \\
&= 2 \cdot g_{i,r}(p) \cdot |X_{\geqslant e}(i+1)| \,.
\end{aligned} \tag{2.36}$$

We need to compute the weight of δz inside $X_{\geqslant e}(i+1)$. Recall that for $z \in C^i(X)$ and $s \in X_{\geqslant e}(i+1)$ the coboundary map is defined as

$$\delta z(s) \;=\; \sum_{s' \prec s} \text{co-res}_{s',s}(z) \;=\; \sum_{s' \prec_j s} (I_{-j} \otimes h_j^T)(z(s'))[..., s_j, ...] \,.$$

Separating i-faces into $s' \succ e$ and $s' \not\succ e$,

$$\delta z(s) = \sum_{e \not\prec s' \prec s} \text{co-res}_{s',s} z + \sum_{e \prec s' \prec s} \text{co-res}_{s',s} z = \sum_{e \not\prec s' \prec s} \text{co-res}_{s',s} z + \delta_e z_e(s) \,. \tag{2.37}$$

We claim that the first summand on the right hand side is non-zero only for few $s \in X_{\geqslant e}(i)$. The key is that for each fixed $s' \not\succ e$, co-res$_{s',s}(z)$ can be non-zero for exactly one $s \in X_{\geqslant e}(i+1)$, namely, the one that is above both s' and e. Each such s' can be 'charged' to e', the $(r-1)$-face directly below both e and s' (this is the empty face in case $r = 1$). However, by assumption on e, no $e' \prec e$ is $f_{i,r}(p)$-heavy for z. So the number of faces s' that are charged to e' is at most $|z_{e'}| < f_{i,r}(p)|X_{\geqslant e'}(i)|$. Summing over all r possible $e' \prec e$, this gives at most

$$r \cdot f_{i,r}(p) \cdot |X_{\geqslant e'}(i)| = |X_{\geqslant e}(i+1)| \cdot g_{i,r}(p)$$

non-zero faces that come from the first summand in (2.37). Observe that we are comparing two numbers of the same order of magnitude: the number of $(i+1)$-faces above e, $|X_{\geqslant e}(i+1)| = \binom{\ell-r}{i+1-r} n^{i+1-r}$, and the number of i-faces above e', $|X_{\geqslant e'}(i)| = \binom{\ell-r+1}{i-r} n^{i+1-r}$.

Plugging in (2.36) we deduce

$$\begin{aligned} |(\delta z)(X_{\geqslant e}(i+1))| &\geqslant |\delta_e(z(X_{\geqslant e}(i)))| - r f_{i,r}(p)|X_{\geqslant e'}(i)| \\ &\geqslant 2g_{i,r}(p)|X_{\geqslant e}(i+1)| - r f_{i,r}(p)|X_{\geqslant e'}(i)| \\ &= g_{i,r}(p) \cdot |X_{\geqslant e}(i+1)| \,, \end{aligned}$$

so e is $g_{i,r}(p)$-heavy for δz. □

2.6 Expansion properties of the geometric complex

In this section we we define some random walks on X that will play an important role in the analysis, and show that all these random walks have good expansion properties, assuming only that the "1-dimensional" random walks on the Cayley graphs $\mathrm{Cay}(G, A_i)$ are expanding. Before proceeding, we start by formulating some useful incidence properties of the geometric complex X.

2.6.1 Incidence properties

Lemma 2.44. *Let $0 \leqslant i \leqslant t$. Let $f \in X(i)$. Then*

$$|\mathfrak{d}(f)| = 2i \qquad \text{and} \qquad |\mathfrak{u}(f)| = (t-i)n\ . \tag{2.38}$$

More generally, if $0 \leqslant \ell < i < k \leqslant t$ then

$$|X_{\geqslant f}(k)| = \binom{t-i}{k-i} n^{k-i} \qquad \text{and} \qquad |X_{\leqslant f}(\ell)| = \binom{i}{\ell} 2^{i-\ell}\ . \tag{2.39}$$

As a consequence,

$$|X(k)| \;=\; \frac{1}{2^k} \sum_{v \in X(0)} |X_{\geqslant v}(k)| \;=\; \binom{t}{k} 2^{t-k} n^k |G|\ .$$

Proof. An i-face is $f = [g; a, b]$ of type S such that $|S| = i$. To specify an element in $\mathfrak{d}(f)$ we specify (i) an index $j \in S$ (i choices) and (ii) a value b_i (2 choices). To specify an element in $\mathfrak{u}(f)$ we specify (i) an index $j \notin S$ ($t - i$ choices) and (ii) a value for a_i (n choices).

More generally, to specify an element of $X_{\geqslant f}(k)$ we specify (i) a subset $S' \subseteq \overline{S}$ of size $|S'| = k - \ell$, and (ii) a value for $a_{S'}$. To specify an element of $X_{\leqslant f}(\ell)$ we specify (i) a subset $S' \subseteq S$ of size $|S'| = i - \ell$ and (ii) an element $b \in \{0,1\}^{S'}$.

Finally for the "consequence", for the first equality note that each $f \in X(k)$ lies in $X_{\geqslant v}(k)$ for exactly 2^k vertices $v \in X(0)$. For the second equality, we have $|X(0)| = 2^t |G|$. □

2.6.2 Markov operators

We recall some basic notation and facts.

Definition 2.45. Let V be a finite set. Let $\mathbb{R}^V$ denote the vector space of real functions on V and $\langle \cdot, \cdot \rangle$ the standard inner product. For $M : \mathbb{R}^V \to \mathbb{R}^V$ linear we say that

1. M is *symmetric* if $\langle M\phi, \psi\rangle = \langle \phi, M\psi\rangle$ for any $\phi, \psi \in \mathbb{R}^V$;

2. M is *Markov* if $M\phi \geqslant 0$ whenever $\phi \geqslant 0$ (entrywise) and $M1_V = 1_V$, with 1_V the constant one function;

3. M is λ-*expanding*, for some $0 < \lambda < 1$, if $\langle M\phi, \phi\rangle \leqslant \lambda\langle \phi, \phi\rangle$ for all $\phi \in \mathbb{R}^V$ such that $\langle \phi, 1_V\rangle = 0$.

We will use the following well-known lemma.

Lemma 2.46. *Let $M : \mathbb{R}^V \to \mathbb{R}^V$ be a symmetric, Markov, λ-expanding linear operator. Then for any $x \in \mathbb{R}^V$,*

$$\langle x, Mx\rangle \leqslant \lambda\|x\|_2^2 + \frac{\|x\|_1^2}{|V|} ,$$

where $\|x\|_1 = \sum_{i\in V} |x_i|$ and $\|x\|_2 = (\sum_{i\in V} x_i^2)^{1/2}$ denote the 1- and 2-norm of x respectively.

Proof. Decompose x as $x = \alpha u + \beta v$, where u, v are orthogonal vectors of norm 1 such that u is colinear to the all-1 vector. Then

$$\begin{aligned}\langle x, Mx\rangle &= \alpha^2 + \beta^2\langle v, Mv\rangle \\ &\leqslant \alpha^2 + \lambda\beta^2 \\ &= \alpha^2(1-\lambda) + \lambda\|x\|_2^2 ,\end{aligned}$$

where the inequality uses that $v \perp u$ and the assumption that M is λ-expanding. Moreover,

$$\begin{aligned}\alpha^2 &= |\langle x, u\rangle|^2 \\ &\leqslant \frac{\|x\|_1^2}{|V|} ,\end{aligned}$$

because $u_i = 1/\sqrt{|V|}$ for all i. Combining both inequalities completes the proof. □

Given a graph $G = (V, E)$, we say that G is λ-expanding if the Markov operator that is given by the normalized adjacency matrix of G is λ-expanding. This notion is only

suitable for graphs that are connected. However, the constant degree complex described in Section 2.3.5 is not connected, so we have to generalize the notion. We say G is λ-expanding up to size $r|V|$ if the graph is a disjoint union of graphs of size $\geqslant r|V|$ where each graph is λ-expanding. In particular, the constant degree complex satisfies $r = \Omega((\log |G|)^{-(t-1)})$. For each set A_j, we let $G_j = \mathrm{Cay}(G, A_j)$ denote the graph whose vertex set is $V_j = G$, and such that there is an edge between (g, g') if and only if $g' = g \cdot a_j$ for some $a_j \in A_j$. Since we assumed that A_j is closed under inverse, $\mathrm{Cay}(G, A_j)$ is an n-regular undirected graph (which for simplicity one may assume does not contain any self-loops, although our arguments extend verbatim if there are).

Definition 2.47. For a set of permutations A on G, we say that the pair $(G; A)$ is λ-*expanding* if the Markov operator $M : \mathbb{R}^G \to \mathbb{R}^G$ such that

$$M(e_g) \;=\; \frac{1}{|A|} \sum_{a \in A} e_{g \cdot a} \tag{2.40}$$

is λ-expanding, where e_g denotes the g-th canonical basis vector of $\mathbb{R}^G$.

More generally, for $0 < r \leqslant 1$ we say that $(G; A)$ is λ-*expanding up to size* $r|G|$ if the graph is a disjoint union of graphs of size $\geqslant r|G|$ where each graph is λ-expanding.

Let M_j be the Markov operator associated with the normalized adjacency matrix of $G_j = \mathrm{Cay}(G, A_j)$ as in (2.40). Hereafter, we assume that G_j is λ-expanding up to size $r|G|$ for all j.

Let $D : \mathbb{R}^X \to \mathbb{R}^X$ be the normalized unsigned boundary operator, i.e.

$$D\phi(f) \;=\; \frac{1}{|\mathfrak{u}\{f\}|} \sum_{f' \succcurlyeq f} \phi(f') \,.$$

Note that $|\mathfrak{u}\{f\}|$ is independent of f (see Lemma 2.44), and so this is a linear operator. Let U be the normalized unsigned coboundary operator, i.e.

$$U\phi(f) \;=\; \frac{1}{|\mathfrak{d}\{f\}|} \sum_{f' \preccurlyeq f} \phi(f') \,.$$

Claim 2.48. *The operators D and U satisfy that for any ℓ and $\phi : X(\ell) \to \mathbb{R}$, $\phi' : X(\ell-1) \to \mathbb{R}$,*

$$\mathop{\mathbb{E}}_{f \in X(\ell-1)} \phi'(f) D\phi(f) = \mathop{\mathbb{E}}_{f' \in X(\ell)} U\phi'(f')\phi(f') \,,$$

where both expectations are uniform.

Proof. To verify this, for any ϕ, ϕ' we have

$$\begin{aligned}
\frac{1}{X(\ell-1)} \sum_{f \in X(\ell-1)} \phi'(f) D\phi(f) &= \frac{1}{X(\ell-1)} \sum_{f \in X(\ell-1)} \frac{1}{\mathfrak{u}\{f\}} \sum_{f' \succ f} \phi'(f)\phi(f') \\
&= \frac{1}{X(\ell)} \sum_{f' \in X(\ell)} \frac{1}{\mathfrak{d}\{f'\}} \sum_{f \prec f'} \phi'(f)\phi(f') \\
&= \frac{1}{X(\ell)} \sum_{f' \in X(\ell)} U\phi'(f')\phi(f') \,.
\end{aligned}$$

Here the second line is because for any $f \prec f'$, $\frac{X(\ell-1)}{X(\ell)} = \frac{\mathfrak{d}\{f'\}}{\mathfrak{u}\{f\}}$. □

2.6.3 Expanding random walks

We define and analyze some random walks on the vertex set $X(k)$.

Definition 2.49. For $0 \leqslant \ell \leqslant k$ define the following walk $W^{(k,\ell)}$ on $X(k)$. Starting at $f \in X(k)$,

1. Choose a uniformly random $v \in X(\ell)$ such that $v \in \mathfrak{d}^{(k-\ell)}\{f\}$. (If $\ell = k$ then set $v = f$.)
2. Write $v = [g; a, b]$ and $S = \mathtt{type}(v)$. Choose a uniformly random $i \in \overline{S}$ and $a_i \in A_i$. Let $g' = ga_i$, $a'_j = a_j$ for $j \in S$, $b'_i = 1 - b_i$, and $b'_j = b_j$ for $j \in \overline{S} - \{i\}$. Let $v' = [g'; a', b']$.
3. Return a uniformly random $f' \in \mathfrak{u}^{(k-\ell)}\{v'\}$.

We write $W^{(k,\ell)} : \mathbb{R}^{X(k)} \to \mathbb{R}^{X(k)}$ for the Markov operator associated with this walk and $W^{(k,\ell)}_{adj} : \mathbb{R}^{X(k)} \to \mathbb{R}^{X(k)}$ for the adjacency matrix associated with this walk.

Definition 2.50. Let $0 \leqslant \ell \leqslant k < t$ and $v \in X(\ell)$. Recall the definition of the set

$$(\mathfrak{d} \circ \mathfrak{u}) X_{\geqslant v}(k) = \left\{ f \in X(k) \;\middle|\; \exists f' \in X_{\geqslant v}(k+1)\,,\ f \prec f' \right\}.$$

In words, these are k-faces that are included in a $(k+1)$-face that contains v. Such faces can either contain v, or intersect but not contain v, or not intersect v. We write the set of k-faces that intersect but do not contain v as

$$\mathrm{Nb}_v(k) = \left\{ f \in X(k) \;\middle|\; \exists f' \in X_{\geqslant v}(k+1)\,,\ f \prec f'\,,\ f \cap v \neq \emptyset, v \not\prec f \right\},$$

and those that do not intersect v as

$$\mathrm{Op}_v(k) = \left\{ f \in X(k) \;\middle|\; \exists f' \in X_{\geqslant v}(k+1)\,,\ f \prec f'\,,\ f \cap v = \emptyset \right\}.$$

From the definition, it is clear that

$$(\mathfrak{d} \circ \mathfrak{u}) X_{\geqslant v}(k) \;=\; X_{\geqslant v}(k) \cup \mathrm{Nb}_v(k) \cup \mathrm{Op}_v(k)\,. \tag{2.41}$$

The following lemma states a useful "covering" property for the neighborhoods $\mathrm{Nb}_v(k)$. First, for any $0 \leqslant \ell < k < t$ we define

$$a_{k,\ell} \;=\; \max_{\substack{v_\ell \in X(\ell),\, v_k, v_k' \in X(k):\\ v_\ell \prec v_k,\, v_\ell \prec v_k'}} \left| \left\{ v_{\ell+1} \in X(\ell+1) :\ v_\ell \prec v_{\ell+1} \preceq v_k \ \text{ and } \ v_k' \in \mathrm{Nb}_{v_{\ell+1}}(k) \right\} \right|. \tag{2.42}$$

We observe the crude bound

$$a_{k,\ell} \;\leqslant\; |X_{\leqslant v_k}(\ell+1)| \;\leqslant\; 2^t \tag{2.43}$$

for all ℓ, k. Here the first inequality is because $v_{\ell+1}$ in (2.42) must be in $X_{\geqslant v_k}(\ell+1)$, and the second inequality is by Lemma 2.44. Furthermore, we observe the following claim (which will not be used in the proof).

Claim 2.51. *If $k = t-1$ then $a_{t-1,\ell} = 1$ for all $\ell < t-1$.*

Proof. Fix $v_\ell \in X(\ell)$ and $v_k, v'_k \in X(k)$ such that $v_\ell \prec v_k$ and $v_\ell \prec v'_k$. We first show that any $v_{\ell+1} \in X(\ell+1)$ such that $v_\ell \prec v_{\ell+1} \preceq v_k$ and $v'_k \in \mathrm{Nb}_{v_{\ell+1}}(k)$ has $\mathtt{type}(v_{\ell+1}) = \mathtt{type}(v_\ell) \cup ([t] \setminus \mathtt{type}(v'_k))$. This is because using the first condition necessarily $\mathtt{type}(v_{\ell+1}) \supseteq \mathtt{type}(v_\ell)$ and has exactly one more element; furthermore, using the second condition that element must be the unique element of $[t]$ that is not in $\mathtt{type}(v'_k)$. Once its type has been fixed, $v_{\ell+1}$ is uniquely determined by the condition $v_\ell \prec v_{\ell+1} \preceq v_k$. □

Lemma 2.52. *For all $v_k \in X(k)$ and $\ell < k$,*

$$\bigsqcup_{v_{\ell+1} \preceq v_k} \mathrm{Nb}_{v_{\ell+1}}(k) \subseteq a_{k,\ell} \cdot \bigsqcup_{v_\ell \prec v_k} X_{\geqslant v_\ell}(k) \,,$$

where $\sqcup$ denotes the union of multisets, i.e. elements are counted with multiplicity, and $a \cdot S$ is all elements of the multiset S with multiplicity multiplied by a.

Proof. Let $v_k \in X(k)$. For $v_{\ell+1} \in X(\ell+1)$, each face $v'_k \in \mathrm{Nb}_{v_{\ell+1}}(k)$ intersects $v_{\ell+1}$ on some ℓ-dimensional face. This is because, since $\mathrm{Nb}_{v_{\ell+1}}(k) \subseteq (\mathfrak{d} \circ \mathfrak{u}) X_{\geqslant v_{\ell+1}}(k)$ by definition, the types of v_k and v'_k must differ by at most one element. Additionally, the intersection could not be an $\ell+1$-dimensional face, otherwise $v_{\ell+1} \preceq v'_k$ which implies $v'_k \notin \mathrm{Nb}_{v_{\ell+1}}(k)$. Therefore,

$$\begin{aligned}
\bigsqcup_{v_{\ell+1} \preceq v_k} \mathrm{Nb}_{v_{\ell+1}}(k) &= \bigsqcup_{v_{\ell+1} \preceq v_k} \bigsqcup_{v'_k \in \mathrm{Nb}_{v_{\ell+1}}(k)} \{v'_k\} \\
&= \bigsqcup_{v_{\ell+1} \preceq v_k} \bigsqcup_{v'_k \in \mathrm{Nb}_{v_{\ell+1}}(k)} \bigsqcup_{v_\ell \prec v_{\ell+1}} 1_{v_{\ell+1} \cap v'_k = v_\ell} \cdot \{v'_k\} \\
&= \bigsqcup_{v_\ell \prec v_k} \bigsqcup_{v'_k \succ v_\ell} \bigsqcup_{v_{\ell+1} : v_\ell \prec v_{\ell+1} \preceq v_k} 1_{v_{\ell+1} \cap v'_k = v_\ell} \cdot 1_{v'_k \in \mathrm{Nb}_{v_{\ell+1}}(k)} \cdot \{v'_k\} \\
&\subseteq \bigsqcup_{v_\ell \prec v_k} \bigsqcup_{v'_k \succ v_\ell} a_{k,\ell} \cdot \{v'_k\} \\
&= a_{k,\ell} \cdot \bigsqcup_{v_\ell \prec v_k} X_{\geqslant v_\ell}(k) \,.
\end{aligned}$$

□

Definition 2.53. For $0 \leqslant \ell \leqslant k$ define the following walk $\mathrm{Op}^{(k,\ell)}$ on $X(k)$. Starting at $f \in X(k)$,

1. Choose a uniformly random $v \in X(\ell)$ such that $v \in \mathfrak{d}^{(k-\ell)}\{f\}$. (If $\ell = k$ then set $v = f$.)

2. Return a uniformly random $f' \in \mathrm{Op}_v(k)$.

We write $\mathrm{Op}^{(k,\ell)} : \mathbb{R}^{X(k)} \to \mathbb{R}^{X(k)}$ for the Markov operator associated with this walk and $\mathrm{Op}^{(k,\ell)}_{adj} : \mathbb{R}^{X(k)} \to \mathbb{R}^{X(k)}$ for the adjacency matrix associated with this walk.

The random walk $\mathrm{Op}^{(k,\ell)}$ is used in the proof of co-systolic distance. For the analysis, it will be convenient to relate it to the walk $W^{(k,\ell)}$. This is done through the following claim.

Claim 2.54. *For any $\mathcal{A} \subseteq X(k)$,*

$$\langle 1_{\mathcal{A}}, \mathrm{Op}^{(k,\ell)}_{adj} 1_{\mathcal{A}} \rangle \leqslant \langle 1_{\mathcal{A}}, W^{(k,\ell)}_{adj} 1_{\mathcal{A}} \rangle = \binom{k}{\ell}\binom{t-\ell}{k-\ell}(t-\ell)2^{k-\ell}n^{k+1-\ell}\langle 1_{\mathcal{A}}, W^{(k,\ell)} 1_{\mathcal{A}} \rangle .$$

Proof. It is not hard to check that $\langle 1_{\mathcal{A}}, \mathrm{Op}^{(k,\ell)}_{adj} 1_{\mathcal{A}} \rangle \leqslant \langle 1_{\mathcal{A}}, W^{(k,\ell)}_{adj} 1_{\mathcal{A}} \rangle$ from the definitions. To show that $W^{(k,\ell)}_{adj} = \binom{k}{\ell}\binom{t-\ell}{k-\ell}(t-\ell)2^{k-\ell}n^{k+1-\ell}W^{(k,\ell)}$, we observe that the normalization factor for $W^{(k,\ell)}$ is, using notations from Definition 2.49,

$$\begin{aligned} |X_{\leqslant f}(\ell)||\overline{S}||A_i||X_{\geqslant v'}(k)| &= \left(\binom{k}{\ell}2^{k-\ell}\right)(t-\ell)(n)\left(\binom{t-\ell}{k-\ell}n^{k-\ell}\right) \\ &= \binom{k}{\ell}\binom{t-\ell}{k-\ell}(t-\ell)2^{k-\ell}n^{k+1-\ell} , \end{aligned}$$

where the first equality uses bounds from Lemma 2.44. □

One would hope that $W^{(k,\ell)}$ has spectral expansion. However, $W^{(k,\ell)}$ shares similarities to the random walk on a hypercube, which is not a good spectral expander. Nevertheless, $W^{(k,\ell)}$ has small-set expansion, as the next lemma shows.

Lemma 2.55. *Assume that for each $j \in \{1, \dots, t\}$ the pair $(G; A_j)$ is λ-expanding. Let $0 \leqslant \ell \leqslant k$ and $\mathcal{A} \subseteq X(k)$. Then*

$$\langle 1_{\mathcal{A}}, W^{(k,\ell)} 1_{\mathcal{A}} \rangle \leqslant \lambda |\mathcal{A}| + \binom{t}{\ell} 2^{t-\ell-1} n^{\ell} \frac{|\mathcal{A}|^2}{r|X(k)|} . \tag{2.44}$$

Note that $\lambda|\mathcal{A}|$ is the dominating term when $\frac{|\mathcal{A}|}{|X(k)|}$ is a small constant.

Proof. Let M be the Markov operator associated to step 2 of the walk $W^{(k,\ell)}$ introduced in Definition 2.49. We first discuss the structural properties of M. Notice that the type of the face does not change after a step of the walk M. Thus, M is a direct sum of Markov operators M_S, for S a subset of size ℓ, each of which corresponds to the restriction of M to faces of type S. M_S can be further decomposed into a sum of operators $M_{S,i}$, for $i \in \overline{S}$, associated to the walk in the i-th direction. This decomposition is not a direct sum, but an average. Finally, $M_{S,i}$ is isomorphic to the direct sum of $n^{\ell} 2^{t-\ell-1}$ disjoint copies of the walk on the double cover of $\mathrm{Cay}(G, A_i)$. This is because a_S and $b_{\overline{S}-i}$ are not changed by $M_{S,i}$; only g and b_i are. Denote the random walk on the double cover of $\mathrm{Cay}(G, A_i)$ as P_i. To summarize,

$$M = \sum_{|S|=\ell} M_S , \qquad M_S = \frac{1}{t-\ell} \sum_{i \in \overline{S}} M_{S,i} , \qquad \text{and} \qquad M_{S,i} = P_i^{\oplus n^{\ell} 2^{t-\ell-1}} .$$

We now prove the lemma. For P_i, by applying the expander mixing lemma, Lemma 2.46, on the double cover of $\mathrm{Cay}(G, A_i)$ we have

$$\langle x, P_i x \rangle \leqslant \lambda \|x\|_2^2 + \frac{\|x\|_1^2}{2r|G|}$$

for all real functions x on the vertices. The factor $2r|G|$ appears in the denominator because each connected component is λ-expanding with size at least $2r|G|$. Because $M_{S,i}$ is a direct sum of P_i, the same inequality carries over:

$$\langle x, M_{S,i} x \rangle \leqslant \lambda \|x\|_2^2 + \frac{\|x\|_1^2}{2r|G|} .$$

Because $M_S = \frac{1}{t-\ell}\sum_{i\in\overline{S}} M_{S,i}$,

$$\langle x, M_S x\rangle \;=\; \frac{1}{t-\ell}\sum_{i\in\overline{S}}\langle x, M_{S,i}x\rangle \;\leqslant\; \lambda\|x\|_2^2 + \frac{\|x\|_1^2}{2r|G|}\;.$$

Finally, since M is a direct sum of M_S, we obtain that for any real function x on the faces $X(\ell)$,

$$\langle x, Mx\rangle \;\leqslant\; \lambda\|x\|_2^2 + \frac{\|x\|_1^2}{2r|G|}\;. \tag{2.45}$$

We can now evaluate

$$\begin{aligned}
\langle 1_{\mathcal{A}}, W^{(k,\ell)}1_{\mathcal{A}}\rangle &= \langle 1_{\mathcal{A}}, U^{\circ(k-\ell)}MD^{\circ(k-\ell)}1_{\mathcal{A}}\rangle\\
&= \frac{|X(k)|}{|X(\ell)|}\langle D^{\circ(k-\ell)}1_{\mathcal{A}}, MD^{\circ(k-\ell)}1_{\mathcal{A}}\rangle\\
&\leqslant \frac{|X(k)|}{|X(\ell)|}\lambda\langle D^{\circ(k-\ell)}1_{\mathcal{A}}, D^{\circ(k-\ell)}1_{\mathcal{A}}\rangle + \frac{|X(k)|}{|X(\ell)|}\frac{\langle 1_{X(\ell)}, D^{\circ(k-\ell)}1_{\mathcal{A}}\rangle^2}{2|G|}\\
&= \frac{|X(k)|}{|X(\ell)|}\lambda\langle D^{\circ(k-\ell)}1_{\mathcal{A}}, D^{\circ(k-\ell)}1_{\mathcal{A}}\rangle + \frac{|X(k)|}{|X(\ell)|}\frac{\langle D^{\circ(k-\ell)}1_{X(k)}, D^{\circ(k-\ell)}1_{\mathcal{A}}\rangle^2}{2r|G|}\\
&\leqslant \lambda\langle 1_{\mathcal{A}}, 1_{\mathcal{A}}\rangle + \frac{|X(\ell)|}{|X(k)|}\frac{\langle 1_{X(k)}, 1_{\mathcal{A}}\rangle^2}{2r|G|}\\
&\leqslant \lambda|\mathcal{A}| + \frac{|X(\ell)|}{|X(k)|}\frac{|\mathcal{A}|^2}{2r|G|}\\
&\leqslant \lambda|\mathcal{A}| + \frac{\binom{t}{\ell}2^{t-\ell}n^\ell|G|}{2r|G|}\frac{|\mathcal{A}|^2}{|X(k)|}\;.
\end{aligned}$$

Here the second line uses Claim 2.48. The third line uses (2.45) and the fact that $D^{\circ(k-\ell)}1_{\mathcal{A}}$ has nonnegative entries, which implies that $\|D^{\circ(k-\ell)}1_{\mathcal{A}}\|_1 = \langle 1_{X(\ell)}, D^{\circ(k-\ell)}1_{\mathcal{A}}\rangle$. The fourth and fifth lines use that D is an averaging operator, and the regularity properties of the complex. The last line uses $|X(\ell)| = \binom{t}{\ell}2^{t-\ell}n^\ell|G|$ by Lemma 2.44. □

2.7 Locally co-minimal distance

The main result of this section is that for any $0 \leqslant k < t$, the co-chain complex $C^*(X,\mathcal{F})$ is a co-systolic expander in dimension k according to Definition 2.3, assuming

the expansion parameter λ (see 2.47) is small enough as a function of t and the robustness parameters introduced in Section 2.5.2.

Quantitatively, we first prove a lower bound on the locally co-minimal distance of $C^k(X)$. This is introduced in Section 2.2.1. For convenience we recall the definition

$$\mathsf{d}_{\mathrm{coloc}}(k) \;=\; \min\Big\{|x| \;:\; x \in \ker \delta_k - \{0\},\, x \text{ is locally co-minimal}\Big\}\,. \tag{2.46}$$

Bounds on the co-cycle expansion parameters follow immediately from a lower bound on $\mathsf{d}_{\mathrm{coloc}}(k)$, as shown in Lemma 2.5.

Proposition 2.56. *Let $0 \leqslant k < t$. Assume that there exists $\kappa_{t-i,k-i} > 0$, for each $0 \leqslant i \leqslant k$, such that for each $S \subseteq \{1,\dots,t\}$ with $|S| = t-i$ the local chain complex $C(L_S)$ is $\kappa_{t-i,k-i}$-robust (at level $k-i$) according to Definition 2.31. Then*

$$\mathsf{d}_{\mathrm{coloc}}(k) \;\geqslant\; \frac{1-\lambda C_1}{C_2} r|X(k)|\,,$$

where

$$C_1 \;=\; \frac{t^2 2^{t^2+3t}}{\prod_{i=0}^{k}\kappa_{t-i,k-i}} \qquad \textit{and} \qquad C_2 \;=\; \frac{t^2 2^{t^2+5t} n^t}{\prod_{i=0}^{k}\kappa_{t-i,k-i}}\,.$$

Remark 2.57. The constants C_1 and C_2 are simplifications of constants C_1' and C_2' which provide a slightly tighter bound, see (2.51) in the proof of Proposition 2.56, where the coefficients $a_{k,\ell}$ are defined in (2.42). To simplify these bounds, we we apply the estimate $a_{k,\ell} \leqslant 2^t$ from (2.43).[10] Then,

$$C_1' = \sum_{\ell=0}^{k} \frac{\binom{k}{\ell}\binom{t-\ell}{k-\ell}(t-\ell)2^{k-\ell}\prod_{i=\ell}^{k-1} a_{k,i}}{\prod_{i=\ell}^{k}\kappa_{t-i,k-i}} \leqslant t\frac{2^t 2^t t 2^t (2^t)^t}{\prod_{i=0}^{k}\kappa_{t-i,k-i}} = C_1\,,$$

and

$$C_2' = \sum_{\ell=0}^{k} \frac{\binom{t}{\ell}2^{t-\ell-1}n^\ell\binom{k}{\ell}\binom{t-\ell}{k-\ell}(t-\ell)2^{k-\ell}\prod_{i=\ell}^{k-1} a_{k,i}}{\prod_{i=\ell}^{k}\kappa_{t-i,k-i}} \leqslant t\frac{2^t 2^t n^t 2^t 2^t t 2^t (2^t)^t}{\prod_{i=0}^{k}\kappa_{t-i,k-i}} = C_2\,.$$

[10] The bound is often much tighter, for example $a_{t-1,\ell} = 1$ for all ℓ as shown in Claim 2.51. Such tighter bounds could be of interest if one cares about the dependence of C_1', C_2' on the dimension t.

Remark 2.58. We explicitly work out the requirements on λ and the $\kappa_{\ell,k}$ that suffice to guarantee a linear (in $X(k)$) co-systolic distance, i.e. such that $1 - \lambda C_1' > 0$, for certain cases of interest. We use the shorthand $\kappa_i = \kappa_{i,i-1}$ for the robustness at the last level of the chain complex. Note that when $k = t - 1$, by Claim 2.51 we have that $a_{k,i} = 1$.

When $t = 2, k = 1$ the sufficient condition is

$$\frac{1}{\kappa_1} + \frac{8}{\kappa_1 \kappa_2} < \frac{1}{\lambda} . \tag{2.47}$$

This is the sufficient condition for obtaining good qLDPC using our construction. When $t = 4, k = 3$ the sufficient condition is

$$\frac{1}{\kappa_1} + \frac{24}{\kappa_1 \kappa_2} + \frac{108}{\kappa_1 \kappa_2 \kappa_3} + \frac{128}{\kappa_1 \kappa_2 \kappa_3 \kappa_4} < \frac{1}{\lambda} . \tag{2.48}$$

This is the sufficient condition for obtaining good (again, up to polylog factors) qLTC using our construction.

Let $k < t$ and $x \in C^k(X)$ be such that $\delta(x) = 0$. For $f \in X(k)$ we say that f is *active* if the projection $x(f)$ of x on V_f is nonzero. Let $\mathcal{A} \subseteq X(k)$ be the set of active faces. At a high level, the proof of Proposition 2.56 amounts to showing that, whenever x is locally co-minimal, the set $\mathcal{A}$ only expands a little under a suitable expanding random walk, which is based on the robustness of the local codes. However, when $\mathcal{A} \neq \emptyset$, the spectral expansion of the graph implies to have such small expansion, the set must be large, leading to a lower bound on $|x|$.

Proof of Proposition 2.56. Let $x \in C^k(X)$ be such that $\delta(x) = 0$ and x is locally co-minimal. We start with a preliminary claim, which shows that the locally co-minimal condition implies the following.

Claim 2.59. *Let $\ell < k$ and $v \in X(\ell)$. Using the isomorphism from Lemma 2.27, the restriction $x(X_{\geqslant v}(k))$ can be seen as a $(k-\ell)$-chain of the dimension $(t-\ell)$ complex $C(L_{\overline{S}})$ where $S = \texttt{type}v$. Let z be that element. Then z is minimal (according to Definition 2.30).*

Proof. It suffices to show that

$$\forall y \text{ supported on } X_{\geqslant v}(k-1)\ , \qquad \Big|x\Big(X_{\geqslant v}(k)\Big) + \delta_{\overline{S}}(y)\Big| \ \geqslant\ \Big|x\Big(X_{\geqslant v}(k)\Big)\Big|\ .$$

Let y be supported on $X_{\geqslant v}(k-1)$. Then $\delta_{\overline{S}}(y) = \delta(y)(X_{\geqslant v}(k))$. Thus $x + \delta(y)$ decomposes as a sum

$$x + \delta(y)\ =\ \Big(x\Big(X_{\geqslant v}(k)\Big) + \delta_{\overline{S}}(y)\Big) + \Big(x - x\Big(X_{\geqslant v}(k)\Big)\Big)\ ,$$

where the two elements on the right-hand side are supported on a disjoint set of faces: $X_{\geqslant v}(k)$ for the first and its complement in $X(k)$ for the second. Thus

$$\begin{aligned}\Big|x\Big(X_{\geqslant v}(k)\Big) + \delta_{\overline{S}}(y)\Big| &= \Big|x + \delta(y)\Big| - \Big|x - x\Big(X_{\geqslant v}(k)\Big)\Big| \\ &\geqslant \Big|x\Big| - \Big|x - x\Big(X_{\geqslant v}(k)\Big)\Big| \\ &= \Big|x\Big(X_{\geqslant v}(k)\Big)\Big|\ ,\end{aligned}$$

where the second line uses local minimality of x (according to Definition 2.4). □

The key step in the proof is the following claim, which uses the robustness assumption on the local chain complex $C(L_S)$.

Claim 2.60. *Let $\ell \leqslant k$. For all $v \in X(\ell)$,*

$$\kappa_{t-\ell,k-\ell} n\ |x(X_{\geqslant v}(k))|\ \leqslant\ |x(\mathrm{Op}_v(k))| + |x(\mathrm{Nb}_v(k))|\ . \tag{2.49}$$

Proof. We apply the robustness condition to $z = x(X_{\geqslant v}(k))$, seen as a $(k-\ell)$-chain of $C(L_{\overline{S}})$ where $S = \mathtt{type}\, v$. By Claim 2.59, z is minimal. Therefore,

$$\kappa_{t-\ell,k-\ell} n\ |x(X_{\geqslant v}(k))|\ \leqslant\ |\delta_{\overline{S}}(x(X_{\geqslant v}(k)))|\ .$$

Because $\delta x = 0$ and $(\mathfrak{d} \circ \mathfrak{u}) X_{\geqslant v}(k) = X_{\geqslant v}(k) \cup \mathrm{Nb}_v(k) \cup \mathrm{Op}_v(k)$ by (2.41), each nonzero entry in $\delta_{\overline{S}}(x(X_{\geqslant v}(k)))$ has to be canceled by some nonzero entries in $x(\mathrm{Op}_v(k))$ or $x(\mathrm{Nb}_v(k))$.

Moreover, each nonzero entry of $x(\mathrm{Op}_v(k))$ or $x(\mathrm{Nb}_v(k))$ cancels at most one nonzero entry of $\delta_{\overline{S}}(x(X_{\geqslant v}(k)))$. Hence,

$$|\delta_{\overline{S}}(x(X_{\geqslant v}(k)))| \leqslant |x(\mathrm{Op}_v(k))| + |x(\mathrm{Nb}_v(k))| \,.$$

Combining the two inequalities above gives (2.49). □

Informally, Claim 2.60 for $\ell = k$ states that whenever a face $v \in X(k)$ is active, i.e. $|x(X_{\geqslant v}(k))| = |x(v)| = 1$, it must have either many neighbors that are active, or many opposite neighbors that are active. In the latter case, we can make a step according to the random walk $\mathrm{Op}^{(k,k)}$ introduced in Defition 2.53, and use the expansion properties of that walk to conclude. However, in the former case a step to $\mathrm{Nb}_v(k)$ would not be expanding (because neighbors share a lower-dimensional face). In that case, we instead first sample a lower-dimensional sub-face of v, and move to an opposite neighbor of the latter face. The following claim analyses this recursive process. Recall the definition of the coefficients $a_{k,\ell}$ in (2.42).

Claim 2.61. *Let $v_k \in X(k)$. Then*

$$|x(v_k)| \leqslant \sum_{\ell=0}^{k} \sum_{v_\ell \preceq v_k} \frac{\prod_{i=\ell}^{k-1} a_{k,i}}{\prod_{i=\ell}^{k}(\kappa_{t-i,k-i}n)} |x(\mathrm{Op}_{v_\ell}(k))| \,. \tag{2.50}$$

Proof. We apply Claim 2.60 recursively, yielding the following sequence of inequalities.

$$\begin{aligned}
|x(v_k)| &\leqslant \frac{|x(\mathrm{Op}_{v_k}(k))|}{\kappa_{t-k,0}n} + \frac{|x(\mathrm{Nb}_{v_k}(k))|}{\kappa_{t-k,0}n} \\
&\leqslant \frac{|x(\mathrm{Op}_{v_k}(k))|}{\kappa_{t-k,0}n} + a_{k,k-1} \sum_{v_{k-1} \prec v_k} \frac{|x(X_{\geqslant v_{k-1}}(k))|}{\kappa_{t-k,0}n} \\
&\leqslant \frac{|x(\mathrm{Op}_{v_k}(k))|}{\kappa_{t-k,0}n} + a_{k,k-1} \sum_{v_{k-1} \prec v_k} \left(\frac{|x(\mathrm{Op}_{v_{k-1}}(k))|}{\kappa_{t-k,0}n \cdot \kappa_{t-k+1,1}n} + \frac{|x(\mathrm{Nb}_{v_{k-1}}(k))|}{\kappa_{t-k,0}n \cdot \kappa_{t-k+1,1}n} \right) \\
&\leqslant \frac{|x(\mathrm{Op}_{v_k}(k))|}{\kappa_{t-k,0}n} + a_{k,k-1} \sum_{v_{k-1} \prec v_k} \frac{|x(\mathrm{Op}_{v_{k-1}}(k))|}{\kappa_{t-k,0}n \cdot \kappa_{t-k+1,1}n}
\end{aligned}$$

$$+ a_{k,k-1}a_{k,k-2} \sum_{v_{k-2} \prec v_k} \frac{|x(X_{\geqslant v_{k-2}}(k))|}{\kappa_{t-k,0} n \cdot \kappa_{t-k+1,1} n}$$

$$\leqslant \cdots$$

$$\leqslant \sum_{\ell=0}^{k} \sum_{v_\ell \prec v_k} \frac{\prod_{i=\ell}^{k-1} a_{k,i}}{\prod_{i=\ell}^{k} (\kappa_{t-i,k-i} n)} |x(\mathrm{Op}_{v_\ell}(k))| \,.$$

Here the second inequality, and every other subsequent step thereafter, uses Lemma 2.52. Note that the recursion terminates because $\mathrm{Nb}_{v_0}(k) = \emptyset$ for all $v_0 \in X(0)$. □

The previous claim draws the consequences of the robustness assumption on the local chains $C(L_S)$, which can be understood as a "local" expansion property. We now conclude the proof by combining this with the global expansion properties of the cubical complex, which are manifest in the expansion properties of the walks $W^{(k,\ell)}$ introduced in Section 2.6.3, see Lemma 2.55. For any $v_k \in X(k)$, applying Claim 2.61 we have

$$\begin{aligned}
\langle 1_{\mathcal{A}}, 1_{v_k} \rangle = |x(v_k)| &\leqslant \sum_{\ell=0}^{k} \sum_{v_\ell \preceq v_k} \frac{\prod_{i=\ell}^{k-1} a_{k,i}}{\prod_{i=\ell}^{k} (\kappa_{t-i,k-i} n)} |x(\mathrm{Op}_{v_\ell}(k))| \\
&= \sum_{\ell=0}^{k} \langle 1_{\mathcal{A}}, \mathrm{Op}_{adj}^{(k,\ell)} 1_{v_k} \rangle \frac{\prod_{i=\ell}^{k-1} a_{k,i}}{\prod_{i=\ell}^{k} (\kappa_{t-i,k-i} n)} \\
&\leqslant \sum_{\ell=0}^{k} \langle 1_{\mathcal{A}}, W^{(k,\ell)} 1_{v_k} \rangle \frac{\binom{k}{\ell}\binom{t-\ell}{k-\ell}(t-\ell) 2^{k-\ell} \prod_{i=\ell}^{k-1} a_{k,i}}{\prod_{i=\ell}^{k} \kappa_{t-i,k-i}} \,,
\end{aligned}$$

where the third line follows from Claim 2.54. By averaging over $v_k \in X(k)$, we obtain

$$\begin{aligned}
\frac{|\mathcal{A}|}{|X(k)|} = \frac{1}{|X(k)|} \langle 1_{\mathcal{A}}, 1_{\mathcal{A}} \rangle &\leqslant \frac{1}{|X(k)|} \sum_{\ell=0}^{k} \langle 1_{\mathcal{A}}, W^{(k,\ell)} 1_{\mathcal{A}} \rangle \frac{\binom{k}{\ell}\binom{t-\ell}{k-\ell}(t-\ell) 2^{k-\ell} \prod_{i=\ell}^{k-1} a_{k,i}}{\prod_{i=\ell}^{k} \kappa_{t-i,k-i}} \\
&\leqslant \frac{|\mathcal{A}|}{|X(k)|} \left(\lambda \sum_{\ell=0}^{k} \frac{\binom{k}{\ell}\binom{t-\ell}{k-\ell}(t-\ell) 2^{k-\ell} \prod_{i=\ell}^{k-1} a_{k,i}}{\prod_{i=\ell}^{k} \kappa_{t-i,k-i}} \right) \\
&\quad + \frac{1}{r} \left(\frac{|\mathcal{A}|}{|X(k)|} \right)^2 \left(\sum_{\ell=0}^{k} \frac{\binom{t}{\ell} 2^{t-\ell-1} n^{\ell} \binom{k}{\ell}\binom{t-\ell}{k-\ell}(t-\ell) 2^{k-\ell} \prod_{i=\ell}^{k-1} a_{k,i}}{\prod_{i=\ell}^{k} \kappa_{t-i,k-i}} \right) \,,
\end{aligned}$$

where the last line follows from Lemma 2.55. Let

$$C_1' = \sum_{\ell=0}^{k} \frac{\binom{k}{\ell}\binom{t-\ell}{k-\ell}(t-\ell) 2^{k-\ell} \prod_{i=\ell}^{k-1} a_{k,i}}{\prod_{i=\ell}^{k} \kappa_{t-i,k-i}} \,,$$

$$C_2' = \sum_{\ell=0}^{k} \frac{\binom{t}{\ell} 2^{t-\ell-1} n^\ell \binom{k}{\ell} \binom{t-\ell}{k-\ell} (t-\ell) 2^{k-\ell} \prod_{i=\ell}^{k-1} a_{k,i}}{\prod_{i=\ell}^{k} \kappa_{t-i,k-i}} \ . \tag{2.51}$$

It follows that either $|\mathcal{A}| = 0$ or $\frac{1-\lambda C_1'}{C_2'} \leqslant \frac{|\mathcal{A}|}{r|X(k)|}$, i.e.

$$|x| \ \geqslant \ \frac{1-\lambda C_1'}{C_2'} r|X(k)| \ .$$

The proof follows using $C_1' \leqslant C_1$ and $C_2' \leqslant C_2$, as discussed in Remark 2.57. □

Remark 2.62. The proof of Proposition 2.56 can be generalized to show that the chain complex has small-set co-boundary expansion, (sequential) co-decoder, and co-decoder with syndrome error by modifying Eq. (2.49). The rest of the proof follows similarly.

For small-set co-boundary expansion, we want to show that for x with small weight, $|\delta x| \geqslant \beta \, dist(x, \operatorname{im}\delta)$ for some β. We can reduce the problem to the case of locally co-minimal x in a way simile to the proof of $\mu_{\text{cosyst}}(i) \geqslant \mathsf{d}_{\text{coloc}}(i)$ in Lemma 2.5. Here we want to show that for locally co-minimal x with small weight, $|\delta x| \geqslant \beta \, |x|$. This follows from the same argument as in the proof of Proposition 2.56. The only modification needed is to replace Eq. (2.49) with

$$\kappa_{t-\ell,k-\ell} n |x(X_{\geqslant v_\ell}(k))| \ \leqslant \ |x(\mathrm{Op}_{v_\ell}(k))| + |x(\mathrm{Nb}_{v_\ell}(k))| + |y(X_{\geqslant v_\ell}(k+1))| \ , \tag{2.52}$$

where $y = \delta x$. The reason is that

$$\kappa_{t-\ell,k-\ell} n |x(X_{\geqslant v_\ell}(k))| \ \leqslant \ |\delta_{\overline{S}}(x(X_{\geqslant v_\ell}(k)))| \ \leqslant \ |x(\mathrm{Op}_{v_\ell}(k))| + |x(\mathrm{Nb}_{v_\ell}(k))| + |y(X_{\geqslant v_\ell}(k+1))| \ .$$

The first inequality follows from the local co-minimality of x as before. The second inequality is modified because now $y = \delta x \neq 0$; therefore, the remaining nonzero entries in $|\delta_{\overline{S}}(x(X_{\geqslant v_\ell}(k)))|$ can now be canceled from $x(\mathrm{Op}_{v_\ell}(k))$, $x(\mathrm{Nb}_{v_\ell}(k))$, or $|y(X_{\geqslant v_\ell}(k+1))|$.

For (sequential) co-decoder, we want to show that for a syndrome $z = \delta x$ obtained from a (qu)bit error x with small weight, one can efficiently find $\tilde{x}$ which is homologous to the actual error, $\tilde{x} - x \in \operatorname{im}\delta$. A natural co-decoder is the generalized small-set flip discussed

in [15]. The main challenge is to show that the generalized small-set flip co-decoder only terminates when $z = 0$ and the key lemma is to show that a flip that reduces the syndrome always exists when the error has small weight. To show such lemma, we suppose the flip does not exist, which implies

$$\kappa_{t-\ell,k-\ell} n |x(X_{\geqslant v_\ell}(k))| \leqslant 2(|x(\mathrm{Op}_{v_\ell}(k))| + |x(\mathrm{Nb}_{v_\ell}(k))|) \ . \tag{2.53}$$

By replacing Eq. (2.49) with the inequality above, the same argument in the proof of Proposition 2.56 implies that $x = 0$.

The inequality (2.53) follows from

$$\kappa_{t-\ell,k-\ell} n |x(X_{\geqslant v_\ell}(k))| \ \leqslant \ \delta_{\overline{S}}(x(X_{\geqslant v_\ell}(k))) \leqslant 2(|x(\mathrm{Op}_{v_\ell}(k))| + |x(\mathrm{Nb}_{v_\ell}(k))|) \ .$$

The first inequality follows from the local co-minimality of x as before. If the second inequality does not hold, then flipping $x(X_{\geqslant v_\ell}(k))$ reduces the weight of $z = \delta x$. In particular, flipping $x(X_{\geqslant v_\ell}(k))$ only affects $z(X_{\geqslant v_\ell}(k+1))$. Additionally, $z(X_{\geqslant v_\ell}(k+1))$ is induced from $\delta_{\overline{S}}(x(X_{\geqslant v_\ell}(k)))$, $x(\mathrm{Op}_{v_\ell}(k))$, and $x(\mathrm{Nb}_{v_\ell}(k))$. Therefore, the current syndrome on $X_{\geqslant v_\ell}(k+1)$ has weight

$$|z(X_{\geqslant v_\ell}(k+1))| \geqslant |\delta_{\overline{S}}(x(X_{\geqslant v_\ell}(k)))| - (|x(\mathrm{Op}_{v_\ell}(k))| + |x(\mathrm{Nb}_{v_\ell}(k))|) \ ,$$

while the new syndrome after the flip of $x(X_{\geqslant v_\ell}(k))$ has weight

$$|z(X_{\geqslant v_\ell}(k+1)) + \delta_{\overline{S}}(x(X_{\geqslant v_\ell}(k)))| \leqslant |x(\mathrm{Op}_{v_\ell}(k))| + |x(\mathrm{Nb}_{v_\ell}(k))| \ .$$

This means that if $|\delta_{\overline{S}}(x(X_{\geqslant v_\ell}(k)))| > 2(|x(\mathrm{Op}_{v_\ell}(k))| + |x(\mathrm{Nb}_{v_\ell}(k))|)$ then $|z(X_{\geqslant v_\ell}(k+1)) + \delta_{\overline{S}}(x(X_{\geqslant v_\ell}(k)))| < |z(X_{\geqslant v_\ell}(k+1))|$, which implies that the flip exists. This concludes the proof of Eq. (2.53). More details of the co-decoder can be found in [15].

For co-decoder with noisy syndrome, we want to show that given the noisy syndrome $z = \delta x + y$ from (qu)bit error x and syndrome error y with small weights, one can efficiently find $\tilde{x}$ which is homologous to the actual (qu)bit error, $\tilde{x} - x \in \operatorname{im} \delta$. We again use the

generalized small-set flip decoder. To guarantees a flip, we apply a similar argument but replace Eq. (2.53) with

$$\kappa_{t-\ell,k-\ell} n |x(X_{\geqslant v_\ell}(k))| \leqslant 2(|x(\mathrm{Op}_{v_\ell}(k))| + |x(\mathrm{Nb}_{v_\ell}(k))| + |y(X_{\geqslant v_\ell}(k+1))|) \tag{2.54}$$

to account for the syndrome error.

2.8 Distance

For integer $1 \leqslant k \leqslant t$ we let $\mu_{\mathrm{syst}}(k)$ be the systolic distance of $C_k(X)$, that is

$$\mu_{\mathrm{syst}}(k) \;=\; \min\left\{|x| \;:\; x \in \ker \partial_k - \mathrm{im}\, \partial_{k+1}\right\}.$$

We also let $\varepsilon_{\mathrm{cyc}}(k)$ be the cycle expansion,

$$\varepsilon_{\mathrm{cyc}}(k) \;=\; \min\left\{\frac{|\partial(x)|}{\min_{y\in\ker\partial_k}|x-y|} \;:\; x \in C_k(X) - \ker\partial_k\right\}.$$

We show a lower bound on the systolic distance, and the cycle expansion, of $C_k(X)$ by reduction to the co-systolic distance at level $k' = t-k$ of a t-dimensional complex $C(\tilde{X})$ which we now define. The complex $C(\tilde{X})$ is defined exactly as $C(X)$, except that the matrices h_i^T and h_i that are used in the definition of the co-boundary and boundary maps for $C(X)$ respectively are replaced by matrices $(h_i^\perp)^T$ and $h_i^\perp$, where $h_i^\perp \in \mathbb{F}_q^{k_i \times n_i}$ is a parity check matrix for the dual code $\mathcal{C}_i^\perp$ (with linearly independent rows).

Proposition 2.63. *For every $0 \leqslant k' \leqslant t-1$, let $\mu_{\mathrm{cosyst}}(k')$ and $\varepsilon_{\mathrm{cocyc}}(k')$ denote the co-systolic distance and co-cycle expansion of $C^{k'}(\tilde{X})$ respectively. Then for any $1 \leqslant k \leqslant t$,*

$$\mu_{\mathrm{syst}}(k) \;\geqslant\; \frac{1}{(2nt)^t} \cdot \mu_{\mathrm{cosyst}}(t-k).$$

Furthermore, for $2 \leqslant k \leqslant t$,

$$\varepsilon_{\mathrm{cyc}}(k) \;\geqslant\; \min\left\{\frac{1}{2(t^2 2^{2t} n^{t+1})^t} \cdot \varepsilon_{\mathrm{cocyc}}(t-k),\; \frac{1}{(2nt)^t} \cdot \frac{\mu_{\mathrm{cosyst}}(t-k+1)}{|X(k)|}\right\}.$$

The proof of Proposition 2.63 is given in Section 2.8.2. Before giving it, we introduce a new sheaf on X, and associated linear maps. The sheaf will play an important role in the proof, facilitating a reduction from the distance back to the case of co-distance.

2.8.1 Preliminaries

We begin by introducing a new sheaf $\mathcal{F}_k$ over X.

Definition 2.64. Let $0 \leqslant k \leqslant t$. We define a sheaf over X, $\mathcal{F}_k$, by letting $\mathcal{F}_k(f) = C_k(X_{\geqslant f})$ for each $f \in X$. If $\dim(f) > k$ we have $\mathcal{F}_k(f) = \{0\}$. We define co-restriction maps in the natural way, by restriction. Namely, for any $f \prec f'$, we let $\text{co-res}_{f,f'} : \mathcal{F}_k(f) \to \mathcal{F}_k(f')$ be defined by

$$\text{co-res}_{f,f'}(x) := x|_{X_{\geqslant f'}(k)} \, .$$

The coboundary map $\Delta_k : C^i(X, \mathcal{F}_k) \to C^{i+1}(X, \mathcal{F}_k)$ is defined as follows. For $y \in C^i(X, \mathcal{F}_k)$ define $\Delta_k y \in C^{i+1}(X, \mathcal{F}_k)$ by setting, for each $f' \in X(i+1)$ and $u \in X_{\geqslant f'}(k)$,

$$\Delta_k y(f')[u] \;=\; \sum_{f \prec\!\!\prec f'} \text{co-res}_{f,f'}(y(f))[u] = \sum_{f \prec\!\!\prec f'} y(f)[u] \, , \tag{2.55}$$

where we have used that $\text{co-res}_{f,f'}(y(f))[u] = y(f)[u]$ for all $f \prec f' \prec u$.

Note that we define the (block) norm for $x \in C^i(X, \mathcal{F}_k)$ as usual by $|x| = \sum_{f \in X(i)} |x(f)|$ and $|x(f)| = 1_{x(f) \neq 0}$.

The coefficient space assigned by $\mathcal{F}_k$ to each face of X is itself a space of k-chains. One possible way to think about $C^i(X, \mathcal{F}_k)$ is that it is similar to the space $C_k(X)$, except that now each face $u \in X(k)$ may have different values, or "opinions", about it that are obtained from different i-faces $f \preceq u$. The map Δ_k measures the amount of inconsistency among the local views. In particular,

Claim 2.65. *Let $z \in C^0(X, \mathcal{F}_k)$ be such that $\Delta_k z = 0$. Then there is $z' \in C_k(X)$ such that $z'|_{X_{\geqslant v}(k)} = z_v$ for all $v \in X(0)$. Furthermore, we can choose such a z' which satisfies $|z'| \leqslant 2^t n^t |z|$.*

Proof. For each $u \in X(k)$ we choose some vertex $v \prec u$ and define $z'(u) := z_v(u)$. We next show that the definition is independent of the choice of v. Let v, v' be vertices such that $v, v' \prec u$. If $e = \{v, v'\} \in X(1)$ then necessarily $z_v(u) = z_{v'}(u)$, since

$$0 = \Delta_k z(e)[u] = z_v(u) + z_{v'}(u) \ .$$

The equality holds true also for any pair $v, v' \prec u$ that are not neighbors, because the graph $(X(0), X(1))$ induced on $\{v : v \prec u\}$ is connected.

The factor $2^t n^t$ originates from the difference in the definition $|z'| = \sum_{u \in X(k)} |z'(u)|$ and $|z| = \sum_{v \in X(0)} |z(v)|$. Indeed,

$$\begin{aligned} |z'| = \sum_{u \in X(k)} |z'(u)| &= \sum_{v \in X(0)} \sum_{u \in X(k) : u \succ v} |z(v)[u]| \\ &\leqslant \sum_{v \in X(0)} |X_{\geqslant v}(k)||z(v)| \\ &\leqslant 2^t n^t \sum_{v \in X(0)} |z(v)| = 2^t n^t |z| \end{aligned}$$

where the last inequality uses that by Lemma 2.44, $X_{\geqslant v}(k) = \binom{t}{k} n^k \leqslant 2^t n^t$. □

The next lemma verifies that the sequence of spaces $C^i(X, \mathcal{F}_k)$ for $i = 0, \ldots, k$, together with the coboundary map Δ_k, forms a co-chain complex.

Lemma 2.66. *The map Δ_k satisfies $\Delta_k \circ \Delta_k = 0$. Moreover, Δ_k is exact on i-chains for $0 < i \leqslant k$, i.e. for any $0 < i \leqslant k$ and $y \in C^i(X, \mathcal{F}_k)$ such that $\Delta_k(y) = 0$ there is a $z \in C^{i-1}(X, \mathcal{F}_k)$ such that $\Delta_k(z) = y$. Furthermore, there exists such a z that satisfies $|z| \leqslant 2^{2t} n^t |y|$.*[11]

[11] In reality, the bound is much tighter $|z| \leqslant |y|$. The proof is however more complicated, so we adopt the loose bound.

Proof. To show the lemma, we first observe that Δ_k is the coboundary map of a local chain complex induced from $X_{\leqslant u}$. In particular,

$$C^i(X, \mathcal{F}_k) = \bigoplus_{f \in X(i)} \bigoplus_{u \in X(k), u \succeq f} V_u = \bigoplus_{u \in X(k)} C^i(X_{\leqslant u}, V_u)$$

where the set $X_{\leqslant u} = \{f \in X, f \preceq u\}$ is a graded incidence poset, consisting of all faces that have an opinion about the k-face u. The map Δ_k then decomposes as the direct sum of maps $\Delta_u : C^i(X_{\leqslant u}, V_u) \to C^{i+1}(X_{\leqslant u}, V_u)$ for each $u \in X(k)$. This means that to check $\Delta_k \circ \Delta_k = 0$, it suffices to check $\Delta_u \circ \Delta_u = 0$ for each $u \in X(k)$.

The chain complex $C^*(X_{\leqslant u})$ is isomorphic to the following "hypercube" construction. Let $S = \texttt{type}(u)$. For $i \in S$, define a 1-dimensional chain complex $C^*(L_i)$ over $\mathbb{F}_q$, where $L_i(0) = \{0, 1\}$, $L_i(1) = \{e_i\}$ and $\delta : C^0(L_i) \to C^1(L_i)$ is defined by $\delta x(e_1) = x(0) + x(1) \in \mathbb{F}_q$. Let $C^*(L_S)$ be obtained by taking the homological product of all $C^*(L_i)$ for $i \in S$. Then $C^*(L_S)$ is an $|S|$-dimensional complex, with a single $|S|$-dimensional face that can be seen as a hypercube, the $(|S| - 1)$-dimensional faces are the $(|S| - 1)$-dimensional faces of the hypercube, etc. Moreover, because $\dim H^0(L_i) = 1$ and $\dim H^1(L_i) = 0$ then by the Künneth formula, $\dim H^0(L_S) = 1$ and $\dim H^j(L_S) = 0$ for all $0 < j \leqslant |S|$. In particular, $C^*(L_S)$ is exact at all $i > 0$. Since each of the maps Δ_u is isomorphic to a copy of $\delta_{\texttt{type}(u)}$, it follows that Δ_k is exact at all levels $i > 0$.

Finally, we have the trivial size bound $|z[u]| \leqslant 2^t |y[u]|$, which follows from $|C^i(L_S)| \leqslant 2^t$ for all i. Therefore,

$$|z| \leqslant \sum_{u \in X(k)} |z[u]| \leqslant \sum_{u \in X(k)} 2^t |y[u]| \leqslant \sum_{f \in X(i)} \sum_{u \in X(k) : u \succeq f} 2^t |y(f)[u]| \leqslant \sum_{f \in X(i)} 2^t |X_{\succcurlyeq f}(k)||y(f)| \leqslant 2^{2t} n^t |y|$$

where the fifth inequality follows from Lemma 2.44, $|X_{\succcurlyeq f}(k)| = \binom{t-i}{k-i} n^{k-i} \leqslant 2^t n^t$. □

We have defined the coboundary maps Δ_k. We now move to define boundary maps ∂_L (where the subscript L stands for "local"). For $S \subseteq \{1, \ldots, t\}$ recall the local chain complex $C(L_S)$ introduced in Section 2.5, and its boundary map ∂_S. Fix $f \in X(i)$

and let $S = \texttt{type}(f)$, so that $|S| = i$ and $|\overline{S}| = t - i$. By Lemma 2.27, the space $\mathcal{F}_k(f) = C_k(X_{\geqslant f})$ is naturally isomorphic to $C_k(L_{\overline{S}})$. Hence there is a naturally defined map $\partial_L : \mathcal{F}_k(f) \to \mathcal{F}_{k-1}(f)$ given by

$$\mathcal{F}_k(f) = C_k(X_{\geqslant f}) \cong C_{k-i}(L_{\overline{S}}) \xrightarrow{\partial_L} C_{k-i-1}(L_{\overline{S}}) \cong C_{k-1}(X_{\geqslant f}) = \mathcal{F}_{k-1}(f) \ . \tag{2.56}$$

This map extends to $C^i(X, \mathcal{F}_k) \xrightarrow{\partial_L} C^i(X, \mathcal{F}_{k-1})$, i-face by i-face.

The next lemma shows that the maps Δ_k and ∂_L commute.

Lemma 2.67. *For any $0 \leqslant i < j \leqslant t$ and $y \in C^i(X, \mathcal{F}_k)$,*

$$\partial_L \circ \Delta_k(y) \ = \ \Delta_{k-1} \circ \partial_L(y) \ \in C^{i+1}(X, \mathcal{F}_{k-1}) \ .$$

In other words, the following diagram is commutative:

$$\begin{array}{ccc} x \in C^{i+1}(X, \mathcal{F}_k) & \xrightarrow{\partial_L} & C^{i+1}(X, \mathcal{F}_{k-1}) \ni z, z' \\ {\scriptstyle \Delta_k}\uparrow & & {\scriptstyle \Delta_{k-1}}\uparrow \\ y \in C^i(X, \mathcal{F}_k) & \xrightarrow{\partial_L} & C^i(X, \mathcal{F}_{k-1}) \ni x' \end{array}$$

Proof. Let $y \in C^i(X, \mathcal{F}_k)$, and define the following chains as in the figure: $x = \Delta_k y \in C^{i+1}(X, \mathcal{F}_k)$, $x' = \partial_L y \in C^i(X, \mathcal{F}_{k-1})$, and let $z = \partial_L x = \partial_L \Delta_k y$ and $z' = \Delta_{k-1} x' = \Delta_{k-1} \partial_L y$ both in $C^{i+1}(X, \mathcal{F}_{k-1})$. The goal is to prove that $z = z'$. For any faces $f \prec v$ of dimension $i+1, k-1$ respectively, we have, (using the definition of ∂_L and (2.55))

$$z(f)[v] = \partial_L x(f)[v] = \sum_{u \succ v} x(f)[u] = \sum_{u \succ v} \Delta_k y(f)[u] = \sum_{u \succ v} \sum_{g \prec f} y(g)[u] \ ,$$

which equals

$$z'(f)[v] = \Delta_{k-1} x'(f)[v] = \sum_{g \prec f} x'(f)[v] = \sum_{g \prec f} \partial_L y(g)[v] = \sum_{g \prec f} \sum_{u \succ v} y(g)[u] \ .$$

□

2.8.2 Proof of Proposition 2.63

Let $1 \leqslant k \leqslant t$ and $x \in C_k(X)$ be such that $\partial(x) = 0$ and $|x| < (2nt)^{-t} \cdot \mu_{\text{cosyst}}(t-k)$. For any vertex $v \in X(0)$, let

$$x_v^{(0)} \;=\; x(X_{\geqslant v})$$

be the local view of x restricted to the k-faces above the vertex v. The set of local views $\left\{x_v^{(0)}\right\}$ can be seen as a cochain $x^{(0)} \in C^0(X, \mathcal{F}_k)$ where $\mathcal{F}_k$ is introduced in Definition 2.64. By construction

$$|x^{(0)}| \leqslant 2^k |x| \leqslant 2^t |x| \;, \tag{2.57}$$

because each value of x on a k-face is copied 2^k times, to the face's 2^k vertices, to form their local views. (Note that here and as always, $|\cdot|$ denotes the block norm taken in the corresponding coefficient space.) Furthermore, from the definition it follows that

$$\Delta_k(x^{(0)}) \;=\; 0 \;, \tag{2.58}$$

where Δ_k is introduced in Definition 2.64.

Using the isomorphism from Lemma 2.27, $x_v^{(0)}$ is naturally seen as an element of $C_k(L_{[t]})$. The condition $\partial(x) = 0$ implies that for all $v \in X(0)$,

$$\partial_{[t]}(x_v^{(0)}) = 0 \;, \tag{2.59}$$

where $\partial_{[t]}$ is the "local" boundary map defined in Section 2.5.

If $k = t$ then the two conditions (2.58) and (2.59) already suffice to conclude the argument. This is done in Claim 2.70 below. If $k < t$ then we reduce to this case by identifying an element $x^{(t-k)}$ that satisfies both conditions and is, in some sense, a "lift" of $x^{(0)}$ to $C^{t-k}(X, \mathcal{F}_t)$. The details follow.

Assume $k < t$. Using exactness for $\partial_{[t]}$ (Lemma 2.28), (2.59) implies that there is a $z^{(0)} = \{z_v^{(0)}\}_v$ where $z_v^{(0)} \in C_{k+1}(L_{[t]})$ for all vertices v, such that

$$\partial_{[t]}(z_v^{(0)}) \;=\; x_v^{(0)} \;. \tag{2.60}$$

Furthermore, if $x_v^{(0)} = 0$ we take $z_v^{(0)} = 0$. Note that $z^{(0)}$ is an element of $C^0(X, \mathcal{F}_{k+1})$, and (2.60) can be rewritten as

$$\partial_L(z^{(0)}) \,=\, x^{(0)}\,, \tag{2.61}$$

where ∂_L is the map introduced around (2.56). We observe the size bound

$$|z^{(0)}| \,=\, \sum_{v\in X(0)} |z_v^{(0)}| \,\leqslant\, \sum_{v\in X(0)} |x_v^{(0)}| \,=\, |x^{(0)}|\,. \tag{2.62}$$

If all $z_v^{(0)}$ are consistent, meaning that whenever $v, v' \prec u$ then $z_v^{(0)}(u) = z_{v'}^{(0)}(u)$, then the $\left\{z_v^{(0)}\right\}$ can be "stitched" together into one $z \in C_{k+1}(X)$ such that for every $v \in X(0)$, $z|_{X_{\geqslant v}} = z_v^{(0)}$. More precisely, by Claim 2.65, if $\Delta_{k+1}(z^{(0)}) = 0$ then there is a $z \in C_{k+1}(X)$ such that for every $v \in X(0)$, $z_{|X_{\geqslant v}(k+1)} = z_v^{(0)}$, and moreover z satisfies $\partial(z) = x$ with $|z| \leqslant 2^t n^t |z^{(0)}|$ which implies the size bound $|z| \leqslant 2^t n^t |z^{(0)}| \leqslant 2^t n^t |x^{(0)}| \leqslant 2^{2t} n^t |x|$ using (2.57) and (2.62).

However, the different $z_v^{(0)}$ may be inconsistent with each other. In such a case we will have to look into the differences between local views along edges of X, and then continue to move into higher and higher dimensions. Towards this we make the following definition.

Definition 2.68. Let $0 \leqslant r \leqslant t - k - 1$. A sequence $z^{(0)}, \ldots, z^{(r)}$, where for each $i \in \{1, \ldots, r\}$, $z^{(i)} \in C^i(X, \mathcal{F}_{k+i+1})$, is said to *explain* $x^{(0)}$ if the following hold. There are $x^{(1)}, \ldots, x^{(r+1)}$ such that $x^{(i)} \in C^i(X, \mathcal{F}_{k+i})$, $x^{(r+1)} = 0$, and for all $i = 0, \ldots, r-1$,

$$\Delta_{k+i+1}(z^{(i)}) = x^{(i+1)}, \qquad x^{(i)} = \partial_L(z^{(i)})\,. \tag{2.63}$$

Furthermore, we require the size bounds

$$|z^{(i)}| \,\leqslant\, |x^{(i)}|, \qquad |x^{(i+1)}| \,\leqslant\, nt|z^{(i)}|\,. \tag{2.64}$$

for all $i = 0, \ldots, r$.

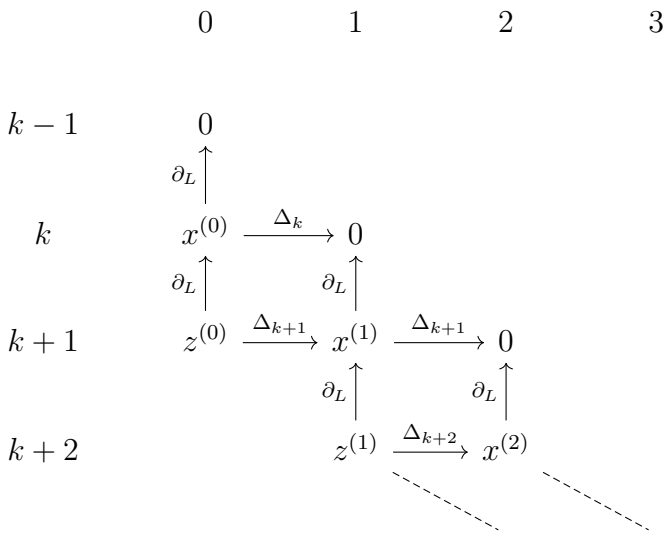

Figure 2.3: An illustration of the requirements on the sequences $x^{(i)}$ and $z^{(i)}$.

Figure 2.3 gives an illustration of the conditions given in Definition 2.68. We first show that having such a sequence lets us find a small chain z whose boundary $\partial(z) = x$.

Claim 2.69. *Suppose that $1 \leqslant k < t$ and there is some $r \leqslant t-k-1$ such that $z^{(0)}, \ldots, z^{(r)}$ explains $x^{(0)}$. Then there is some $z \in C_{k+1}(X)$ such that $\partial(z) = x$ and $|z| \leqslant (t2^{2t}n^{t+1})^t|x|$.*

Proof. The special case where $r = 0$ was handled directly above, relying on Claim 2.65. We reduce the general case to this case. For general $r > 0$, $x^{(r+1)} = \Delta_{k+r+1}(z^{(r)})$ by (2.63), so $x^{(r+1)} = 0$ implies (by exactness, Lemma 2.66) that there is an $u^{(r-1)} \in C^{r-1}(X, \mathcal{F}_{k+r+1})$ such that $\Delta_{k+r+1}(u^{(r-1)}) = z^{(r)}$. Furthermore, we can choose $u^{(r-1)}$ in such a way that $|u^{(r-1)}| \leqslant 2^{2t}n^t|z^{(r)}|$.

Let $w^{(r-1)} = \partial_L(u^{(r-1)})$. Then $w^{(r-1)} \in C^{r-1}(X, \mathcal{F}_{k+r})$, and $|w^{(r-1)}| \leqslant |u^{(r-1)}|$ (because both vectors are labeled by faces in $X(r-1)$). Furthermore,

$$\begin{aligned}\Delta_{k+r}(w^{(r-1)}) &= \Delta_{k+r}(\partial_L(u^{(r-1)})) \\ &= \partial_L(\Delta_{k+r+1}(u^{(r-1)})) \\ &= \partial_L(z^{(r)})\end{aligned}$$

$$= x^{(r)} \ ,$$

where the second equality is by Lemma 2.67 and the last is by (2.63). Now let $\tilde{z}^{(r-1)} = z^{(r-1)} + w^{(r-1)}$. We easily verify that this element satisfies $\Delta_{k+r+1}(\tilde{z}^{(r-1)}) = 0$ and $\partial_L(\tilde{z}^{(r-1)}) = x^{(r-1)}$. This means that $z^{(0)}, \ldots, z^{(r-2)}, \tilde{z}^{(r-1)}$ explains $x^{(0)}$. We went from the assumption that $x^{(r+1)} = 0$ to the same situation, but now $x^{(r)} = 0$, and furthermore

$$|\tilde{z}^{(r-1)}| \leqslant |z^{(r-1)}| + 2^{2t} n^t |z^{(r)}| \leqslant (1 + nt2^{2t} n^t)|z^{(r-1)}| \ , \tag{2.65}$$

where the second inequality uses (2.64). Iterating, we reduce the case of general r to the case where $r = 0$, which as shown above (see Claim 2.65) completes the proof of Proposition 2.63. The weight of $\tilde{z}^{(0)}$ before applying Claim 2.65 satisfies

$$|\tilde{z}^{(0)}| \leqslant \Big(1 + (nt2^{2t} n^t) + ... + (nt2^{2t} n^t)^r\Big)|x^{(0)}| \leqslant (nt2^{2t} n^t)^{r+1}|x^{(0)}| \leqslant (t2^{2t} n^{t+1})^{t-1}|x^{(0)}|$$

where the second inequality follows from $2 \leqslant nt2^{2t} n^t$ and the third inequality follows from $r \leqslant t - k - 1 \leqslant t - 2$. Thus,

$$|z| \leqslant 2^t n^t |\tilde{z}^{(0)}| \leqslant 2^t n^t (t2^{2t} n^{t+1})^{t-1}|x^{(0)}| \leqslant 2^t 2^t n^t (t2^{2t} n^{t+1})^{t-1}|x| \leqslant (t2^{2t} n^{t+1})^t |x| \ .$$

□

Now we show how to construct a sequence that satisfies the conditions of Definition 2.68. We do this inductively. Suppose that $z^{(j)}$ and $x^{(j)}$ have been defined for all $0 \leqslant j < r$, such that they satisfy both conditions in the definition. Suppose first that $r < t - k$. We show how to extend the sequence.

Let $x^{(r)} = \Delta_{k+r}(z^{(r-1)})$. We wish to define $z^{(r)}$. By assumption, $\partial_L(z^{(r-1)}) = x^{(r-1)}$, which implies

$$\begin{aligned} \partial_L(x^{(r)}) &= \partial_L \circ \Delta_{k+r}(z^{(r-1)}) \\ &= \Delta_{k+r-1} \circ \partial_L(z^{(r-1)}) \end{aligned}$$

$$= \Delta_{k+r-1}(x^{(r-1)})$$
$$= 0\,, \tag{2.66}$$

where the first equality is because $x^{(r-1)} = \Delta_{k+r}(z^{(r-2)})$, the second equality is by Lemma 2.67 and the last because $\Delta^2 = 0$ by Lemma 2.66. Therefore, using the exactness condition for the local complex we find $z^{(r)} \in C^r(X, \mathcal{F}_{k+r+1})$ such that $\partial_L(z^{(r)}) = x^{(r)}$, as desired.

Moreover, because $|z^{(r)}|$ and $|x^{(r)}|$ are labeled by faces in $X(r)$ we have $|z^{(r)}| \leqslant |x^{(r)}|$. And because each nonzero face $f \in X(r)$ in $z^{(r)}$ results in at most $|\mathfrak{u}(f)|$ nonzero faces in $x^{(r+1)} = \Delta_k z^{(r)}$, $|x^{(r+1)}| \leqslant |\mathfrak{u}(f)||z^{(r)}| \leqslant nt|z^{(r)}|$ where we apply Lemma 2.44, $|\mathfrak{u}(f)| = (t-i)n \leqslant nt$.

It remains to treat the case where we have defined $x^{(0)}, \ldots, x^{(t-k)}$, but $x^{(t-k)} \neq 0$. This is handled by our last claim.

Claim 2.70. *Suppose that $x^{(t-k)} \neq 0$ and $|x^{(t-k)}| < \mu_{\text{cosyst}}(t-k)$. Then there is some $v \in C_{k+1}(X)$ such that $\partial(v) = x$ and $|v| \leqslant 2(t^2 2^{2t} n^{t+1})^t \frac{1}{\varepsilon_{\text{cocyc}}(t-k-1)}|x|$.*

Proof. If $k = t$, then we have that $\Delta_t(x^{(0)}) = 0$ as shown in (2.58), and $\partial_L(x^{(0)}) = 0$ which follows from (2.59). If $k < t$, then both equalities hold as well, as we now show. Using $x^{(t-k)} = \Delta_t(z^{(t-k-1)})$ and $\Delta_t \circ \Delta_t = 0$ by Lemma 2.66, we get that

$$\Delta_t(x^{(t-k)}) \;=\; 0\,. \tag{2.67}$$

Moreover, using a similar chain of equalities as for (2.66),

$$\begin{aligned}\partial_L(x^{(t-k)}) &= \partial_L \circ \Delta_t(z^{(t-k-1)}) \\ &= \Delta_{t-1} \circ \partial_L(z^{(t-k-1)}) \\ &= \Delta_{t-1}(x^{(t-k-1)}) \\ &= 0\,.\end{aligned}$$

Note that for each $f \in X(t-k)$, $x^{(t-k)}(f) \in \mathbb{F}_q^{\prod_{j \notin \mathtt{type}(f)} A_j}$. So $\partial_L(x^{(t-k)}) = 0$ means that $x^{(t-k)}(f) \in \otimes_{j \notin \mathtt{type}(f)} \mathcal{C}_j$ by Lemma 2.29. In particular, by replacing the tensor codeword $x^{(t-k)}(f)$ by a decoding $\tilde{x}^{(t-k)}(f) \in \mathbb{F}_q^{\prod_{j \notin \mathtt{type}(f)} k_j}$ of it, we obtain a $(t-k)$-cochain $\tilde{x}^{(t-k)} \in C^{t-k}(\tilde{X})$ such that, for any $f \in \tilde{X}(t-k+1)$,

$$\begin{aligned}
\Big(\bigotimes_{j \notin \mathtt{type}(f)} (h_j^\perp)^T \Big) \tilde{\delta}(\tilde{x}^{(t-k)})(f) &= \Big(\bigotimes_{j \notin \mathtt{type}(f)} (h_j^\perp)^T \Big) \Big(\sum_{f' \prec f} \widetilde{\text{co-res}}_{f',f}(\tilde{x}^{(t-k)}(f')) \Big) \\
&= \sum_{f' \prec f} \Big(\Big(\bigotimes_{j \notin \mathtt{type}(f')} (h_j^\perp)^T \Big) \tilde{x}^{(t-k)}(f') \Big) \\
&= \sum_{f' \prec f} x^{(t-k)}(f') \\
&= \Delta_t x^{(t-k)}(f) \\
&= 0 \, .
\end{aligned}$$

Here, the first equality is by definition of $\tilde{\delta}$, where $\widetilde{\text{co-res}}$ is defined as co-res in (2.10) but with the map h_i^T replaced by $(h_i^\perp)^T$. The second equality is by definition of $\widetilde{\text{co-res}}$. The third equality is by the definition of $(h_i^\perp)^T$, which re-encodes $\tilde{x}^{(t-k)}$ in $x^{(t-k)}$. The fourth equality is by the definition of Δ_t in (2.55). The last equality is by (2.67). Finally, because $(h_i^\perp)^T$ are injective, this implies $\tilde{\delta}(\tilde{x}^{(t-k)})(f) = 0$ for all f which means $\tilde{\delta}(\tilde{x}^{(t-k)}) = 0$.

Because $\tilde{\delta}(\tilde{x}^{(t-k)}) = 0$, as long as $|\tilde{x}^{(t-k)}| < \mu_{\text{cosyst}}(t-k)$, $\tilde{x}^{(t-k)} \in \text{im } \tilde{\delta}$, and we can find an $\tilde{u}^{(t-k-1)} \in C^{t-k-1}(\tilde{X})$ such that $\tilde{\delta}(\tilde{u}^{(t-k-1)}) = \tilde{x}^{(t-k)}$ and furthermore

$$|\tilde{u}^{(t-k-1)}| \leqslant \frac{1}{\varepsilon_{\text{cocyc}}(t-k-1)} |\tilde{x}^{(t-k)}| = \frac{1}{\varepsilon_{\text{cocyc}}(t-k-1)} |x^{(t-k)}|$$

where the last equality follows from the construction of $\tilde{x}^{(t-k)}$, which is the decoding of $x^{(t-k)}$. Note that $|\tilde{x}^{(t-k)}| < \mu_{\text{cosyst}}(t-k)$ is satisfied when $|x| < (2nt)^{-t} \mu_{\text{cosyst}}(t-k)$, because $|\tilde{x}^{(t-k)}| = |x^{(t-k)}| \leqslant (nt)^{t-k} |x^{(0)}| \leqslant (nt)^{t-k} 2^t |x|$ where the last inequality follows from (2.57).

By re-encoding $\tilde{u}$, i.e. applying $\otimes_{j \notin \mathtt{type}(f)} \tilde{h}_j^T$ to each $\tilde{u}^{(t-k-1)}(f)$, we obtain an element $u^{(t-k-1)} \in C^{t-k-1}(X, \mathcal{F}_t)$ such that $\Delta_t(u^{(t-k-1)}) = x^{(t-k)}$ (because $\tilde{\delta}(\tilde{u}^{(t-k-1)}) = \tilde{x}^{(t-k)}$) and

$\partial_L(u^{(t-k-1)}) = 0$ (because $u^{(t-k-1)}$ is obtained by applying $\otimes_{j\notin \mathtt{type}(f)}\tilde{h}_j^T$). From there we can reduce the problem to Claim 2.69, by setting $\tilde{z}^{(t-k-1)} = z^{(t-k-1)} + u^{(t-k-1)}$ and observing that it satisfies the conditions of the claim, where $\tilde{x}^{(t-k)} = \Delta_t(\tilde{z}^{(t-k-1)}) = 0$ and $\partial_L(\tilde{z}^{(t-k-1)}) = x^{(t-k-1)}$.

Similar to Claim 2.69, for $r = t-k-1$ we have

$$\begin{aligned}
|\tilde{z}^{(r-1)}| &\leqslant |z^{(r-1)}| + 2^{2t}n^t|\tilde{z}^{(r)}| \\
&\leqslant |z^{(r-1)}| + 2^{2t}n^t(|z^{(r)}| + |u^{(r)}|) \\
&\leqslant (1 + nt2^{2t}n^t)|z^{(r-1)}| + \frac{2^{2t}n^t}{\varepsilon_{\mathrm{cocyc}}(t-k)(r)}|x^{(r+1)}| \\
&\leqslant (1 + nt2^{2t}n^t)|z^{(r-1)}| + \frac{2^{2t}n^t}{\varepsilon_{\mathrm{cocyc}}(r)}(nt)^t 2^t|x| \ ,
\end{aligned}$$

and

$$|\tilde{z}^{(0)}| \leqslant (t2^{2t}n^{t+1})^{t-1}|x^{(0)}| + \frac{(t2^{2t}n^{t+1})^{t-2}}{\varepsilon_{\mathrm{cocyc}}(r)}(nt)^t 2^t|x| \ .$$

Finally,

$$\begin{aligned}
|z| \leqslant 2^t n^t|\tilde{z}^{(0)}| \leqslant (t2^{2t}n^{t+1})^t|x| + \frac{(t2^{2t}n^{t+1})^{t-2}}{\varepsilon_{\mathrm{cocyc}}(r)}t^t n^{2t}2^{2t}|x| &\leqslant (t^2 2^{2t}n^{t+1})^t\left(1 + \frac{1}{\varepsilon_{\mathrm{cocyc}}(r)}\right)|x| \\
&\leqslant \frac{2(t^2 2^{2t}n^{t+1})^t}{\varepsilon_{\mathrm{cocyc}}(r)}|x| \ .
\end{aligned}$$

□

We now conclude. First, we show the lower bound on $\mu_{\mathrm{syst}}(k)$. It suffices to show that whenever $x \in C_k(X)$ satisfies $\partial(x) = 0$ and $|x| \leqslant (2nt)^{-t}\mu_{\mathrm{cosyst}}(t-k)$, there is $z \in C_{k+1}(X)$ such that $\partial(z) = x$. If $k = t$ this follows immediately from Claim 2.70 and the bound (2.57). If $k < t$, we construct a sequence $z^{(0)}, \ldots, z^{(r)}$ that explains $x^{(0)}$. If such a sequence is found with $r \leqslant t-k-1$ then Claim 2.69 gives the desired conclusion. If not, then our assumption on $|x|$ together with (2.64) shows that the assumption of Claim 2.70 is satisfied, allowing us to conclude the argument.

Finally, we show the "Furthermore" part of Proposition 2.63. Let $x \in C_k(X)$ be the vector with the smallest cycle expansion where $x' := \partial(x) \neq 0$,

$$|x| \;=\; \min_{y \in \ker \partial_k} |x - y| \tag{2.68}$$

and $|\partial(x)| = \varepsilon_{\mathrm{cyc}}(k)|x|$. If $|x'| \geqslant (2nt)^{-t}\mu_{\mathrm{cosyst}}(t-(k-1))$ then because $|x| \leqslant |X(k)|$ we deduce $\varepsilon_{\mathrm{cyc}}(k) \geqslant (2nt)^{-t}\mu_{\mathrm{cosyst}}(t-(k-1))/|X(k)|$. If $|x'| < (2nt)^{-t}\mu_{\mathrm{cosyst}}(t-(k-1))$, then following the arguments above, applied to x', which satisfies $\partial(x') = 0$, we find a $z' \in C_k(X)$ such that $\partial(z) = x'$ and $|z'| \leqslant 2(t^2 2^{2t} n^{t+1})^t \frac{1}{\varepsilon_{\mathrm{cocyc}}(t-k)}|x'|$. By (2.68), $|z'| \geqslant |x|$, which concludes the proof.

Dissertation Note

Chapter 2 is based on joint work with Irit Dinur and Thomas Vidick, published as "Expansion of high-dimensional cubical complexes with application to quantum locally testable codes," *Proc. IEEE FOCS*, 2024.

Bibliography

[1] Irit Dinur, Shai Evra, Ron Livne, Alexander Lubotzky, and Shahar Mozes. Good locally testable codes. *arXiv preprint arXiv:2207.11929*, 2022.

[2] Pavel Panteleev and Gleb Kalachev. Asymptotically good quantum and locally testable classical LDPC codes. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 375–388, 2022.

[3] Fernando Granha Jeronimo, Tushant Mittal, Ryan O'Donnell, Pedro Paredes, and Madhur Tulsiani. Explicit abelian lifts and quantum LDPC codes. In Mark Braverman, editor, *13th Innovations in Theoretical Computer Science Conference, ITCS 2022, January 31 - February 3, 2022, Berkeley, CA, USA*, volume 215 of *LIPIcs*, pages 88:1–88:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.

[4] Pavel Panteleev and Gleb Kalachev. Personal communication, 2024.

[5] Adam Wills, Ting-Chun Lin, and Min-Hsiu Hsieh. Tradeoff constructions for quantum locally testable codes. *arXiv preprint arXiv:2309.05541*, 2023.

[6] Matthew B Hastings. Quantum codes from high-dimensional manifolds. In *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.

[7] Anthony Leverrier, Vivien Londe, and Gilles Zémor. Towards local testability for quantum coding. *Quantum*, 6:661, 2022.

[8] Michael Sipser and Daniel A. Spielman. Expander codes. *IEEE Trans. Inform. Theory*, 42(6, part 1):1710–1722, 1996. Codes and complexity.

[9] Pavel Panteleev and Gleb Kalachev. Quantum LDPC codes with almost linear minimum distance. *IEEE Transactions on Information Theory*, pages 1–1, 2021.

[10] Uriya A First and Tali Kaufman. On good 2-query locally testable codes from sheaves on high dimensional expanders. *arXiv preprint arXiv:2208.01778*, 2022.

[11] Tali Kaufman, David Kazhdan, and Alexander Lubotzky. Ramanujan complexes and bounded degree topological expanders. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 484–493, 2014.

[12] Shai Evra and Tali Kaufman. Bounded degree cosystolic expanders of every dimension. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 36–48, 2016.

[13] Gleb Kalachev. High-dimensional expansion of product codes is stronger than robust and agreement testability, 2023.

[14] Gleb Kalachev and Pavel Panteleev. Two-sided robustly testable codes. *arXiv preprint arXiv:2206.09973*, 2022.

[15] Irit Dinur, Min-Hsiu Hsieh, Ting-Chun Lin, and Thomas Vidick. Good quantum LDPC codes with linear time decoders. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 905–918, 2023.

[16] G. Kalachev and P. Panteleev. Maximally extendable product codes are good coboundary expanders, 2024.

[17] Nathan Linial and Roy Meshulam. Homological connectivity of random 2-complexes. *Combinatorica*, 26(4):475–487, 2006.

[18] Mikhail Gromov. Singularities, expanders and topology of maps. part 2: from combinatorics to topology via algebraic isoperimetry. *Geometric and Functional Analysis*, 20(2):416–526, 2010.

[19] Daniel Gottesman. Fault-tolerant quantum computation with constant overhead. *Quantum Information & Computation*, 14(15-16):1338–1372, 2014.

[20] Anurag Anshu, Nikolas Breuckmann, and Chinmay Nirkhe. NLTS hamiltonians from good quantum codes. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1090–1096, 2023.

[21] Shai Evra, Tali Kaufman, and Gilles Zémor. Decodable quantum LDPC codes beyond the n distance barrier using high-dimensional expanders. *SIAM Journal on Computing*, (0):FOCS20–276, 2022.

[22] Matthew B Hastings, Jeongwan Haah, and Ryan O'Donnell. Fiber bundle codes: breaking the $n^{1/2}\,\mathrm{poly}\log(n)$ barrier for quantum LDPC codes. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1276–1288, 2021.

[23] Nikolas P. Breuckmann and Jens N. Eberhardt. Balanced product quantum codes. *IEEE Transactions on Information Theory*, 67(10):6653–6674, 2021.

[24] Ting-Chun Lin and Min-Hsiu Hsieh. c^3-locally testable codes from lossless expanders. In *IEEE International Symposium on Information Theory, ISIT 2022, Espoo, Finland, June 26 - July 1, 2022*, pages 1175–1180. IEEE, 2022.

[25] Anthony Leverrier and Gilles Zémor. Quantum tanner codes. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 872–883. IEEE, 2022.

[26] Dorit Aharonov and Lior Eldar. Quantum locally testable codes. *SIAM Journal on Computing*, 44(5):1230–1262, 2015.

[27] Lior Eldar and Aram W Harrow. Local Hamiltonians whose ground states are hard to approximate. In *2017 IEEE 58th annual symposium on foundations of computer science (FOCS)*, pages 427–438. IEEE, 2017.

[28] Dorit Aharonov, Itai Arad, and Thomas Vidick. Guest column: the quantum PCP conjecture. *Acm sigact news*, 44(2):47–79, 2013.

[29] Adam Wills, Ting-Chun Lin, and Min-Hsiu Hsieh. General distance balancing for quantum locally testable codes. *arXiv preprint arXiv:2305.00689*, 2023.

[30] Shouzhen Gu, Eugene Tang, Libor Caha, Shin Ho Choe, Zhiyang He, and Aleksander Kubica. Single-shot decoding of good quantum LDPC codes. *arXiv preprint arXiv:2306.12470*, 2023.

[31] Tali Kaufman, David Kazhdan, and Alexander Lubotzky. Isoperimetric inequalities for ramanujan complexes and topological expanders. *Geometric and Functional Analysis*, 26(1):250–287, 2016.

[32] Michael Freedman and Matthew Hastings. Building manifolds from quantum codes. *Geometric and Functional Analysis*, 31(4):855–894, 2021.

[33] Matthew B Hastings. Quantum codes from high-dimensional manifolds. *arXiv preprint arXiv:1608.05089*, 2016.

[34] Elia Portnoy. Local quantum codes from subdivided manifolds. *arXiv preprint arXiv:2303.06755*, 2023.

[35] Alexander Lubotzky. High dimensional expanders. In *Proceedings of the International Congress of Mathematicians: Rio de Janeiro 2018*, pages 705–730. World Scientific, 2018.

[36] Bruce Jordan and Ron Livne. The Ramanujan property for regular cubical complexes. *Duke Mathematical Journal*, 105:85–103, 1999.

[37] A Robert Calderbank and Peter W Shor. Good quantum error-correcting codes exist. *Physical Review A*, 54(2):1098, 1996.

[38] Andrew M Steane. Error correcting codes in quantum theory. *Physical Review Letters*, 77(5):793, 1996.

[39] Noga Alon and Yuval Roichman. Random Cayley graphs and expanders. *Random Structures & Algorithms*, 5(2):271–284, 1994.

[40] N. Alon, O. Goldreich, J. Hrastad, and R. Peralta. Simple constructions of almost k-wise independent random variables. *Random Structures and Algorithms*, 3, 1992. See also Addendum in *Random Structures and Algorithms* 4 pages ?? 1993.

[41] Ting-Chun Lin. Transversal non-clifford gates for quantum ldpc codes on sheaves. *arXiv preprint arXiv:2410.14631*, 2024.

[42] Sergey Bravyi and Matthew B Hastings. Homological product codes. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 273–282, 2014.

[43] Tali Kaufman and Ran J Tessler. New cosystolic expanders from tensors imply explicit quantum ldpc codes with $\Omega(\sqrt{n}\log^k n)$ distance. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1317–1329, 2021.

Chapter 3

Transversal Non-Clifford Gates for Quantum LDPC Codes on Sheaves

Chapter 3 is sole-authored work (arXiv:2410.14631), with a companion paper with Louis Golowich (arXiv:2410.14662).

3.1 Introduction

A major goal in quantum computing is to build a fault-tolerant quantum computer. Without fault tolerance, running quantum algorithms (quantum circuits) directly leads to error propagation, and the resulting output becomes unreliable. In contrast, a fault-tolerant version can achieve arbitrarily low error, provided that every gate has error below a certain threshold. It has been proven that any quantum algorithm can be implemented in a fault-tolerant manner. This result is known as the threshold theorem and was first demonstrated using concatenation of quantum codes [1, 2, 3].

Although fault-tolerant circuits have reliable output, they often require more qubits and longer runtime. These costs, known as the space and time overheads, are not only of

theoretical importance but also poses practical challenges. As a result, researchers have been exploring various fault-tolerant schemes each with different trade-offs and advantages. Some notable techniques include code concatenation [1, 2, 3], magic state distillation [4], gate teleportation [5], gauge fixing or code switching [6, 7], transversal gates [8], and quantum low-density parity-check (qLDPC) codes [9]. See [10] for a review.

One of those fault-tolerant schemes relies on qLDPC codes with transversal gates. In this approach, the quantum state in the original circuit $|\psi\rangle$ is encoded into a new state $\mathrm{Enc}(|\psi\rangle)$ using a qLDPC code. Additionally, gate operations from the original circuit U are implemented as transversal operations $\mathrm{Enc}(U) = \bigotimes_i U_i$ where U_i acts on the i-th qubit, such that $\mathrm{Enc}(U|\psi\rangle) = \mathrm{Enc}(U)\,\mathrm{Enc}(|\psi\rangle)$. The advantage of transversal operations is that they do not spread errors. The advantage of qLDPC codes is that the syndromes can be extracted rather reliably in constant depth. Therefore, every layer of gates in the original circuit is transformed into two layers where the first layer applies $\mathrm{Enc}(U)$ on $\mathrm{Enc}(|\psi\rangle)$, and the second layer runs the decoder to correct the error. In many cases, this process keeps the errors under control which leads to a fault tolerant quantum circuit. (In actuality, this scheme often requires code switching to achieve a universal gate set, but we will not discuss that aspect here.)

The challenge in implementing this idea lies in constructing qLDPC codes that support transversal gates. While it is known that any quantum CSS code supports transversal CNOT gates, additional gates are required to obtain a universal gate set. Some examples include: 2D color codes which support transversal H gates, 3D color codes which support transversal T gates [11] and 3D toric codes which support transversal CCZ gates [12]. More recently, there have been efforts to construct codes from manifolds [13] and codes formed by gluing color codes [14, 15]. However, for these constructions, a key code parameter k_{CCZ}, which quantifies the number of logical CCZ gates that can be applied, is currently unknown. This highlights the need for further construction of qLDPC codes that

support non-Clifford gates.

In this work, we provide a large family of qLDPC codes that supports transversal CCZ gates. These are known as sheaf codes [10, 16], and they encompass all recent constructions of qLDPC codes [17, 25, 18, 19, 20]. In a companion paper [21], we present an instance of this family that achieves $[[n, k \geqslant n^{1-\varepsilon}, k_{\mathrm{CCZ}} \geqslant n^{1-\varepsilon}, d \geqslant n^{1/3}/\operatorname{poly}\log n]]$ over qubits with $\operatorname{poly}\log n$ locality for arbitrarily small $\varepsilon > 0$. In the context of magic state distillation, this means that, given n noisy CCZ resource states, the code can output k_{CCZ} more reliable CCZ states. This leads to a magic state distillation protocol with $\gamma = \log(n/k_{\mathrm{CCZ}})/\log(d) \to 0$. In comparison, recent magic state distillation protocols with $\gamma \to 0$ [22, 23, 24] rely on codes with stabilizers that have linear weight.

Remark 3.1. We note an upcoming paper with overlapping results [25]. We will discuss the similarities and differences in the next version.

3.1.1 Main result

To explain the result, we first briefly review the notion of sheaves over cellular complexes. Given a cellular complex X of dimension t, the traditional cochain complex $\mathcal{C}^\bullet(X, \mathbb{F}_q)$, is defined by functions (also known as cochains) which assign a value in $\mathbb{F}_q$ to each cell of the complex. In contrast, a cochain complex based on a sheaf, assigns a value in a vector space $\mathcal{F}_\sigma$ to each cell σ, where the vector space $\mathcal{F}_\sigma$ depends on the cell. The vector spaces $\mathcal{F}_\sigma$ are specified by $\mathcal{F}_\tau$ for the $(t-1)$-cells $\tau \in X(t-1)$. This additional structure $\{\mathcal{F}_\tau\}_{\tau \in X(t-1)}$ is known as the local codes in Tanner code and plays a crucial role in constructing good qLDPC codes.

Given the sheaf $\mathcal{F}$, we can define a cochain complex $\mathcal{C}^\bullet(X, \mathcal{F})$, analogous to the standard cochain complex. The corresponding quantum code is then obtained by choosing a particular dimension ℓ where the X checks are placed at the $(\ell-1)$-cochains, the qubits

are placed at the ℓ-cochains, and the Z checks are placed at the $(\ell+1)$-cochains.

In algebraic topology, it is well-known that the cup product can be defined on the cochains $\mathcal{C}^\bullet(X, \mathbb{F}_q)$. One of our main results is to extend the cup product to cochain complexes based on sheaves $\mathcal{C}(X, \mathcal{F})$. The cup product is defined on the cochains

$$\mathcal{C}^i(X, \mathcal{F}_1) \times \mathcal{C}^j(X, \mathcal{F}_2) \xrightarrow{\cup} \mathcal{C}^{i+j}(X, \mathcal{F}_1 \odot \mathcal{F}_2) \tag{3.1}$$

and induces a corresponding cup product on cohomology classes

$$H^i(X, \mathcal{F}_1) \times H^j(X, \mathcal{F}_2) \xrightarrow{\cup} H^{i+j}(X, \mathcal{F}_1 \odot \mathcal{F}_2). \tag{3.2}$$

For each $(t-1)$-cell $\tau \in X(t-1)$, $(\mathcal{F}_1 \odot \mathcal{F}_2)_\tau \coloneqq \{c_1 \odot c_2 : c_1 \in (\mathcal{F}_1)_\tau, c_2 \in (\mathcal{F}_2)_\tau\}$ where $c_1 \odot c_2$ denotes the entrywise product of c_1 and c_2.

Using the cup product, we derive a sufficient condition for constructing qLDPC codes that support transversal CCZ gates. In particular, if three sheaves $\mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3$ satisfy the condition that for every set of vectors $c_1 \in (\mathcal{F}_1)_\tau, c_2 \in (\mathcal{F}_2)_\tau, c_3 \in (\mathcal{F}_3)_\tau$, we have $\sum_{i=1}^{\Delta} (c_1)_i (c_2)_i (c_3)_i = 0$, then this induces a transversal CCZ operation on the three quantum codes $\mathcal{C}(X, \mathcal{F}_1), \mathcal{C}(X, \mathcal{F}_2), \mathcal{C}(X, \mathcal{F}_3)$.

This general framework enables us to explore a wide family of codes, including the new qLDPC code with $\gamma \to 0$ studied in [21]. Additionally, it clarifies the sufficient conditions to realizing transversal non-Clifford gates, which paves the way for future construction of qLDPC codes with improved parameters.

In addition to the result on the cup product we also study various aspects of sheaves over cellular complexes. Specifically, we propose an alternative definition of sheaves that differs from [26] and prove an analogous statement of the Poincaré duality in the context of cellular complexes.

3.1.2 Proof overview

Our construction is inspired by various observations and previous works. The first observation is that the triorthogonal code introduced in [27] heavily relies on the trilinear function $f(\alpha_1, \alpha_2, \alpha_3) = \sum_{i=1}^{n} (\alpha_1)_i (\alpha_2)_i (\alpha_3)_i$, where $\alpha_1, \alpha_2, \alpha_3$ are vectors in $\mathbb{F}_2^n$. This suggests that a degree 3 function is necessary to access the non-Clifford gates at the third level of the Clifford hierarchy. This also raises the possibility that other degree 3 functions, such as more general trilinear functions, may be sufficient to capture the non-Clifford features. This idea is formalized into the CCZ formalism discussed in Section 3.3. In particular, all we need is a trilinear function f defined on the cohomology classes, which we mean by the condition

$$f(\zeta_1, \zeta_2, \zeta_3) = f(\zeta_1 + \beta_1, \zeta_2 + \beta_2, \zeta_3 + \beta_3) \tag{3.3}$$

for any ζ_i in the cocycle (corresponding to the logical operators) and any β_i in the coboundary (corresponding to the stabilizers). This modification provides more flexibility to the construction of codes with transversal non-Clifford gates.

The second observation is that many known codes with non-Clifford operations can be linked to the triple intersection number of their corresponding topological spaces. For example, the $[[8, 3, 2]]$ code [28] and 3D toric codes support transversal CCZ gates [12]. This is related to the fact that the 3D torus has three homology clases $H_2(\mathbb{T}^3, \mathbb{Z}/2\mathbb{Z}) = (\mathbb{Z}/2\mathbb{Z})^3$ with a nontrivial triple intersection. Furthermore, the topological nature of the triple intersection number implies that it is defined on the homology classes. This observation suggests an approach to study quantum code with a topological nature and define f as the triple intersection. This strategy has also been observed recently and applied to 3D manifolds in [13].

This brings us to the final observation that the good qLDPC codes in the recent breakthrough [17, 25, 18] do have a topological interpretation, including their higher

dimensional analogue [20]. This topological perspective is inspired by a series of work on geometrically local codes [34, 29, 30, 31] which implicitly uses the topological nature of the qLDPC codes. For qLDPC codes constructed from 3D cubical complexes, the X logical operators can be interpreted as surface configurations on the cubical complex. This geometric interpretation allows us to define the triple intersection number, which leads to the transversal CCZ gates on these codes. Together, these observations provide the foundation for our main result regarding the cup product on chain complexes induced from sheaves.

3.1.3 Further directions

Quantum LDPC codes with transversal non-Clifford gates and nearly linear distance. An immediate question is to improve the code parameters obtained in the companion work [21]. It would be desireable to obtain qLDPC codes with linear distance (or up to polylogs). Achieving such a result should also lead to quantum locally testable codes (qLTCs) with transversal non-Clifford gates with inverse polylog soundness, by extending the arguments from [20].

Fully-fledged fault-tolerant scheme. With the improved qLDPC codes and qLTCs, the next goal is to construct a complete fault-tolerant quantum computing scheme based on such codes. A key question here is to develop an efficient code-switching protocols between these codes. This probably relies on the relation between the cellular complexes these codes are based on. In particular, the structural property that a cubical complex contains lower-dimensional cubical complexes, may facilitate code-switching.

Quantum PCP conjecture. The fault-tolerant scheme based on quantum LDPC codes with transversal gates may also contribute to progress on the quantum PCP conjecture [28].

It has been pointed in [32] that an adversarially robust fault-tolerant scheme can be used to construct local Hamiltonians for which it is QMA-hard to decide whether the energy density is $< a$ or $> a + 1/\operatorname{poly}\log n$. If the scheme using qLDPC codes with transversal gates proves to be adversarially robust, it would be a major step toward addressing the qPCP conjecture.

3.2 Preliminary

3.2.1 Chain complex

We review some definitions for chain complexes, which are helpful for explaining the technical parts of the paper. In our context, all vector spaces are equipped with a chosen basis. Additionally, we place significant emphasis on the LDPC condition, which is essentially about sparsity of matrices. A matrix is sparse if each row and column contains at most a constant number of nonzero entries. Note that the definition requires a chosen basis on the vector spaces.

Definition 3.2 (Chain complex)**.** A chain complex $\mathcal{C}$ is collection of vector spaces over $\mathbb{F}_q$, $\mathcal{C}^i$, and linear maps $\delta^i : \mathcal{C}^i \to \mathcal{C}^{i+1}$, such that $\delta^{i+1}\delta^i = 0$.

A chain complex with basis $\mathcal{C}$, additionally have a chosen basis for every vector space $\mathcal{C}^i$.

A chain complex with basis $\mathcal{C}$ is sparse if every coboundary δ^i is sparse under the chosen basis.

Definition 3.3 (Chain map)**.** A chain map f from $\mathcal{C}$ to $\mathcal{D}$ is a collection of maps $f^i : \mathcal{C}^i \to \mathcal{D}^i$, such that $\delta^i_{\mathcal{D}} f^i = f^{i+1}\delta^i_{\mathcal{C}}$.

A chain map f is sparse if every f^i is sparse under the chosen basis.

Definition 3.4 (Chain homotopy)**.** Given two chain maps f, g from $\mathcal{C}$ to $\mathcal{D}$, a chain homotopy from f to g is a collection of maps $h^i : \mathcal{C}^i \to \mathcal{D}^{i-1}$ such that $f^i - g^i = \delta^{i-1}_{\mathcal{D}} h^i + h^{i+1}\delta^i_{\mathcal{C}}$.

A chain homotopy h is sparse if every h^i is sparse under the chosen basis.

Two chain complexes $\mathcal{C}, \mathcal{D}$ are homotopy equivalent if there exists chain maps f from $\mathcal{C}$ to $\mathcal{D}$ and g from $\mathcal{D}$ to $\mathcal{C}$ such that $g \circ f$ is homotopic to $\mathrm{id}_{\mathcal{C}}$ and $f \circ g$ is homotopic to $\mathrm{id}_{\mathcal{D}}$.

Such homotopy equivalence is sparse if the chain homotopies are sparse.

From now on, we only consider the sparse version of every definition above. Therefore, we often drop the word sparse.

Given a chain complex, we can change the direction of the linear maps and form the dual chain complex consists of the boundary operators. Let $\mathcal{C}_i$ be the dual vector space $\mathcal{C}^i \to \mathbb{F}_q$. In our context, since each vector space has a chosen basis, we can identify $\mathcal{C}_i \cong \mathcal{C}^i$. Furthermore, we can define the associate boundary operators $\partial_i : \mathcal{C}_i \to \mathcal{C}_{i-1}$ as the matrix transpose of δ^{i-1}, $\partial_i \coloneqq (\delta^{i-1})^T$. The boundary operators automatically satisfy

$$\partial_{i-1}\partial_i = 0$$

which again forms a chain complex. For most of the discussion, we will focus on cochains rather than chains.

We introduce some standard terminologies. The elements in $\mathcal{C}^i$ ($\mathcal{C}_i$) are called the i-cochains (i-chains). The elements of the kernel of the coboundary (boundary) operators are called cocycles (cycles)

$$Z^i \coloneqq \ker \delta^i = \{\alpha \in \mathcal{C}^i : \delta^i \alpha = 0\}, \qquad Z_i \coloneqq \ker \partial_i = \{\alpha \in \mathcal{C}_i : \partial_i \alpha = 0\}. \tag{3.4}$$

The elements of the image of the coboundary (boundary) operators are called coboundaries (boundaries)

$$B^i \coloneqq \operatorname{im} \delta^{i-1} = \{\delta^{i-1}\alpha : \alpha \in \mathcal{C}^{i-1}\}, \qquad B_i \coloneqq \operatorname{im} \partial_{i+1} = \{\partial_{i+1}\alpha : \alpha \in \mathcal{C}_{i+1}\}. \tag{3.5}$$

Because $\delta^i\delta^{i-1} = 0$ ($\partial_i\partial_{i+1} = 0$), it follows that $B^i \subseteq Z^i$ ($B_i \subseteq Z_i$). The cohomology (homology) group is the quotient of the coboundaries by the cocycles (boundaries by the cycles)

$$H^i \coloneqq Z^i/B^i, \qquad H_i \coloneqq Z_i/B_i. \tag{3.6}$$

The elements are known as the cohomology classes which will be denoted as the equivalence class of a coboundary $[\zeta] = \{\zeta + \beta : \beta \in B^i\}$ for $\zeta \in Z^i$.

3.2.2 Quantum operators over finite fields

We briefly review qudits over finite fields [33]. Let $q = p^r$, be a power of primes. The basis of qudits over $\mathbb{F}_q$ is generated by $\{|x\rangle : x \in \mathbb{F}_q\}$. The Pauli operators over $\mathbb{F}_q$ are

$$X^a|x\rangle = |x + a\rangle, \qquad Z^a|x\rangle = e^{\frac{2\pi i}{p}\,\mathrm{tr}_{\mathbb{F}_q/\mathbb{F}_p}(ax)}|x\rangle \tag{3.7}$$

for $a \in \mathbb{F}_q$. Notice that Z^bX^a and X^aZ^b differ by a phase factor

$$Z^bX^a = e^{\frac{2\pi i}{p}\,\mathrm{tr}_{\mathbb{F}_q/\mathbb{F}_p}(ab)}X^aZ^b. \tag{3.8}$$

With the Pauli operators defined, the Clifford hierarchy can be generalized to the setting of qudits. Let C_1 be the Pauli group generated by the Pauli operators. We define the Clifford hierarchy iteratively, where the i-th level C_i is defined as

$$C_i = \{U : UC_1U^\dagger \subseteq C_{i-1}\}. \tag{3.9}$$

The second level C_2 are also known as the Clifford group, and its elements are known as the Clifford gates. The hierarchy characterize some form of computing power. In particular, the gates in C_2 do not form a universal gate set and can be simulated efficiently using classical computers. In contrast, the gates in C_3 form a universal gate set. These gates outside of C_2 are known as the non-Clifford gates and are essential for building universal quantum computers.

We illustrate some of the gates in the hierarchy. At the second level C_2, for $a \in \mathbb{F}_q$, we have the following operators

$$M^a|x\rangle = |ax\rangle, \tag{3.10}$$

$$CZ^a|x,y\rangle = e^{\frac{2\pi i}{p}\,\mathrm{tr}_{\mathbb{F}_q/\mathbb{F}_p}(axy)}|x,y\rangle, \tag{3.11}$$

$$CNOT^a|x,y\rangle = |x, ax+y\rangle. \tag{3.12}$$

At the third level C_3, for $a \in \mathbb{F}_q$, we have the following operators

$$T^a|x\rangle = e^{\frac{2\pi i}{p}\,\mathrm{tr}_{\mathbb{F}_q/\mathbb{F}_p}(ax^3)}|x\rangle, \tag{3.13}$$

$$CCZ^a|x,y,z\rangle = e^{\frac{2\pi i}{p}\,\mathrm{tr}_{\mathbb{F}_q/\mathbb{F}_p}(axyz)}|x,y,z\rangle. \tag{3.14}$$

In particular, since M^a is a Clifford gate, in the context of magic state distillation, CCZ^a and CCZ^1 are of the same resources. This follows from the relation

$$CCZ^a = M^{a^{-1}} CCZ^1 M^a. \tag{3.15}$$

It is common to perform alphabet reduction which reduces qudits over $\mathbb{F}_q$ into r qudits over $\mathbb{F}_p$. This allows us to obtain codes over qubits ($p = 2$) while working with large finite fields ($q = 2^r$). By viewing $\mathbb{F}_q$ as a vector space over $\mathbb{F}_p$, $\mathbb{F}_q \cong (\mathbb{F}_p)^r$, we can express $|x\rangle$ as $|x_1\rangle \otimes |x_2\rangle \otimes ... \otimes |x_r\rangle$, where $x_1, ..., x_r \in \mathbb{F}_p$. One can verify that the Clifford hierarchies over $\mathbb{F}_q$ and $\mathbb{F}_p$ agree after alphabet reduction. Conceptually, this means that the computation model only depends on the characteristic of the field. More concretely, there are several methods to perform alphabet reduction while preserving the properties of the stabilizer codes [23, 24]. Therefore, there is no need to be intimidated by the large field size $\mathbb{F}_q$.

Notice that several notations in this section overlap with those used in the chain complex. From now on, we will only use CCZ, while X and Z will mainly refer to concepts in the context of the chain complex.

3.2.3 Quantum CSS codes

A quantum CSS code of length n over $\mathbb{F}_q$ is specified by a three term chain complex

$$Q : \mathbb{F}_q^{m_x} \xrightarrow{\delta^0} \mathbb{F}_q^n \xrightarrow{\delta^1} \mathbb{F}_q^{m_z} \tag{3.16}$$

where m_x is the number of X checks, n is the number of qudits, and m_z is the number of Z checks. The coboundary maps δ^0 and δ^1 specify how the checks act on the qudits.

It is understood that the X and Z logical operators correspond to cocycles Z^1 and cycles Z_1, respectively. Additionally, the X and Z stabilizers correspond to coboundaries B^1 and boundaries B_1, respectively.

The dimension $k = \dim H^1$ is the number of logical qubits. The distance is $d = \min(d^1, d_1)$ where

$$d^1 = \min_{\alpha \in Z^1 \setminus B^1} |\alpha|, \qquad d_1 = \min_{\alpha \in Z_1 \setminus B_1} |\alpha|. \tag{3.17}$$

These are the minimal weights of the nontrivial X and Z logical operators.

We say the quantum code is a low-density parity-check (LDPC) code if the chain complex is sparse, i.e. each check interacts with a bounded number of qudits and each qudit interacts with a bounded number of checks.

To simplify notation in later sections, where many indices are needed to label variables, we will sometimes drop the indices and use Z for Z^1 and B for B^1 as the discussion mostly focuses on the properties of cochains.

3.2.4 Cubical complexes and cubical codes

The cubical complexes are formed by vertices, edges, squares, cubes, and their high dimensional counterparts. Their first application in error correcting codes was introduced in [20] to construct quantum locally testable codes. Prior to this, their 1D and 2D counterparts

had been applied in [34, 35] and [36, 37]. Our discussion will focus on labeled cubical complexes because it is more straightforward to apply local codes.

Labeled cubical complexes

A t-dimensional labeled cubical complex $X(V; A_1, A_2, ..., A_t)$ can be constructed from a set V of size $N = |V|$ and subsets of permutations $A_1, ..., A_t \subset \mathrm{Aut}(V)$ of size $\Delta = |A_i|$. These permutations should pairwise commute, i.e. $a_i a_j = a_j a_i$ for all $a_i \in A_i$ and $a_j \in A_j$. The complex is also denoted as $X(V; \{A_i\})$.

Each k-dimensional cell ($0 \leqslant k \leqslant t$) is associated with a type, which is a subset $S \subseteq \{1, ..., t\}$ of size $|S| = k$. Denote $\overline{S}$ as the complement of S in $\{1, ..., t\}$. A cell σ of type S is specified by

$$\sigma = (v; (a_j)_{j \in S}, (b_j)_{j \in \overline{S}}), \tag{3.18}$$

where $v \in V$, $a_j \in A_j$ for each $j \in S$, and $b_j \in \{0, 1\}$ for each $j \in \overline{S}$. Let $X(S)$ be the faces of type S. Then, $|X(S)| = |V|\Delta^k 2^{t-k}$. The type of a cell σ is denoted as $\texttt{type}(\sigma)$.

Geometrically, a k-cell of type S, $\sigma = (v; a, b)$ contains the following 2^k vertices

$$\left\{ \left(\prod_{j: b'_j = 1} a_j \cdot v; b' || b \right) : b' \in \{0, 1\}^S \right\}, \tag{3.19}$$

where $b'||b$ denotes string concatenation with re-ordering of the indices in the natural way. We use the notation $a_j \cdot v$ to denote the element $a_j(v) \in V$ for $a_j \in A_j$ is a permutation of V. Note that the order of the product in $\prod_j a_j \cdot v$ does not matter since the permutations a_j from different sets A_j pairwise commute.

A family of labeled cubical complexes to keep in mind is to take V to be an abelian group G, and A_i to be a collection of left actions $A_i = \{(g \mapsto \tilde{a}_i g) : \tilde{a}_i \in \tilde{A}_i\}$ for a generating subset $\tilde{A}_i \subset G$. Since the group is abelian, the elements from different sets pairwise commute. Because of this example, from now on, we will use G instead of V to denote the vertex set.

When $t = 2$, we can take G to be a non-abelian group, A_1 to be a collection of left actions $A_1 = \{(g \mapsto \tilde{a}_1 g) : \tilde{a}_1 \in \tilde{A}_1\}$, and A_2 to be a collection of right actions $A_2 = \{(g \mapsto g\tilde{a}_2) : \tilde{a}_2 \in \tilde{A}_2\}$. The permutations $a_1 \in A_1$ and $a_2 \in A_2$ commutes, because left and right actions commute. This is known as the left-right Cayley complex [37] which is an essential ingredient to the recent breakthroughs in classical locally testable codes and quantum low-density parity check codes [37, 17].

Another family of labeled cubical complexes comes from a Cartesian product of three bipartite graphs. These are used for constructing qLDPC codes with $\gamma \to 0$ in the companion paper [21].

We introduce some additional notations. Let $X(i)$ be the set of i-cells. Given two cells σ, τ, we write $\sigma \preceq \tau$ if σ is a subcell of τ, and write $\sigma \prec \tau$ if σ is an 'immediate' subcell of τ where their dimensions differ by 1.

$X_{\geqslant \sigma}$ denotes the sub-complex consists of all cells above σ

$$X_{\geqslant \sigma} = \{\tau \in X : \tau \succeq \sigma\}. \tag{3.20}$$

$X_{\leqslant \sigma}$ denotes the sub-complex consists of all cells below σ

$$X_{\leqslant \sigma} = \{\tau \in X : \tau \preceq \sigma\}. \tag{3.21}$$

We define $X_{\geqslant \sigma}(k) = X_{\geqslant \sigma} \cap X(k)$ and $X_{\leqslant \sigma}(k) = X_{\leqslant \sigma} \cap X(k)$.

Chain complexes from cubical complexes

To generate a chain complex, we attach each cell with local coefficients similar to a Tanner code construction. Let $C_1, ..., C_t \subseteq \mathbb{F}_q^\Delta$ be classical codes of length Δ with dimension m_i and generating matrix $h_i^T : \mathbb{F}_q^{m_i} \to \mathbb{F}_q^\Delta$. The local coefficient space associated with a k-cell σ of type S is $\mathcal{F}_\sigma = \bigotimes_{i \in \overline{S}} \mathbb{F}_q^{m_i}$.

The local coefficient can be identified with a codeword in the tensor code $\bigotimes_{i \in \overline{S}} C_i$. Given $c \in \mathcal{F}_\sigma$, it is clear that $(\bigotimes_{i \in \overline{S}} h_i^T)c$ is a codeword in $\bigotimes_{i \in \overline{S}} C_i \subseteq (\mathbb{F}_q^\Delta)^{\otimes \overline{S}}$. The codeword

can be further interpreted as a function that assigns every t-cells above σ with a value in $\mathbb{F}_q$, $(\bigotimes_{i\in\overline{S}} h_i^T)c : X_{\geqslant\sigma}(t) \to \mathbb{F}_q$. Since each t-cell above σ is obtained by extending the directions in $\overline{S}$ and each direction has $|A_i| = \Delta$ labels to choose from, we can identify $\Delta^{\overline{S}}$ with $X_{\geqslant\sigma}(t)$ and obtain $(\mathbb{F}_q^{\Delta})^{\otimes\overline{S}} \cong \mathbb{F}_q^{X_{\geqslant\sigma}(t)} \cong (X_{\geqslant\sigma}(t) \to \mathbb{F}_q)$. This induces a map from $\mathcal{F}_\sigma$ to $\mathbb{F}_q^{X_{\geqslant\sigma}}$ and we define $\mathcal{G}_\sigma \subseteq \mathbb{F}_q^{X_{\geqslant\sigma}(t)}$ as the image. It is clear that $\mathcal{G}_\sigma \cong \mathcal{F}_\sigma$. These $\mathcal{G}_\sigma$ will be helpful when discussing the generalization to sheaf code in Section 3.6 and defining the co-restriction map below.

We first define the vector spaces for the desired chain complex. The vector space $\mathcal{C}^i$ is the space of functions α, that assigns a value $\alpha(\sigma) \in \mathcal{F}_\sigma$ for each i-cell σ. In particular,

$$\mathcal{C}^i = \bigoplus_{\sigma\in X(i)} \mathcal{F}_\sigma. \tag{3.22}$$

To define the coboundary maps, we first describe the co-restriction maps. For every pair of cells $\sigma \preceq \tau$, we have an embedding $X_{\geqslant\tau}(t) \subseteq X_{\geqslant\sigma}(t)$. This induces the map $\mathcal{G}_\sigma \to \mathcal{G}_\tau$, which maps $c : X_{\geqslant\sigma}(t) \to \mathbb{F}_q$ to $c|_{X_{\geqslant\tau}(t)} : X_{\geqslant\tau}(t) \to \mathbb{F}_q$. Because $\mathcal{G}_\sigma \cong \mathcal{F}_\sigma$, this induces a map from $\mathcal{F}_\sigma \to \mathcal{F}_\tau$, which we denote as co-res$_{\sigma,\tau}$.

More explicitly, co-res can be defined by mapping through h_j^T among certain directions and restrict to certain slices. Given a pair of cells $\sigma \prec \tau$ with $\texttt{type}(\tau)\backslash\texttt{type}(\sigma) = \{j\}$,

$$\text{co-res}_{\sigma,\tau}(c) = \Big((I_{-j} \otimes h_j^T)c\Big)[..., a_i, ...], \tag{3.23}$$

where $a_j \in A_j$ is the label for the j-th direction of τ. For a general pair of cells $\sigma \preceq \tau$, we again apply h_j^T and restrict to $a_j \in A_j$ for all $j \in \texttt{type}(\tau)\backslash\texttt{type}(\sigma)$. One can verify that this agrees with the previous definition of co-res.

The coboundary maps $\delta^i : \mathcal{C}^i \to \mathcal{C}^{i+1}$ is defined by

$$\delta^i\alpha(\tau) = \sum_{\sigma\prec\tau} \text{co-res}_{\sigma,\tau}(\alpha(\sigma)). \tag{3.24}$$

(In general, there will be signs, $\sum_{\sigma\prec\tau}(-)^{(\cdots)}$ co-res$_{\sigma,\tau}(\alpha(\sigma))$. We evade this technical discussion by studying finite fields with characteristic 2 where $-1 = 1$.)

We can verify that $\delta \circ \delta = 0$, where we have

$$\delta^{i+1}\delta^i\alpha(\pi) = \sum_{\sigma \prec \tau \prec \pi} \text{co-res}_{\sigma,\pi}(\alpha(\sigma)) = 0. \tag{3.25}$$

The last equality holds because for each pair $\sigma \in X(i), \pi \in X(i+2)$, there exist two $\tau \in X(i+1)$ such that $\sigma \prec \tau \prec \pi$.

We denote the chain complex as $\mathcal{C}(X, \{C_j\}_{j=1}^t)$ or $\mathcal{C}(G, \{A_j\}_{j=1}^t, \{C_j\}_{j=1}^t)$.

Cubical quantum codes

Recall that a quantum CSS code is equivalent to a three-term chain complex. The previous section provides a chain complex $\mathcal{C}(G, \{A_i\}, \{C_i\})$. So the only additional data is an integer ℓ which specifies which of the three terms to select. This leads to the quantum code $Q(G, \{A_i\}, \{C_i\}, \ell)$ defined by

$$\mathcal{C}^{\ell-1}(G, \{A_i\}, \{C_i\}) \to \mathcal{C}^{\ell}(G, \{A_i\}, \{C_i\}) \to \mathcal{C}^{\ell+1}(G, \{A_i\}, \{C_i\}). \tag{3.26}$$

In particular, the X checks are placed on the $(\ell-1)$-cells, the qubits are placed on the ℓ-cells, and the Z checks are placed on the $(\ell+1)$-cells,

Later, we will mainly discuss quantum codes based on 2D square complexes and 3D cubical complexes with $\ell = 1$.

3.3 CCZ formalism

This section introduces the CCZ formalism, where the goal is to establish a condition on quantum codes that implies a transversal operation of CCZ gate. This framework should be compared to the triorthogonal condition introduced in [27], which specifies a condition on quantum codes that implies a transversal operation for the T gate. Our formalism can be generalized to C^{r-1}Z gates or subsystem codes, although we will not discuss them explicitly in this work.

3.3.1 Definition of CCZ codes

In the CCZ formalism, the data consists of three quantum codes Q_1, Q_2, Q_3 of length n_1, n_2, n_3, respectively, and a trilinear map $f : \mathbb{F}_q^{n_1} \times \mathbb{F}_q^{n_2} \times \mathbb{F}_q^{n_3} \to \mathbb{F}_q$. We say a function is trilinear if the function is linear when any of the two inputs are fixed. For example, $f(\alpha_1, \alpha_2, \alpha_3) + f(\alpha_1', \alpha_2, \alpha_3) = f(\alpha_1 + \alpha_1', \alpha_2, \alpha_3)$ and $f(a\alpha_1, \alpha_2, \alpha_3) = af(\alpha_1, \alpha_2, \alpha_3)$ for $a \in \mathbb{F}_q$. The trilinear map f has to satisfy an additional condition which we will explain shortly.

We associate the trilinear map f with an unitary operator CCZ^f defined by

$$\mathrm{CCZ}^f |\alpha_1, \alpha_2, \alpha_3\rangle = e^{\frac{2\pi i}{p} \mathrm{tr}_{\mathbb{F}_q/\mathbb{F}_p}(f(\alpha_1,\alpha_2,\alpha_3))} |\alpha_1, \alpha_2, \alpha_3\rangle \tag{3.27}$$

where $\alpha_i \in \mathbb{F}_q^{n_i}$ and $|\alpha_1, \alpha_2, \alpha_3\rangle$ is the state on $n_1 + n_2 + n_3$ qudits labeled by strings $\alpha_1, \alpha_2, \alpha_3$ in the Z basis. (Note that strings in the Z basis correspond to Pauli X operators.) This is analogous to the definition of a CCZ gate. In fact, we will see that CCZ^f can be implemented by applying multiple CCZ gates.

To ensure that CCZ^f maps codewords to codewords, we need to impose a condition on f. The codewords in $Q_1 \otimes Q_2 \otimes Q_3$ are spanned by

$$|[\zeta_1], [\zeta_2], [\zeta_3]\rangle \coloneqq \sum_{\beta_1 \in B_1} \sum_{\beta_2 \in B_2} \sum_{\beta_3 \in B_3} |\zeta_1 + \beta_1, \zeta_2 + \beta_2, \zeta_3 + \beta_3\rangle \tag{3.28}$$

for $\zeta_i \in Z_i$. Because CCZ^f only changes the phase under the Z basis, $\mathrm{CCZ}^f |[\zeta_1], [\zeta_2], [\zeta_3]\rangle$ is a codeword iff

$$\mathrm{CCZ}^f |[\zeta_1], [\zeta_2], [\zeta_3]\rangle \propto |[\zeta_1], [\zeta_2], [\zeta_3]\rangle \tag{3.29}$$

which is equivalent to

$$\mathrm{tr}_{\mathbb{F}_q/\mathbb{F}_p}(f(\zeta_1, \zeta_2, \zeta_3)) = \mathrm{tr}_{\mathbb{F}_q/\mathbb{F}_p}(f(\zeta_1 + \beta_1, \zeta_2 + \beta_2, \zeta_3 + \beta_3)) \tag{3.30}$$

for all $\zeta_i \in Z_i$ and $\beta_i \in B_i$. Therefore, we require f to satisfy

$$f(\zeta_1, \zeta_2, \zeta_3) = f(\zeta_1 + \beta_1, \zeta_2 + \beta_2, \zeta_3 + \beta_3) \tag{3.31}$$

which guarantees that CCZ^f maps codewords to codewords. This condition is appealing as it is saying that f is uniquely defined on the cohomology classes. This is how we will refer to this condition.

Remark 3.5. The condition we imposed (3.31) may seem stronger than what is necessary (3.30). However, they are in fact equivalent. For all $a \in \mathbb{F}_q$, because $a\zeta_1 \in Z_1$ and $a\beta_1 \in B_1$, from (3.30) we have

$$\mathrm{tr}_{\mathbb{F}_q/\mathbb{F}_p}(f(a\zeta_1, \zeta_2, \zeta_3)) = \mathrm{tr}_{\mathbb{F}_q/\mathbb{F}_p}(f(a\zeta_1 + a\beta_1, \zeta_2 + \beta_2, \zeta_3 + \beta_3)). \tag{3.32}$$

Because f is trilinear, we have

$$\mathrm{tr}_{\mathbb{F}_q/\mathbb{F}_p}(af(\zeta_1, \zeta_2, \zeta_3)) = \mathrm{tr}_{\mathbb{F}_q/\mathbb{F}_p}(af(\zeta_1 + \beta_1, \zeta_2 + \beta_2, \zeta_3 + \beta_3)) \tag{3.33}$$

for all $a \in \mathbb{F}_q$.

Since finite extensions of a finite field is separable, we have $\mathbb{F}_q/\mathbb{F}_p$ is separable. Furthermore, for separable extensions $\mathbb{F}_q/\mathbb{F}_p$, the bilinear trace form $\mathrm{tr}_{\mathbb{F}_q/\mathbb{F}_p}(ax)$ is nondegenerate, meaning that if $\mathrm{tr}_{\mathbb{F}_q/\mathbb{F}_p}(ax) = 0$ for all $a \in \mathbb{F}_q$, then $x = 0$. Thus, $f(\zeta_1, \zeta_2, \zeta_3) = f(\zeta_1 + \beta_1, \zeta_2 + \beta_2, \zeta_3 + \beta_3)$ which is exactly (3.31).

This leads to the definition of a CCZ code.

Definition 3.6. A CCZ code consists of three quantum codes Q_1, Q_2, Q_3 and a trilinear map $f : \mathbb{F}_q^{n_1} \times \mathbb{F}_q^{n_2} \times \mathbb{F}_q^{n_3} \to \mathbb{F}_q$ such that

$$f(\zeta_1, \zeta_2, \zeta_3) = f(\zeta_1 + \beta_1, \zeta_2 + \beta_2, \zeta_3 + \beta_3) \tag{3.34}$$

for all $\zeta_i \in Z_i$ and $\beta_i \in B_i$.

Generally, n_1, n_2, n_3 are of the same order.

Example 3.7 (Triorthogonal code)**.** Three copies of the same triorthogonal code [27] form a CCZ code. Let Q be a triorthogonal code of length n. Let $\{\beta^a\}$ be the vectors correspond

to stabilzer generators and let $\{\zeta^a\}$ be vectors corresponds to logical generators. Recall that a triorthogonal code is defined over $\mathbb{F}_2$ and satisfies

1. $\sum_{i=1}^n \beta_i^a = 0$,
2. $\sum_{i=1}^n \zeta_i^a = 1$,
3. $\sum_{i=1}^n \alpha_i^a \alpha_i^b = 0$ for any distinct $\alpha^a, \alpha^b \in \{\beta^a\} \cup \{\zeta^a\}$,
4. $\sum_{i=1}^n \alpha_i^a \alpha_i^b \alpha_i^c = 0$ for any distinct $\alpha^a, \alpha^b, \alpha^c \in \{\beta^a\} \cup \{\zeta^a\}$.

Set $f(\alpha_1, \alpha_2, \alpha_3) = \sum_{i=1}^n \alpha_{1,i}\alpha_{2,i}\alpha_{3,i}$, where $\alpha_{1,i}$ is the i-th entry of α_1. It is straightforward to check that the conditions for triorthogonal codes imply the condition for CCZ codes, Equation (3.34).

Example 3.8 (3D surface code)**.** Three 3D surface codes form a CCZ code. Since similar facts have been observed in [12, 13], we will not construct f explicitly. We will instead make some high level remarks.

The key idea is that the X-logical operators of the 3D surface codes can be viewed as 1-forms. A natural way to form a trilinear operator among three 1-forms is through the cup product. Schematically, we can write $f(\alpha_1, \alpha_2, \alpha_3) = \int_M \alpha_1 \cup \alpha_2 \cup \alpha_3$. It is well-known in algebraic topology that cup product is well-defined on the cohomology classes. Therefore, the construction automatically satisfy the condition for CCZ codes.

An alternative perspective, which may be more visually appealing, is to note that 1-forms are Poincaré dual to 2D surfaces. In this dual perspective, f becomes the intersection number of the three 2D surfaces. This can be visualized in Figure 3.1. In particular, the fact that there is "one" nontrivial triple intersection corresponds to having "one" logical CCZ gate.

In this example, we see that the reason f is defined on the cohomology classes is largely induced from the topological nature of the construction. In particular, the

intersection number is a topological invariant. This is the main idea that allows us to construct transversal CCZ gates on qLDPC codes. Later, we will first provide a topological interpretation of the qLDPC codes, which naturally induces the construction of f through the intersection number.

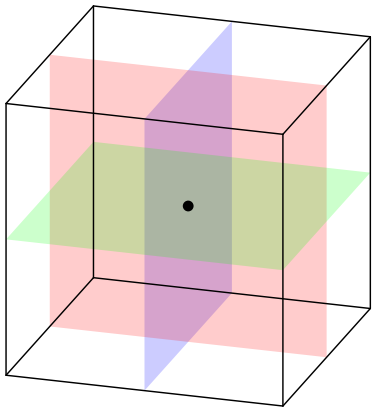

Figure 3.1: The X-logical operators of three 3D surface codes, each oriented along a different axis. These logical operators are represented by the corresponding surfaces. A crucial feature of these surfaces is that their triple intersection number is always 1.

Remark 3.9. Codes with transversal T gates appear to be related to triple self-intersection. An example with nontrivial triple self-intersection is the 3D real projective space $\mathbb{RP}^3$ which has a homology $H_2(\mathbb{RP}^3, \mathbb{Z}/2\mathbb{Z}) = \mathbb{Z}/2\mathbb{Z}$ with a nontrivial triple self-intersection.

3.3.2 Extra code parameters of CCZ codes: $n_{\text{CCZ}}, k_{\text{CCZ}}$

Besides the traditional code parameters, such as n, k, d, a CCZ code has additional parameters n_{CCZ}, k_{CCZ}, which are dependent on f. Intuitively, n_{CCZ} is the number of CCZ gates needed to implement CCZ^f, and k_{CCZ} is the number of logical CCZ gates applied by acting CCZ^f. In the context of magic state distillation, this means, by consuming n_{CCZ} noisy CCZ states, the error correcting code can output k_{CCZ} cleaner CCZ states. Therefore, the corresponding magic state distillation protocol has overhead $\log^\gamma(1/\varepsilon)$ for $\gamma = \log(n_{\text{CCZ}}/k_{\text{CCZ}})/\log(d)$, where ε is the target error rate.

We first discuss n_{CCZ}. Notice that the operator CCZ^f defined in (3.27) can be rewritten as

$$\mathrm{CCZ}^f = \prod_{j_1\in[n_1],j_2\in[n_2],j_3\in[n_3]} \mathrm{CCZ}_{j_1,j_2,j_3}^{f(e_{j_1},e_{j_2},e_{j_3})} = \prod_{j_1\in[n_1],j_2\in[n_2],j_3\in[n_3],f(e_{j_1},e_{j_2},e_{j_3})\neq 0} \mathrm{CCZ}_{j_1,j_2,j_3}^{f(e_{j_1},e_{j_2},e_{j_3})} \tag{3.35}$$

where $\mathrm{CCZ}_{j_1,j_2,j_3}^{f(e_{j_1},e_{j_2},e_{j_3})}$ is the $\mathrm{CCZ}^{f(e_{j_1},e_{j_2},e_{j_3})}$ gate acting on the j_1, j_2, j_3-th qudits in codes Q_1, Q_2, Q_3, respectively. Recall that, CCZ^a gates for different $a \neq 0$ are equivalent resources up to Clifford operations, as discussed in Section 3.2.2.

One may notice that this application of CCZ gates is not strictly transversal. There are two perspective to address this concern. First, for the purpose of constructing fault-tolerant circuits, the encoded unitary does not need to be strictly transversal. In fact, the same fault-tolerant analysis holds as long as the circuit has constant depth, which is primarily the setting we consider. Since "transversal" is a well-established term, we will slightly abuse its meaning. Similarly, constant depth circuits are also acceptable for magic state distillation.

Alternatively, constant depth applications of CCZ gates can be transformed into strictly transversal operation by enlarging the quantum code, [21, Lemma 2.45]. Conceptually, this can be compared to the weight reduction techniques, which also implies a certain equivalence between different locality constants.

The above discussion leads to the following definition.

Definition 3.10. Given a CCZ code, we define

$$n_{\mathrm{CCZ}} \coloneqq |\{(j_1, j_2, j_3) : f(e_{j_1}, e_{j_2}, e_{j_3}) \neq 0\}| \tag{3.36}$$

where e_j is the standard basis vector with all entries equal to 0 except for the j-th entry which is 1.

It is clear that $n_{\mathrm{CCZ}} \leqslant n_1 n_2 n_3$. However, in the context of fault tolerance, we want f to be sparse, meaning that each qubit is only involved in a bounded number of CCZ gates.

This sparsity condition is analogous to the LDPC condition. Since each j_1 only appears a bounded number of times on the RHS of (3.36), in this context, we have $n_{\mathrm{CCZ}} = O(n_1)$. In fact, we often have $n_{\mathrm{CCZ}} = \Theta(n_1) = \Theta(n_2) = \Theta(n_3)$.

Remark 3.11. In magic state distillation with perfect Clifford gates, n_{CCZ} is a key code parameter as it quantifies the number of resource state consumed. However, in the context of transversal gates for building fault-tolerant circuits, it is crucial for the implementation of CCZ^f to have a low depth. That means it is important to bound the maximal number of CCZ gates that a qubit participates in, which leads to the definition of the code parameter $w_{\mathrm{CCZ}} = \max(w_1, w_2, w_3)$, where

$$w_1 = \max_{j_1 \in [n_1]} |\{(j_2, j_3) \in [n_2] \times [n_3] : f(e_{j_1}, e_{j_2}, e_{j_3}) \neq 0\}|, \tag{3.37}$$

and w_2 and w_3 are defined similarly.

We now discuss k_{CCZ}, which is the number of logical CCZ gates applied by CCZ^f. Let k_{CCZ} be the largest integer such that there exist k_{CCZ} X logical operators in each of the codes Q_1, Q_2, Q_3, denoted as $\zeta_{1,j}, \zeta_{2,j}, \zeta_{3,j}$ for $j \in [k_{\mathrm{CCZ}}]$, with the property that

$$f(\zeta_{1,j_1}, \zeta_{2,j_2}, \zeta_{3,j_3}) = 1_{j_1=j_2=j_3} \tag{3.38}$$

for $j_1, j_2, j_3 \in [k_{\mathrm{CCZ}}]$, where $1_{j_1=j_2=j_3}$ is equal to 0 except when $j_1 = j_2 = j_3$ which is 1. Once we found those X logical operators, we treat them as the data qubits and set all other qubits to be in the 0 state. One can see that CCZ^f act on the data qubits as k_{CCZ} copies of CCZ gates.

More conceptually, k_{CCZ} can be defined as the subrank of the 3-tensor obtained by restricting f to the logical operators. We first describe the 3-tensor, then review the definition of subrank. Let $\{\zeta_{1,j_1'}\}_{j_1'=1}^{k_1}$ be a set of basis for the X logical operators, i.e. the representatives of a basis for the cohomology group $H(Q_1)$. Define $\{\zeta_{2,j_2'}\}_{j_2'=1}^{k_2}$ and

$\{\zeta_{3,j_3'}\}_{j_3'=1}^{k_3}$ similarly. Let T be the 3-tensor in $\mathbb{F}_q^{k_1} \otimes \mathbb{F}_q^{k_2} \otimes \mathbb{F}_q^{k_3}$, where

$$T_{j_1',j_2',j_3'} = f(\zeta_{1,j_1'}, \zeta_{2,j_2'}, \zeta_{3,j_3'}). \tag{3.39}$$

One can see that T contains all the information about how f acts on the codespace.

We now review the definition of subrank. The subrank of a 3-tensor $T \in \mathbb{F}_q^{k_1} \otimes \mathbb{F}_q^{k_2} \otimes \mathbb{F}_q^{k_3}$, denoted as $Q(T)$, is the largest integer r such that there exist linear maps $M_i : \mathbb{F}_q^{k_i} \to \mathbb{F}_q^r$ $(r \leqslant k_i)$ with the property that

$$(M_1 \otimes M_2 \otimes M_3)T = \sum_{j=1}^{r} e_j \otimes e_j \otimes e_j. \tag{3.40}$$

It is straightforward to check that the earlier definition of k_{CCZ} is equal to $Q(T)$. In particular, the choice of the X logical operators $\zeta_{1,j}, \zeta_{2,j}, \zeta_{3,j}$ for $j \in k_{\mathrm{CCZ}}$ corresponds to the choice of the linear maps M_1, M_2, M_3.

The above discussion leads to the following definition.

Definition 3.12. Given a CCZ code, we define

$$k_{\mathrm{CCZ}} \coloneqq Q(T) \tag{3.41}$$

where T is the 3-tensor obtained by restricting f to the logical operators and $Q(T)$ is the subrank of T.

Remark 3.13. In the definition of k_{CCZ}, we treat CCZ gates as the desired resource and discard everything else. It is possible that the remnant contains nontrivial resource which may be recaptured and utilize for better efficiency. We leave this possibility for future exploration.

To summarize, the discussion of CCZ codes generally contains 3 parts:

1. Construct three codes Q_1, Q_2, Q_3 together with a trilinear map f.

2. Show that the trilinear map f is uniquely defined on the cohomology classes, i.e. f satisfies Equation (3.34).

3. Bound n_{CCZ} and k_{CCZ}.

In this work, we focus on the first two aspects. The third aspect will be discussed in the companion paper [21].

3.4 Interpret logical operators of qLDPC codes as strings and surfaces

It is well-known that in a 2D surface code, logical operators can be viewed as 1D strings. Similarly, in a 3D surface code, the Z logical operator corresponds to a 1D string and the X logical operator corresponds to a 2D surface. These geometric interpretations were crucial for understanding the intersection number that induces the CCZ gate on a 3D surface code, as discussed in Theorem 3.8. In this section, we would like to obtain a similar geometric interpretation for qLDPC codes, which will allow us to construct f through the intersection number.

This geometric interpretation for qLDPC codes is largely motivated by the recent works on subdivided codes [29, 31]. One key step in these works is to impose a 2D structure on qLDPC codes and attach 2D surface codes on these 2D regions. See Figure 3.2. This process provides a family of codes with different levels of subdivision characterized by L, where the original code has $L = 1$. Since the new code is mostly formed by 2D surface codes, the logical operators can be naturally interpreted as 1D strings that form a tree, which branches at the place where the surfaces meet. See Figure 3.3.

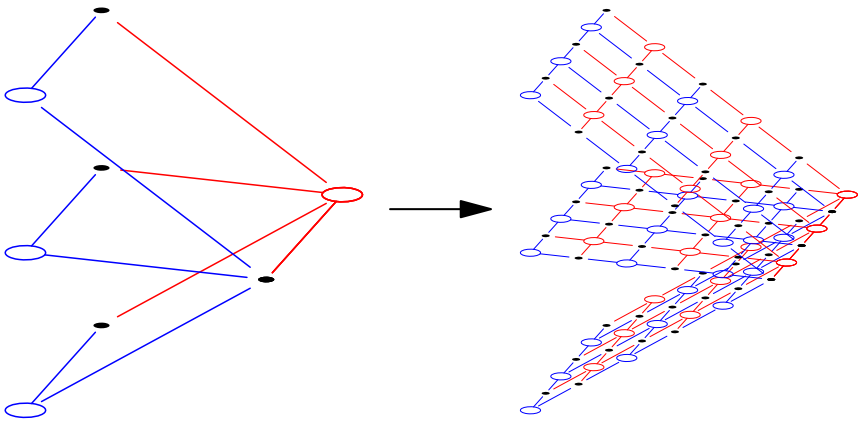

Figure 3.2: The figure illustrates the subdivision process for part of a quantum code.

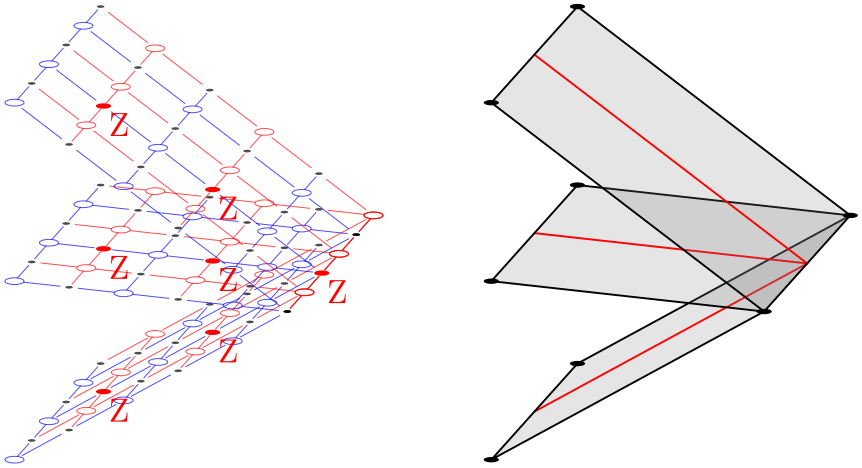

Figure 3.3: The left figure shows part of a Z logical operator, which has a natural string structure illustrated on the right.

In a physics language, the boundary edge in the figure is condensed by $e_1e_2e_3, m_1m_2, m_2m_3$ where e_i and m_i are the electric and magnetic particle on the i-th layer. See Figure 3.4. We say a set of particles (e.g. $e_1e_2e_3$) condenses on a boundary if this set of particles can be created without extra energy near the boundary. This is related to the fact that three logical Z strings ending on this boundary do not violate any stabilizers.

There are certain criteria on what sets can be condensed. For general TQFTs, we need to make sure every condensing set is 'bosonic' and different condensing sets mutually commute [38, 39]. For $\mathbb{Z}/2\mathbb{Z}$-gauge theories (i.e. toric codes), by using the braiding property

of the e and m particles, the condition can be rephrased as the following. If we map e_i to Z_i and m_i to X_i, then as Pauli operators, they commutes. Indeed $Z_1Z_2Z_3, X_1X_2, X_2X_3$ mutually commutes. Because of this mapping, we will refer to these boundary condensing patterns as the boundary stabilizers. If we further require that the boundary is 'gapped', which is the case for stabilizer codes with large distances, every Pauli operator that commutes with the boundary stabilizers has to be a boundary stabilizer, i.e. the quantum code induced by the boundary stabilizers has $k = 0$.

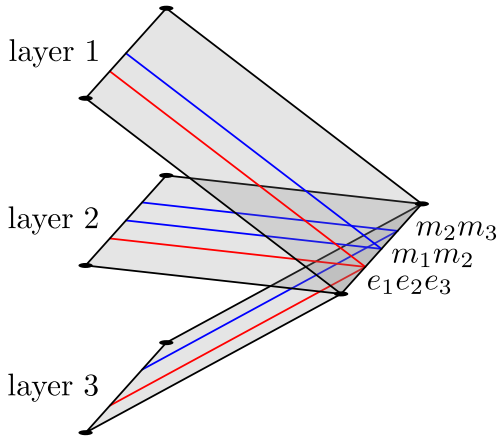

Figure 3.4: The figure shows three condensing patterns for the code example above.

If the code is a CSS code, then each condensing set is either all formed by e or m. That means the boundary stabilizers also form a CSS code and we can specify the condensing pattern using two classical linear codes $C_e, C_m \subseteq \mathbb{F}_q^\Delta$, where Δ is the number of squares meeting at the boundary. For example, the condensing pattern $e_1e_2e_3, m_1m_2, m_2m_3$ has $C_e = \text{span}\{(1,1,1)\}$, $C_m = \text{span}\{(1,1,0),(0,1,1)\}$. From the discussion in the last paragraph, we observe that when the boundary is 'gapped', C_e and C_m are the dual to each other, $C_m = C_e^\perp$.

In fact, the two classical codes C_e, C_m can be identified with the local codes used in the Tanner code constructions [17, 25, 18]. In particular, for local codes C_A, C_B, the horizontal edges have condensing data $C_e = C_A, C_m = C_A^\perp$ and the vertical edges have

condensing data $C_e = C_B, C_m = C_B^\perp$.

As a fun fact, this observation can be used to prove that all three constructions [17, 25, 18] are (sparse) chain homotopy equivalent as defined in Section 3.2.1. This implies that if one of the three codes has certain properties, such as large distance or having a decoder, then other codes also have such properties.

This observation also elucidate the appearance of dual codes when analyzing Z distance versus X distance of the qLDPC codes. In the proof of the papers above, showing Z distance boils down to certain properties of the local codes C_A, C_B, while showing X distance boils down to certain properties of the local codes $C_A^\perp, C_B^\perp$. We can now explain this phenomenon as follows. The proof of distance is essentially about analyzing the structure of the string patterns. Since the Z logical operators correspond to strings that ends on the boundaries according to C_e, proving Z distance only involves C_e, i.e. C_A, C_B. Similarly, since the X logical operators correspond to strings that ends on the boundaries according C_m, proving X distance only involves C_m, i.e. $C_A^\perp, C_B^\perp$.

More generally, the discussion about boundary conditions applies to codes based on 2D simplicial complexes such as the quantum variant of [19], to codes based on higher dimensional cubical complexes [20], and to sheaf codes [10, 16]. The formal discussion for sheaf codes will appear in Section 3.6.

In this work, the geometric interpretation of the logical operators is used to define f as the intersection number. To achieve this, we will formalize this relation by providing an explicit mapping for quantum codes constructed from 2D and 3D cubical complexes.

3.4.1 An explicit map from Pauli X operators to geometric objects for 2D and 3D cubical codes

For the purpose of constructing CZ and CCZ gates, all we need is a mapping for the Pauli X operators. This makes the task especially simple if we place the checks and qubits geometrically, which is what we did in Section 3.2.4.

We start with quantum codes based on square complexes $\mathcal{C}(G, \{A_1, A_2\}, \{C_1, C_2\})$. Given a vector on the qudits $\alpha \in \mathcal{C}^1(G, \{A_1, A_2\}, \{C_1, C_2\})$ that represents a Pauli X operator. The local coefficient on each edge e, $\alpha(e) \in \mathcal{F}_e$, corresponds to a function on the Δ adjacent faces $h_i^T \alpha(e) \in \mathcal{G}_e \subseteq \mathbb{F}_q^\Delta$. These values can be attached to the segments between the midpoint of the edge e and the center of the corresponding face f. (This geometric object should be thought of as the chain in algebraic topology.) See Figure 3.5.

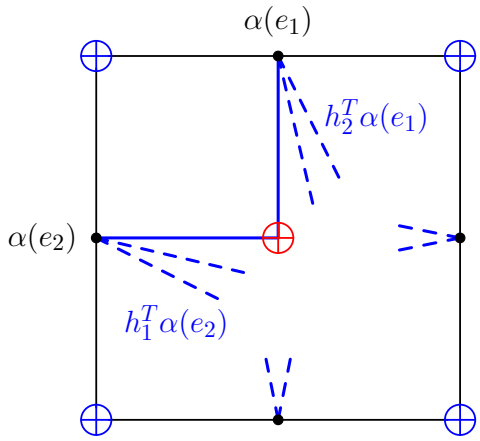

Figure 3.5: Each edge has a local coefficient $\alpha(e) \in V_e$. Such a local coefficient induces a vector $h_i^T \alpha(e) \in \mathbb{F}_q^\Delta$. The vector $h_i^T \alpha(e)$ is then used to induce segments between the midpoints of the edges and the centers of the faces. Notice that because $\delta\alpha = 0$, the sum of the segment values on a face is 0.

We observe three properties of this mapping. First, the string pattern near the boundary belongs to the local codes C_1 or C_2. Second, the condition of a X logical operator $\delta\alpha = 0$, corresponds to the condition that the sum of the segment values on each face is

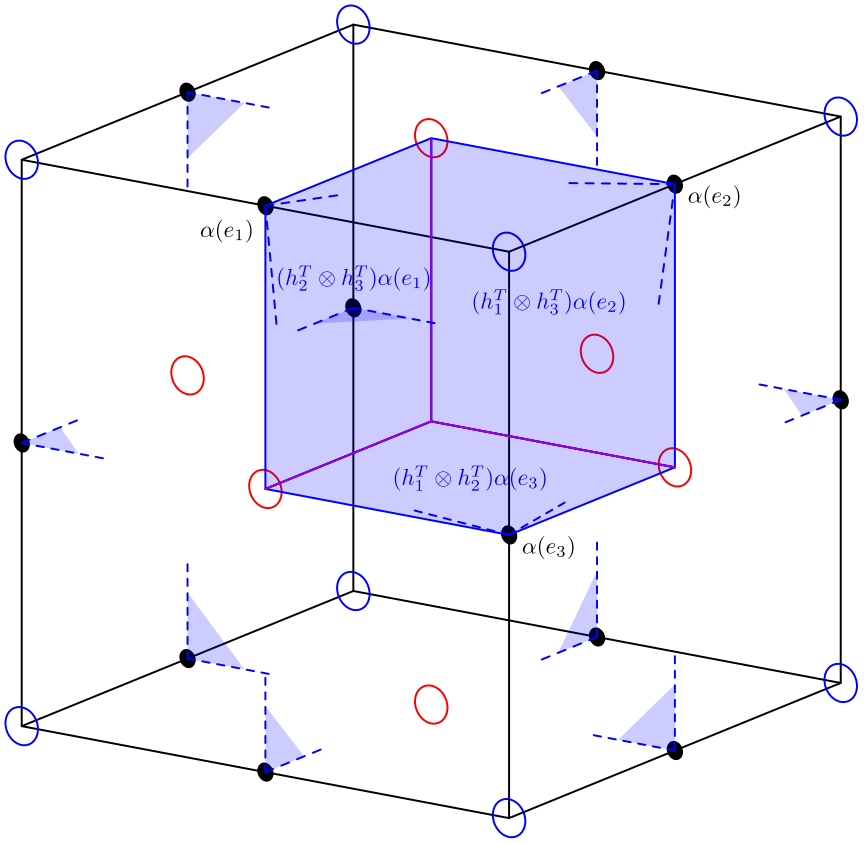

Figure 3.6: Each edge has a local coefficient $\alpha(e) \in V_e$. Such a local coefficient induces a vector $(h_i^T \otimes h_j^T)\alpha(e) \in \mathbb{F}_q^{\Delta^2}$. The vector $(h_i^T \otimes h_j^T)\alpha(e)$ is then used to induce the blue surfaces. Notice that because $\delta\alpha = 0$, the sum of the surface values near a purple line is 0.

0, because $\delta\alpha$ is obtained by applying the co-res map on $\alpha(e)$, which is exactly applying h_i^T on $\alpha(e)$. Finally, the X stabilizer generators correspond to logical operators that are supported locally on the segments around a vertex.

A similar mapping exists for 3D cubical codes $\mathcal{C}(G, \{A_1, A_2, A_3\}, \{C_1, C_2, C_3\})$. Given a vector on the qudits $\alpha \in \mathcal{C}^1(G, \{A_1, A_2, A_3\}, \{C_1, C_2, C_3\})$. The local coefficient on each edge e, $\alpha(e) \in \mathcal{F}_e$, corresponds to a function on the Δ^2 adjacent cubes $(h_i^T \otimes h_j^T)\alpha(e) \in \mathcal{G}_e \subseteq \mathbb{F}_q^{\Delta^2}$. These values can be attached to the faces spanned by the midpoint of the edge e, the two center of the corresponding neighboring faces, and the center of the corresponding cube. See Figure 3.6.

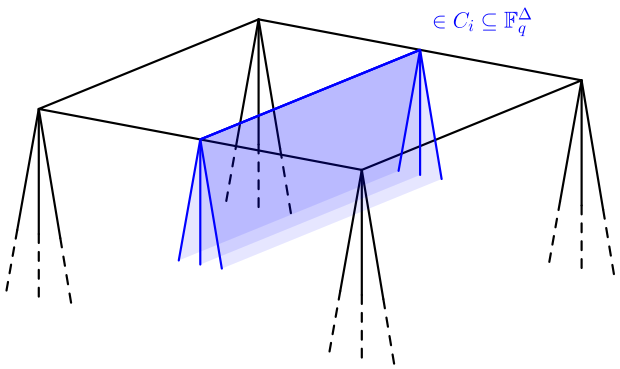

Figure 3.7: There are Δ cubes incident to a top boundary face. Each such cube carries a value. These values form a codeword in C_i.

Similarly, we observe three properties of this mapping. First, the surface pattern near the boundary belongs to the local codes C_i. See Figure 3.7. The Δ^2 values induced from an edge belongs to $C_i \otimes C_j$. When viewed from the 2D boundaries perpendicular to direction i, this splits into Δ vectors in C_i on Δ boundaries. Second, the condition of a X logical operator $\delta\alpha = 0$, corresponds to the condition that the sum of the surface values near an edge connecting the center of a face (Z check) and the center of the cube is 0. See Figure 3.6. Finally, the X stabilizer generators correspond to logical operators that are supported locally on the surfaces around a vertex. This description leads to an explicit way to view X logical operators as surfaces, which will be used to define the trilinear function f in the next section.

3.5 Transversal CCZ operations on 3D cubical codes

The goal of this section is to construct a trilinear function f for the 3D cubical code, which as discussed in Section 3.3, leads to transversal CCZ operations. This construction is based on the geometrical interpretation of logical operators described in the previous section. We first present a simpler case for the 2D square code, where we construct a bilinear function that is defined on the cohomology classes. We then generalize this idea

to the 3D cubical code. Since the discussion involves three codes, we will label the three codes using $1, 2, 3$, and label the directions using u, v, w in this section.

3.5.1 Warm up: a bilinear map on a 2D square code

Let X be the square complex $X(G, \{A_u, A_v\})$. Consider two quantum codes based on the square complex

$$Q_1 = Q(X, \{C_u^1, C_v^1\}, \ell = 1), \quad Q_2 = Q(X, \{C_u^2, C_v^2\}, \ell = 1). \tag{3.42}$$

The goal is to construct a bilinear map $f : \mathcal{C}^1(X, \{C_u^1, C_v^1\}) \times \mathcal{C}^1(X, \{C_u^2, C_v^2\}) \to \mathbb{F}_q$ such that

$$f(\zeta_1, \zeta_2) = f(\zeta_1 + \beta_1, \zeta_2 + \beta_2) \tag{3.43}$$

for all $\zeta_i \in Z^1(X, \{C_u^i, C_v^i\})$ and $\beta_i \in B^1(X, \{C_u^i, C_v^i\})$.

To achieve this, we require that $C_u^1 \odot C_u^2$ and $C_v^1 \odot C_v^2$ are both contained within the parity check code, where $C^1 \odot C^2 = \text{span}\{c_1 \odot c_2 : c_1 \in C^1, c_2 \in C^2\}$ is the vector space generated by $c_1 \odot c_2$ and $c_1 \odot c_2$ is the entrywise product of c_1, c_2. That means, for every $c_1 \in C_u^1, c_2 \in C_u^2$, we have $\sum_{i=1}^{\Delta} (c_1)_i (c_2)_i = 0$ where $(c_1)_i$ is the i-th entry of c_1.

We now construct the bilinear map $f : \mathcal{C}^1(X, \{C_u^1, C_v^1\}) \times \mathcal{C}^1(X, \{C_u^2, C_v^2\}) \to \mathbb{F}_q$. The goal is to map the vectors $\alpha_1 \in \mathcal{C}^1(X, \{C_u^1, C_v^1\})$ and $\alpha_2 \in \mathcal{C}^1(X, \{C_u^2, C_v^2\})$ into geometric strings, as described in Section 3.4.1, and then define f as the intersection number of these strings. The appeared difficulty is that the strings overlap, making it unclear how to define their intersection number directly. The solution is well-understood in mathematics, which is to perturb the strings into a "generic position". In our case, we can achieve this by shifting the string as shown in Figure 3.8. Note that these strings have associated values attached. When computing the intersection, we multiply the attached

values. By summing up the contributions from all faces, we have the definition

$$f(\alpha_1, \alpha_2) := \sum_{g\in G, a_u\in A_u, a_v\in A_v} \text{co-res}_{e_{*1},sq}(\alpha_1(e_{*1}))\ \text{co-res}_{e_{0*},sq}(\alpha_2(e_{0*})) + \text{co-res}_{e_{1*},sq}(\alpha_1(e_{1*}))\ \text{co-res}_{e_{*0},sq}(\alpha_2(e_{*0})) \tag{3.44}$$

where $e_{*0} = (g; a_u, 0)$, $e_{0*} = (g; 0, a_v)$, $e_{*1} = (a_v g; a_u, 1)$, $e_{1*} = (a_u g; 1, a_v)$, $sq = (g; a_u, a_v)$.[1]

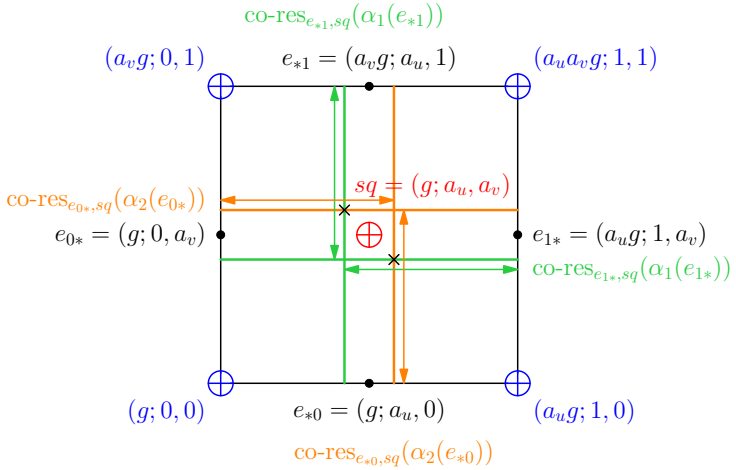

Figure 3.8: The figure shows two quantum codes constructed based on a square complex, $Q_1 = Q(X, \{C_u^1, C_v^1\}, \ell = 1)$ and $Q_2 = Q(X, \{C_u^2, C_v^2\}, \ell = 1)$. The associated Pauli X operators are colored in green and in orange. On each face, there are two contributions to the intersection number marked with a cross. These corresponds to the two terms $\text{co-res}_{e_{*1},sq}(\alpha_1(e_{*1}))\,\text{co-res}_{e_{0*},sq}(\alpha_2(e_{0*}))$ and $\text{co-res}_{e_{1*},sq}(\alpha_1(e_{1*}))\,\text{co-res}_{e_{*0},sq}(\alpha_2(e_{*0}))$.

What remains to be done is to verify that f is well-defined on the cohomology classes, meaning that it satisfies Equation (3.43). This is where the assumption that both $C_a^1 \odot C_a^2$ and $C_b^1 \odot C_b^2$ are contained within the parity check code comes into play. The following discussion aims to provide intuition for this result, while the formal proof is deferred to Section 3.6.

[1]For general finite fields $\mathbb{F}_q$ with characteristic other than 2, the definition has to be modified to incorporate certain negative signs.

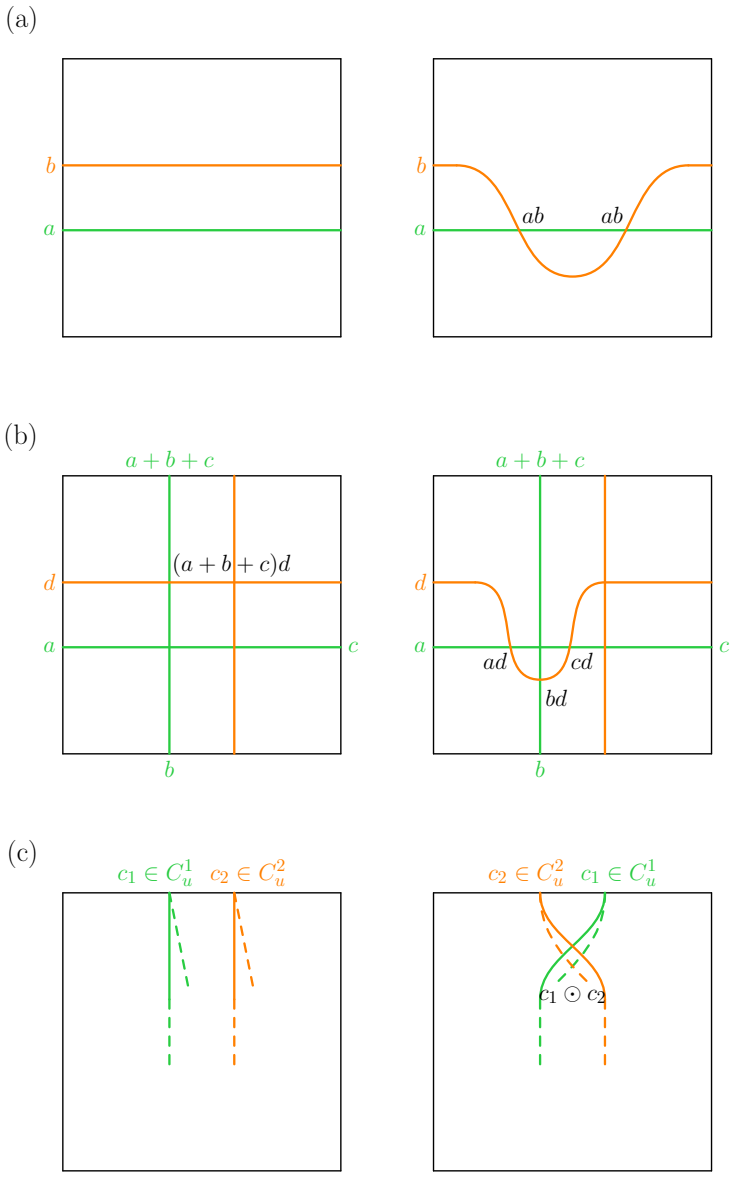

Figure 3.9: The three main types of topological deformations. We observe that the intersection number is invariant under those deformations. This explains why f is well-defined on the cohomology classes.

In the language of topology, the condition in Equation (3.43) is saying that the intersection number of two logical operators ζ_1, ζ_2 is invariant under topological deformations. There are mainly three types of deformations, as illustrated in Figure 3.9. The invariance of the first type holds for the simple reason that the intersections have the same value. The invariance of the second type follows from the property that, for logical operators, the four string values on a face sums to 0. The invariance of the third type is guaranteed by the fact that the string patterns near the boundary are elements of the local codes, $c_1 \in C_u^1, c_2 \in C_u^2$. The change in the intersection number is given by $\sum_{i=1}^{\Delta}(c_1)_i(c_2)_i$, which equals to 0 because $C_u^1 \odot C_u^2$ is contained within the parity check code. We also require $C_v^1 \odot C_v^2$ to be contained within the parity check code to ensure invariance for similar topological deformations along the vertical boundaries.

3.5.2 A trilinear map on a 3D cubical code

The same idea of using intersections to define f can be extended to 3D cubical codes. Let X be the cubical complex $X(G, \{A_u, A_v, A_w\})$. Consider three quantum codes based on the cubical complex

$$\begin{aligned} Q_1 &= Q(X, \{C_u^1, C_v^1, C_w^1\}, \ell = 1), \\ Q_2 &= Q(X, \{C_u^2, C_v^2, C_w^2\}, \ell = 1), \\ Q_3 &= Q(X, \{C_u^3, C_v^3, C_w^3\}, \ell = 1). \end{aligned} \tag{3.45}$$

The goal is to construct a trilinear map $f : \mathcal{C}^1(X, \{C_u^1, C_v^1, C_w^1\}) \times \mathcal{C}^1(X, \{C_u^2, C_v^2, C_w^2\}) \times \mathcal{C}^1(X, \{C_u^3, C_v^3, C_w^3\}) \to \mathbb{F}_q$ such that

$$f(\zeta_1, \zeta_2, \zeta_3) = f(\zeta_1 + \beta_1, \zeta_2 + \beta_2, \zeta_3 + \beta_3) \tag{3.46}$$

for all $\zeta_i \in Z^1(X, \{C_u^i, C_v^i, C_w^i\})$ and $\beta_i \in B^1(X, \{C_u^i, C_v^i, C_w^i\})$. As discussed in Section 3.3, these data lead to a transversal CCZ operation on the codes.

To achieve this, we require that $C_u^1 \odot C_u^2 \odot C_u^3$, $C_v^1 \odot C_v^2 \odot C_v^3$, $C_w^1 \odot C_w^2 \odot C_w^3$, are all contained within the parity check code. That means, for every $c_1 \in C_u^1, c_2 \in C_u^2, c_3 \in C_u^3$, we have $\sum_{i=1}^{\Delta}(c_1)_i(c_2)_i(c_3)_i = 0$ where $(c_1)_i$ is the i-th entry of c_1. This can be satisfied, for example, by taking low-degree Reed-Solomon codes over $\mathbb{F}_q$ with $\Delta = q$.

We can similarly construct the trilinear map f by mapping the vectors $\alpha_i \in \mathcal{C}^i(X, \{C_u^i, C_v^i, C_w^i\})$ into geometric surfaces, as described in Section 3.4.1, and then define f as the triple intersection number of these surfaces. We again shift the surfaces slightly to obtain the triple intersection number as illustrated in Figure 3.10. Overall, we have

$$
\begin{aligned}
f(\alpha_1, \alpha_2, \alpha_3) = \sum_{g \in G, a_u \in A_u, a_v \in A_v, a_w \in A_w} & \\
\text{co-res}_{e_{\dots},cu}(\alpha_1(a_v a_w g; a_u, 1, 1))\ \text{co-res}_{e_{\dots},cu}(\alpha_2(a_w g; 0, a_v, 1))\ \text{co-res}_{e_{\dots},cu}(\alpha_3(g; 0, 0, a_w)) & \\
+ \text{co-res}_{e_{\dots},cu}(\alpha_1(a_v a_w g; a_u, 1, 1))\ \text{co-res}_{e_{\dots},cu}(\alpha_2(a_v g; 0, 1, a_w))\ \text{co-res}_{e_{\dots},cu}(\alpha_3(g; 0, a_v, 0)) & \\
+ \text{co-res}_{e_{\dots},cu}(\alpha_1(a_u a_w g; 1, a_v, 1))\ \text{co-res}_{e_{\dots},cu}(\alpha_2(a_u g; 1, 0, a_w))\ \text{co-res}_{e_{\dots},cu}(\alpha_3(g; a_u, 0, 0)) & \\
+ \text{co-res}_{e_{\dots},cu}(\alpha_1(a_u a_w g; 1, a_v, 1))\ \text{co-res}_{e_{\dots},cu}(\alpha_2(a_w g; a_u, 0, 1))\ \text{co-res}_{e_{\dots},cu}(\alpha_3(g; 0, 0, a_w)) & \\
+ \text{co-res}_{e_{\dots},cu}(\alpha_1(a_u a_v g; 1, 1, a_w))\ \text{co-res}_{e_{\dots},cu}(\alpha_2(a_v g; a_u, 1, 0))\ \text{co-res}_{e_{\dots},cu}(\alpha_3(g; 0, a_v, 0)) & \\
+ \text{co-res}_{e_{\dots},cu}(\alpha_1(a_u a_v g; 1, 1, a_w))\ \text{co-res}_{e_{\dots},cu}(\alpha_2(a_u g; 1, a_v, 0))\ \text{co-res}_{e_{\dots},cu}(\alpha_3(g; a_u, 0, 0)) &
\end{aligned}
$$

where $e_{\dots}$ is the edge in the argument of α_i, and $cu = (g; a_u, a_v, a_w)$. For example, in the first term, $e_{\dots} = (a_v a_w g; a_u, 1, 1)$. Although these expressions are explicit, they are cumbersome to work with. Therefore, we will reformulate the definition of f more succinctly using the language of sheaves in Section 3.6.

Analogous to the discussion in the previous section, for f to be well-defined on the cohomology classes, we want the intersection number to remain invariant under of the topological deformations. One of the nontrivial deformations occurs near the boundary, where we deform the surface as shown in Figure 3.11. With some imagination, we observe that each cube introduces an extra triple intersection with the value $(c_1)_i(c_2)_i(c_3)_i$, where $c_1 \in C_u^1, c_2 \in C_u^2, c_3 \in C_u^3$ correspond to the patterns near the boundary. This means that

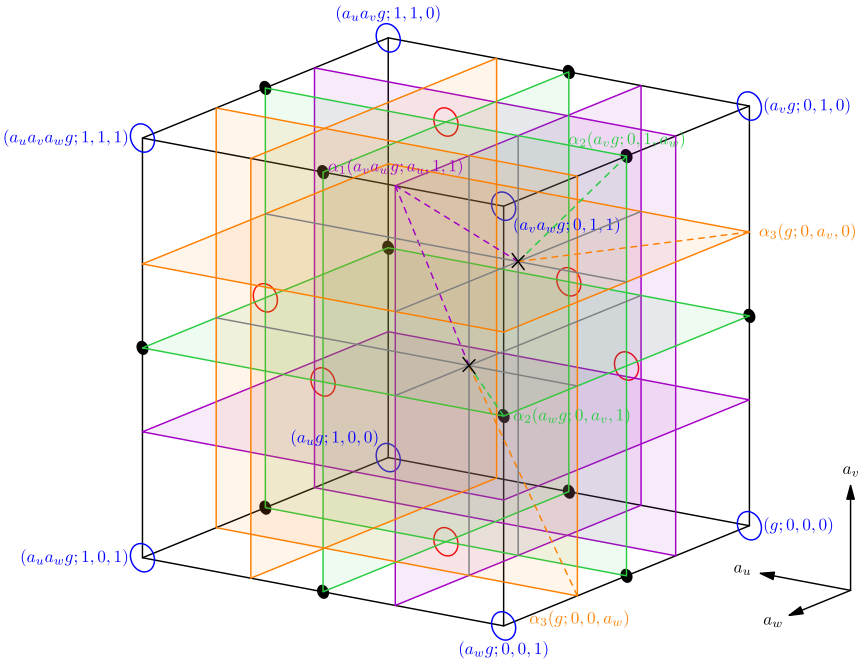

Figure 3.10: The figure shows three quantum codes constructed based on a cubical complex, Q_1, Q_2, Q_3. The associated Pauli X operators are colored in purple, green, and orange respectively. In each cube, there are a total of 6 intersections which contribute to the 6 terms in the expression of f. We draw the first two contributions in the figure marked by the crosses.

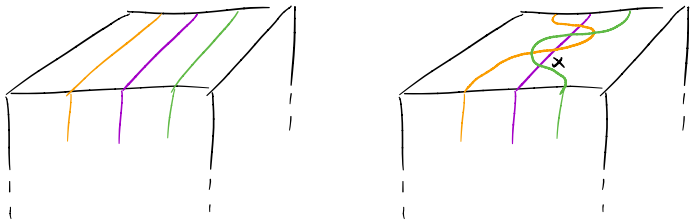

Figure 3.11: A topological deformation where the invariance of f relies on the assumption that $C_u^1 \odot C_u^2 \odot C_u^3$ is contained within the parity check code.

the intersection number is changed by $\sum_{i=1}^{\Delta}(c_1)_i(c_2)_i(c_3)_i$, which evaluates to 0 due to the assumption that $C_u^1 \odot C_u^2 \odot C_u^3$ is contained within the parity check code.

3.6 Transversal CCZ operations on sheaf codes

In this section, we prove the technical result that enables transversal CCZ operations on sheaf codes. This includes the cubical code discussed in the previous section. We begin by reviewing cellular complexes and sheaf complexes, This is followed by a discussion on how to map Pauli X operators geometrically. Finally, we formalize these intuitions by defining the cup product on all sheaf codes.

We remark that there have are several variants of sheaves that have been considered [26, 10, 16]. There is the vanilla sheaf code in [26, 10] and the Tanner sheaf code and maximally extendable sheaf code defined in [16]. In this work, we suggest a narrower definition of sheaf codes compared to the definitions in [26]. In fact, our definition almost agrees with the definition of Tanner sheaf code in [16], except a technical difference that will be explained later.

3.6.1 Cellular complexes

Let X be a t-dimensional cellular complex, where $X(i)$ denotes the set of i-cells. Given two cells σ, τ, we write $\sigma \preceq \tau$ if σ is a subcell of τ, and write $\sigma \prec \tau$ if σ is an 'immediate' subcell of τ where their dimensions differ by 1.

Mathematically speaking, it is easier to formalize the notion of a simplicial complexes, then obtain a cellular complex by merging the simplices. However, we will skip those details.

In the later discussion we assume that for every pair of i-cell σ and $(i+2)$-cell τ, such that $\sigma \prec \tau$, there are exactly two $(i+1)$-cells π such that $\sigma \prec \pi \prec \tau$. This simplifies the discussion on the gluability of a sheaf. This property generally holds with enough subdivision of a cellular complex and is true for the cubical complexes discussed earlier. In fact, this property can be shown rigorously for simplicial complexes. (Note that it is generally assumed for simplicial complexes that the vertices of each simplex are distinct.)

Claim 3.14. *Given a simplicial complex. For every pair of i-simplex σ and $(i+2)$-simplex τ, such that $\sigma \prec \tau$, there are exactly two $(i+1)$-simplices π such that $\sigma \prec \pi \prec \tau$.*

Proof. Let V be the vertices of σ and W be the vertices of τ. Since the vertices of each cell are distinct, $W \backslash V$ has size 2. Thus, the two $(i+1)$-cells π are specified by the vertices $V \cup \{v\}$ for $v \in W \backslash V$. □

We introduce some notations. $X_{\succcurlyeq \sigma}$ denotes the sub-complex consists of all cells above σ

$$X_{\succcurlyeq \sigma} = \{\tau \in X : \tau \succeq \sigma\}. \tag{3.47}$$

$X_{\preccurlyeq \sigma}$ denotes the sub-complex consists of all cells below σ

$$X_{\preccurlyeq \sigma} = \{\tau \in X : \tau \preceq \sigma\}. \tag{3.48}$$

We define $X_{\succcurlyeq \sigma}(k) = X_{\succcurlyeq \sigma} \cap X(k)$ and $X_{\preccurlyeq \sigma}(k) = X_{\preccurlyeq \sigma} \cap X(k)$.

For the purpose of constructing sheaves in the next section, it will be useful to keep the following analogies in mind. We think of each t-cell as a point, and each i-cell σ as an "open" set that contains the points $X_{\geqslant\sigma}(t)$. Therefore, the relation of set inclusion $X_{\geqslant\sigma}(t) \supseteq X_{\geqslant\tau}(t)$, corresponds to the relation $\sigma \preceq \tau$. Notice that these "open" sets do not form a standard topology, as they do not satisfy the property that the union of every family of open sets is open. Instead, they form the Alexandrov topology, which only requires that the intersection of every family of open sets is open. This perspective of Alexandrov topology is discussed in [26] and [16], although we do not use this perspective in the later discussions.

3.6.2 Sheaves on cellular complexes

Sheaf theory is a well-established topic in mathematics. Recently, chain complexes induced from sheaves have emerged as a natural framework for constructing codes [10, 19, 20, 16]. For readers unfamiliar with sheaf theory, it can be helpful to consult chapter 2 in [40] to understand the motivation behind sheaves, which are traditionally defined over topological spaces. It is known that sheaves over topological spaces have several categorical properties and is the right object to study in various settings.

It is important to note, however, that the definition of sheaves over cellular complexes differs somewhat from the traditional definition of sheaves over topological spaces. Therefore, one should be careful when drawing parallels between established results on sheaves over topological spaces and similar results for cellular complexes. In fact, the definition of sheaves on cellular complex introduced in [26], seems more similar to the traditional definition of presheaves in some respects[2], and one might wonder whether it is the right object to study

[2]The reason the author refers to it as a sheaf is because the definition satisfies certain categorical properties. As of my current understanding, the author may have implicitly applied certain sheafification to make the object more well-behaved.

in the context of error correcting codes.

These concerns motivate us to introduce additional conditions on sheaves over cellular complexes. The conditions we impose have certain analogs in the traditional definition of sheaves and could be useful for certain applications.

We begin by reviewing the definition of sheaves over cellular complexes as introduced in [26]. To distinguish this from our definition of sheaves, we refer to this structure as a "presheaf" with quotation marks. A "presheaf" $\mathcal{F}$ on a cellular complex X is the following data.

- For each i-cell $\sigma \in X(i)$, we have a vector space over $\mathbb{F}_q$, $\mathcal{F}_\sigma$. Recall that each cell corresponds to a set $X_{\geqslant\sigma}(t)$, and we call the elements of $\mathcal{F}_\sigma$ as sections of $\mathcal{F}$ over $X_{\geqslant\sigma}(t)$.
- For each inclusion $\sigma \preceq \tau$ (i.e. $X_{\geqslant\sigma}(t) \supseteq X_{\geqslant\tau}(t)$), we have a restriction map $\mathrm{res}_{\sigma,\tau} : \mathcal{F}_\sigma \to \mathcal{F}_\tau$.

The data is required to satisfy the following two conditions.

- The map $\mathrm{res}_{\sigma,\sigma}$ is the identity: $\mathrm{res}_{\sigma,\sigma} = \mathrm{id}_{\mathcal{F}_\sigma}$.
- If $\sigma \preceq \pi \preceq \tau$ are inclusions, then the restriction maps commute, i.e. $\mathrm{res}_{\sigma,\tau} = \mathrm{res}_{\pi,\tau}\,\mathrm{res}_{\sigma,\pi}$.

When the context is clear, we write $|_\tau$ in place of $\mathrm{res}_{\sigma,\tau}$.

A useful way to think about an element in $\mathcal{F}_\sigma$ is to view it as a function defined on $X_{\geqslant\sigma}(t)$. In particular, given $c \in \mathcal{F}_\sigma$, we can define the value of c at $\tau \in X_{\geqslant\sigma}(t)$ as $c(\tau) = \mathrm{res}_{\sigma,\tau}(c)$. This explains why we call the elements of $\mathcal{F}_\sigma$ as sections of $\mathcal{F}$ over $X_{\geqslant\sigma}(t)$. However, there is a caveat to this interpretation. In particular, the map from $\mathcal{F}_\sigma$ to functions on $X_{\geqslant\sigma}(t)$ is not necessarily injective for "presheaves". That means that some information in $\mathcal{F}_\sigma$ may be lost when viewed as a function on $X_{\geqslant\sigma}(t)$. This is one reason why

additional conditions are required to guarantee desireable behaviors for "presheaves". As we will see, the conditions we impose imply that the map is injective, so that the element in $\mathcal{F}_\sigma$ one-one corresponds to a function on $X_{\geqslant\sigma}(t)$.

Even though we refer to the object as a "presheaf", it differs from traditional presheaves over topological spaces in several ways. While there appears to be a dictionary between open sets in topological spaces and the regions $X_{\geqslant\sigma}(t)$ for some cell σ, this analogy has some subtleties. For example, in topology, the union of open sets is an open set, but the union of regions $X_{\geqslant\sigma}(t)$ is generally not of the form $X_{\geqslant\tau}(t)$. The presents a problem: a presheaf should assign a vector space to each open set, yet the current definition of a "presheaf" over a cellular complex does not assign a vector space to the union of such regions. This indicates that the definition of "presheaves" over cellular complexes lacks some nonlocal data compared to presheaves over topological spaces. This is why we put the quotation marks around "presheaves".

The definition of sheaves is designed to address these two issues. The first issue is resolved by the *identity axiom*, which ensures that sections are uniquely determined by their local values. The second issue is resolved by the *gluability axiom*, which states that local sections can be glued together to form sections over larger sets. We first review these two conditions for sheaves over topological spaces.

- Identity axiom: If $\{U_i\}_{i\in I}$ is an open cover of U, and $c_1, c_2 \in \mathcal{F}_U$, and $\mathrm{res}_{U,U_i} c_1 = \mathrm{res}_{U,U_i} c_2$ for all i, then $c_1 = c_2$.

- Gluability axiom: If $\{U_i\}_{i\in I}$ is an open cover of U, then given $c_i \in \mathcal{F}_{U_i}$ for all i, such that $\mathrm{res}_{U_i,U_i\cap U_j} c_i = \mathrm{res}_{U_j,U_i\cap U_j} c_j$ for all i, j, then there is some $c \in \mathcal{F}_U$ such that $\mathrm{res}_{U,U_i} c = c_i$ for all i.

These two conditions can be interpreted as the exactness of the following chain complex

$$0 \to \mathcal{F}_U \to \prod_{i \in I} \mathcal{F}_{U_i} \to \prod_{i,j \in I, i \neq j} \mathcal{F}_{U_i \cap U_j}. \tag{3.49}$$

The second map is $c \mapsto \{\mathrm{res}_{U,U_i}(c)\}_{i \in I}$ and the last map is $\{c_i\}_{i \in I} \mapsto \{c_i - c_j\}_{i,j \in I, i \neq j}$, where we impose an ordering on the elements in I. Identity is exactness at $\mathcal{F}_U$ and gluability is exactness at $\prod_i \mathcal{F}_{U_i}$.

In the context of topological spaces, these conditions must be satisfied by all possible open sets. However, for cellular complexes, the analogous notion of open sets is a bit different. So instead we impose the following condition: for every i-cell $\sigma \in X$ with $i \leqslant t-2$, the following chain complex is exact

$$0 \to \mathcal{F}_\sigma \to \prod_{\sigma_{i+1} \in X_{\geqslant \sigma}(i+1)} \mathcal{F}_{\sigma_{i+1}} \to \prod_{\sigma_{i+2} \in X_{\geqslant \sigma}(i+2)} \mathcal{F}_{\sigma_{i+2}} \tag{3.50}$$

and for i-cell $\sigma \in X$ with $i = t-1$, the following chain complex is exact

$$0 \to \mathcal{F}_\sigma \to \prod_{\sigma_{i+1} \in X_{\geqslant \sigma}(i+1)} \mathcal{F}_{\sigma_{i+1}}. \tag{3.51}$$

We will refer to the exactness at $\mathcal{F}_\sigma$ as the identity axiom on σ and the exactness at $\prod_{\sigma_{i+1} \in X_{\geqslant \sigma}(i+1)} \mathcal{F}_{\sigma_{i+1}}$ as the gluability axiom on σ.

The conditions we impose follow closely with the definition of sheaves over topological spaces, with the difference being that we now restrict ourselves to a specific family of "open covers" and a particular set of intersections between the "open sets". It is worth noting that in the traditional setting each intersection $U_i \cap U_j$ corresponds to exactly two open sets U_i and U_j, which ensures that the sections on U_i and U_j are consistent. Similarly, in our context, for every $(i+2)$-cell $\sigma_{i+2} \in X_{\geqslant \sigma}(i+2)$, there are exactly two $(i+1)$-cells σ_{i+1} such that $\sigma \prec \sigma_{i+1} \prec \sigma_{i+2}$ (discussed near Theorem 3.14), which guarantees that the section on the two $(i+1)$-cells are consistent. In fact, when X is a simplicial complex, and we take $\{\sigma_{i+1}\}_{\sigma_{i+1} \in X_{\geqslant \sigma}(i+1)}$ as an open cover of σ, the two expressions in Equation (3.50) and Equation (3.49) are exactly the same.

We now describe an explicit characterization of $\mathcal{F}$, as a consequence of the sheaf conditions.

Lemma 3.15. *For every cell $\sigma \in X$, the map $\mathcal{F}_\sigma \xrightarrow{\iota_\sigma} \prod_{\tau \in X_{\geqslant \sigma}(t)} \mathcal{F}_\tau$, which sends $c \mapsto \prod_{\tau \in X_{\geqslant \sigma}(t)} \mathrm{res}_{\sigma,\tau}(c)$, is injective.*

Furthermore, the restriction map $\mathrm{res}_{\sigma,\pi}$ corresponds to the restriction map on $X(t)$

$$\begin{array}{ccc} \mathcal{F}_\sigma & \xrightarrow{\iota_\sigma} & \prod_{\tau \in X_{\geqslant \sigma}(t)} \mathcal{F}_\tau \\ {\scriptstyle \mathrm{res}_{\sigma,\pi}}\downarrow & & \downarrow{\scriptstyle \mathrm{res}_{X_{\geqslant \sigma}(t), X_{\geqslant \pi}(t)}} \\ \mathcal{F}_\pi & \xrightarrow{\iota_\pi} & \prod_{\tau \in X_{\geqslant \pi}(t)} \mathcal{F}_\tau \end{array} \tag{3.52}$$

where $\mathrm{res}_{X_{\geqslant \sigma}(t), X_{\geqslant \pi}(t)}$ is the map that restricts the domain of the function from $X_{\geqslant \sigma}(t)$ to $X_{\geqslant \pi}(t)$.

Proof. We first show the injectivity of ι_σ through induction. For σ with dimension $i = t$, the result is straightforward. For σ with dimension $i = t-1$, the injectivity is implied from the identity axiom on σ.

Assuming the statement holds for cells with dimension greater than i, we want to show that the statement also holds for σ with dimension i. By the identity and gluability axioms on σ, we have

$$\mathcal{F}_\sigma \overset{g_\sigma}{\cong} \{\{c_\pi\}_{\pi \in X_{\geqslant \sigma}(i+1)} : \forall \pi, \pi' \in X_{\geqslant \sigma}(i{+}1), \rho \in X_{\geqslant \sigma}(i{+}2),\ c_\pi \in \mathcal{F}_\pi, \mathrm{res}_{\pi,\rho}(c_\pi) = \mathrm{res}_{\pi',\rho}(c_{\pi'})\} \tag{3.53}$$

where $c \mapsto \{\mathrm{res}_{\sigma,\pi}(c)\}_{\pi \in X_{\geqslant \sigma}(i+1)}$. The condition $\mathrm{res}_{\pi,\rho}(c_\pi) = \mathrm{res}_{\pi',\rho}(c_{\pi'})$ is imposed only when $\pi \prec \rho$ and $\pi' \prec \rho$. Denote the RHS as S.

To show that ι_σ is injective, we construct the inverse map $\mathrm{im}\,(\iota_\sigma) \to \mathcal{F}_\sigma$. Given $h \in \mathrm{im}\,(\iota_\sigma) \subseteq \prod_{\tau \in X_{\geqslant \sigma}(t)} \mathcal{F}_\tau$. Because $h(\tau) = \mathrm{res}_{\pi,\tau}(\mathrm{res}_{\sigma,\pi}(c))$ for all $\tau \succeq \pi$, we have $\prod_{\tau \in X_{\geqslant \pi}(t)} h(\tau) \in \iota_\pi(\mathcal{F}_\pi)$. By the induction hypothesis, ι_π is injective, which allows us to define $c_\pi = \iota_\pi^{-1}(\prod_{\tau \in X_{\geqslant \pi}(t)} h(\tau)) \in \mathcal{F}_\pi$ for every $\pi \in X_{\geqslant \sigma}(i+1)$. It is straightforward to check that $\{c_\pi\}_{\pi \in X_{\geqslant \sigma}(i+1)} \in S$. Therefore, we can define the candidate inverse map ι_σ^{-1} as

$\iota_\sigma^{-1}(h) = g_\sigma^{-1}(\{c_\pi\}_{\pi \in X_{\geqslant \sigma}(i+1)})$. It is straightforward to check that the composition $\iota_\sigma^{-1} \circ \iota_\sigma$ is the identity map on $\mathcal{F}_\sigma$.

The second statement about the restriction map is a direct consequence of the presheaf condition, $\mathrm{res}_{\pi,\tau} \mathrm{res}_{\sigma,\pi} = \mathrm{res}_{\sigma,\tau}$ for $\tau \in X(t)$. □

Therefore, from now on, we will ocacasionally abuse the notation by viewing $\mathcal{F}_\sigma$ as a subset of $\prod_{\tau \in X_{\geqslant \sigma}(t)} \mathcal{F}_\tau$ and interpreting elements of $\mathcal{F}_\sigma$ as functions on $X_{\geqslant \sigma}(t)$.

Lemma 3.16. *For every cell $\sigma \in X$,*

$$\mathcal{F}_\sigma = \{c \in \prod_{\tau \in X_{\geqslant \sigma}(t)} \mathcal{F}_\tau : \forall \tau' \in X_{\geqslant \sigma}(t-1),\ \mathrm{res}_{\sigma,\tau'}(c) \in \mathcal{F}_{\tau'}\}. \tag{3.54}$$

In particular, the sheaf $\mathcal{F}$ is completely determined by $\{\mathcal{F}_{\tau'}\}_{\tau' \in X(t-1)}$.

Proof. We again proof by induction. For σ with dimension $i = t-1$ or t the statement is clearly true.

Assuming the statement holds for cells with dimension greater than i, we want to show that the statement also holds for σ with dimension i. Denote the RHS of Equation (3.54) as T_σ. By the identity and gluability axioms on σ

$$\mathcal{F}_\sigma \overset{g_\sigma}{\cong} \{\{c_\pi\}_{\pi \in X_{\geqslant \sigma}(i+1)} : \forall \pi, \pi' \in X_{\geqslant \sigma}(i+1), \rho \in X_{\geqslant \sigma}(i+2),\ c_\pi \in \mathcal{F}_\pi, \mathrm{res}_{\pi,\rho}(c_\pi) = \mathrm{res}_{\pi',\rho}(c_{\pi'})\}. \tag{3.55}$$

Denote the RHS as S.

It is clear that $\mathcal{F}_\sigma \subseteq T_\sigma$, which is a direct consequence of the presheaf condition. What remains is to show that $T_\sigma \subseteq \mathcal{F}_\sigma$. Given $h \in T_\sigma$, we define $c_\pi = \prod_{\tau \in X_{\geqslant \pi}(t)} h(\tau)$ for every $\pi \in X_{\geqslant \sigma}(i+1)$. It is clear that $c_\pi \in T_\pi$. By the induction hypothesis, $T_\pi = \mathcal{F}_\pi$. Therefore, $\{c_\pi\}_{\pi \in X_{\geqslant \sigma}(i+1)} \in S$, which implies $h \in \mathcal{F}_\sigma$, as desired. □

In fact, the statements of Theorems 3.15 and 3.16 together is equivalent to the identity and gluability axioms. The means that we could have defined sheaves over

cellular complexes using the statements in the lemmas. We denote the unique sheaf with $\iota_\sigma(\mathcal{F}_\sigma) = C_\sigma \subseteq \mathbb{F}_q^{X_{\geqslant\sigma}(t)}$ as $\mathcal{F}(X, \{C_\sigma\}_{\sigma\in X(t-1)})$, We refer to C_σ as the local codes, which serves a role similar to the local codes used in Tanner codes.

Conceptually, these $\mathcal{F}_{\tau'}$ for $\tau' \in X(t-1)$ should be regarded as analogs of stalks. In the context of sheaf over topological spaces, it is known that "many properties of the sheaf are determined by its stalks". Similarly, in our context of cellular complex, we can adopt the same principle: "many properties of the global code are determined by its local codes".

This characterization of sheaves allows us to define the concept of a dual sheaf $\mathcal{F}^\perp$, where $\iota_\sigma((\mathcal{F}^\perp)_\sigma) = (\iota_\sigma(\mathcal{F}_\sigma))^\perp$ for $\sigma \in X(t-1)$. (This notion was also proposed in [16] using the notation $\mathcal{F}^T$.) As we will see, under suitable conditions, there exists a Poincaré duality between $\mathcal{F}$ and $\mathcal{F}^\perp$. This will allow us to identify the X and Z logical operators with the cohomologies of $\mathcal{F}$ and $\mathcal{F}^\perp$, respectively.

Interestingly, the two statements in the lemmas were used in [16] to define the Tanner sheaves. The only technical difference is that we allow $\mathcal{F}_\tau \neq \mathbb{F}_q$ for $\tau \in X(t)$. In [16], the statement of Theorem 3.15 is assumed as part of the definition of sheaves, and the statement of Theorem 3.16 is treated as an additional requirement specific for Tanner sheaves. In contrast, we argue that these two conditions probably should be included right from the beginning in the definition of sheaves over cellular complexes, as they are closely related to the identity and gluability axioms that are central to the definition of sheaves over topological spaces.

Remark 3.17. Our definition of sheaf depends on the value of t. In particular, we apply different conditions for $i = t-1$ and $t \leqslant t-2$. As a result, if we take the discrete union of a t-dimensional cellular complex X with sheaf $\mathcal{F}$ and a t'-dimensional cellular complex X' with sheaf $\mathcal{F}'$, the resulting structure is generally not a sheaf when $t \neq t'$, even though the discrete union of X and X' reamins a cellular complex. This contrasts with the definition of sheaves in [26]

We argue that this behavior is a feature, not a bug. In our context, even though we describe the objects through t-dimensional cellular complexes, the goal is to model t-dimensional manifolds. Similarly, the discrete union of two manifolds of different dimensions is no longer a manifold. In particular, this dependence on t is manifest in the statement of Poincaré duality, which we will establish in Section 3.8.

3.6.3 Sheaf cohomology and chain complexes

Given a sheaf $\mathcal{F}$ on a cellular complex X, we can construct a chain complex $\mathcal{C}(X, \mathcal{F})$ analogous to the construction of the cellular cohomology. This sheaf-induced chain complex has recently become a valuable tool for constructing classical and quantum codes [10, 19, 20, 16].

We first define the vector spaces. For $0 \leqslant i \leqslant t$, let $\mathcal{C}^i(X, \mathcal{F})$ be the space of cochains defined on $\sigma \in X(i)$ with values in $\mathcal{F}_\sigma$

$$\mathcal{C}^i(X, \mathcal{F}) = \bigoplus_{\sigma \in X(i)} \mathcal{F}_\sigma. \tag{3.56}$$

Otherwise, $\mathcal{C}^i(X, \mathcal{F}) = 0$. It is clear that $\mathcal{C}^i(X, \mathcal{F})$ forms a vector space, since each $\mathcal{F}_\sigma$ is a vector space.

We now define the coboundary maps $\delta^i : \mathcal{C}^i(X, \mathcal{F}) \to \mathcal{C}^{i+1}(X, \mathcal{F})$. Let $\alpha \in C^i$ be a cochain. To define $\delta^i(\alpha) \in \mathcal{C}^{i+1}(X, \mathcal{F})$, it is sufficient to specify $(\delta^i \alpha)(\tau) \in \mathcal{F}_\tau$ for every $\tau \in X(i+1)$. In particular, for $0 \leqslant i \leqslant t-1$ we define

$$\delta^i \alpha(\tau) = \sum_{\sigma \prec \tau} \mathrm{res}_{\sigma,\tau}(\alpha(\sigma)). \tag{3.57}$$

Otherwise, δ^i is defined as the zero map. (In general, the definition contains signs, $\sum_{\sigma \prec \tau} (-)^{(\dots)} \mathrm{res}_{\sigma,\tau}(\alpha(\sigma))$. We evade this technical discussion by studying finite fields with characteristic 2 where $-1 = 1$.)

We finally check that $\delta \circ \delta = 0$. We have

$$\delta^{i+1}\delta^i \alpha(\pi) = \sum_{\sigma \prec \tau \prec \pi} \mathrm{res}_{\sigma,\pi}(\alpha(\sigma)) = 0. \tag{3.58}$$

The last equality holds because for each pair $\sigma \in X(i), \pi \in X(i+2)$, there exist two $\tau \in X(i+1)$ such that $\sigma \prec \tau \prec \pi$, as discussed near Theorem 3.14.

Note that the standard cellular complex of X can also be viewed as a chain complex induced from a sheaf, where the sheaf consists of constant sections, i.e. each $\mathcal{F}_\sigma \cong \mathbb{F}_q$ forms a repetition code. So it makes sense to overload the notation and express the cellular complex as $\mathcal{C}(X, \mathbb{F}_q)$.

Using the chain complex $\mathcal{C}^\bullet(X, \mathcal{F})$, we can define the dual chain complex $\mathcal{C}_\bullet(X, \mathcal{F})$. These two chain complexes further induce the cohomology and homology groups, $H^\bullet(X, \mathcal{F})$ and $H_\bullet(X, \mathcal{F})$.

We observe a relation between the homology and cohomology of the sheaf and its dual sheaf.

Claim 3.18. *Given a sheaf $\mathcal{F}$ on a cellular complex X, we have*

$$H_t(X, \mathcal{F}) \cong H^0(X, \mathcal{F}^\perp). \tag{3.59}$$

Proof. Following from the characterization of sheaves, let $\mathcal{F} = \mathcal{F}(X, \{C_\sigma\}_{\sigma \in X(t-1)})$ and $\mathcal{F}^\perp = \mathcal{F}(X, \{C_\sigma^\perp\}_{\sigma \in X(t-1)})$. By definition, $H_t(X, \mathcal{F})$ is induced from the following part of the chain complex $0 \xrightarrow{\partial_{t+1}} \mathcal{C}_t(X, \mathcal{F}) \xrightarrow{\partial_t} \mathcal{C}_{t-1}(X, \mathcal{F})$ and $H^0(X, \mathcal{F}^\perp)$ is induced from the following part of the chain complex $0 \xrightarrow{\delta^{-1}} \mathcal{C}^0(X, \mathcal{F}^\perp) \xrightarrow{\delta^0} \mathcal{C}^1(X, \mathcal{F}^\perp)$. That means

$$H_t(X, \mathcal{F}) = \ker \partial_t = \{\alpha \in \prod_{\tau \in X(t)} \mathcal{F}_\tau : \forall \sigma \in X(t-1), \sum_{\tau \succ \sigma} \mathrm{res}^T_{\sigma,\tau}(\alpha(\tau)) = 0\} \tag{3.60}$$

and

$$H^0(X, \mathcal{F}^\perp) = \ker \delta^0 = \{\beta \in \prod_{\pi \in X(0)} \mathcal{F}_\pi : \forall \rho \in X(1), \sum_{\pi \prec \rho} \mathrm{res}_{\pi,\rho}(\beta(\pi)) = 0\}. \tag{3.61}$$

Recall that, in our context $\mathcal{F}_\sigma$ has a choice of basis, which allows us to define the transpose $\mathrm{res}^T_{\sigma,\tau} : \mathcal{F}_\tau \to \mathcal{F}_\sigma$. We would like to show that both are isomorphic to the set of global sections of $\mathcal{F}^\perp$

$$S = \{\alpha \in \prod_{\tau \in X(t)} \mathcal{F}_\tau : \forall \sigma \in X(t-1), \{\alpha(\tau)\}_{\tau \in X_{\geqslant \sigma}(t)} \in C^\perp_\sigma\}. \tag{3.62}$$

We begin by showing $H_t(X, \mathcal{F}) \cong S$. The key is to understand the meaning of $\sum_{\tau \succcurlyeq \sigma} \mathrm{res}^T_{\sigma,\tau}(\alpha(\tau)) = 0$. Let $\iota_\sigma : \mathcal{F}_\sigma \to \prod_{\tau \in X_{\geqslant \sigma}(t)} \mathcal{F}_\tau$ be the map such that $c \mapsto \prod_{\tau \in X_{\geqslant \sigma}(t)} \mathrm{res}_{\sigma,\tau}(c)$. Notice that

$$\sum_{\tau \succcurlyeq \sigma} \mathrm{res}^T_{\sigma,\tau}(\alpha(\tau)) = 0 \iff \iota^T_\sigma(\{\alpha(\tau)\}_{\tau \in X_{\geqslant \sigma}(t)}) = 0 \iff \{\alpha(\tau)\}_{\tau \in X_{\geqslant \sigma}(t)} \in C^\perp_\sigma. \tag{3.63}$$

The last iff holds because $C_\sigma = \mathrm{im}\, \iota_\sigma$.

We now show $H^0(X, \mathcal{F}^\perp) \cong S$. For each $\rho \in X(1)$, there are exactly two $\pi_0, \pi_1 \in X(0)$ such that $\pi_0, \pi_1 \prec \rho$. This implies that the restrictions on ρ agree, $\mathrm{res}_{\pi_0,\rho}(\beta(\pi_0)) = \mathrm{res}_{\pi_1,\rho}(\beta(\pi_1))$. Similar to the proof in Theorem 3.16, one can show that for each $\tau \in X(t)$, the value $\mathrm{res}_{\pi,\tau}(\beta(\pi))$ is independent of $\pi \in X(0)$. This allows us to map β to the function $\alpha \in \prod_{\tau \in X(t)} \mathcal{F}_\tau$, where $\alpha(\tau) = \mathrm{res}_{\pi,\tau}(\beta(\pi))$ for any choice of $\pi \in X(0)$. The image α satisfies $\{\alpha(\tau)\}_{\tau \in X_{\geqslant \sigma}(t)} \in C^\perp_\sigma$ for all $\sigma \in X(t-1)$ because $\beta(\pi) \in \mathcal{F}^\perp_\pi$ and its restriction $\mathrm{res}_{\pi,\sigma} \in \mathcal{F}^\perp_\sigma$. It is straightforward to check that this is an isomorphism. (This relation between $H^0(X, \mathcal{F}^\perp)$ and the global sections holds generally for sheaves.) □

The equation $H_t(X, \mathcal{F}) \cong H^0(X, \mathcal{F}^\perp)$ is reminiscent of Poincaré duality. In fact, as we will show in Section 3.8, under suitable assumptions detailed in Section 3.7, this relation holds more generally, with $H_{t-i}(X, \mathcal{F}) \cong H^i(X, \mathcal{F}^\perp)$ for $0 \leqslant i \leqslant t-1$.

3.6.4 Quantum LDPC codes based on sheaves

We are ready to describe a family of qLDPC codes based on Tanner sheaf complexes. We will later show that they have a natural geometric interpretation and a natural cup

product.

Recall that a quantum CSS code is equivalent to a three-term chain complex. The previous section provides a chain complex $\mathcal{C}(X, \{C_\sigma\}_{\sigma \in X(t-1)})$. So the only additional data is an integer ℓ which specifies which of the three terms to select. This leads to the quantum code $Q(X, \{C_\sigma\}_{\sigma \in X(t-1)}, \ell)$ defined by

$$\mathcal{C}^{\ell-1}(X, \{C_\sigma\}_{\sigma \in X(t-1)}) \to \mathcal{C}^{\ell}(X, \{C_\sigma\}_{\sigma \in X(t-1)}) \to \mathcal{C}^{\ell+1}(X, \{C_\sigma\}_{\sigma \in X(t-1)}). \tag{3.64}$$

The X distance is related to the small set coboundary expansion [41] of the above chain complex.

Following the discussions in Section 3.4, and more rigorously in Section 3.8[3], the Z distance is related to the small set coboundary expansion of another chain complex

$$\mathcal{C}^{t-\ell-1}(X, \{C_\sigma^\perp\}_{\sigma \in X(t-1)}) \to \mathcal{C}^{t-\ell}(X, \{C_\sigma^\perp\}_{\sigma \in X(t-1)}) \to \mathcal{C}^{t-\ell+1}(X, \{C_\sigma^\perp\}_{\sigma \in X(t-1)}). \tag{3.65}$$

We will not discuss how to prove small set coboundary expansion in this paper, but the tools in [20, Sec 6] are well-suited for this task.

3.6.5 A map from cochains to chains for sheaf complexes

Similar to Section 3.4.1, we map Pauli X operators to geometric objects which can be used to define the triple intersection number. More specifically, we will map cochains to chains. Strictly speaking, this discussion is not necessary, since we will take a different route and define the cup product directly. However, we still provide this discussion because these geometric objects are easier to visualized and are helpful to keep in mind. For simplicity, in this and the following sections, we will primarily focus on the case where $\mathcal{F}_\tau = \mathbb{F}_q$ for

[3]The discussion in Section 3.8 is not complete, as we have only show the correspondece at the level of homology. To establish the claimed result, we must also show a correspondence that includes a bound between the weight of the chain and cochain. Nevertheless, a similar setting for cubical complexes was addressed in [20, Sec 7] and it is possible to generalize that proof to general cellular complexes.

all $\tau \in X(t)$. In Section 3.6.9, we will extend the discussion to the case where $\mathcal{F}_\tau$ is not necessarily $\mathbb{F}_q$.

For $0 \leqslant i \leqslant t$, we will map every cochain $\alpha \in \mathcal{C}^i(X, \mathcal{F})$, to a chain on the "interior" of the nerve of X, denoted as $\tilde{\alpha} \in \mathcal{C}_{t-i}(\mathcal{N}_{int}(X), \mathbb{F}_q)$. The nerve of X, denoted as $\mathcal{N}(X)$, is a simplicial complex where the vertices correspond to the cells in X, and the edges correspond to pairs of cells related by inclusion $\{(\sigma, \tau) : \sigma \precneqq \tau\}$. More generally, the i-simplices of $\mathcal{N}(X)$ correspond to sequences of cell inclusions of length $i+1$, $\{(\sigma_0, ..., \sigma_i) : \sigma_0 \precneqq ... \precneqq \sigma_i\}$. The boundary of the nerve $\mathcal{N}_\partial(X)$ consists of simplices that lie on the $(t-1)$-skeleton of X. The interior of the nerve $\mathcal{N}_{int}(X)$ consists of simplices in $\mathcal{N}(X) \backslash \mathcal{N}_\partial(X)$. Formally, the chains in $\mathcal{C}_\bullet(\mathcal{N}_{int}(X), \mathbb{F}_q)$ can be thought of as the quotient $\mathcal{C}_\bullet(\mathcal{N}(X), \mathbb{F}_q)/\mathcal{C}_\bullet(\mathcal{N}_\partial(X), \mathbb{F}_q)$, as in the construction of relative homology.

We are ready to define the map from a cochain $\alpha \in \mathcal{C}^i(X, \mathcal{F})$ to a chain $\tilde{\alpha} \in \mathcal{C}_{t-i}(\mathcal{N}_{int}(X), \mathbb{F}_q)$. Given a pair of cells $\sigma \in X(i), \tau \in X(t)$ with $\sigma \preceq \tau$. Let $S(\sigma, \tau)$ be the collection of $(t-i)$-simplices, $\{(\sigma_0, ..., \sigma_{t-i}) : \sigma_0 \precneqq ... \precneqq \sigma_{t-i}\}$ with $\sigma_0 = \sigma$ and $\sigma_{t-i} = \tau$. We assign the value $\text{res}_{\sigma,\tau}(\alpha(\sigma)) \in \mathbb{F}_q$ to every simplex in $S(\sigma, \tau)$. Formally,

$$\tilde{\alpha} = \sum_{\sigma \in X(i), \tau \in X(t)} \sum_{\pi \in S(\sigma,\tau)} \text{res}_{\sigma,\tau}(\alpha(\sigma))\pi. \tag{3.66}$$

Notice that $\tilde{\alpha}$ do not involve any simplices that lie on the boundary $\mathcal{N}_\partial(X)$.

This construction is consistent with the one used for cubical complexes discussed in Section 3.4.1. Similar to the case of a cubical complex, the corresponding $(t-i)$-dimensional surface $\tilde{\alpha}$ ends on the boundaries in a specific way, where the values near a boundary $\tau' \in X(t-1)$ form a vector in $C_{\tau'} \subseteq \mathbb{F}_q^{X_{\succcurlyeq \tau'}(t)}$.

This geometric perspective helps clarify the upcoming discussion of the cup product. The dual of the cup product is the intersection. Consider a $(t-i)$-dimensional surface $\tilde{\alpha}_1$ associated with $\alpha_1 \in \mathcal{C}^i(X, \{C_{1,\tau'}\}_{\tau' \in X(t-1)})$ and a $(t-j)$-dimensional surface $\tilde{\alpha}_2$ associated with $\alpha_2 \in \mathcal{C}^j(X, \{C_{2,\tau'}\}_{\tau' \in X(t-1)})$. The intersection of $\tilde{\alpha}_1$ and $\tilde{\alpha}_2$ is a $(t-i-j)$-dimensional

surface $\tilde{\alpha}$. Furthermore, the values near the boundary of $\tilde{\alpha}$ are derived from the product of the boundary values of $\tilde{\alpha}_1$ and $\tilde{\alpha}_2$, which lie in $C_{1,\tau'} \odot C_{2,\tau'}$, where $C_1 \odot C_2 = \{c_1 \odot c_2 : c_1 \in C_1, c_2 \in C_2\}$ and $c_1 \odot c_2$ is the entrywise product of c_1 and c_2. By taking the dual from the surface $\tilde{c}$ back to the cochain c, this suggests that the the cup product of c_1 and c_2 should lie in $\mathcal{C}^{i+j}(X, \{C_{1,\tau'} \odot C_{2,\tau'}\}_{\tau' \in X(t-1)})$. This is exactly what we will show in the remaining sections.

3.6.6 Cup product on a sheaf complex over a simplicial complex

We first consider a simpler case where X is a simplicial complex. In the next section, we will triangulate general cellular complex, which allows us to reduce the problem to the case of simplicial complex.

The goal is to define a cup product for the cohomology classes

$$H^i(X, \mathcal{F}_1) \times H^j(X, \mathcal{F}_2) \xrightarrow{\cup} H^{i+j}(X, \mathcal{F}_1 \odot \mathcal{F}_2) \tag{3.67}$$

where $\mathcal{F}_1 \odot \mathcal{F}_2 = \mathcal{F}(X, \{C_{1,\sigma} \odot C_{2,\sigma}\}_{\sigma \in X(t-1)})$ where $C_1 \odot C_2 = \text{span}\{c_1 \odot c_2 : c_1 \in C_1, c_2 \in C_2\}$ is the vector space spanned by $c_1 \odot c_2$ and $c_1 \odot c_2$ is the entrywise product of c_1, c_2. To do so, we will first define a cup product for the cochains

$$\mathcal{C}^i(X, \mathcal{F}_1) \times \mathcal{C}^j(X, \mathcal{F}_2) \xrightarrow{\cup} \mathcal{C}^{i+j}(X, \mathcal{F}_1 \odot \mathcal{F}_2) \tag{3.68}$$

and show that such map interacts well with the coboundary map.

Notice that the "presheaf" constructed from $\{\mathcal{F}_{1,\sigma} \odot \mathcal{F}_{2,\sigma}\}_{\sigma \in X}$ is generally different from the sheaf $\mathcal{F}_1 \odot \mathcal{F}_2$ defined earlier. In particular, while $\mathcal{F}_{1,\sigma} \odot \mathcal{F}_{2,\sigma} \subseteq (\mathcal{F}_1 \odot \mathcal{F}_2)_\sigma$, they are not necessarily equal. This parallels the situation the traditional sheaf theory, where one often needs to apply sheafification after performing elementary operations.

Recall that for two cochains on a simplicial complex $\alpha \in \mathcal{C}^i(X, \mathbb{F}_q), \beta \in \mathcal{C}^j(X, \mathbb{F}_q)$,

the cup product $\alpha \cup \beta \in \mathcal{C}^{i+j}(X, \mathbb{F}_q)$ is defined by

$$(\alpha \cup \beta)([v_0, ..., v_i, ..., v_{i+j}]) = \alpha([v_0, ..., v_i])\beta([v_i, ..., v_{i+j}]) \tag{3.69}$$

for every $(i+j)$-simplex $[v_0, ..., v_i, ..., v_{i+j}]$ in X, where $[v_0, ..., v_i]$ denotes the simplex with vertices $v_0, ..., v_i$. Notice that the definition assumes an ordering of the vertices.

The idea behind defining the new cup product is simple. We combine the contributions from each simplex analogous to how we combine the values from each intersection. Since the cup product on each simplex depends on the vertex ordering, we assign a global ordering to all vertices in X from the start. This ensures a consistent ordering among all simplices, which will be important later on.

Given $\alpha \in \mathcal{C}^i(X, \mathcal{F})$, let $\alpha_{\leqslant\tau} : X_{\leqslant\tau}(i) \to \mathbb{F}_q$ be the map where $\sigma \mapsto \alpha(\sigma)|_\tau$. (Recall that $c|_\tau$ is a shorthand for $\mathrm{res}_{\sigma,\tau}(c)$) This local component $\alpha_{\leqslant\tau}$ captures the $(t-i)$-dimensional surface structure in the t-simplex τ. It is clear that there is a bijection between $\alpha \in \mathcal{C}^i(X, \mathcal{F}^{full})$ and $\{\alpha_{\leqslant\tau}\}_{\tau \in X(t)} \in \prod_{\tau \in X(t)}(X_{\leqslant\tau}(i) \to \mathbb{F}_q)$ where $\mathcal{F}^{full}_\sigma = (X_{\geqslant\sigma}(t) \to \mathbb{F}_q)$ could be any local section. That means given the values $\{\alpha_{\leqslant\tau}\}_{\tau \in X(t)}$, we can combine them into an element in $\mathcal{C}^i(X, \mathcal{F}^{full})$.

Notice that $\alpha_{\leqslant\tau} : X_{\leqslant\tau}(i) \to \mathbb{F}_q$ can be interpreted as a i-cochain on the simplex $X_{\leqslant\tau}$, since $(X_{\leqslant\tau}(i) \to \mathbb{F}_q) \cong \mathcal{C}^i(X_{\leqslant\tau}, \mathbb{F}_q)$. Thus, we define $\alpha \cup \beta$ as

$$(\alpha \cup \beta)_{\leqslant\tau} = (\alpha_{\leqslant\tau}) \cup (\beta_{\leqslant\tau}), \tag{3.70}$$

where $(\alpha_{\leqslant\tau}) \cup (\beta_{\leqslant\tau})$ is the usual cup product $\mathcal{C}^i(X_{\leqslant\tau}, \mathbb{F}_q) \times \mathcal{C}^j(X_{\leqslant\tau}, \mathbb{F}_q) \xrightarrow{\cup} \mathcal{C}^{i+j}(X_{\leqslant\tau}, \mathbb{F}_q)$.

At this moment, we only know that $\alpha \cup \beta \in \mathcal{C}^{i+j}(X, \mathcal{F}^{full})$. We now show that $\alpha \cup \beta \in \mathcal{C}^{i+j}(X, \mathcal{F}_1 \odot \mathcal{F}_2)$. All we need to show is that $(\alpha \cup \beta)|_\sigma \in \mathcal{F}_{1,\sigma} \odot \mathcal{F}_{2,\sigma}$, since $\mathcal{F}_{1,\sigma} \odot \mathcal{F}_{2,\sigma} \subseteq (\mathcal{F}_1 \odot \mathcal{F}_2)_\sigma$. We have

$$\begin{aligned}(\alpha\cup\beta)(\sigma) &= \{(\alpha\cup\beta)(\sigma)|_\tau\}_{\tau \in X_{\geqslant\sigma}(t)} = \{(\alpha\cup\beta)_{\leqslant\tau}(\sigma)\}_{\tau \in X_{\geqslant\sigma}(t)} = \{((\alpha_{\leqslant\tau})\cup(\beta_{\leqslant\tau}))(\sigma)\}_{\tau \in X_{\geqslant\sigma}(t)} \\ &= \{\alpha_{\leqslant\tau}(\sigma_1)\beta_{\leqslant\tau}(\sigma_2)\}_{\tau \in X_{\geqslant\sigma}(t)} = \{\alpha_{\leqslant\tau}(\sigma_1)\}_{\tau \in X_{\geqslant\sigma}(t)} \odot \{\beta_{\leqslant\tau}(\sigma_2)\}_{\tau \in X_{\geqslant\sigma}(t)}\end{aligned} \tag{3.71}$$

where the equalities in the first line follow from definitions. The first equality in the second line follows from the definition of the usual cup product. For $\sigma = [v_0, ..., v_i, ..., v_{i+j}]$, we define $\sigma_1 = [v_0, ..., v_i], \sigma_2 = [v_i, ..., v_{i+j}]$. Note that this step relies on having a global ordering of the vertices, which ensures the vertex ordering is independent of τ.

Notice that $\{\alpha_{\leqslant\tau}(\sigma_1)\}_{\tau\in X_{\geqslant\sigma}(t)}$ is a restriction of $\{\alpha_{\leqslant\tau}(\sigma_1)\}_{\tau\in X_{\geqslant\sigma_1}(t)} = \alpha(\sigma_1)$. Because $\alpha(\sigma_1) \in \mathcal{F}_{1,\sigma_1}$ and by the definition of the sheaf, the restriction to $X_{\geqslant\sigma}(t)$ is a vector in $\mathcal{F}_{1,\sigma}$, i.e. $\{\alpha_{\leqslant\tau}(\sigma_1)\}_{\tau\in X_{\geqslant\sigma}(t)} \in \mathcal{F}_{1,\sigma}$. Similarly, $\{\beta_{\leqslant\tau}(\sigma_2)\}_{\tau\in X_{\geqslant\sigma}(t)} \in \mathcal{F}_{2,\sigma}$. Therefore, $(\alpha \cup \beta)(\sigma) \in \mathcal{F}_{1,\sigma} \odot \mathcal{F}_{2,\sigma}$.

We have established a cup product for the cochains $\mathcal{C}^i(X, \mathcal{F}_1) \times \mathcal{C}^j(X, \mathcal{F}_2) \xrightarrow{\cup} \mathcal{C}^{i+j}(X, \mathcal{F}_1 \odot \mathcal{F}_2)$. We now show that it interacts well with the coboundary operator and induces a cup product for the cohomology classes.

First, a simple lemma.

Lemma 3.19. *$(\delta\alpha)_{\leqslant\tau} = \delta(\alpha_{\leqslant\tau})$ for $\alpha \in \mathcal{C}^i(X, \mathcal{F})$.*

Proof. We simply apply the definitions

$$(\delta\alpha)_{\leqslant\tau}(\sigma) = (\delta\alpha)(\sigma)|_\tau = \sum_{\sigma'\prec\sigma} \alpha(\sigma')|_\tau = \sum_{\sigma'\prec\sigma} \alpha_{\leqslant\tau}(\sigma') = \delta(\alpha_{\leqslant\tau}) \tag{3.72}$$

where the first and third equalities follow from the definition of $(*)_{\leqslant\tau}$. The second equality follows from the definition of δ on $\mathcal{C}^i(X, \mathcal{F})$ and the forth equality follows from the definition of δ on $\mathcal{C}^i(X_{\leqslant\tau}, \mathbb{F}_q)$. □

This is the key formula.

Lemma 3.20. *$\delta(\alpha \cup \beta) = (\delta\alpha \cup \beta) + (\alpha \cup \delta\beta)$ for $\alpha \in \mathcal{C}^i(X, \mathcal{F}_1)$ and $\beta \in \mathcal{C}^j(X, \mathcal{F}_2)$.*

Proof. All we need to show is that $(\delta(\alpha \cup \beta))_{\leqslant\tau} = (\delta\alpha \cup \beta)_{\leqslant\tau} + (\alpha \cup \delta\beta)_{\leqslant\tau}$ for all $\tau \in X(t)$.

We have

$$(\delta(\alpha \cup \beta))_{\leqslant \tau} = \delta((\alpha \cup \beta)_{\leqslant \tau}) = \delta(\alpha_{\leqslant \tau} \cup \beta_{\leqslant \tau}) = (\delta(\alpha_{\leqslant \tau}) \cup \beta_{\leqslant \tau}) + (\alpha_{\leqslant \tau} \cup \delta(\beta_{\leqslant \tau}))$$
$$= ((\delta\alpha)_{\leqslant \tau} \cup \beta_{\leqslant \tau}) + (\alpha_{\leqslant \tau} \cup (\delta\beta)_{\leqslant \tau}) = (\delta\alpha \cup \beta)_{\leqslant \tau} + (\alpha \cup \delta\beta)_{\leqslant \tau} \quad (3.73)$$

where the first and fifth equalities follow from Lemma 3.19, and the second and forth equalities follow from the definition of the new cup product. The third equality is the nontrivial result that holds for the usual cup product (see, for example, [42, Lemma 3.6]). □

It is clear that if α and β are cocycles, then $\alpha \cup \beta$ is again a cocycle. Also, if one of α, β is a cocycle and the other is a coboundary, say $\delta\alpha = 0$ and $\beta = \delta\gamma$, then $\alpha \cup \beta = \alpha \cup \delta\gamma = \delta(\alpha \cup \gamma)$ is again a coboundary. This implies that there is an induced cup product on the cohomology classes $H^i(X, \mathcal{F}_1) \times H^j(X, \mathcal{F}_2) \xrightarrow{\cup} H^{i+j}(X, \mathcal{F}_1 \odot \mathcal{F}_2)$.

3.6.7 Cup product on a sheaf complex

For the purpose of constructing quantum code, we can actually stop here and use the result from the last section on simplicial complexes. The reason is that for any cellular complex X with sheaf $\mathcal{F}$, we can always subdivide it into a simplicial complex $\tilde{X}$ and extend the sheaf into $\tilde{\mathcal{F}}$ (see below). The codes defined on $(X, \mathcal{F})$ and $(\tilde{X}, \tilde{\mathcal{F}})$ are essentially equivalent. Therefore, the result from the last section is sufficient for all practical purpose. Nevertheless, since our discussion in Section 3.5 was on cubical complex, we will use this opportunity to demonstrate the tools to convert between simplicial and cellular complexes.

Given a cellular complex X and a sheaf $\mathcal{F}$ on X. Let $\tilde{X}$ be a simplicial complex that triangulates X. The goal is to use the cup products defined on $\mathcal{C}(\tilde{X}, \tilde{\mathcal{F}})$ and $H(\tilde{X}, \tilde{\mathcal{F}})$ in the previous section to induce cup products on $\mathcal{C}(X, \mathcal{F})$ and $H(X, \mathcal{F})$. The idea fairly simple: since X and $\tilde{X}$ have the same topology, the chain complexes defined on X and $\tilde{X}$ will be chain homotopy equivalent. However, explicitly constructing the maps between them is somewhat involved, and will form the bulk of this section.

Define the sheaf $\tilde{\mathcal{F}}$: We begin by defining $\tilde{\mathcal{F}}$. To achieve this, we need to specify a vector space $\tilde{\mathcal{F}}_\sigma \subseteq \mathbb{F}_q^{\tilde{X}_{\geqslant\sigma}(t)}$ for every $\sigma \in \tilde{X}$. The idea is that each t-cell in X is decomposed into many t-simplices in $\tilde{X}$. The vector space $\tilde{\mathcal{F}}_\sigma$ is induced from such a decomposition by duplicating the value in a t-cell to the t-simplices it contains. Let τ_σ be the smallest cell in X that contains σ. Notice that every t-simplices above σ is contained in some t-cells above τ_σ. Let $I : \tilde{X}_{\geqslant\sigma}(t) \to X_{\geqslant\tau_\sigma}(t)$ be the map that maps every t-simplex to the unique t-cell it is contained in. This induces the map $I_* : (X_{\geqslant\tau_\sigma}(t) \to \mathbb{F}_q) \to (\tilde{X}_{\geqslant\sigma}(t) \to \mathbb{F}_q)$, which allows us to define $\tilde{\mathcal{F}}_\sigma$ as the image $I_*(\mathcal{F}_{\tau_\sigma})$. It is not hard to check that $\mathcal{F}_{\tau_\sigma} \cong \tilde{\mathcal{F}}_\sigma$. Essentially, the map duplicates the values in $X_{\geqslant\tau_\sigma}(t)$ to the values $\tilde{X}_{\geqslant\sigma}(t)$. It is also not hard to check that this duplication operation commutes with the entrywise product, $\widetilde{\mathcal{F}_1 \odot \mathcal{F}_2} = \tilde{\mathcal{F}}_1 \odot \tilde{\mathcal{F}}_2$.

Define the maps g and $\tilde{g}$ between the topological spaces X and $\tilde{X}$: Let X_i and $\tilde{X}_i$ be the i-skeleton of X and $\tilde{X}$. We now view X and $\tilde{X}$ as topological spaces and construct two maps $g : X \to \tilde{X}$ and $\tilde{g} : \tilde{X} \to X$, such that $\tilde{g}g$ and $g\tilde{g}$ are homotopic to the identity maps on X and $\tilde{X}$, respectively. Furthermore, we require that the i-skeleton is map into the i-skeleton, $g(X_i) \subseteq \tilde{X}_i, \tilde{g}(\tilde{X}_i) \subseteq X_i$. This allows us to use g and $\tilde{g}$ to naturally induce the maps between the chains (cochains) of X and $\tilde{X}$. Additionally, for each simplex $\sigma \in \tilde{X}$, we want $\tilde{g}(\sigma) \subseteq \tau_\sigma$ where $\tau_\sigma \in X$ is the smallest cell in X that contains σ. This guarantees the correspondence between the local sections.

Since $X_i \subseteq \tilde{X}_i$, we can naturally define g as such an embedding. Defining $\tilde{g}$ is more complicated and we will construct it iteratively. We begin by defining $\tilde{g}$ on the 0-cells. For each $\sigma \in \tilde{X}(0)$, we map σ to a vertex in τ_σ. Now, assume that $\tilde{g} : \tilde{X}_i \to X_i$ is defined. Our goal is to extend the definition of $\tilde{g}$ to every $(i+1)$-simplex $\sigma \in \tilde{X}(i+1)$. Given the image on the boundary $\tilde{g}(\partial\sigma)$, we want to ensure that $\tilde{g}(\sigma)$ lies within τ_σ. By the induction hypothesis, it is known that $\tilde{g}(\partial\sigma)$ is supported on τ_σ. Since τ is contractable, $\tilde{g}(\partial\sigma)$ is contractable on the $(i+1)$-skeleton of τ_σ, $(\tau_\sigma)_{i+1}$. Such a contraction induces a map from

σ into $(\tau_\sigma)_{i+1}$, thereby defining $\tilde{g}$ on σ. By continuing this process, we obtain $\tilde{g} : \tilde{X} \to X$.

We check that $\tilde{g}g$ and $g\tilde{g}$ are homotopic to the identity maps. It is clear that $\tilde{g}g$ is the the identity map on X. For $g\tilde{g}$, one can again build the homotopy map iteratively, which we will skip the details.

Define the induced chain and cochain maps: With g and $\tilde{g}$, we can define the corresponding chain maps $g_i : \mathcal{C}_i(X, \mathbb{Z}) \to \mathcal{C}_i(\tilde{X}, \mathbb{Z})$ and $\tilde{g}_i : \mathcal{C}_i(\tilde{X}, \mathbb{Z}) \to \mathcal{C}_i(X, \mathbb{Z})$. Notice that, by construction, the images of g and $\tilde{g}$ for each cell is a combination of cells, which leads to the map on the chains. In particular, g_i can be explicitly expressed as, $g_i(\sigma) = \sum_{\tilde{\sigma} \in \tilde{X}(i), \tilde{\sigma} \subseteq \sigma} \tilde{\sigma}$, where $\tilde{\sigma}$ are the simplices of the same dimension that form the cell σ. For $\tilde{g}_i(\tilde{\sigma})$, it is harder to express it explicitly, but we have the property that $\tilde{g}_i(\tilde{\sigma})$ is only supported on the chain within $\tau_{\tilde{\sigma}}$. Since these chain maps are induced from maps on the topological spaces, they naturally commute with the boundary operators.

The chain maps $g_i, \tilde{g}_i$ naturally induce the maps on the cochains $g^i : \mathcal{C}^i(\tilde{X}, \tilde{\mathcal{F}}) \to \mathcal{C}^i(X, \mathcal{F})$, $\tilde{g}^i : \mathcal{C}^i(X, \mathcal{F}) \to \mathcal{C}^i(\tilde{X}, \tilde{\mathcal{F}})$. For g^i, we define

$$g^i(\tilde{\alpha})(\sigma) = \sum_{\tilde{\sigma} \in g_i(\sigma)} I_*^{-1}(\tilde{\alpha}(\tilde{\sigma})). \tag{3.74}$$

Recall $I_* : \mathcal{F}_{\tau_{\tilde{\sigma}}} \to \mathcal{F}_{\tilde{\sigma}}$ is an isomorphism. Because $\tilde{\sigma} \subseteq \sigma$ and both have the same dimension, we have $\tau_{\tilde{\sigma}} = \sigma$. Therefore, $I_*^{-1}(\tilde{\alpha}(\tilde{\sigma})) \in \mathcal{F}_\sigma$ and the expression is well-defined. For $\tilde{g}^i$, we define

$$\tilde{g}^i(\alpha)(\tilde{\sigma}) = I_*\Big(\sum_{\sigma \in \tilde{g}_i(\tilde{\sigma})} \alpha(\sigma)|_{\tau_{\tilde{\sigma}}} \Big). \tag{3.75}$$

By the construction of $\tilde{g}$, σ is a cell in $\tau_{\tilde{\sigma}}$. This allows us to restrict $\alpha(\sigma) \in \mathcal{F}_\sigma$ to $\alpha(\sigma)|_{\tau_{\tilde{\sigma}}} \in \mathcal{F}_{\tau_{\tilde{\sigma}}}$. Note that the sum accounts for the multiplicities: if σ is covered twice in $\tilde{g}_i(\tilde{\sigma})$, then we consider two contributions from σ. Since these cochain maps are induced from maps on the topological spaces, they naturally commute with the coboundary operators.

Define the cup product: With the cochains maps $g^i, \tilde{g}^i$, we are ready to compose them together with the cup product on $\mathcal{C}(\tilde{X}, \tilde{F})$ to form the cup product on $\mathcal{C}(X, F)$,

$$\mathcal{C}^i(X,\mathcal{F}_1)\times\mathcal{C}^j(X,\mathcal{F}_2) \xrightarrow{\tilde{g}^{1,i}\times\tilde{g}^{2,j}} \mathcal{C}^i(\tilde{X},\tilde{\mathcal{F}}_1)\times\mathcal{C}^j(\tilde{X},\tilde{\mathcal{F}}_2) \xrightarrow{\cup} \mathcal{C}^{i+j}(\tilde{X},\tilde{\mathcal{F}}_1\odot\tilde{\mathcal{F}}_2) \xrightarrow{g^{i+j}} \mathcal{C}^{i+j}(X,\mathcal{F}_1\odot\mathcal{F}_2). \tag{3.76}$$

We use the property that $\tilde{\mathcal{F}}_1 \odot \tilde{\mathcal{F}}_2 = \widetilde{\mathcal{F}_1 \odot \mathcal{F}_2}$ in the last map. Since the cochain maps commute with the coboundary operators, this induces the cup product for the cohomology classes $H^i(X, \mathcal{F}_1) \times H^j(X, \mathcal{F}_2) \xrightarrow{\cup} H^{i+j}(X, \mathcal{F}_1 \odot \mathcal{F}_2)$.

3.6.8 A local condition that allows transversal CCZ operations for sheaf codes

We now use the cup product to prove our main theorem.

Theorem 3.21. *Let $Q(X, \{C_{1,\sigma}\}_{\sigma\in X(t-1)}, \ell_1)$, $Q(X, \{C_{2,\sigma}\}_{\sigma\in X(t-1)}, \ell_2)$, $Q(X, \{C_{3,\sigma}\}_{\sigma\in X(t-1)}, \ell_3)$ be three quantum codes Q_1, Q_2, Q_3 such that $\ell_1 + \ell_2 + \ell_3 = t$. Let $\mathcal{F}_i$ be the sheaf $\mathcal{F}(X, \{C_{i,\sigma}\}_{\sigma\in X(t-1)})$. We assume that the local sections satisfy $(\mathcal{F}_i)_\tau = \mathbb{F}_q$ for $\tau \in X(t)$. We define a trilinear map $f : \mathcal{C}^{\ell_1}(X, \mathcal{F}_1) \times \mathcal{C}^{\ell_2}(X, \mathcal{F}_2) \times \mathcal{C}^{\ell_3}(X, \mathcal{F}_3) \to \mathbb{F}_q$ by*

$$f(\alpha_1, \alpha_2, \alpha_3) = \sum_{\tau\in X(t)} ((\alpha_1 \cup \alpha_2) \cup \alpha_3)(\tau). \tag{3.77}$$

If $C_{1,\sigma} \odot C_{2,\sigma} \odot C_{3,\sigma}$ is contained within the parity check code for every $\sigma \in X(t-1)$, then (Q_1, Q_2, Q_3, f) forms a CCZ code.

Note that $(\alpha_1 \cup \alpha_2) \cup \alpha_3$ lies in $\mathcal{C}^{\ell_1+\ell_2+\ell_3}(X, \mathcal{F}_1 \odot \mathcal{F}_2 \odot \mathcal{F}_3)$, which can be viewed as a function in $X(t) \to \mathbb{F}_q$.

Proof. All we need to show is that f is defined on the cohomology classes. Given $\zeta_1, \zeta_2, \zeta_3$ and $\beta_1, \beta_2, \beta_3$ with $\zeta_i \in Z^{\ell_i}(X, \mathcal{F}_i)$ and $\beta_i \in B^{\ell_i}(X, \mathcal{F}_i)$. Because the cup product

is defined on the cohomology classes, we have $(\zeta_1' \cup \zeta_2') \cup \zeta_3' - (\zeta_1 \cup \zeta_2) \cup \zeta_3 = \beta$ for some $\beta \in B^t(X, \mathcal{F}_1 \odot \mathcal{F}_2 \odot \mathcal{F}_3)$, where $\zeta_i' = \zeta_i + \beta_i$. This implies

$$f(\zeta_1', \zeta_2', \zeta_3') - f(\zeta_1, \zeta_2, \zeta_3) = \sum_{\tau \in X(t)} \beta(\tau). \tag{3.78}$$

Let $\beta = \delta\alpha$ with $\alpha \in \mathcal{C}^{t-1}(X, \mathcal{F}_1 \odot \mathcal{F}_2 \odot \mathcal{F}_3)$. We have

$$\sum_{\tau \in X(t)} \beta(\tau) = \sum_{\tau \in X(t)} \delta\alpha(\tau) = \sum_{\tau \in X(t)} \sum_{\sigma \prec \tau} \alpha(\sigma)|_\tau = \sum_{\sigma \in X(t-1)} \sum_{\tau \succ \sigma} \alpha(\sigma)|_\tau = 0. \tag{3.79}$$

In the last equality, we use $\{\alpha(\sigma)|_\tau\}_{\tau \succ \sigma} \in C_{1,\sigma} \odot C_{2,\sigma} \odot C_{3,\sigma}$ and $\sum_{\tau \succ \sigma} \alpha(\sigma)|_\tau = 0$ which follows from the assumption that $C_{1,\sigma} \odot C_{2,\sigma} \odot C_{3,\sigma}$ is contained within the parity check code. □

More generally, we can obtain multiple CCZ operations, by taking four codes, $Q(X, \{C_{i,\sigma}\}_{\sigma \in X(t-1)}, \ell_i)$ for $i = 1, 2, 3, 4$, where $\ell_1 + \ell_2 + \ell_3 + \ell_4 = t$. (It is ok for the last code to be classical with $\ell_4 = 0$.) We assume that for every $\sigma \in X(t-1)$, $C_{1,\sigma} \odot C_{2,\sigma} \odot C_{3,\sigma} \odot C_{4,\sigma}$ is contained within the parity check code. We can then define trilinear maps f_{ζ_4} for every $\zeta_4 \in Z^{\ell_4}(X, \{C_{4,\sigma}\}_{\sigma \in X(t-1)})$ as

$$f_{\zeta_4}(\alpha_1, \alpha_2, \alpha_3) = \sum_{\tau \in X(t)} (((\alpha_1 \cup \alpha_2) \cup \alpha_3) \cup \zeta_4)(\tau). \tag{3.80}$$

Similarly, one can show that f_{ζ_4} is defined on the cohomology classes. This is related to a known fact that a non-targeted CCCZ operation can induce targeted CCZ operations.

3.6.9 Abstract discussions

We make some abstract remarks and extend the cup product to the case where $\mathcal{F}_\tau$ for $\tau \in X(t)$ is not necessarily $\mathbb{F}_q$.

In mathematics, it is understood that given sheaves $\mathcal{F}_1, \mathcal{F}_2$ of vector spaces over $\mathbb{F}_q$ on a topological space X, there is the cup product

$$H^i(X, \mathcal{F}_1) \times H^j(X, \mathcal{F}_2) \xrightarrow{\cup} H^{i+j}(X, \mathcal{F}_1 \otimes \mathcal{F}_2). \tag{3.81}$$

In the context where the the value at each point is in $\mathbb{F}_q$, this corresponds to the entrywise product we considered above, $(f, g) \mapsto fg$ where $fg(x) = f(x)g(x)$. Thus, our result on cellular complexes can be seen as an analog of this well-known result from traditional sheaf theory.

This indicates that our result can be extended. In particular, given two sheaves $\mathcal{F}_1, \mathcal{F}_2$ on a cellular complex X, we can construct the sheaf $\mathcal{F}_1 \otimes \mathcal{F}_2$ defined by $(\mathcal{F}_1 \otimes \mathcal{F}_2)_\sigma(\tau) = \mathcal{F}_{1,\sigma}(\tau) \otimes \mathcal{F}_{2,\sigma}(\tau)$ for every $\sigma \in X(t-1)$ and $\tau \in X_{\geqslant\sigma}(t)$. By following a similar proof strategy, we can derive the cup product $H^i(X, \mathcal{F}_1) \times H^j(X, \mathcal{F}_2) \xrightarrow{\cup} H^{i+j}(X, \mathcal{F}_1 \otimes \mathcal{F}_2)$.

Acknowledgements

TCL was supported in part by funds provided by the U.S. Department of Energy (D.O.E.) under the cooperative research agreement DE-SC0009919 and by the Simons Collaboration on Ultra-Quantum Matter, which is a grant from the Simons Foundation (652264 JM).

3.7 Appendix: Locally acyclic sheaves

For the purpose of studying sheaf cohomology, it is natural to demand that the cohomology vanishes on each cell. This is essential to obtain Poincaré duality which will be discussed in the next section. We denote $X_{\leqslant\sigma}$ as the cellular complex formed by σ, and $\mathcal{F}_{\leqslant\sigma}$ as the restriction of the sheaf where $\mathcal{F}_{\leqslant\sigma}(\rho) = \mathcal{F}_\rho$ for all $\rho \preceq \sigma$.

Definition 3.22. Given a t dimensional cellular complex X and and a sheaf $\mathcal{F}$ on X. We say $\mathcal{F}$ is locally acyclic if for each i-cell σ, $H^j(X_{\leqslant\sigma}, \mathcal{F}_{\leqslant\sigma}) = 0$ for $0 < j \leqslant t - i$.

For simplicity, we will denote $H^j(X_{\leqslant\sigma}, \mathcal{F}_{\leqslant\sigma})$ as $H^j(\sigma, \mathcal{F})$. More explicitly, $\mathcal{F}$ is locally acyclic iff for each i-cell σ, the chain complex

$$\prod_{\sigma_{i+1}\in X_{\geqslant\sigma}(i+1)} \mathcal{F}_{\sigma_{i+1}} \to \prod_{\sigma_{i+2}\in X_{\geqslant\sigma}(i+2)} \mathcal{F}_{\sigma_{i+2}} \to ... \to \prod_{\sigma_t\in X_{\geqslant\sigma}(t)} \mathcal{F}_{\sigma_t} \tag{3.82}$$

is exact.

The acyclic condition is often essential for estabilishing a meaningful cohomology theory, as it determine which structures are considered trivial. For example, in the case of topological spaces, if X is contractible, then $H^i(X, \mathbb{F}_q) = 0$ for all $i > 0$. Similarly, for schemes, if X is an affine scheme and $\mathcal{F}$ is a quasicoherent sheaf, then $H^i(X, \mathcal{F}) = 0$ for all $i > 0$.

In the context of quantum codes, the acyclic condition naturally arises because it implies that every cocycle ζ supported on $X_{\leqslant\sigma}$ must be a coboundary $\zeta = \delta\alpha$ for some α supported on $X_{\leqslant\sigma}$. If this condition is not satisfied, it is almost saying that ζ is a nontrivial logical operator with a constant weight. This is undesireable, since our goal is to construct LDPC codes with large code distances.

Another way to view the locally acyclic condition is that this is analogous to the condition of a Leray cover over topological spaces. The Leray cover implies that the associate Čech cohomology is isomorphic to the sheaf cohomology. This is not essential and we will not dwell on this point further.

Currently, all sheaves used for code constructions, such as those in [17, 37, 25, 18, 19, 20], are locally acyclic. Most of these constructions achieve local acyclicity by employing tensor codes. The only exception is [19], which uses Reed-Solomon and Reed-Muller codes.

3.8 Appendix: Poincaré duality for cellular complexes

Theorem 3.23. *Given a t dimensional cellular complex X and and a sheaf $\mathcal{F}$ on X. If $\mathcal{F}$*

is locally acyclic, for $0 \leqslant i \leqslant t-1$*, we have*

$$H_{t-i}(X, \mathcal{F}) \cong H^i(X, \mathcal{F}^\perp). \tag{3.83}$$

Proof. Our proof utilizes a spectral sequence, similar to the proof strategy used in Thm 8.9 of Section 14 in [43]. This strategy often arises when working with Čech cohomology which is closely related to our setting.

The idea is to construct a double complex $(K^{\bullet,\bullet}, \delta, d)$ where $\delta : K^{p,q} \to K^{p+1,q}, d : K^{p,q} \to K^{p,q+1}$ and $d^2 = 0, \delta^2 = 0, d\delta = \delta d$. Let $D = d + \delta$. Notice that $D^2 = 0$[4], which defines a chain complex $(K^\bullet, D)$, where $K^k = \bigoplus_{p+q=k} K^{p,q}$.

It is known that one can compute a sequence of complexes $\{(E_r, d_r)\}_{r=0}^{\infty}$, where each page E_r is the cohomology of the previous page E_{r-1}. The key result of spectral sequences is that, if the sequence converges, the resulting complex corresponds to the cohomology group of $(K^\bullet, D)$. For a more detailed statement, we refer to Section 14 in [43].

In our use of the spectral sequence, we alter the roles of δ and d, which provides an alternative way to compute the cohomology group of $(K^\bullet, D)$. The desired result follows from the fact that both computations lead to the same cohomology group.

We now describe the double complex. For $0 \leqslant p \leqslant r \leqslant t$, let

$$K^{p,q} = \prod_{\sigma \in X(p)} \mathcal{C}^r(\sigma, \mathcal{F}) \tag{3.84}$$

where $r = t - q$ and $\mathcal{C}^r(\sigma, \mathcal{F}) = \prod_{\tau \in X_{\geqslant \sigma}(r)} \mathcal{F}_\tau$. Otherwise, $K^{p,q} = 0$.

For $0 \leqslant p < r \leqslant t$, let $\delta : K^{p,q} \to K^{p+1,q}$ be the map

$$(\delta\alpha)(\tau) = \sum_{\sigma \prec \tau} \alpha(\sigma) \tag{3.85}$$

where $\tau \in X(p)$. Otherwise, δ is the zero map.

[4]We again avoid the discussion of the negative signs by working in characteristic 2.

For $0 \leqslant p < r \leqslant t$, let $d : K^{p,q} \to K^{p,q+1}$ be the map

$$(d\alpha)(\tau) = \partial(\alpha(\tau)) \tag{3.86}$$

where $\partial : \mathcal{C}^r(\tau, \mathcal{F}) \to \mathcal{C}^{r-1}(\tau, \mathcal{F})$ is the boundary map where $(\partial\beta)(\pi) = \sum_{\rho \succ - \pi} \beta(\rho)$. Otherwise, d is the zero map.

It is straightforward that $d^2 = 0, \delta^2 = 0$. For $0 \leqslant p < r \leqslant t$ with $r - p \geqslant 2$ and $\alpha \in K^{p,q}$ we have

$$(d\delta\alpha)(\tau)(\pi) = \partial(d\alpha(\tau))(\pi) = \sum_{\rho \succ - \pi} d\alpha(\tau)(\rho) = \sum_{\rho \succ - \pi} \sum_{\sigma \prec - \tau} \alpha(\sigma)(\rho) \tag{3.87}$$

and

$$(\delta d\alpha)(\tau)(\pi) = \sum_{\sigma \prec - \tau} (d\alpha)(\sigma)(\pi) = \sum_{\sigma \prec - \tau} \partial(\alpha(\sigma))(\pi) = \sum_{\sigma \prec - \tau} \sum_{\rho \succ - \pi} \alpha(\sigma)(\rho) \tag{3.88}$$

which implies $d\delta\alpha = \delta d\alpha$. Otherwise $d\delta = \delta d$ as they are both zero maps. We have now established that $(K^{\bullet,\bullet}, \delta, d)$ forms a double complex.

We are ready to discuss the spectral sequence. We will state the resulting pages in the spectral sequence, and then justify each step. The double complex forms the zero-th page

$$E_0 = \begin{array}{c} q \\ \uparrow \\ \left| \begin{array}{ccccc} \prod_{\sigma \in X(0)} \mathcal{C}^0(\sigma, \mathcal{F}) & 0 & 0 & \cdots & 0 \\ \prod_{\sigma \in X(0)} \mathcal{C}^1(\sigma, \mathcal{F}) & \prod_{\sigma \in X(1)} \mathcal{C}^1(\sigma, \mathcal{F}) & 0 & \cdots & 0 \\ \prod_{\sigma \in X(0)} \mathcal{C}^2(\sigma, \mathcal{F}) & \prod_{\sigma \in X(1)} \mathcal{C}^2(\sigma, \mathcal{F}) & \prod_{\sigma \in X(2)} \mathcal{C}^2(\sigma, \mathcal{F}) & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \prod_{\sigma \in X(0)} \mathcal{C}^t(\sigma, \mathcal{F}) & \prod_{\sigma \in X(1)} \mathcal{C}^t(\sigma, \mathcal{F}) & \prod_{\sigma \in X(2)} \mathcal{C}^t(\sigma, \mathcal{F}) & \cdots & \prod_{\sigma \in X(t)} \mathcal{C}^t(\sigma, \mathcal{F}) \end{array} \right| \\ \longrightarrow \\ p \end{array} \tag{3.89}$$

The first spectral sequence has

$$E_1 = H_\delta E_0 = \begin{array}{c} q \\ \uparrow \\ \\ \\ \\ \\ \\ \end{array} \begin{array}{|ccccc} \mathcal{C}_0(X,\mathcal{F}) & 0 & 0 & \cdots & 0 \\ \mathcal{C}_1(X,\mathcal{F}) & ? & 0 & \cdots & 0 \\ \mathcal{C}_2(X,\mathcal{F}) & 0 & ? & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \mathcal{C}_t(X,\mathcal{F}) & 0 & 0 & \cdots & ? \end{array} \begin{array}{l} \\ \\ \\ \\ \\ \\ \rightarrow p \end{array} \tag{3.90}$$

and

$$E_\infty = E_2 = H_d H_\delta E_0 = \begin{array}{c} q \\ \uparrow \\ \\ \\ \\ \\ \\ \end{array} \begin{array}{|ccccc} H_0(X,\mathcal{F}) & 0 & 0 & \cdots & 0 \\ H_1(X,\mathcal{F}) & ? & 0 & \cdots & 0 \\ H_2(X,\mathcal{F}) & 0 & ? & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ H_t(X,\mathcal{F}) & 0 & 0 & \cdots & ? \end{array} \begin{array}{l} \\ \\ \\ \\ \\ \\ \rightarrow p \end{array} \tag{3.91}$$

Even though some of the values on the diagonals is unknown, we have $H^i(K,D) = H_{t-i}(X,\mathcal{F})$ for $0 \leqslant i \leqslant t-1$.

The second spectral sequence has

$$E_1 = H_d E_0 = \begin{array}{c} q \\ \uparrow \\ \end{array} \begin{array}{ccccc} 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \mathcal{C}^0(X,\mathcal{F}^\perp) & \mathcal{C}^1(X,\mathcal{F}^\perp) & \mathcal{C}^2(X,\mathcal{F}^\perp) & \cdots & \mathcal{C}^t(X,\mathcal{F}^\perp) \end{array} \rightarrow p \tag{3.92}$$

and

$$E_\infty = E_2 = H_\delta H_d E_0 = \begin{array}{c} q \\ \uparrow \\ \end{array} \begin{array}{ccccc} 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ H^0(X,\mathcal{F}^\perp) & H^1(X,\mathcal{F}^\perp) & H^2(X,\mathcal{F}^\perp) & \cdots & H^t(X,\mathcal{F}^\perp) \end{array} \rightarrow p \tag{3.93}$$

We have $H^i(K,D) = H^i(X,\mathcal{F}^\perp)$ for $0 \leqslant i \leqslant t$.

By comparing the results from the two spectral sequences, we have the desired result $H_{t-i}(X,\mathcal{F}) \cong H^i(X,\mathcal{F}^\perp)$ for $0 \leqslant i \leqslant t-1$.

It remains to show the following statements:

1. $0 \to \mathcal{C}_0(X,\mathcal{F}) \xrightarrow{f} \prod_{\sigma\in X(0)} \mathcal{C}^0(\sigma,\mathcal{F}) \to 0$ is exact.
2. $0 \to \mathcal{C}_i(X,\mathcal{F}) \xrightarrow{g} \prod_{\sigma\in X(0)} \mathcal{C}^i(\sigma,\mathcal{F}) \xrightarrow{\delta} \prod_{\sigma\in X(1)} \mathcal{C}^i(\sigma,\mathcal{F})$ is exact for $1 \leqslant i \leqslant t$.
3. $\prod_{\sigma\in X(i)} \mathcal{C}^{i+1}(\sigma,\mathcal{F}) \xrightarrow{d} \prod_{\sigma\in X(i)} \mathcal{C}^i(\sigma,\mathcal{F}) \to 0$ is exact for $0 \leqslant i \leqslant t-1$.
4. $0 \to \mathcal{C}^t(X,\mathcal{F}^\perp) \to \prod_{\sigma\in X(t)} \mathcal{C}^t(\sigma,\mathcal{F}) \to 0$ is exact.

5. $0 \to \mathcal{C}^i(X, \mathcal{F}^\perp) \xrightarrow{h} \prod_{\sigma \in X(i)} \mathcal{C}^t(\sigma, \mathcal{F}) \xrightarrow{d} \prod_{\sigma \in X(i)} \mathcal{C}^{t-1}(\sigma, \mathcal{F})$ is exact for $0 \leqslant i \leqslant t-1$.

We will specify the maps f, g, h.

We show the first statement. Notice that $\mathcal{C}^0(\sigma, \mathcal{F}) = \mathcal{F}_\sigma$. Let f be the map $\alpha \mapsto \prod_{\sigma \in X(0)} \alpha(\sigma)$. It is straightforward to check that f is an isomorphism. (Recall that in our context $\mathcal{C}_0(X, \mathcal{F}) = \mathcal{C}^0(X, \mathcal{F})$.)

We show the second statement. Let g be the map $\alpha \mapsto \prod_{\sigma \in X(0)} \{\alpha(\tau)\}_{\tau \in X_{\geqslant \sigma}(i)}$. It is clear that g is injective. It remains to show $\operatorname{im} g = \ker \delta$, specifically that $\operatorname{im} g \supseteq \ker \delta$. Recall that for $\beta \in \prod_{\sigma \in X(0)} \mathcal{C}^i(\sigma, \mathcal{F})$, $(\delta\beta)(\rho) = \sum_{\sigma \prec \rho} \beta(\sigma)$. In this case, ρ is a 1-cell, so there are exactly two 0-cells σ, σ' with $\sigma, \sigma' \prec \rho$. This implies that the functions $\beta(\sigma)$ defined on $\tau \in X_{\geqslant \sigma}(i)$ and $\beta(\sigma')$ defined on $\tau \in X_{\geqslant \sigma'}(i)$ agree on their overlap $\tau \in X_{\geqslant \rho}(i)$. Since the functions $\{\beta(\sigma)\}_{\sigma \in X(0)}$ are consistent, we can glue them together to form a global section α defined on all of $X(i)$. It is clear that $\beta = f(\alpha)$, and thus $\operatorname{im} g \supseteq \ker \delta$.

We show the third statement. It suffices to show that $\mathcal{C}^i(\sigma, \mathcal{F}) \to \mathcal{C}^{i+1}(\sigma, \mathcal{F})$, i.e. $\mathcal{F}_\sigma \to \prod_{\sigma_{i+1} \in X_{\geqslant \sigma}(i+1)} \mathcal{F}_{\sigma_{i+1}}$, is injective. This follows directly from the injectivity axiom on σ.

We show the fourth statement. The statement is true as both $\mathcal{C}^t(X, \mathcal{F}^\perp)$ and $\prod_{\sigma \in X(t)} \mathcal{C}^t(\sigma, \mathcal{F})$ are the space of functions defined on $X(t)$, $\mathcal{C}^t(X, \mathcal{F}^\perp) = \prod_{\sigma \in X(t)} \mathcal{F}_\sigma^\perp \cong \prod_{\sigma \in X(t)} \mathcal{F}_\sigma = \prod_{\sigma \in X(t)} \mathcal{C}^t(\sigma, \mathcal{F})$.

We show the fifth statement. Notice that $\mathcal{C}^i(X, \mathcal{F}^\perp) = \prod_{\sigma \in X(i)} \mathcal{F}_\sigma^\perp$. So all we need is to show is the chain complex $0 \to \mathcal{F}_\sigma^\perp \xrightarrow{e_\sigma} C^t(\sigma, \mathcal{F}) \xrightarrow{\partial} C^{t-1}(\sigma, \mathcal{F})$ is exact for each $\sigma \in X(i)$, where $e = \prod_{\sigma \in X(i)} e_\sigma$ and e_σ maps c to $\{\operatorname{res}_{\sigma,\tau} c\}_{\tau \in X_{\geqslant \sigma}(t)}$. In fact, the exactness of such chain complex was discussed in Theorem 3.18, where we take σ as the cellular complex X. This concludes the proof for the case where $\mathcal{F}$ is locally acyclic.

For statements 3, 4, 5, it is helpful to keep in mind the following long exact sequence

for each i-cell σ:

$$0 \to \mathcal{F}_\sigma \to \prod_{\sigma_{i+1} \in X_{\geqslant \sigma}(i+1)} \mathcal{F}_{\sigma_{i+1}} \to \prod_{\sigma_{i+2} \in X_{\geqslant \sigma}(i+2)} \mathcal{F}_{\sigma_{i+2}} \to ... \to \prod_{\sigma_t \in X_{\geqslant \sigma}(t)} \mathcal{F}_{\sigma_t} \to \mathcal{F}_\sigma^\perp \to 0 \tag{3.94}$$

which can also be written as

$$0 \to \mathcal{C}^i(\sigma, \mathcal{F}) \to \mathcal{C}^{i+1}(\sigma, \mathcal{F}) \to \mathcal{C}^{i+2}(\sigma, \mathcal{F}) \to ... \to \mathcal{C}^t(\sigma, \mathcal{F}) \to \mathcal{F}_\sigma^\perp \to 0. \tag{3.95}$$

□

Dissertation Note

Chapter 3 is based on the sole-authored work "Transversal non-Clifford gates for quantum LDPC codes on sheaves," *arXiv:2410.14631*, 2024.

Bibliography

[1] Dorit Aharonov and Michael Ben-Or. Fault-tolerant quantum computation with constant error. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pages 176–188, 1997.

[2] A Yu Kitaev. Quantum computations: algorithms and error correction. *Russian Mathematical Surveys*, 52(6):1191–1249, 1997.

[3] Emanuel Knill, Raymond Laflamme, and Wojciech H Zurek. Resilient quantum computation: error models and thresholds. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 454(1969):365–384, 1998.

[4] Sergey Bravyi and Alexei Kitaev. Universal quantum computation with ideal clifford gates and noisy ancillas. *Physical Review A—Atomic, Molecular, and Optical Physics*, 71(2):022316, 2005.

[5] Daniel Gottesman and Isaac L Chuang. Quantum teleportation is a universal computational primitive. *arXiv preprint quant-ph/9908010*, 1999.

[6] Héctor Bombín. Gauge color codes: optimal transversal gates and gauge fixing in topological stabilizer codes. *New Journal of Physics*, 17(8):083002, 2015.

[7] Jonas T Anderson, Guillaume Duclos-Cianci, and David Poulin. Fault-tolerant conversion between the steane and reed-muller quantum codes. *Physical review letters*, 113(8):080501, 2014.

[8] Peter W Shor. Fault-tolerant quantum computation. In *Proceedings of 37th conference on foundations of computer science*, pages 56–65. IEEE, 1996.

[9] Daniel Gottesman. Fault-tolerant quantum computation with constant overhead. *arXiv preprint arXiv:1310.2984*, 2013.

[10] Daniel Gottesman. Opportunities and challenges in fault-tolerant quantum computation. *arXiv preprint arXiv:2210.15844*, 2022.

[11] Hector Bombin and Miguel-Angel Martin-Delgado. Topological computation without braiding. *Physical review letters*, 98(16):160502, 2007.

[12] Michael Vasmer and Dan E Browne. Three-dimensional surface codes: Transversal gates and fault-tolerant architectures. *Physical Review A*, 100(1):012312, 2019.

[13] Guanyu Zhu, Shehryar Sikander, Elia Portnoy, Andrew W Cross, and Benjamin J Brown. Non-clifford and parallelizable fault-tolerant logical gates on constant and almost-constant rate homological quantum ldpc codes via higher symmetries. *arXiv preprint arXiv:2310.16982*, 2023.

[14] Christophe Vuillot and Nikolas P Breuckmann. Quantum pin codes. *IEEE Transactions on Information Theory*, 68(9):5955–5974, 2022.

[15] Thomas R Scruby, Arthur Pesah, and Mark Webster. Quantum rainbow codes. *arXiv preprint arXiv:2408.13130*, 2024.

[16] Pavel Panteleev and Gleb Kalachev. Maximally extendable sheaf codes. *arXiv preprint arXiv:2403.03651*, 2024.

[17] Pavel Panteleev and Gleb Kalachev. Asymptotically good quantum and locally testable classical ldpc codes. *arXiv preprint arXiv:2111.03654*, 2021.

[18] Irit Dinur, Min-Hsiu Hsieh, Ting-Chun Lin, and Thomas Vidick. Good quantum ldpc codes with linear time decoders. *arXiv preprint arXiv:2206.07750*, 2022.

[19] Irit Dinur, Siqi Liu, and Rachel Yun Zhang. New codes on high dimensional expanders. *arXiv preprint arXiv:2308.15563*, 2023.

[20] Irit Dinur, Ting-Chun Lin, and Thomas Vidick. Expansion of higher-dimensional cubical complexes with application to quantum locally testable codes. *arXiv preprint arXiv:2402.07476*, 2024.

[21] Louis Golowich and Ting-Chun Lin. Quantum ldpc codes with transversal non-clifford gates via products of algebraic codes, 2024.

[22] Adam Wills, Min-Hsiu Hsieh, and Hayata Yamasaki. Constant-overhead magic state distillation. *arXiv preprint arXiv:2408.07764*, 2024.

[23] Louis Golowich and Venkatesan Guruswami. Asymptotically good quantum codes with transversal non-clifford gates. *arXiv preprint arXiv:2408.09254*, 2024.

[24] Quynh T Nguyen. Good binary quantum codes with transversal ccz gate. *arXiv preprint arXiv:2408.10140*, 2024.

[25] Nikolas P. Breuckmann, Margarita Davydova, Jens N. Eberhardt, and Nathanan Tantivasadakarn. Cups and gates, pt. i, 2024.

[26] Justin Michael Curry. *Sheaves, cosheaves and applications*. University of Pennsylvania, 2014.

[27] Sergey Bravyi and Jeongwan Haah. Magic-state distillation with low overhead. *Physical Review A*, 86(5):052329, 2012.

[28] E. Campbell. The smallest interesting colour code, 2016.

[29] Ting-Chun Lin, Adam Wills, and Min-Hsiu Hsieh. Geometrically local quantum and classical codes from subdivision. *arXiv preprint arXiv:2309.16104*, 2023.

[30] Dominic J. Williamson and Nouédyn Baspin. Layer codes. *arXiv preprint arXiv:2309.16503*, 2023.

[31] Xingjian Li, Ting-Chun Lin, and Min-Hsiu Hsieh. Transform arbitrary good quantum ldpc codes into good geometrically local codes in any dimension. *arXiv preprint arXiv:2408.01769*, 2024.

[32] Anurag Anshu, Nikolas P Breuckmann, and Quynh T Nguyen. Circuit-to-hamiltonian from tensor networks and fault tolerance. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 585–595, 2024.

[33] Alexei Ashikhmin and Emanuel Knill. Nonbinary quantum stabilizer codes. *IEEE Transactions on Information Theory*, 47(7):3065–3072, 2001.

[34] Michael Sipser and Daniel A Spielman. Expander codes. *IEEE transactions on Information Theory*, 42(6):1710–1722, 1996.

[35] R Tanner. A recursive approach to low complexity codes. *IEEE Transactions on information theory*, 27(5):533–547, 1981.

[36] Pavel Panteleev and Gleb Kalachev. Asymptotically good quantum and locally testable classical ldpc codes. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 375–388, 2022.

[37] Irit Dinur, Shai Evra, Ron Livne, Alexander Lubotzky, and Shahar Mozes. Locally testable codes with constant rate, distance, and locality. *arXiv preprint arXiv:2111.04808*, 2021.

[38] Alexei Kitaev and Liang Kong. Models for gapped boundaries and domain walls. *Communications in Mathematical Physics*, 313(2):351–373, 2012.

[39] Fiona J Burnell. Anyon condensation and its applications. *Annual Review of Condensed Matter Physics*, 9(1):307–327, 2018.

[40] Ravi Vakil. The rising sea: Foundations of algebraic geometry. *preprint*, 2017. http://math.stanford.edu/~vakil/216blog/FOAGaug2922public.pdf.

[41] Max Hopkins and Ting-Chun Lin. Explicit lower bounds against $\omega(n)$-rounds of sum-of-squares. *arXiv preprint arXiv:2204.11469*, 2022.

[42] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2000.

[43] Raoul Bott and Loring W Tu. *Differential forms in algebraic topology*, volume 82. Springer Science & Business Media, 2013.

Chapter 4

CFT Ground States as Critical Points of an Entropy Function

Chapter 4 is based on joint work with John McGreevy (arXiv:2303.05444).

swimming in a sea of field theories

searching for shape

the RG flow reveals a hub—
roads threading outward

to grasp the center
i reach with the tweezer
of the vector fixed point equation

Quantum entanglement and quantum information have played important roles in the study of quantum matter. For 2+1D gapped phases, this includes topological entanglement entropy [1, 2], entanglement spectrum [3], and the recent work on chiral central charge [4] and minimal total central charge [5]. For 1+1D conformal field theories (CFTs), the

entanglement entropy [6, 7, 8] of the ground state is related to the central charge. These tools are useful to distinguish quantum phases analytically and numerically.

More ambitiously, one may wonder if the reverse holds, namely, could there be entanglement conditions that are satisfied and only satisfied by ground states for certain quantum phases? One proposal is given by Shi, Kato, and Kim [9] where they stated two conditions that are conjectured to be satisfied and only satisfied by the ground states of topological quantum field theories[1]. Remarkably, using the two conditions, they are able to derive many known properties of 2+1D topological orders. Another nice feature is that because the two conditions only involve the entropies of local regions, the conditions can be checked easily. This program is called entanglement bootstrap and can be applied to other settings, including gapped domain walls [10] and higher dimensions [11].

Given the success of entanglement bootstrap for gapped topological orders, one may wonder if a similar set of entropy conditions exists for gapless states. In this work, we propose a new set of ultra-violet (UV)-independent entropy conditions that apply to the ground states of 1+1D unitary CFTs. The conditions additionally hold for 1+1D gapped phases at RG fixed points. In the spirit of entanglement bootstrap, we conjecture that these conditions characterize the ground states of 1+1D RG fixed points with Lorentz symmetry.[2]

[1]More precisely, any state that satisfies the conditions is the ground state of a Hamiltonian that belongs to a topological phase. And any topological phase has a Hamiltonian whose ground state satisfies the conditions.

[2]Notice that in this statement we make a distinction between CFT and RG fixed point. In fact we know examples of lattice wavefunctions that satisfy Equation (4.1) that are not CFT ground states. Nevertheless, they are renormalization group fixed points.

To be clear, for the more specific case of CFT, this paper shows that Equation (4.1) is a necessary condition, and we conjecture that Equation (4.1) can also be viewed as a sufficient condition for RG fixed points. Furthermore, the following argument suggests that demanding Equation (4.1) for multiple choices of regions may be a sufficient condition for CFT.

Let A, B, C be three consecutive intervals; see Figure 4.1. Our essential observation is that the ground state of a 1+1D unitary CFT $|\psi\rangle$ satisfies

$$K_\Delta \propto I \qquad \text{and} \qquad K_\Delta|\psi\rangle \propto |\psi\rangle^{34} \tag{4.1}$$

where

$$\begin{aligned} K_\Delta \coloneqq (K_{AB} + K_{BC}) - \eta(K_A + K_C) \\ - (1-\eta)(K_B + K_{ABC}) \end{aligned} \tag{4.2}$$

and $K_X \coloneqq -\log \rho_X$[5] is the entanglement Hamiltonian of the reduced density matrix $\rho_X \coloneqq \mathrm{Tr}_{\bar{X}} |\psi\rangle\langle\psi|$, I is the identity operator, and η is the cross ratio of the intervals. Moreover, when the central charge c is known, the proportionality constant is given by

$$K_\Delta = \frac{c}{3} h(\eta) \qquad \text{and} \qquad K_\Delta|\psi\rangle = \frac{c}{3} h(\eta)|\psi\rangle \tag{4.3}$$

The condition Equation (4.1) is a strong constraint; the fact that there are solutions to Equation (4.1) is itself surprising. Notice that each vector equation has roughly the same number of constraints as the number of variables in $|\psi\rangle$, and therefore imposing the vector equation for multiple choices of intervals is overconstrained. When the state has $\geqslant 6$ intervals, if we require the vector equation for *every* choice of three contiguous intervals, there are at least 2 independent vector equations, generically. That means, in general, we do not expect any $|\psi\rangle$ that satisfy every vector equation to exist. However, as we have shown, each CFT ground state is a solution to the vector equations. This is suggesting that these vector equations have intricate relations that are yet to be understood, and it is possible that this relation between equations has something to do with the Virasoro algebra.

[3]Even though the operator equation $K_\Delta \propto I$ implies the vector equation $K_\Delta|\psi\rangle \propto |\psi\rangle$, we find the vector equation to be more robust than the operator equation. Therefore, we write both.

[4]In [12], we give an argument that Eq. (4.1) for different choices of A, B, C exhausts the linear relations amongst single-interval entanglement Hamiltonians in a CFT groundstate.

[5]The reader might worry that $-\log \rho_X$ is not well-defined when ρ_X has zero eigenvalues. Here, we note that the equation $K_\Delta|\psi\rangle \propto |\psi\rangle$ remains well-defined, because it projects out the problematic eigenvectors. This point is discussed in further detail after the derivation of the main results.

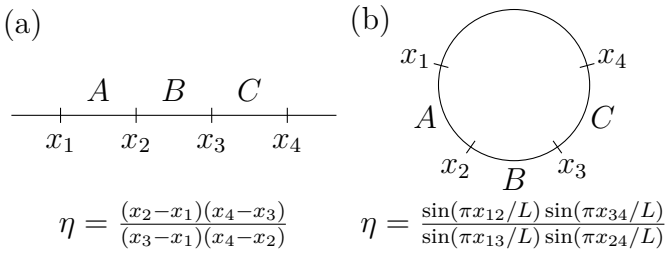

Figure 4.1: (a) Three consecutive intervals on an infinite system with the corresponding cross ratio η. (b) Three consecutive intervals on a circle with circumference L. Here, $x_{ij} \coloneqq x_j - x_i$ denotes the distance between x_i and x_j.

where $h(\eta) \coloneqq -\eta \log \eta - (1-\eta)\log(1-\eta)$ is the binary entropy function. Finally, we observe that Eq. (4.1) is the condition for $|\psi\rangle$ to be a critical point of the following function:

$$\begin{aligned} S_\Delta(|\psi\rangle) \coloneqq (S_{AB} + S_{BC}) - \eta(S_A + S_C) \\ - (1-\eta)(S_B + S_{ABC}) \end{aligned} \tag{4.4}$$

where $S_X \coloneqq S(\rho_X)$ is the von Neumann entanglement entropy between X and $\bar{X}$. The function is nonnegative because it is a convex combination of two nonnegative quantities, by weak monotonicity $S_{AB} + S_{BC} - S_A - S_C \geqslant 0$ and strong subadditivity $S_{AB} + S_{BC} - S_B - S_{ABC} \geqslant 0$ [13, 14]. All the statements above can be extended to the ground state on a finite circle and the thermal state on an infinite line.

The formulae discussed above were derived in continuum CFT. In many physical realizations, CFT arises as an approximation to a lattice model. Importantly, we show numerically that the equations hold approximately for various lattice models and the error $\left| K_\Delta|\psi\rangle - \frac{c}{3}h(\eta)|\psi\rangle \right|$ decays as the number of sites increases (as a power law).

Therefore, after obtaining the basic formulas, we explore their implications on states supported on a discrete lattice, by the following device: divide up the spatial direction into a collection of line segments, and regard each segment as a site of a lattice model with infinite-dimensional local Hilbert space. We thereby identify new equations that could

potentially hold for lattice models with finite-dimensional local Hilbert space, which can then be validated numerically. Specifically, we express the entanglement Hamiltonians of intervals and the six global conformal generators as a combination of 1-site and 2-site entanglement Hamiltonians of the CFT ground state. This has two important implications. First, when the state satisfies the entropy formula, the six global conformal generators have the expected scaling properties under real-space RG. This provides evidence that a state satisfying the entropy formula enjoys conformal symmetry. Second, the result yields an algorithm that can reconstruct the CFT local Hamiltonian from the CFT ground state. This is remarkable since the parent Hamiltonian construction [15] that applies to matrix product states cannot be applied to gapless phases due to their long-range correlations. This algorithm offers an alternative method for recovering the local Hamiltonians from states near RG fixed points, including gapped and gapless phases. This result corroborates the general principle that a single representative wavefunction contains all the universal data about a state of matter.

We want to highlight an implicit theme of this work, which aims to go beyond the traditional algebraic formulation of CFT and provide an alternative analytic formulation. Typically, CFTs are defined using algebras with equalities, which do not allow for the discussion of *approximate CFTs*. However, many examples exist that we would like to categorize as approximate CFTs, such as QFTs that are slight perturbations of CFTs or critical lattice models that are CFTs in the IR limit. Because these models have different Hilbert spaces, finding a criterion that applies to all cases is challenging. To address this issue, we utilize the entanglement entropy and entanglement Hamiltonian, which are agnostic to the Hilbert space. We propose to define approximate CFTs as systems whose ground state approximately satisfies Eq. 4.1. Further research is needed to determine the usefulness of this proposal.

4.1 CFTs and their entanglement properties

Conformal field theories (CFTs) are field theories with scale invariance. Such scale invariance often implies a larger invariance called the conformal symmetries, which include transformations that locally look like rescalings. These conformal symmetries appear naturally in many physical systems, including the fixed points of the RG flow and the critical points of statistical models.

For 1+1D CFT, the entanglement entropy [8] and the entanglement Hamiltonian (EH) [16] of the ground state on an interval $[x_1, x_2]$ are known to be

$$S_{[x_1,x_2]} = \frac{c}{3} \log \frac{x_2 - x_1}{\varepsilon} \tag{4.5}$$

and

$$K_{[x_1,x_2]} = 2\pi \int_{x_1}^{x_2} dx \frac{(x - x_1)(x_2 - x)}{x_2 - x_1} T_{00}(x) + \text{const.} \tag{4.6}$$

where c is the central charge, ε is the uniform UV cutoff, T is the stress-energy tensor, and const. is a number that depends on the UV cutoff ε. These two equations are the key to showing the main results. As we will see, even though both equations suffer from UV divergences, the combinations in Eq. (4.1) and Eq. (4.3) are free from UV divergences.

4.2 Derivation of the main results

We first show Eq. (4.1) and Eq. (4.3) for the ground state of a 1+1D CFT on an infinite system. Then we show the equivalence of Eq. (4.1) to the critical point condition of S_Δ defined in Eq. (4.4). Let $A = [x_1, x_2]$, $B = [x_2, x_3]$, and $C = [x_3, x_4]$ be three consecutive intervals. For convenience, we define the function $f_{[x',x'']}(x) = \frac{(x-x')(x''-x)}{x''-x'} 1_{[x',x'']}$ where $1_{[x',x'']}$ is the indicator function. By Eq. (4.6),

$$(K_{AB} + K_{BC}) - \eta(K_A + K_C) - (1 - \eta)(K_B + K_{ABC})$$

$$= \int_{-\infty}^{\infty} dx \Big((f_{AB} + f_{BC}) - \eta(f_A + f_C) \\ - (1-\eta)(f_B + f_{ABC}) \Big) T_{00}(x) + \text{const.} \tag{4.7}$$

A straightforward calculation shows $(f_{AB} + f_{BC}) - \eta(f_A + f_C) - (1-\eta)(f_B + f_{ABC}) = 0$ from which Equation (4.1) follows.

To obtain the ratio in Eq. (4.3), we multiply $\langle\psi|$ on both sides. Because $\langle\psi|K_A|\psi\rangle = S_A$, the left hand side becomes $(S_{AB} + S_{BC}) - \eta(S_A + S_C) - (1-\eta)(S_B + S_{ABC})$ which evaluates to $\frac{c}{3}h(\eta)$ using Eq. (4.5).

Now we show that Eq. (4.1) is precisely the condition for vanishing variation of S_Δ with respect to the state. Recall that $|\psi\rangle$ is a critical point if the gradient is 0 subject to $(\langle\psi| + \langle d\psi|)(|\psi\rangle + |d\psi\rangle) = 1$, i.e., $\langle d\psi|\psi\rangle + \langle\psi|d\psi\rangle = 0$. To compute the gradient of $S_\Delta = (S_{AB} + S_{BC}) - \eta(S_A + S_C) - (1-\eta)(S_B + S_{ABC})$ we use the first order derivative of the entanglement entropy:

$$dS_X(|\psi\rangle) = \langle d\psi|K_X|\psi\rangle + \langle\psi|K_X|d\psi\rangle \tag{4.8}$$

as reviewed in the Supplementary Material. Therefore, the critical point $|\psi\rangle$ has

$$dS_\Delta = \langle d\psi|K_\Delta|\psi\rangle + \langle\psi|K_\Delta|d\psi\rangle = 0 \tag{4.9}$$

for all $|d\psi\rangle$ satisfying $\langle d\psi|\psi\rangle + \langle\psi|d\psi\rangle = 0$. Recall $K_\Delta = (K_{AB} + K_{BC}) - \eta(K_A + K_C) - (1-\eta)(K_B + K_{ABC})$. This is equivalent to $K_\Delta|\psi\rangle \propto |\psi\rangle$ in Eq. (4.1). Note that the left equation in Eq. (4.1), $K_\Delta \propto I$, does not follow from $K_\Delta|\psi\rangle \propto |\psi\rangle$ and the critical point condition in general, but does follow under certain assumptions that will be explained later in this section.

More generally, the result can be extended to the ground state on a finite circle and the thermal state on an infinite line. In both cases, the entanglement entropy and the entanglement Hamiltonian are known [8, 16], so the proof strategy still works. The difference is to replace the cross-ratio η with the effective cross-ratio η_{eff} where $\eta_{\text{eff}}^L = \frac{\sin(\pi x_{12}/L)\sin(\pi x_{34}/L)}{\sin(\pi x_{13}/L)\sin(\pi x_{24}/L)}$

for the ground state on a circle of length L and $\eta_{\text{eff}}^{\beta} = \frac{\sinh(\pi x_{12}/\beta)\sinh(\pi x_{34}/\beta)}{\sinh(\pi x_{13}/\beta)\sinh(\pi x_{24}/\beta)}$ for the thermal state on an infinite line of inverse temperature β.

We now discuss the subtle relation between the operator equation $K_\Delta = \frac{c}{3}h(\eta)$ and the vector equation $K_\Delta|\psi\rangle = \frac{c}{3}h(\eta)|\psi\rangle$. We first present an argument that shows their equivalence for quantum field theories, then discuss the difference in their approximate versions. It is clear that the operator equation implies the vector equation and typically the other way does not hold for finite dimensional Hilbert spaces. One extreme example for finite dimensional Hilbert spaces is that $K_X = -\log \rho_X$ could have singularities when ρ_X has zero eigenvalues, while $K_X|\psi\rangle$ is still well-defined with continuity, because $\lim_{\lambda\to 0} \lambda \log \lambda = 0$. On the other hand, for quantum field theories, due to the Reeh-Schlieder theorem [17], $|\psi\rangle \in \mathcal{H}_{ABCD}$ is cyclic, which means $O_D|\psi\rangle$ is dense in $\mathcal{H}_{ABC}$ for O_D supported on D. Because K_Δ is only supported on $\mathcal{H}_{ABC}$, K_Δ and O_D commutes, and we have $K_\Delta O_D|\psi\rangle = \frac{c}{3}h(\eta)O_D|\psi\rangle$. Since $O_D|\psi\rangle$ is dense, this implies the operator equation $K_\Delta = \frac{c}{3}h(\eta)$.

In the approximate case, one might hope that $K_\Delta \approx \frac{c}{3}h(\eta)$ in the operator norm and $K_\Delta|\psi\rangle \approx \frac{c}{3}h(\eta)|\psi\rangle$ in the vector norm. Numerically, we observe that the approximate operator equation does not hold, while the approximate vector equation holds. This suggests that the vector equation is more valid because it is stable against perturbation. Nevertheless, we suspect that the operator equation remains stable under a different norm which requires further investigation.

4.3 Local-to-global implications

In the last section, we established that $K_\Delta = \frac{c}{3}h(\eta)$ for any consecutive intervals A, B, C. In this section, we apply this relation to intervals with integer-valued endpoints which can be associated with a lattice model. The benefit of having a lattice model is that

it allows for concrete numerical verifications. We demonstrate that several quantities can be expressed as a sum of local operators. First, we prove that the EH of an interval can be expressed as a sum of 1-site and 2-site EHs from Eq. (4.1). This implies that the reduced density matrix on an arbitrary interval can be reconstructed from the 2-site reduced density matrix.

We first introduce the notation. To simplify the presentation, we define the reduced EH as $\widetilde{K}_X \coloneqq K_X - \langle\psi|K_X|\psi\rangle$, which is the EH shifted by a constant so that $\langle\psi|\widetilde{K}_X|\psi\rangle = 0$. This removes the constant term in Eq. (4.6), which implies that $\widetilde{K}_\Delta = 0$. The reduced EH of an interval $[a, b]$ is denoted as $\widetilde{K}_{[a,b]}$. In the following discussion, we only consider intervals with $a, b \in \mathbb{Z}$, so that they can be associated with a lattice model. The corresponding lattice model regards each interval of length 1, $[a, a+1]$, as a single site[6]. Therefore, the interval $[a, b]$ corresponds to $b - a$ sites. From now on, we refer to the reduced EH simply as EH.

We now show that any EH of an interval can be written as a sum of 1-site and 2-site EHs, using Eq. (4.1) recursively. We begin with the example of a 3-site EH. Take $A = [0, 1], B = [1, 2], C = [2, 3]$. The equivalent form of Eq. (4.1) $\widetilde{K}_\Delta = 0$ implies

$$\widetilde{K}_{[0,2]} + \widetilde{K}_{[1,3]} - \frac{1}{4}\widetilde{K}_{[0,1]} - \frac{1}{4}\widetilde{K}_{[2,3]} - \frac{3}{4}\widetilde{K}_{[1,2]} - \frac{3}{4}\widetilde{K}_{[0,3]} = 0. \tag{4.10}$$

Therefore,

$$\widetilde{K}_{[0,3]} = -\frac{1}{3}\widetilde{K}_{[0,1]} + \frac{4}{3}\widetilde{K}_{[0,2]} - \widetilde{K}_{[1,2]} + \frac{4}{3}\widetilde{K}_{[1,3]} - \frac{1}{3}\widetilde{K}_{[2,3]}. \tag{4.11}$$

Similarly, for a 4-site EH, take $A = [0, 1], B = [1, 2], C = [2, 4]$, and we have

$$\begin{aligned}\widetilde{K}_{[0,4]} &= -\frac{1}{2}\widetilde{K}_{[0,1]} + \frac{3}{2}\widetilde{K}_{[0,2]} - \widetilde{K}_{[1,2]} + \frac{3}{2}\widetilde{K}_{[1,4]} - \frac{1}{2}\widetilde{K}_{[2,4]} \\ &= -\frac{1}{2}\widetilde{K}_{[0,1]} + \frac{3}{2}\widetilde{K}_{[0,2]} - \frac{3}{2}\widetilde{K}_{[1,2]} + 2\widetilde{K}_{[1,3]}\end{aligned}$$

[6]Of course, these 'sites' have infinite-dimensional Hilbert spaces, and this is just a rewriting of the continuum CFT information. In what follows, we will use the same relation in systems where the Hilbert space of each site is replaced by a finite-dimensional Hilbert space.

$$- \frac{3}{2}\widetilde{K}_{[2,3]} + \frac{3}{2}\widetilde{K}_{[2,4]} - \frac{1}{2}\widetilde{K}_{[3,4]}. \tag{4.12}$$

More generally, for an n-site EH, take $A = [0,1], B = [1,2], C = [2,n]$, and we have

$$\begin{aligned}\widetilde{K}_{[0,n]} = &-\frac{n-2}{n}\widetilde{K}_{[0,1]} + \frac{2(n-1)}{n}\widetilde{K}_{[0,2]} - \widetilde{K}_{[1,2]} \\ &+ \frac{2(n-1)}{n}\widetilde{K}_{[1,n]} - \frac{n-2}{n}\widetilde{K}_{[2,n]}.\end{aligned} \tag{4.13}$$

By recursively expanding $\widetilde{K}_{[1,n]}$ and $\widetilde{K}_{[2,n]}$, we obtain

$$\widetilde{K}_{[0,n]} = \sum_{j=-\infty}^{\infty} f_2(j+1)\widetilde{K}_{[j,j+2]} + f_1(j+\frac{1}{2})\widetilde{K}_{[j,j+1]} \tag{4.14}$$

where $f_2(x) = \frac{2x(n-x)}{n}1_{[0,n]}$, $f_1(x) = \frac{-2x(n-x)+\frac{3}{2}}{n}1_{[0,n]}$, and $1_{[0,n]}$ is the indicator function. Therefore, any EH can be written as a sum of 1-site and 2-site EHs, which means the reduced density matrix on arbitrary intervals can be reconstructed from the 2-site reduced density matrix. This could be viewed as a solution to the quantum marginal problem for CFTs where the Petz recovery map does not apply.

Remark 4.1. We note that the decomposition of EH into a sum of 1-site and 2-site EHs is similar to Eq. (4.6) in field theory, where EH is a sum of local terms $T_{00}(x)$. We also remark that there are different ways to decompose $\widetilde{K}_{[0,n]}$, yet all of them lead to the same expression. This nontrivial property implies certain consistency relations between $\widetilde{K}_\Delta = 0$ across different choices of A, B, C.

The state contains the universal data— The same logic that leads to Eq. (4.1) implies that the CFT Hamiltonian (and hence all of the CFT data) can be extracted from the groundstate. In fact, the same is true of all 6 global conformal generators $H, P, D, M_{10}, C_0, C_1$ which are the Hamiltonian, momentum, dilatation, boost, and special conformal generators. We first find their corresponding expressions for CFT ground states when Eq. (4.6) applies. Because $H = \int_{-\infty}^{\infty} dx\, T_{00}(x)$, $M_{10} = \int_{-\infty}^{\infty} dx\, xT_{00}(x)$, $C_0 =$

$\int_{-\infty}^{\infty} dx\, x^2 T_{00}(x)$, applying Eq. (4.6), we have

$$\begin{aligned} H &= \frac{1}{\pi}\sum_{j=-\infty}^{\infty}\left(\widetilde{K}_{[j,j+2]} - \widetilde{K}_{[j,j+1]}\right) \\ &= \frac{1}{\pi}\sum_{j=-\infty}^{\infty}\widetilde{K}'_{[j,j+2]}, \qquad (4.15) \\ M_{10} &= \frac{1}{2\pi}\sum_{j=-\infty}^{\infty}\left((2j+2)\widetilde{K}_{[j,j+2]} - (2j+1)\widetilde{K}_{[j,j+1]}\right) \\ &= \frac{1}{\pi}\sum_{j=-\infty}^{\infty}(j+1)\widetilde{K}'_{[j,j+2]}, \qquad (4.16) \\ C_0 &= \frac{1}{\pi}\sum_{j=-\infty}^{\infty}\left((j+1)^2\widetilde{K}_{[j,j+2]} - (j^2+j+1)\widetilde{K}_{[j,j+1]}\right) \\ &= \frac{1}{\pi}\sum_{j=-\infty}^{\infty}\left((j+1)^2\widetilde{K}'_{[j,j+2]} - \frac{1}{2}\widetilde{K}_{[j,j+1]}\right), \qquad (4.17) \end{aligned}$$

where $\widetilde{K}'_{[j,j+2]} = \widetilde{K}_{[j,j+2]} - \frac{1}{2}\widetilde{K}_{[j,j+1]} - \frac{1}{2}\widetilde{K}_{[j+1,j+2]}$ is introduced to simplify the equations. Because $P = \mathbf{i}[M_{10}, H]$, $D = \frac{\mathbf{i}}{2}[C_0, H]$, $C_1 = \mathbf{i}[C_0, M_{10}]$, we have

$$\begin{aligned} P &= \frac{\mathbf{i}}{\pi^2}\sum_{j=-\infty}^{\infty}\left[\widetilde{K}'_{[j+1,j+3]}, \widetilde{K}'_{[j,j+2]}\right], \qquad (4.18) \\ D &= \frac{\mathbf{i}}{\pi^2}\sum_{j=-\infty}^{\infty}\Bigg((j+\frac{3}{2})\left[\widetilde{K}'_{[j+1,j+3]}, \widetilde{K}'_{[j,j+2]}\right] \\ &\qquad - \frac{1}{4}\left[\widetilde{K}_{[j,j+1]}, \widetilde{K}'_{[j,j+2]}\right] \\ &\qquad - \frac{1}{4}\left[\widetilde{K}_{[j+1,j+2]}, \widetilde{K}'_{[j,j+2]}\right]\Bigg), \qquad (4.19) \\ C_1 &= \frac{\mathbf{i}}{\pi^2}\sum_{j=-\infty}^{\infty}\Bigg((j+1)(j+2)\left[\widetilde{K}'_{[j+1,j+3]}, \widetilde{K}'_{[j,j+2]}\right] \\ &\qquad - \frac{j+1}{2}\left[\widetilde{K}_{[j,j+1]}, \widetilde{K}'_{[j,j+2]}\right] \\ &\qquad - \frac{j+1}{2}\left[\widetilde{K}_{[j+1,j+2]}, \widetilde{K}'_{[j,j+2]}\right]\Bigg). \qquad (4.20) \end{aligned}$$

This constructs the 6 global conformal generators from 1-site and 2-site EHs. We use the term reconstructed Hamiltonian to refer to the output of the process of taking a state and constructing an associated Hamiltonian as in (4.15). This state could be on a lattice or in the continuum.

We make two comments on the expression for P. First, a similar expression was considered in [18, Sec II.C]. The idea is to define the momentum density $p_j = \mathbf{i}[h_j, h_{j-1}]$ in terms of the Hamiltonian density h_j where $H = \sum h_j$. One challenge in implementing this idea in practice is that the Hamiltonian H is only equal to the CFT Hamiltonian H_{CFT} up to a multiplicative factor. Consequently, the momentum density p_j is determined only up to a multiplicative factor. In our work, we fix this multiplicative factor and provide a better theoretical understanding. Second, there is a no-go theorem that prevents expressing the momentum operator as a sum of local operators [19, Corollary 6.1] which seemingly contradicts our expression above. There are several ways to reconcile this apparent contradiction. One is to note that our expression holds exactly only for lattice models with an infinite local Hilbert space dimension, whereas the no-go theorem applies to models with a finite local Hilbert space dimension. Another way to reconcile this is to recognize that, for models with a finite local Hilbert space dimension, $e^{-\mathbf{i}P}$ is not precisely equal to the lattice translation operator [20]. Instead, they are only approximately equal at low energies, as we verify numerically below.

We further comment that there are additional commutation relations between the generators which are not utilized in this work, such as $[M_{10}, P] = -\mathbf{i}H$ and $[H, P] = 0$. In the context of field theories, this phenomenon where higher order commutator of EHs are related to linear combinations of EHs can be understood from the operator product expansion (OPE) of $T_{zz}(t, x)$. First, because EH is a sum of $T_{zz}(t = 0, x)$ and $T_{\bar{z}\bar{z}}(t = 0, x)$, the commutator of EHs is a sum of the commutators of $T_{zz}(t = 0, x)$ and $T_{\bar{z}\bar{z}}(t = 0, x)$. Second, the commutator of $T_{zz}(t = 0, x)$ is generated by the singular part of the OPE of T_{zz}. Finally, the singular part of the OPE of T_{zz} is generated by T_{zz} and its descendent $\partial_z T_{zz}$. The second and the final part can be summarized into the following expression [21,

Equation (B23)]

$$[T_{zz}(0,x_1),T_{zz}(0,x_2)] = \frac{\mathbf{i}\pi c}{6}\partial_{x_1}^3\delta(x_1-x_2) + 4\pi\mathbf{i}T_{zz}(0,x_2)\partial_{x_1}\delta(x_1-x_2) - 2\pi\mathbf{i}\partial_z T_{zz}(0,x_2)\delta(x_1-x_2). \tag{4.21}$$

We leave these further relations between the higher order commutators of EHs and the linear combinations of EHs for future study.

Having defined the 6 conformal generators using 1-site and 2-site EHs, we now show that they have the expected scaling from real-space RG. We first study the case of the Hamiltonian H. To perform real-space RG, we consider a new system which blocks 2 sites in the original system into 1 site. We then compare the two reconstructed Hamiltonians using Eq. (4.15), which are

$$H_1 = \frac{1}{\pi}\sum_{j=-\infty}^{\infty}\widetilde{K}_{[j,j+2]} - \widetilde{K}_{[j,j+1]} \tag{4.22}$$

$$H_2 = \frac{1}{\pi}\sum_{j=-\infty}^{\infty}\widetilde{K}_{[2j,2j+4]} - \widetilde{K}_{[2j,2j+2]}. \tag{4.23}$$

Note that the state we used for the reconstructions is the same. The only difference is the size of the block. By expanding the 4-site EH as a sum of 1-site and 2-site EH using Eq. (4.12), we have $H_2 = 2H_1$. Similarly, one can define P_2 and show that $P_2 = 2P_1$, $D_2 = D_1$, $(M_{10})_2 = (M_{10})_1$, $(C_0)_2 = \frac{1}{2}(C_0)_1$, and $(C_1)_2 = \frac{1}{2}(C_1)_1$. These are precisely the scalings of the generators for CFT under the transformation $x \to x/2$, $t \to t/2$.

It is perhaps not surprising that the scaling property holds because the construction is motivated by the field theories. However, what is surprising is that the derivation of the scaling property only uses the condition (4.1), which is independent from the field theory description of CFTs. This observation supports the conjecture that states satisfying Eq. (4.1) are the CFT ground states.

4.4 Numerical tests

We now move on to numerics and verify that Eqs. (4.1), (4.3), (4.15) and (4.18) hold approximately for critical lattice models. First, we test the validity of the main Eqs. (4.1) and (4.3) on small sizes of four different critical lattice models and on large sizes of free fermions. We find that these equations hold approximately for all the models we consider and the error decreases as the size increases. Then, we test the reconstruction of the Hamiltonian and momentum operator Eqs. (4.15) and (4.18) on small sizes of the four models.

The critical lattice models we consider are the critical transverse field Ising model, critical three-state Potts model, XX model, and Heisenberg model[7] defined as:

$$H_{\text{Ising}} = \sum_i -Z_i - X_i X_{i+1} \tag{4.24}$$

$$H_{\text{Potts}} = \sum_i -\mathcal{Z}_i - \mathcal{Z}_i^\dagger - \mathcal{X}_i \mathcal{X}_{i+1}^\dagger - \mathcal{X}_i^\dagger \mathcal{X}_{i+1} \tag{4.25}$$

$$H_{\text{XX}} = \sum_i X_i X_{i+1} + Y_i Y_{i+1} \tag{4.26}$$

$$H_{\text{Heisenberg}} = \sum_i X_i X_{i+1} + Y_i Y_{i+1} + Z_i Z_{i+1} \tag{4.27}$$

where X, Y, Z are the Pauli matrices and $\mathcal{X}, \mathcal{Z}$ are the qutrit Pauli matrices

$$\mathcal{X} = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}, \quad \mathcal{Z} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \omega & 0 \\ 0 & 0 & \omega^2 \end{pmatrix}, \quad \omega = e^{2\pi \mathbf{i}/3} \tag{4.28}$$

Table 4.1 lists the errors associated with Eqs. (4.1) and (4.3). The error in Eq. (4.1) is defined as the norm of the component in $K_\Delta|\psi\rangle$ orthogonal to $|\psi\rangle$, $\left| K_\Delta|\psi\rangle - \langle\psi|K_\Delta|\psi\rangle|\psi\rangle \right|$, which is also the standard deviation of K_Δ, $\sqrt{\langle K_\Delta^2\rangle - \langle K_\Delta\rangle^2}$. The error in Eq. (4.3) is

[7]These models are all quantum critical and their universal properties are described respectively by the following conformal field theories: Ising model with $c = 1/2$; Potts model with $c = 4/5$; compact scalar CFT with radius 1 and $1/\sqrt{2}$.

Table 4.1: The error in Eq. (4.1) $\left|K_\Delta|\psi\rangle - \langle\psi|K_\Delta|\psi\rangle|\psi\rangle\right|$ and the error in Eq. (4.3) $\left|K_\Delta|\psi\rangle - \frac{c}{3}h(\eta)|\psi\rangle\right|$ for ground states on circles with circumferences $L = 4, 8, 12$.

L	Error in Eq. (4.1)			Error in Eq. (4.3)		
	4	8	12	4	8	12
Ising Model	0.0282	0.0090	0.0057	0.0288	0.0090	0.0057
Potts Model	0.0422			0.0434		
XX Model	0	0.0399	0.0120	0.0486	0.0416	0.0125
Heisenberg	0	0.0562	0.0279	0.0566	0.0583	0.0286

defined as the norm $\left|K_\Delta|\psi\rangle - \frac{c}{3}h(\eta)|\psi\rangle\right|$. We computed these errors for ground states on circles with circumferences $L = 4, 8, 12$, where (A, B, C) have lengths $(1, 1, 1), (2, 2, 2)$, and $(3, 3, 3)$, respectively. In all cases, $\eta = 1/2$. We observe that except for the accidental case where the XX model and Heisenberg model with circumference 4 have 0 errors, the error decreases as the system size increases. This result is consistent with the intuition that the lattice model approximates the CFT better in the IR as the system size increases.

Figure 4.2 shows the error in Eq. (4.1) for free fermions across various system sizes. We again observe that the error decreases as the system size increases which roughly scales as $1/L^2$. Simulating free fermions on a large system size is feasible due to their lower complexity [22, 23].

We now test the reconstructed Hamiltonian and momentum on a circle in Eqs. (4.3) and (4.4) where

$$H_{rec} = \frac{\sin(2\pi/L)}{2L\sin^2(\pi/L)} \sum_{j=0}^{L-1} \widetilde{K}_{[j,j+2]} - \widetilde{K}_{[j,j+1]} \tag{4.29}$$

$$P_{rec} = \frac{\mathbf{i}}{\pi^2} \sum_{j=0}^{L-1} \left[\widetilde{K}'_{[j+1,j+3]}, \widetilde{K}'_{[j,j+2]}\right]. \tag{4.30}$$

[8] The factor $\frac{\sin(2\pi/L)}{2L\sin^2(\pi/L)}$ in H_{rec} comes from the coefficient obtained when expressing EH

[8] The analogous formula for the exact reconstructed momentum is

$$P_{rec} = \frac{\sin^3(2\pi/L)\pi}{8L\sin^4(\pi/L)} \frac{i}{\pi^2} \sum_{j=0}^{L-1} [\widetilde{K}_{[j+1,j+3]}, \widetilde{K}_{[j,j+2]}]$$

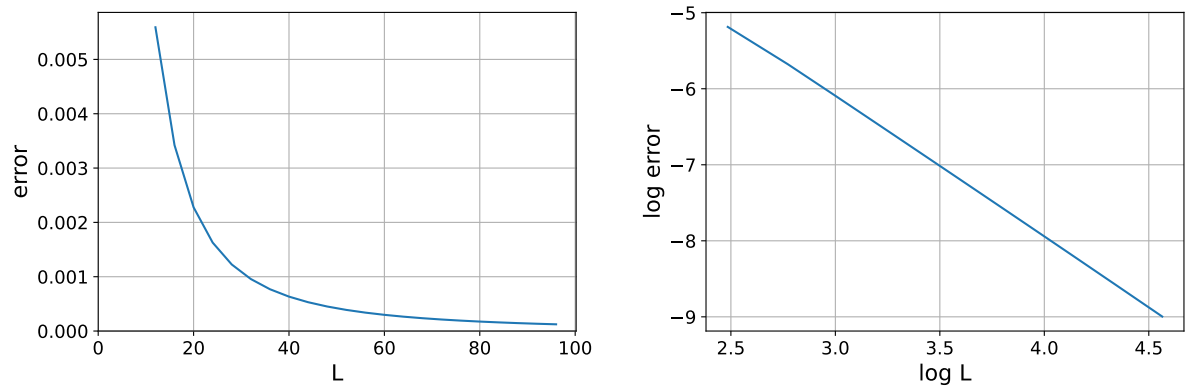

Figure 4.2: The error in Eq. (4.1) $\left|K_\Delta|\psi\rangle - \langle\psi|K_\Delta|\psi\rangle|\psi\rangle\right|$ for the ground state of the critical free fermions. The circle is partitioned into four equal sized intervals, where the lengths of A, B, C are $L/4$ and has cross ratio $\eta = 1/2$. The error roughly scales as $1/L^2$.

as an integral of $T_{00}(x)$. On the other hand, the factor $\frac{\mathbf{i}}{\pi^2}$ in P_{rec} is based solely on the understanding at the infinite-size limit in Eq. (4.18). Despite this, as we will see, the reconstructed Hamiltonian and momentum agree excellently. [9]

Before presenting the result, we first remark that the field theory Hamiltonian and momentum satisfy

$$H_{CFT}\,|\Delta, s\rangle = \frac{2\pi}{L}\left(\Delta - \frac{c}{12}\right)|\Delta, s\rangle \tag{4.31}$$

$$P_{CFT}\,|\Delta, s\rangle = \frac{2\pi}{L}s\,|\Delta, s\rangle, \tag{4.32}$$

where $|\Delta, s\rangle$ is the image under the state-operator correspondence of a scaling operator of dimension Δ and spin s. We emphasize that H_{rec} is equal to H_{CFT} up to a constant shift so that H_{rec} has ground state energy 0, i.e. $H_{rec} = H_{CFT} - E_0$ [10]. In particular,

$$- \left(1 - \frac{2\sin^2(\pi/L)}{\sin^2(2\pi/L)}\right)[\widetilde{K}_{[j,j+2]}, \widetilde{K}_{[j,j+1]}]$$
$$- \left(1 - \frac{2\sin^2(\pi/L)}{\sin^2(2\pi/L)}\right)[\widetilde{K}_{[j+1,j+2]}, \widetilde{K}_{[j,j+2]}]$$

[9] Note that H_{rec} is equal to H_{CFT} since both have the same expression in terms of T. On the contrary, because we put the system on a circle P_{rec} is only equal to P_{CFT} in the thermodynamic limit.

[10] The constant $-\frac{2\pi}{L}\frac{c}{12}$ can be recovered by using the EHs of the ground state on an infinite system

the multiplicative factor is fixed, meaning that the scaling dimensions Δ can be obtained without the rescaling required if one directly uses the spectrum of an arbitrary critical Hamiltonian.

Figure 4.3 compares the spectra of the original Hamiltonian to the spectrum of the reconstructed Hamiltonian at low energy. The reconstructed Hamiltonian H_{rec} is rescaled by $\frac{L}{2\pi}$ and the original Hamiltonian is rescaled to fit H_{rec}. We observe the spectra have excellent agreement even at the small system size where $L = 4$.[11]

Figure 4.4 compares the spectrum of $\mathbf{i} \log T$ to the spectrum of the reconstructed momentum at low energy, where T is the translation operator by 1 site. We test if $T \approx e^{-\mathbf{i}P_{rec}}$. Again both spectra are rescaled by $\frac{L}{2\pi}$ and we expect both to take integer values. We observe the spectra agree at low energy and the agreement improves as the system size increases.

We provide three technical comments on this test. First, we note that in lattice models, it is not generally true that $T = e^{-\mathbf{i}P_{CFT}}$, for example, in the Heisenberg and XX models. This is because $\mathbf{i} \log T$ may have a constant shift s_0, where $\mathbf{i} \log T|\psi\rangle = (s_0 + \frac{2\pi}{L}s)|\psi\rangle$. As an example, if we instead define the transverse field Ising model as $H_{\text{Ising}} = \sum_i -Z_i + X_i X_{i+1}$, then the new ground state is obtained by applying Z on the even sites of the usual ground state. This results with a shift $s_0 = \pi$. Sometimes this

instead of the EHs of the length L finite system, which we have tested for critical free fermions. However, we do not discuss this aspect further because it is generally not possible to know the EHs on an infinite system directly.

[11] We emphasize that this is a pleasant surprise: although we use field theory to derive our formula for the reconstructed Hamiltonian, its form is numerically effective even on quite small lattices. It would be interesting to understand why it works so well.

Locality clearly plays an important role. We note that the ideal Hamiltonian reconstructed from the saddle point of the entropy function S_Δ and the model Hamiltonians we study (Eqs. 24-27) both involve at most 2-site operators. This means that the artificial model Hamiltonian can already be close to the ideal reconstructed CFT Hamiltonian.

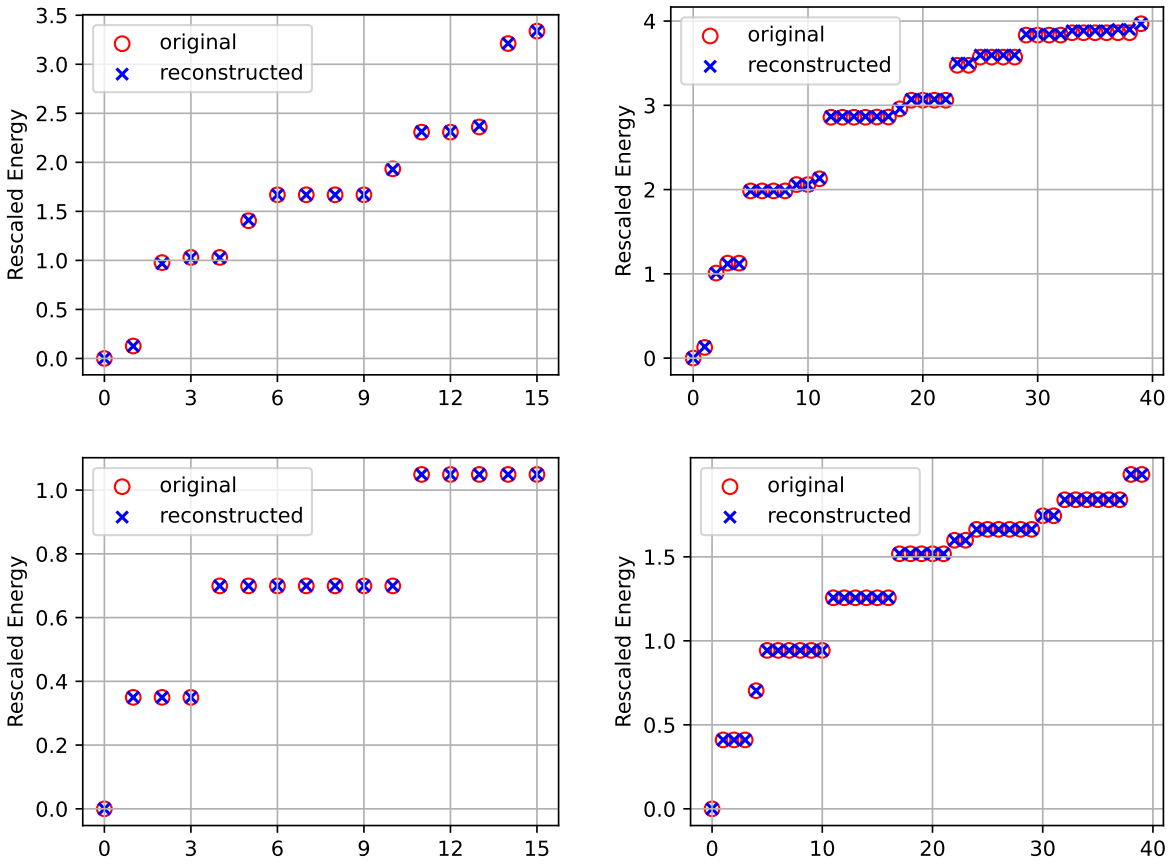

Figure 4.3: Spectra of the original Hamiltonian and the reconstructed Hamiltonian, where the states are ordered by energy along the X-axis. Top left: Ising model with $L = 4$. Top right: Lowest 40 eigenvalues for Ising model with $L = 12$. Bottom left: Heisenberg model with $L = 4$. Bottom right: Lowest 40 eigenvalues for Heisenberg model with $L = 12$.

shift cannot be removed easily due to anomaly [20]. This is why we did not show the reconstruction in the case of Heisenberg and XX models for simplicity. Next, we remark on two surprising properties of P_{rec}. One is that $\frac{L}{2\pi}P_{rec}$ is close to taking integer values at low energy. The other is that it is often the case that $-\pi \leqslant P_{rec} \leqslant \pi$. Both properties are expected when knowing $T \approx e^{-\mathbf{i}P_{rec}}$. However, from Eq. (4.30) alone, it is unclear why this happens; we leave this for future research. Finally, although we lack a full analytic understanding of the factor $\frac{\mathbf{i}}{\pi^2}$ in P_{rec}, it works well numerically. The precise nature of this factor on a finite system, and whether this is a correct choice or merely a coincidence, is left for further exploration.

Further directions— The further directions are sorted based on some subjective measure of attractiveness.

(a) We show that CFT ground states satisfy the condition in Eq. (4.1). Motivated by entanglement bootstrap, we ask whether the converse is also true, i.e. if a state satisfying the condition can be interpreted as a CFT ground state. Additionally, we would like the statement to be robust, useful even when the state only satisfies the condition approximately.

(b) We proposed various formulas that hold exactly when the state is a CFT ground state. However, for the ground state of critical lattice models, the formula only holds approximately. What is the convergence behavior and the finite size scaling?

(c) Relatedly, one may wonder if we allow combinations of not only 1-site and 2-site EH but also 3-site EH, which combination has the best approximation to the Hamiltonian and the momentum. Furthermore, what if we include EH on larger sites? Could this lead to a sequence of reconstructions that converges to the actual operator? Finding a good approximation is crucial for obtaining a better estimation of the OPE based on the method discussed in [24].

(d) Can we reproduce the key results in CFT directly using quantum information?

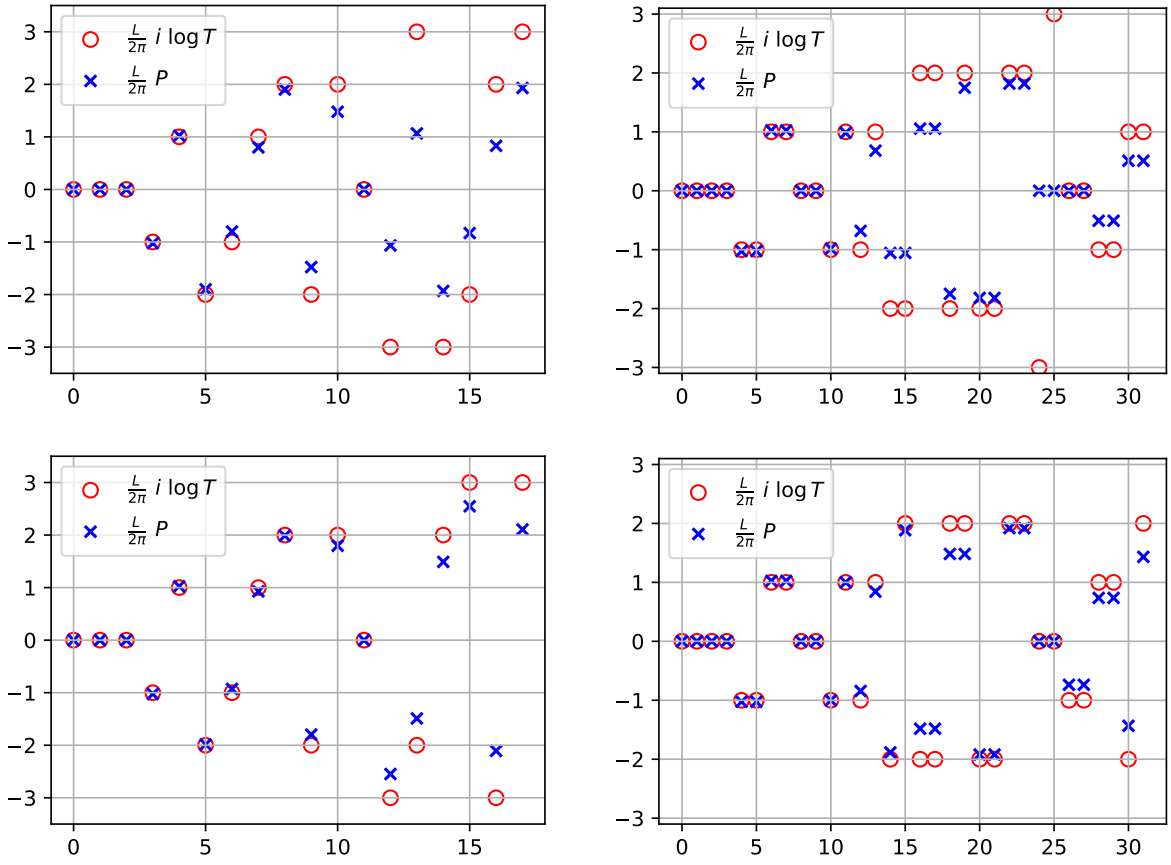

Figure 4.4: Spectra of $\mathbf{i} \log T$ and the reconstructed momentum operator P, where T is the translation operator by 1 site. The states are ordered by energy along the X-axis. In the infinite-size limit, this order corresponds to the order of the scaling dimension of the corresponding operator from the state-operator correspondence. For the small system sizes shown here, these orders are not yet the same. This is why the order changes in the figure when the system size changes. Top left: Ising model with $L = 8$. Bottom left: Ising model with $L = 12$. Top right: Potts model with $L = 6$. Bottom right: Potts model with $L = 8$.

For example, can we show that all nontrivial states must have $c \geqslant \frac{1}{2}$? Can we show the ground state energy is $-\frac{2\pi}{L}\frac{c}{12}$ on a circle with circumference L? Can we show that the reconstructed momentum operator should be approximately integer valued at low energy and has a norm roughly below π?

(e) The appearance of $h(\eta)$ in Eq. (4.3) suggests that η can be interpreted as a probability. Could we find a physical meaning for this observation?

(f) We provide an information theoretic criterion for CFT ground state. Can we transform this criterion into an algorithm that searches for CFTs by screening states that satisfy Eq. (4.1)? We believe the answer is yes and plan to explore this aspect in an upcoming paper.

(g) We provide an entropy criterion that appears to describe ground states for 1+1D unitary CFTs. We suspect a similar formula exists for 1+1D chiral CFTs, d+1D CFTs, and perhaps complex CFTs [25, 26].

(h) We show that CFT ground states are critical points of the function S_Δ, with a value proportional to the central charge. What is its relation to the Zamolodchikov c-function [27] and the entropic c-function [28]? For example, it is known that the second and third-order expansion of the Zamolodchikov c-function near the CFT point contains the CFT data, including the scaling dimensions and the OPEs. Does the same hold for the entropy function we proposed?

(i) We showed that 1+1D phases at RG fixed points with Lorentz symmetry satisfy Eq. (4.1). What happens to other RG fixed points without Lorentz symmetry, such as those with dynamical critical exponent $z \neq 1$. In the case of free fermion models with $z \in \mathbb{Z}^{>0}$, the formula continues to hold. When z is even, the ground state is simply a product state, and when z is odd, the ground state is same as the ground state for $z = 1$, which is a CFT.

4.5 Acknowledgments

We thank Isaac Kim and Hao-Chung Cheng for suggesting a simple proof to compute the first-order derivative of entanglement entropy in Eq. (4.8). We thank Meng Cheng and Yijian Zhou for clarifying why in general $T \neq e^{-\mathbf{i}P}$. We thank Bowen Shi and Daniel Ranard for helpful discussions. This work was supported in part by funds provided by the U.S. Department of Energy (D.O.E.) under the cooperative research agreement DE-SC0009919, and by the Simons Collaboration on Ultra-Quantum Matter, which is a grant from the Simons Foundation (652264, JM). TCL would like to thank the assistance of ChatGPT in improving parts of the paper for clarity, including this sentence.

Dissertation Note

Chapter 4 is based on joint work with John McGreevy, published as "Conformal field theory ground states as critical points of an entropy function," *Phys. Rev. Lett.*, 2023.

Bibliography

[1] Alexei Kitaev and John Preskill. Topological entanglement entropy. *Physical review letters*, 96(11):110404, 2006.

[2] Michael Levin and Xiao-Gang Wen. Detecting topological order in a ground state wave function. *Physical review letters*, 96(11):110405, 2006.

[3] Hui Li and F Duncan M Haldane. Entanglement spectrum as a generalization of entanglement entropy: Identification of topological order in non-abelian fractional quantum hall effect states. *Physical review letters*, 101(1):010504, 2008.

[4] Isaac H Kim, Bowen Shi, Kohtaro Kato, and Victor V Albert. Chiral central charge from a single bulk wave function. *Physical Review Letters*, 128(17):176402, 2022.

[5] Karthik Siva, Yijian Zou, Tomohiro Soejima, Roger S. K. Mong, and Michael P. Zaletel. Universal tripartite entanglement signature of ungappable edge states. *Phys. Rev. B*, 106(4):L041107, 2022.

[6] Christoph Holzhey, Finn Larsen, and Frank Wilczek. Geometric and renormalized entropy in conformal field theory. *Nuclear physics b*, 424(3):443–467, 1994.

[7] Pasquale Calabrese and John Cardy. Entanglement entropy and quantum field theory. *Journal of statistical mechanics: theory and experiment*, 2004(06):P06002, 2004.

[8] Pasquale Calabrese and John Cardy. Entanglement entropy and conformal field theory. *Journal of physics a: mathematical and theoretical*, 42(50):504005, 2009.

[9] Bowen Shi, Kohtaro Kato, and Isaac H Kim. Fusion rules from entanglement. *Annals of Physics*, 418:168164, 2020.

[10] Bowen Shi and Isaac H Kim. Entanglement bootstrap approach for gapped domain walls. *Physical Review B*, 103(11):115150, 2021.

[11] Jin-Long Huang, John McGreevy, and Bowen Shi. Knots and entanglement. *arXiv preprint arXiv:2112.08398*, 2021.

[12] Ting-Chun Lin and John McGreevy. Supplemental material. 2023.

[13] Elliott H Lieb and Mary Beth Ruskai. Proof of the strong subadditivity of quantum-mechanical entropy. *Les rencontres physiciens-mathématiciens de Strasbourg-RCP25*, 19:36–55, 1973.

[14] Michael A Nielsen and Isaac Chuang. Quantum computation and quantum information, 2002.

[15] David Perez-Garcia, Frank Verstraete, Michael M Wolf, and J Ignacio Cirac. Matrix product state representations. *arXiv preprint quant-ph/0608197*, 2006.

[16] John Cardy and Erik Tonni. Entanglement hamiltonians in two-dimensional conformal field theory. *Journal of Statistical Mechanics: Theory and Experiment*, 2016(12):123103, 2016.

[17] Helmut Reeh and Siegfried Schlieder. Bemerkungen zur unitäräquivalenz von lorentz-invarianten feldern. *Il Nuovo Cimento (1955-1965)*, 22(5):1051–1068, 1961.

[18] Ashley Milsted and Guifre Vidal. Extraction of conformal data in critical quantum spin chains using the koo-saleur formula. *Physical Review B*, 96(24):245105, 2017.

[19] Daniel Ranard, Michael Walter, and Freek Witteveen. A converse to lieb–robinson bounds in one dimension using index theory. In *Annales Henri Poincaré*, volume 23, pages 3905–3979. Springer, 2022.

[20] Meng Cheng and Nathan Seiberg. Lieb-schultz-mattis, luttinger, and't hooft–anomaly matching in lattice systems. *arXiv preprint arXiv:2211.12543*, 2022.

[21] Yijian Zou, Bowen Shi, Jonathan Sorce, Ian T Lim, and Isaac H Kim. Modular commutators in conformal field theory. *Physical Review Letters*, 129(26):260402, 2022.

[22] Ingo Peschel and Viktor Eisler. Reduced density matrices and entanglement entropy in free lattice models. *Journal of physics a: mathematical and theoretical*, 42(50):504003, 2009.

[23] Maurizio Fagotti and Pasquale Calabrese. Entanglement entropy of two disjoint blocks in xy chains. *Journal of Statistical Mechanics: Theory and Experiment*, 2010(04):P04016, 2010.

[24] Yijian Zou, Ashley Milsted, and Guifre Vidal. Conformal fields and operator product expansion in critical quantum spin chains. *Physical Review Letters*, 124(4):040604, 2020.

[25] Victor Gorbenko, Slava Rychkov, and Bernardo Zan. Walking, Weak first-order transitions, and Complex CFTs. *JHEP*, 10:108, 2018.

[26] Victor Gorbenko, Slava Rychkov, and Bernardo Zan. Walking, Weak first-order transitions, and Complex CFTs II. Two-dimensional Potts model at $Q > 4$. *SciPost Phys.*, 5(5):050, 2018.

[27] Alexander B Zamolodchikov. Irreversibility of the flux of the renormalization group in a 2d field theory. *JETP lett*, 43(12):730–732, 1986.

[28] Horacio Casini and Marina Huerta. A c-theorem for entanglement entropy. *Journal of Physics A: Mathematical and Theoretical*, 40(25):7031, 2007.

[29] David Sutter, Mario Berta, and Marco Tomamichel. Multivariate trace inequalities. *Communications in Mathematical Physics*, 352:37–58, 2017.

Chapter 5

A Systematic Search for Conformal Field Theories in Very Small Spaces

Chapter 5 is based on joint work with Xiang Li and John McGreevy (arXiv:2509.04596).

with a tweezer in hand, i dive

the water is vast—
more than i can hold

the tweezer opens into a probe
the water lifts into contour

i search for the saddle

in its hollow, glimmers gather
some hidden gem, some ordinary stone

i gather them all,
their virtue to be known

5.1 Introduction

Conformal field theories (CFTs) are the endpoints of renormalization group flows of relativistic quantum field theories, govern the universal information at continuous phase transitions, and in the case of 1+1 dimensions, are a crucial ingredient in perturbative string vacua. The space of CFTs remains an object of mystery. Even in 1+1 dimensions, where conformal symmetry implies that a theory represents an infinite-dimensional Virasoro algebra [1], a classification of unitary CFTs only exists when the Virasoro central charge c is less than 1 [2]. Our knowledge of CFTs with $c \geqslant 1$ relies on ad hoc constructions, and solvable models with extra symmetries (see e.g. [3]). As a result, all well-understood explicit constructions produce CFTs that are rational with respect to some chiral algebra (Virasoro or one of its extensions). Despite a general belief that the generic CFT is irrational, and evidence from holographic duality for the existence of many irrational CFTs, an explicit example of an irrational CFT is not yet known with certainty (but see e.g., [4, 5, 6, 7, 8] and references therein for a candidate construction of a few such theories).

The Entanglement Bootstrap is a program to understand the universal properties of quantum many-body states from their entanglement structure, and it provides a new perspective on quantum field theory (QFT). Let us define a universal property of a state to be a property of a renormalization group fixed point, shared by its whole basin of attraction. The idea is that the universal information in a groundstate is that of its low-energy QFT. In the case of quantum critical lattice models, this low-energy QFT is a CFT.

The thus-far-most-well-developed aspect of Entanglement Bootstrap is its application to liquid gapped bosonic phases [9, 10, 11, 12, 13, 14, 15, 16], where the low-energy field theory is a topological quantum field theory (TQFT). A key ingredient of Entanglement Bootstrap for gapped states is a set of local conditions on the state that guarantee that the state is a good representative of a gapped liquid topological order. From these local

conditions, much of the machinery of TQFT has been constructed.

In [17], we proposed an analogous local condition for 1+1 dimensional CFT groundstates,

$$K_\Delta \left|\psi\right\rangle \propto \left|\psi\right\rangle, \;\; K_\Delta \equiv \eta \hat{\Delta}(A,B,C) + (1-\eta)\hat{I}(A:C|B), \tag{5.1}$$

$$\hat{\Delta}(A,B,C) \equiv K_{AB} + K_{BC} - K_A - K_C, \tag{5.2}$$

$$\hat{I}(A:C|B) \equiv K_{AB} + K_{BC} - K_B - K_{ABC} \tag{5.3}$$

where A, B, C are three adjacent intervals, and η is the cross-ratio specified by their four endpoints. The fact that a CFT groundstate satisfies this condition, which we call the vector fixed-point equation (VFPE), follows from known properties of the CFT groundstate entanglement Hamiltonian of an interval [18, 19]. Conversely, there is strong evidence that a state satisfying the VFPE actually must be a fixed point of the renormalization group[1]. Amongst the evidence for this conjecture, we mention its strong consequences for the form of the entanglement Hamiltonian [20], that a reformulation of this condition [21] can logically result in the emergence of conformal geometry, as well as the ability to extract expressions for all six global conformal generators (including the Hamiltonian) and Virasoro generators [22]. This VFPE, although is derived from continuum CFT, works surprisingly well for critical lattice models, even in very small system sizes with small local Hilbert spaces. This is not true for all results derived from field theory, due to strong finite size effect. The VFPE is designed in such a way that finite-size effects and UV physics are suppressed to some extent. This is the key reason that we can impose this condition on small systems to study universal properties of phases of matter.

The VFPE turns out [17] to be the condition that the state is a critical point of a certain linear combination of entropies of three contiguous regions, $S_\Delta \equiv \langle\psi|K_\Delta|\psi\rangle$. This

[1]We note that not all such fixed points are CFTs; some of them have zero correlation length, rather than infinite correlation length. We will explain below how to distinguish these two classes of states.

suggests the possibility that S_Δ can behave like a Morse function on the space of 1+1d states, similarly to the famous Zamolodchikov c-function [23], which is instead defined (for relativstic QFT) in terms of correlation functions of local operators, and in terms of which the renormalization group is gradient flow. The relationship between our entropy function and the renormalization group is still open, but this picture has an exciting practical consequence: it suggests a systematic algorithm to enumerate CFTs. A nice feature of this algorithm is that it requires no input about symmetry or of a choice of Hamiltonian. In particular, there is no reason that CFTs found this way should be rational for some chiral algebra.

In this paper, we take steps toward implementing this algorithm. A crucial ingredient is the ability [17] to reconstruct a local Hamiltonian from the groundstate. The reconstructed Hamiltonian, on a circle with L sites, is

$$H^\psi_{\text{rec}} = \frac{\cot\frac{\pi}{L}}{L}\sum_i \left(K_{i,i+1} - K_i\right). \tag{5.4}$$

where K_X is the entanglement Hamiltonian for region X from the state $|\psi\rangle$. This Hamiltonian has several virtues. First, it is a local Hamiltonian; in fact it involves only nearest-neighbor interactions. This means that from a small system, we can reconstruct the Hamiltonian for a large system. Second, H_{rec} is a non-negative operator. This can be seen by write it as

$$H^\psi_{\text{rec}} = \frac{\cot\frac{\pi}{L}}{2L}\sum_i \hat{\Delta}(i-1,i,i+1), \tag{5.5}$$

where $\hat{\Delta} \geqslant 0$ by operator weak monotonicity [24]. Third, if the state is the groundstate of a CFT, and each site is identified with an interval in the continuum, it is exactly the Hamiltonian of the CFT on the circle, with the normalization chosen so that the eigenvalues are $2\pi\Delta_i/L$ where $\{\Delta_i\}$ are the CFT scaling dimensions,

$$(H^\psi_{\text{rec}} - E_0)|\Delta_i\rangle = \frac{2\pi}{L}\Delta_i|\Delta_i\rangle, \tag{5.6}$$

where E_0 is the ground state energy. Elsewhere [20] we will show that a state $|\psi\rangle$ satisfying the fixed point equation is the groundstate of its reconstructed Hamiltonian H^{ψ}_{rec}[2]. The spectrum of H_{rec} computed from the groundstates of a critical quantum lattice model matches surprisingly well with that of the original Hamiltonian even for small system sizes. The fixed point equation for $|\psi\rangle$ also implies $z = 1$ scale invariance of H^{ψ}_{rec} [17].

5.2 Strategy for enumeration of CFTs

With confidence from the numerical tests in [17, 25], we are ready to describe an algorithm that potentially enumerates all unitary CFTs. The algorithm contains two parts. The first part generates approximate CFT ground states $|\psi\rangle_{ABCD}$ on a circle of 4 evenly partitioned intervals, A, B, C, D. We replace each interval with a single 'site', a Hilbert space of finite dimension, d, which play the role of a control parameter in our search. We proceed by finding states that are the critical points of S_Δ with $\eta = 1/2$. We search within the set of translation-invariant states on four sites. For now, we also restrict our search to time-reversal-invariant, i.e. real, states (see Figure 5.1).

The second part of the procedure generates approximate CFT Hamiltonians from the ground states using (5.4). Once we have the Hamiltonian, we can obtain the CFT data, including scaling dimensions and OPEs, through known algorithms [24]. Note that the purpose of this algorithm is not to give accurate values for the CFT data. The purpose is to generate as many CFTs as possible, including previously unknown ones. If one wants accurate values, improvements to the Hamiltonians can be made, and conformal bootstrap [26] can be used afterward to turn these approximate values into tighter bounds.

[2]The assumptions required to show the groundstate result are stronger than those used to enumerate CFT in this paper. Specifically, proving the groundstate result requires taking the limit of many sites, $L \to \infty$, and demands that the fixed-point equation hold for all possible three intervals. Whereas in this paper, we rely only on the case $L = 4$.

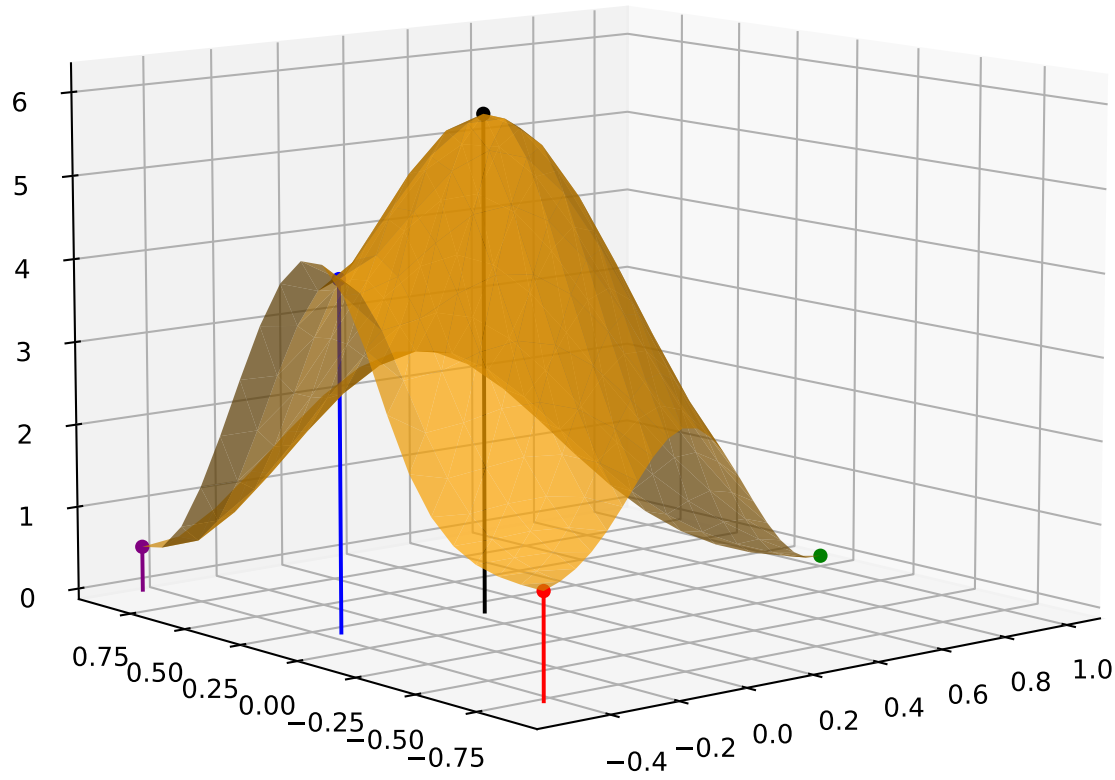

Figure 5.1: A two-dimensional cross section of the function S_Δ on the space of translation-invariant real wavefunctions on four qubits. Remarkably, five of the critical points listed in Table 5.1 can be seen, including trivial (green), Ising (purple), Heisenberg (red), Fibonacci (blue), and maximum (black).

Before continuing, let us introduce the operator $\hat{c}_\Delta$ proportional to H_{rec}:

$$\hat{c}_\Delta(\psi) \equiv \frac{3}{2\log 2}(K_{AB} - K_B + K_{BC} - K_C + K_{CD} - K_D + K_{DA} - K_A) \quad (5.7)$$

with expectation value

$$c_\Delta(\psi) \equiv \langle\psi|\hat{c}_\Delta(\psi)|\psi\rangle. \quad (5.8)$$

Since $\hat{c}_\Delta \propto H_{\text{rec}}$, we have $\hat{c}_\Delta|\psi\rangle \propto H_{\text{rec}}|\psi\rangle \propto |\psi\rangle$ for CFT groundstates. The advantage of this normalization is that when $|\psi\rangle$ is the ground state of the field theory, c_Δ is exactly the central charge of the corresponding CFT. Thus, equivalently, the goal in the first part is to find all the critical points of c_Δ.

Search. The process of finding the critical point consists of two steps. First, we find an initial point close to the critical point. Notice that we cannot just do gradient descent on c_Δ, because we want to find all critical points, not just the local minima.[3] Generally, performing gradient descent only gives local minima. On the other hand, it is indeed possible to perform gradient descent over the error function[4], which is the standard deviation of $\hat{c}_\Delta$[5]:

$$\text{err} \equiv \sigma(\hat{c}_\Delta) \equiv \sqrt{\langle\psi|\left(\hat{c}_\Delta(\psi)\right)^2|\psi\rangle - \langle\psi|\,\hat{c}_\Delta(\psi)\,|\psi\rangle^2}. \quad (5.9)$$

Second, we use a second-order Newton's method to converge to the critical point. Some further details of the search procedure are given in App. 5.5.

Note that c_Δ is invariant under local unitary transformations. Thus, each critical point is actually an orbit of this gauge redundancy.

[3]If instead the goal is to find CFTs with spectral gap > 2 (also known as dead-end CFTs, since they cannot be perturbed by relevant operators), then it may be sufficient to study the local minima.

[4]Note that when acting on a translation-invariant state, H_{rec} and K_Δ are proportional to each other by a purification move (e.g., $K_{ABC}|\psi\rangle = K_D|\psi\rangle$).

[5]Note that our choice of normalization differs from that used in [17].

Analysis. Once the critical points are identified, we perform a series of diagnostics to determine whether their properties are consistent with those of a conformal field theory.

As a preliminary step, we discard states with degenerate ground states, since it is believed that a CFT has a unique ground state.

Next, we check whether the low-energy spectrum of the reconstructed Hamiltonian H_{rec} exhibits universal behavior as the system size L increases. Let $E_{L,i}$ denote the i-th eigenvalue of H_{rec} on L sites. In the absence of relevant perturbations, one expects that for fixed i, the limit $\lim_{L\to\infty} L \cdot E_{L,i}$ exists. In fact this condition is also satisfied even if a small relevant perturbation is present, as illustrated in In Fig. 5.9. Thus, the diagnostic is to test whether the low-lying spectrum appears to decrease at the rate $1/L$ for L that does not exceed the correlation length.

We also examine whether the ground state exhibits the entanglement entropy scaling expected of a CFT, namely $S(\ell) = \frac{c}{3} \log \sin \frac{\pi \ell}{L} + \text{const}$. The diagnostic is to test whether a plot of $S(\ell)$ vs $\log \sin \frac{\pi \ell}{L}$ yields a straight line.

In both cases, we use periodic MPS to obtain the low-energy spectrum and corresponding eigenstates.

5.3 Results

Table 5.1 and Figure 5.2 list the critical states we found over 4 qubits and 4 qutrits respectively. In the case of qutrits, we perform sampling – it might not be a complete list, but a list of solutions we have found. In both cases, we restrict our search to translation-invariant and time-reversal-invariant states. For qubits and qutrits, respectively, this space has 5 and 23 dimensions.

Table 5.1: All critical 4-qubit states. Actual CFTs are in blue. We defined $|\alpha\rangle \equiv (|0011\rangle+|0110\rangle+|1100\rangle+|1001\rangle)/2$ and $|\beta\rangle \equiv (|1010\rangle+|0101\rangle)/\sqrt{2}$ as elements of our basis of translation-invariant states. The coefficients for the Ising solutions are $[a_1, a_2, a_5, a_6] = [0.11036, 0.80152, 0, 0, -0.51103, 0.29023]$. In fact, there are two inequivalent maximum point state with the same H_{rec} and S_Δ that are not related by on-site unitaries.

c_Δ	Description	Explicit form up to on-site unitary
0	cat states	$a\|0000\rangle + b\|1111\rangle$
0.526	Ising CFT	$a_1\|0000\rangle + a_2\|1111\rangle + a_5\|\alpha\rangle + a_6\|\beta\rangle$
1.132	W state	$\frac{1}{2}(\|0001\rangle + \|0010\rangle + \|0100\rangle + \|1000\rangle)$
1.211	XX model	$-\frac{1}{\sqrt{2}}\|\beta\rangle + \frac{1}{\sqrt{2}}\|\alpha\rangle$
1.245	Heisenberg	$-\sqrt{\frac{2}{3}}\|\beta\rangle + \frac{1}{\sqrt{3}}\|\alpha\rangle$
3.510	ferromagnet	$\frac{1}{\sqrt{3}}\|0000\rangle + \frac{1}{\sqrt{6}}(\|0111\rangle + \|1011\rangle + \|1101\rangle + \|1110\rangle)$
4.165	Fibonacci chain	$\frac{1}{\sqrt[4]{5}}\left(\|0000\rangle + \sqrt{\sqrt{5}-1}\,\|\beta\rangle\right)$
6.000	maximum point	$\frac{1}{2}(\|0000\rangle \pm \|1111\rangle + \|0101\rangle + \|1010\rangle)$

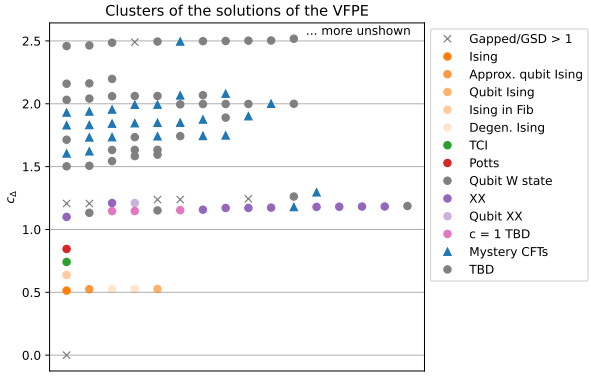

Figure 5.2: Critical 4-qutrit states. We do not display the states with larger values of c_Δ. "Degen. Ising" refers to a state with the spectrum of the critical Ising model, but with two-fold degeneracy. "Ising in Fib." refers to a state whose spectrum on L sites is the first ℓ_L states of the Ising model below a gap, where ℓ_L is the Lth Lucas number. States labelled 'gap' are gapped fixed points. H_{rec} of the states marked with an x have a groundstate degeneracy on 4 sites and were discarded.

Qubits. The entries in the table are all RG fixed points, but some of them are zero-correlation-length fixed points. Actual CFTs are in blue[6]. These two classes of states can be distinguished by reconstructing the local Hamiltonian from the 4-qubit state and studying it on larger system sizes. For the XX and Heisenberg solution, we discover that the reconstructed Hamiltonian are exactly the well-known Hamiltonians $H = XX + YY$ and $H = XX + YY + ZZ$ respectively up to a prefactor and an additive constant. In fact, this prefactor, which gives the exactly correct scaling behavior of the Hamiltonian, is a bonus from using the reconstructed Hamiltonian – one does not need to manually scale the spectrum. The reconstructed Hamiltonians from the gapped fixed point states on qubits all have a groundstate degeneracy.

Some comments about the zero-correlation-length fixed points that appear in the table: The groundstate degeneracy and c_Δ have a peculiar relation, i.e. $\log(\text{degeneracy}) = 0.46 c_\Delta$. There are only two eigenstates with finite energy for the reconstructed Hamiltonian of the W state: one ground state is the W state, and the other is the product state $|0\rangle^{\otimes n}$. This is consistent with the theorem of [27], which forbids a local Hamiltonian whose *unique* groundstate is the W state. The reconstructed Hamiltonians for the states labelled 'ferromagnet' and 'Fibonacci' are (up to an additive constant):

$$H_{\text{Ferromagnet}} \propto -\sum_i \left(X_i X_{i+1} + Y_i Y_{i+1} + Z_i Z_{i+1}\right) \tag{5.10}$$

$$H_{\text{Fibonacci}} \propto \sum_i (|11\rangle\langle 11|)_{ii+1} \tag{5.11}$$

which explains the nomenclature. The latter is well-known to produce the golden chain [28, 29, 30, 8], which at system size $L = 2, 3, 4, 5, 6, 7, ...$, has a ground state degeneracy $\ell_L = 3, 4, 7, 11, 18, 29, ...$, where ℓ_L are the Lucas numbers. The maximum point of S_Δ on

[6]There exist translation-invariant states that are equivalent under an on-site unitary, but not under a translation-invariant on-site unitary. For example, two realizations of the Ising CFT, $a_1|0000\rangle + a_2|1111\rangle + a_5|\alpha\rangle \pm a_6|\beta\rangle$ are related by the unitary $I \otimes Z \otimes I \otimes Z$.

four qubits is something like a perfect tensor, where $S_2 = 2 \log 2$ attains its maximal value for any two contiguous sites.

It is not a coincidence that the three CFTs that fit into the 4-qubit Hilbert space in Table 5.1 are ones for which we know a parent Hamiltonian acting on nearest-neighbor qubits. Our reconstructed Hamiltonian always acts on 2 sites at a time. The question of which CFTs will fit in a given local Hilbert space, with our nearest-neighbor reconstructed Hamiltonian, is closely related in spirit to the c-d conjecture of [31], which says that a lattice model of local Hilbert space dimension d and nearest-neighbor interactions can only realize a CFT with $c \leqslant d - 1$. The results so far are consistent with that conjecture. In fact, our search strategy can in principle provide a rather stringent test of that conjecture: In our search with local dimension d, we should find no CFTs with $c > d - 1$.

Qutrits. Compared to Table 5.1, when searching in the space of translation-invariant, real states on four qutrits, we find more CFTs (Figure 5.2). In the list for qutrits, there are multiple representatives of the Ising model. One is the same state as realized on qubits, making no use of the third register (labelled 'Ising qubits'), and the other is a new state whose value of c_Δ is closer to $1/2$. Interestingly, there is a solution which encodes the Ising spectrum in the Fibonacci chain groundsate subspace.

We find representatives of the next two unitary minimal models after Ising, with $c_\Delta = 0.744$ (tricritical Ising has $c = 7/10$) and $c_\Delta = 0.844$ (3-state Potts has $c = 4/5$), as well as many more points along the $c = 1$ line. At larger c we find many more (of order 25) states whose reconstructed spectrum looks like that of a CFT, possibly with small relevant perturbation. In Appendix 5.6 we give details of the information extracted from the solutions of the VFPE on qutrits.

Further discussion. What determines the order in which CFTs appear in our search, as we increase the local Hilbert space dimension? There is a preference for small values of c, which, since the reconstructed Hamiltonian always acts on precisely two sites, is

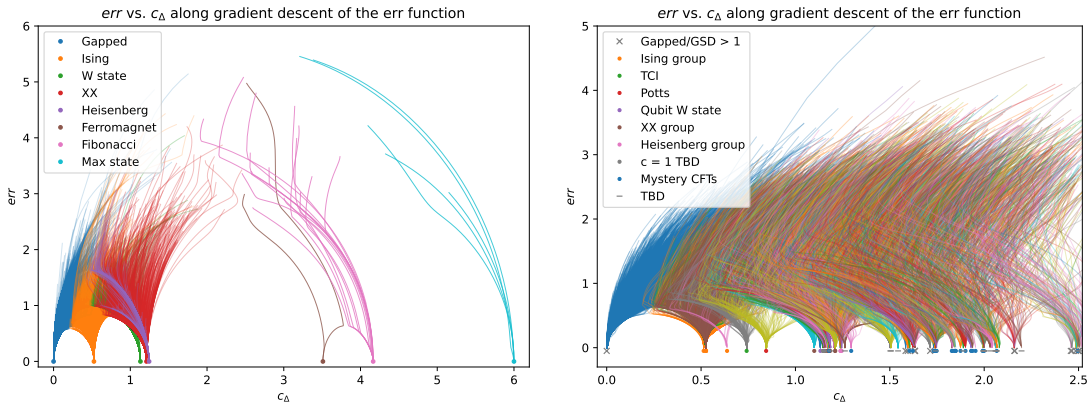

Figure 5.3: Error of the VFPE versus c_Δ along the trajectory of the gradient descent of the error function on 4 qubits (top) and 4 qutrits (bottom). We first do a uniform sampling in the space of translation-invariant and time reversal symmetric states. Then we use the groundstates of their reconstructed Hamiltonians H_{rec} [Eq. (5.4)] as the initial states in the gradient descent. There are forbidden regions for c_Δ when the errors of VFPEs are small. The names in the figure legends will be explained below.

consistent with the spirit of the c-d conjecture of [31]. We observe that CFTs with primaries of small dimension do not fit in small Hilbert spaces: In Figure 5.5 we show the error of the VFPE in the XXZ model as a function of Δ, which determines the radius of the $c = 1$ boson; when the radius takes extreme values (near 0 or ∞) there are light primaries and the error grows. This explains the slow appearance of the unitary minimal models in our search as we increase d.

Approximate solutions of the VFPE. In addition to the exact or numerically-exact solutions found above, it is also interesting to study approximate solutions of the VFPE. One reason we are led to study such approximate solutions is the behavior of CFTs with exactly marginal operators in our search: on four qubits, we find two special points in the $c = 1$ moduli space, as isolated critical points of S_Δ, despite the fact that the XXZ

model

$$H_{\mathrm{XXZ}} = \sum_i \left(X_i X_{i+1} + Y_i Y_{i+1} + \Delta Z_i Z_{i+1}\right). \tag{5.12}$$

is a nearest-neighbor Hamiltonian on qubits that realizes a whole family of $c = 1$ CFTs with different radii. We conclude that the error of the VFPE in these states is a finite-size effect. By relaxing our search criteria to states with *small* error of the VFPE, we can indeed see the whole $c = 1$ moduli space (though we have not yet found the Ginsparg archipelago [32, 33] likely because they have several primaries with small scaling dimensions), see Figure 5.5. The pattern of approximate solutions also contains information about the space of CFTs. In particular, one can see voids in the set of states in the plot of error versus c_Δ around known forbidden regions, see Figures 5.3 and 5.4. There could be some quantum-information-theoretic reason that forbids states in these areas: we conjecture that, when $\sigma(\hat{c}_\Delta) \leqslant \varepsilon$, the values of c_Δ are constrained to take a set of specific values with a small window of deviations. It is also suggestive that we can observe similar voids in regions for $c > 1$ where there is so far no known condition forbidding unitary CFTs.

5.4 Discussion

c_Δ as a c-function? Starting at a saddle point $|\psi\rangle$ of c_Δ identified with a CFT, we can parameterize nearby states as $|\psi\rangle + |d\psi\rangle$. In turn we can expand $|d\psi\rangle$ in a basis of eigenstates of H^ψ_{rec}, which, by the state-operator correspondence of CFT, correspond to operators of definite scaling dimension, $|\Delta_i\rangle$. Thus, we can ask about the change in c_Δ upon deforming the state by a given $|\Delta_i\rangle$. It is natural to ask about a possible relation between the spectrum of H^ψ_{rec} and the spectrum of the Hessian of c_Δ. The example of CFTs with exactly marginal operators shows that the match is not precise. However, we observe numerically in examples that the number of relevant operators is equal to the number of downhill directions of c_Δ, that is, negative eigenvalues of its Hessian. A visceral

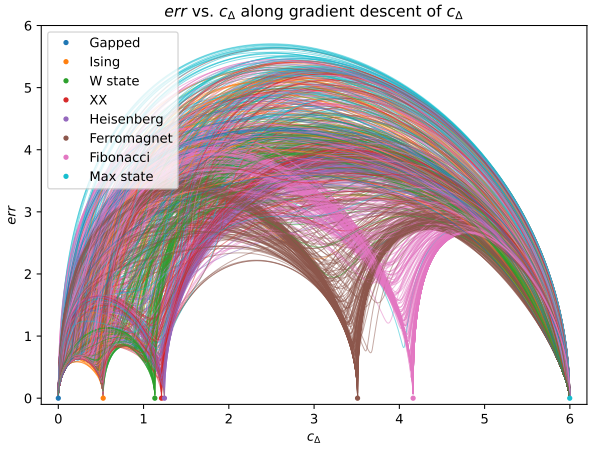

Figure 5.4: Each curve represents a trajectory of gradient descent on c_Δ, flowing from right to left. We generate these trajectories by starting with a solution of the VFPE on 4 qubits, adding a small random perturbation, and then perform gradient descent and ascent on c_Δ to complete each trajectory. Notice, for example, the flow from XX to Ising.

representation of the structure of this function is given in Figure 5.4, where we perturb the solutions of the VFPE and do gradient flow by c_Δ. From the plot one can read out flows between fixed points.

The set of CFTs. Does a solution of the VFPE on four sites encode all the data of the associated CFT? This seems quite believable, though turning this into a universal algorithm remains challenging. An overly optimistic possibility is that the Hamiltonian reconstructed from such a state, when extended to large system size, would automatically realize the CFT in the IR limit. While this is true for some of the states (such as Heisenberg and XX on qubits), most of the states we find produce H_{rec} with a finite correlation length (of order $15-30$ lattice spacing), even when the spectrum of H_{rec} for small sizes agrees with the known list of scaling dimensions. We can systematically improve upon this situation in many ways. A simple method is: start from the groundstate of H^{ψ}_{rec} on 8 sites of dimension d, block the sites into 4 supersites of dimension d^2, and use Newton's method to find a lower-error state on the larger local Hilbert space.

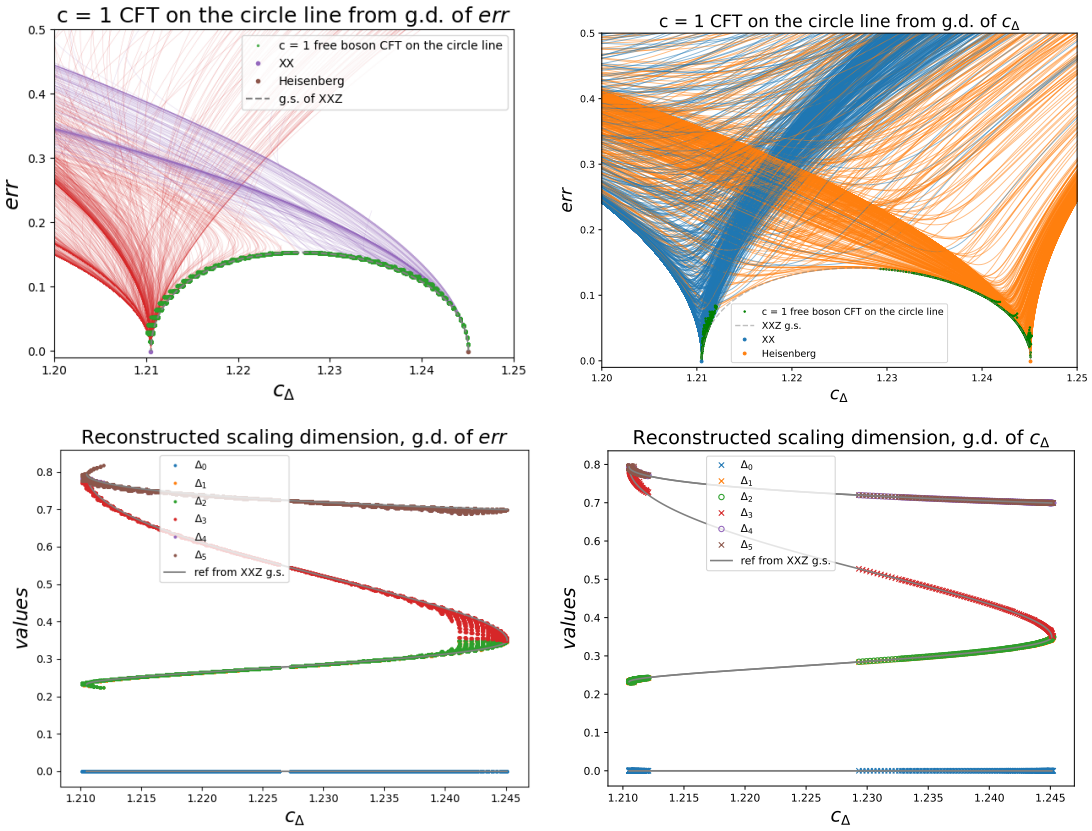

Figure 5.5: Top: Gradient descent near the $c = 1$ fixed points on qubits, XX ($c_\Delta = 1.211$) and Heisenberg ($c_\Delta = 1.245$). Top left is gradient descent on the error, and top right is gradient descent on c_Δ. The void between these two points is a finite-size effect. Also shown (in gray) are (c_Δ, err) in groundstates of the XXZ model as Δ varies from 0 to 1, which interpolate between XX and Heisenberg; for intermediate Δ on 4 sites, there is a finite error. The green dots are obtained by selecting the states whose (c_Δ, err) is close (with deviation less than 10^{-3}) to that of an XXZ groundstate. In this way we match a value of c_Δ with a value of the parameter Δ in H_{XXZ}. Bottom: The spectra of H_{rec} from states at the boundary of the void of each of the two top plots, respectively. These spectra agree with those of the $c = 1$ theory on the circle line and that of the XXZ model (5.12) at the values of Δ determined by the green dots.

In addition to well-understood states with $c < 1$ and $c = 1$, in the qutrit search we identify several candidate CFT states in the range $1 < c < 2$ (see Tables 5.2 to 5.10). These states clearly deserve further attention. It would also be interesting to study further the measure on the set of CFTs defined by the canonical measure on the 4-site Hilbert space. Our search strategy provides a rich source of potential counterexamples to the c-d conjecture of [31] – if we had found a smoking-gun CFT state with $c > 2$ in our qutrit search it would have disproved this conjecture.

Criticality in small spaces. We should emphasize an important practical point. H_{rec} is a useful tool for learning from small systems: If you just study the spectrum of an arbitrary lattice Hamiltonian on four sites it doesn't look like CFT, the normalization is arbitrary. You need to get the stress tensor right to fix the normalization. But this is generally a very high level in the spectrum and costly to find. The fact that $H_{\text{rec}} = \sum_i \lambda_i K_i$ is completely fixed by $\sum_i \lambda_i \beta(x) \stackrel{!}{=} 1$ means that its normalization is already fixed, as in (5.4). Moreover, H_{rec} for the optimized states we find gives the best realization of the continuum CFT in the given Hilbert space, often better than our favorite lattice model. We have thus provided a method to improve upon a given quantum critical Hamiltonian with the given local Hilbert space dimension.

The error of the VFPE also provides a simple practical tool for drawing phase diagrams of a given family of Hamiltonians. Scanning a phase diagram for minimal error of the VFPE detects fixed points of the RG. In Figure 5.6 we give an example of this method. More examples of this application of the VFPE can be found in [34].

Acknowledgments

We thank Paul Fendley, Tarun Grover, Janet Hung, Yannick Meurice, Daniel Parker, Brandon Rayhaun, Shu-Heng Shao, German Sierra, and the participants of the 2025

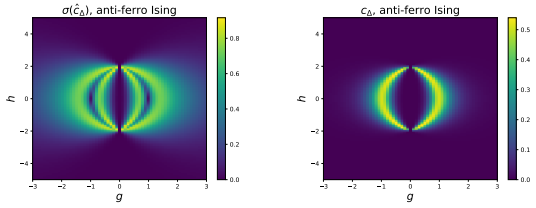

Figure 5.6: The error of the VFPE and the value of c for the *anti*-ferromagnetic Ising chain with transverse and longitudinal fields: the *anti*-ferromagnetic Ising chain with transverse and longitudinal fields: $H = +\sum_{j=1}^{L} Z_j Z_{j+1} - h_X \sum_j X_j - h_Z \sum_j Z_j$ with periodic boundary conditions, $L = 8$, and $A = \{1, 2\}, B = \{3, 4\}, C = \{5, 6\}$. We can see two gapped fixed points with $c = 0$ separated by a wall of Ising transitions with $c \approx 0.5$. After a sublattice rotation to map to a ferromagnetic Ising model in a staggered field, it has a hidden Ising symmetry acting by $\prod_j X_j \otimes$ translation by one site.

Conformal Bootstrap School in São Paulo, especially Marco Meineri, João Penedones, Leonardo Rastelli, Slava Rychkov, Balt van Rees, Pedro Vieira, for useful discussions and comments. We thank Shing-Tung Yao for providing inspiration to finish the paper in a timely manner. DMRG calculations were done using the iTensor library [35, 36] in Julia. This work was supported in part by funds provided by the U.S. Department of Energy (D.O.E.) under cooperative research agreement DE-SC0009919, and by the Simons Collaboration on Ultra-Quantum Matter, which is a grant from the Simons Foundation (652264, JM). JM received travel reimbursement from the Simons Foundation; the terms of this arrangement have been reviewed and approved by the University of California, San Diego in accordance with its conflict of interest policies.

5.5 Appendix: Details of search procedure

The goal is to obtain (ideally all) the critical points of the function $c_\Delta = \langle \hat{c}_\Delta \rangle$ on the space of translation-invariant and time-reversal symmetric states on 4 qudit sites. In

the scope of this paper, we choose each site to be a qubit or a qutrit. After we obtain all the critical points, we then want to single out those that correspond to a CFT, which will be referred to as states of interest.

The search process contains mainly three components: (1) generator, (2) clusterer, (3) analyzer. The flow charts can be found in Figures 5.7 and 5.8.

The generator is responsible for producing critical states with small error to be analyzed. This involves sampling from a prior distribution and using gradient descent and Newton's method to find states with small error.

The clusterer groups states with the same physical properties together, allowing for a more efficient analysis by reducing the number of states that need to be considered.

The analyzer performs a detailed examination of the critical states to determine whether they correspond to CFTs. This is achieved through periodic matrix product state (MPS) simulations. In addition, the procedure enables the estimation of key quantities, including central charges and scaling dimensions. The states can also be used to estimate OPE data [37], though this will be left for future work.

Generator. The generator mainly contains (1) a sampler that produces initial translation invariant and time-reversal symmetric states. (2) a gradient descent optimizer that refines the states to minimize the error until the error is less than 10^{-4}. (3) a Newton's method optimizer that further refines the states to minimize the error until machine precision.

In between, we employ various filters to throw away states that are not desirable. This includes states with $c_\Delta \approx 0$ (trivial states) and states that converge too slowly.

(1) We first produce a basis $\{|s\rangle\}$ for the translation invariant states on 4 sites. A state can be expressed as $|\psi\rangle = \sum_s x_s |s\rangle$, with x_s being real so that the state is time-reversal symmetric. We then sample the coefficients x_s from a normal distribution and normalize x_s, which amounts to uniformly sampling from a hypersphere.

(2) We then perform gradient descent to minimize the error function. Before doing gradient descent directly on the error function $\sigma(\hat{c}_\Delta)$, on the initial states $|\psi\rangle_{\text{init}}$ we first construct its reconstructed Hamiltonian and then compute its groundstate $|\psi_0\rangle$. This step can significantly reduce $\sigma(\hat{c}_\Delta)$ as well as reduce c_Δ, which can improve the frequency with which we obtain states of interest. Starting from $|\psi_0\rangle$, we perform gradient descent on the error function $\sigma(\hat{c}_\Delta)$. To improve computational efficiency, we use numerical differentiation, namely:

$$\partial_s \sigma = \frac{\sigma(|\psi + \varepsilon\rangle_s) - \sigma(|\psi - \varepsilon\rangle_s)}{2\varepsilon}, \tag{5.13}$$

where $|\psi \pm \varepsilon\rangle_s = (|\psi\rangle \pm \varepsilon\,|s\rangle)/\mathcal{N}_\pm$ is the norm-preserving variation with the translation invariant basis $|s\rangle$ and $\mathcal{N}_\pm$ is the normalization factor. In practice, we choose $\varepsilon = 10^{-5}$. We've compared the numerical differentiation with the analytic differentiation, and we found that choosing ε of the magnitude 10^{-5} has the smallest discrepancy. The numerical differentiation is much faster compared with the analytic differentiation. Moreover, since the goal of this step is just to reduce the error to a small amount rather than use it to obtain the exact solution, the numerical error introduced by the numerical differentiation is acceptable. Based on these considerations, we choose to use the numerical differentiation.

With the the gradient of $\sigma(|\psi\rangle) = \sigma[\hat{c}_\Delta(|\psi\rangle)]$, which is of the form $\nabla\sigma(|\psi\rangle) = \sum_s \partial_s \sigma(|\psi\rangle)\,|s\rangle$, we can update the state

$$|\psi\rangle_{\text{new}} = |\psi\rangle - \gamma \cdot \nabla\sigma(|\psi\rangle). \tag{5.14}$$

The step size γ is chosen through backtracking line search algorithm. Briefly speaking, we first choose a (relatively large) step size γ_{test} to update the state, examine the decrease of the error. Then we gradually reduce the step size by a factor α: $\gamma_{\text{test}} \to \alpha\gamma_{\text{test}}$, comparing the magnitude of the error decrease with the previous, until we find the best step size in the sequence of tests. From the tests, we found that $\alpha = 0.5$ gave the best performance. Since the landscape we explore is quite rugged, we need to be careful about the step size

to ensure convergence. This algorithm will choose a larger step size when the error is large, which speeds up the computation. When the error is small, the number of tests to find the optimal step size will increase and a small step size will be selected.

(3) Once the error becomes sufficiently small $\sigma(\hat{c}_\Delta) < 10^{-4}$, the gradient descent method becomes inefficient and we switch to Newton's method to accelerate convergence.[7] Traditionally, Newton's method is applied near local minima to achieve quadratic convergence. However, it is also known that Newton's method can be applied near saddle points to locate critical points. We adopt this approach to our search procedure.

The Hessian is computed through analytic computations. The time cost for analytically computing the Hessian is similar to that using numerical differentiation. Moreover, since the task in this stage is to obtain an (almost) exact solution up to numerical precision, analytic computation has a relatively small numerical error. If required, we can further reduce the numerical error by changing the machine precision, but we did not find that improving the precision of the numerical estimate of the critical points of the VFPE improved the results for physical quantities. Once the Hessian A is computed, we can update the state by

$$|\psi\rangle_{\text{new}} = |\psi\rangle - \beta \cdot A^{-1} \cdot (\nabla c_\Delta), \tag{5.15}$$

where the gradient of c_Δ can be computed as $\nabla c_\Delta = \hat{c}_\Delta |\psi\rangle - c_\Delta |\psi\rangle$. Via several tests, we found that $\beta \in (10^{-1}, 10^{-2})$ is the best choice for efficiency. One potential issue is that A might contain zero eigenvalues. We regulated the computation of the inverse by adding an identity matrix with a small factor: $A \to A + \delta I$, with $\delta \sim 10^{-7}$. At the end of this

[7] Intuitively, the Hessian of c_Δ is believed to be related to the scaling dimensions of the CFT, assuming that c_Δ is related to the Zamolodchikov c-function. This suggests the condition number, the ratio between the largest and smallest eigenvalues of the Hessian, is rather poor. Such ill-conditioning may explain the slowness of the current gradient descent method.

Alternatively, one could employ variants of gradient descent with momentum estimation. Whether these variants perform better is deferred to future study.

step, we typically obtain $\sigma(\hat{c}_\Delta) < 10^{-8}$, which is approaching the machine precision for computing $\sigma(\hat{c}_\Delta)$.

In practice, Newton's method is time-consuming. To address this, we clustered the results from the gradient descent step and applied Newton's method only to a representative state from each cluster. The clustering algorithm is described below.

Clusterer. The clusterer uses c_Δ as the key to group states together and outputs one representative state for each group. When the error $\sigma(\hat{c}_\Delta)$ is small enough, as shown in Fig. 5.3, the values of c_Δ is clustered. In our numerical results, when $\sigma(\hat{c}_\Delta) \sim 10^{-4}$, in each cluster the standard deviation of c_Δ is of order 10^{-7}. States in the same cluster are believed to have similar properties, with deviations decreasing as $\sigma(\hat{c}_\Delta)$ further decreases.

We tested this conjecture by computing the spectrum of $K_{i,i+1} - K_i$ for the states in a cluster and examining their difference. Since $K_{i,i+1} - K_i$ (the reconstructed Hamiltonian density) are the same up to numerical errors, it is plausible that the reconstructed Hamiltonian Eq. (5.4) will have the same properties. In fact, in various cases of states of interest, states in the same cluster are simply related to each other by a product of on-site unitaries (up to numerical error), which correspond to the same CFTs. For example, two states in the Ising cluster on qubits at $c_\Delta = 0.526$ in Table 5.1 with $\alpha_6 \to -\alpha_6$ are related by the on-site but not translation-invariant unitary $\psi \to Z_B Z_D \, |\psi\rangle$.

In extreme settings where c_Δ attains the maximal allowed value, the states are no longer related by on-site unitaries, nevertheless, the reconstructed Hamiltonian and its groundstate still have the same properties. We also tested multiple representatives from clusters at larger c_Δ and found the same spectrum of the extended H_{rec} on $L > 4$ sites.

Analyzer. For each cluster, we pick a representative state and check if its properties are consistent with those of a quantum critical groundstate. We first employ a simple filter that rejects states whose H_{rec} on 4 sites has degenerate ground states.

We then study the low-energy behavior of H_{rec} on many sites to approach the

thermodynamic limit to extract the universal data. In particular, we want the energy and the eigenstates. Due to the complexity in representing quantum states, we use MPS to represent these low-energy states. We were able to find the first 8 eigenstates at system sizes $L = 4, 6, 8, 10, 12, 14$ for the first 88 clusters of states, up to $c_\Delta \approx 2.5$. For certain states, we also compute $L = 16$ when the spectrum depends on $L\text{mod}4$.

Additionally, for each system size, we compute the groundstate entanglement entropy of intervals of size $\ell = 1..L/2$. The computation for each cluster takes about 2.5 hours on a MacBook Pro 2021 with Apple M1 Max Chip, but is much faster for gapped states. We also did the computation for a sample of clusters at larger c_Δ.

5.6 Appendix: Details of qutrit search results

In Tables 5.2 to 5.18 we present an analysis of the first 88 solutions of the VFPE on qutrits that pass the small-system filters (the states themselves can be found in the supplementary material with the arxiv submission). For each of these states, we study H^{ψ}_{rec} on a chain of length $L \in \{4, 6, 8, 10, 12, 14, (16)\}$ with periodic boundary conditions using the DMRG algorithm (or exact diagonalization for some of the smaller systems). We find the first 8 eigenstates, and the entanglement entropy $S(\ell)$ in the groundstate of intervals of size $\ell = 1..L/2$. Based on this information, we group the states into three groups: states that we believe correspond to CFTs (Tables 5.2 to 5.10), states that we believe do not correspond to CFTs (Tables 5.15 to 5.18), and states where we cannot draw a conclusion at the moment (Tables 5.11 to 5.14). For the states that we believe correspond to CFTs, we show the spectrum of the reconstructed dilatation operator, $\frac{L}{2\pi} H_{\text{rec}}$, which (see (5.6)) is normalized so that the numbers should correspond to the list of scaling dimensions of the CFT. In the middle of the plot, we show the spectrum normalized so that $E_i = 1$, where $i = 1, 2$ or 3 is chosen so that the spectra collapse best with system size. On the right, we

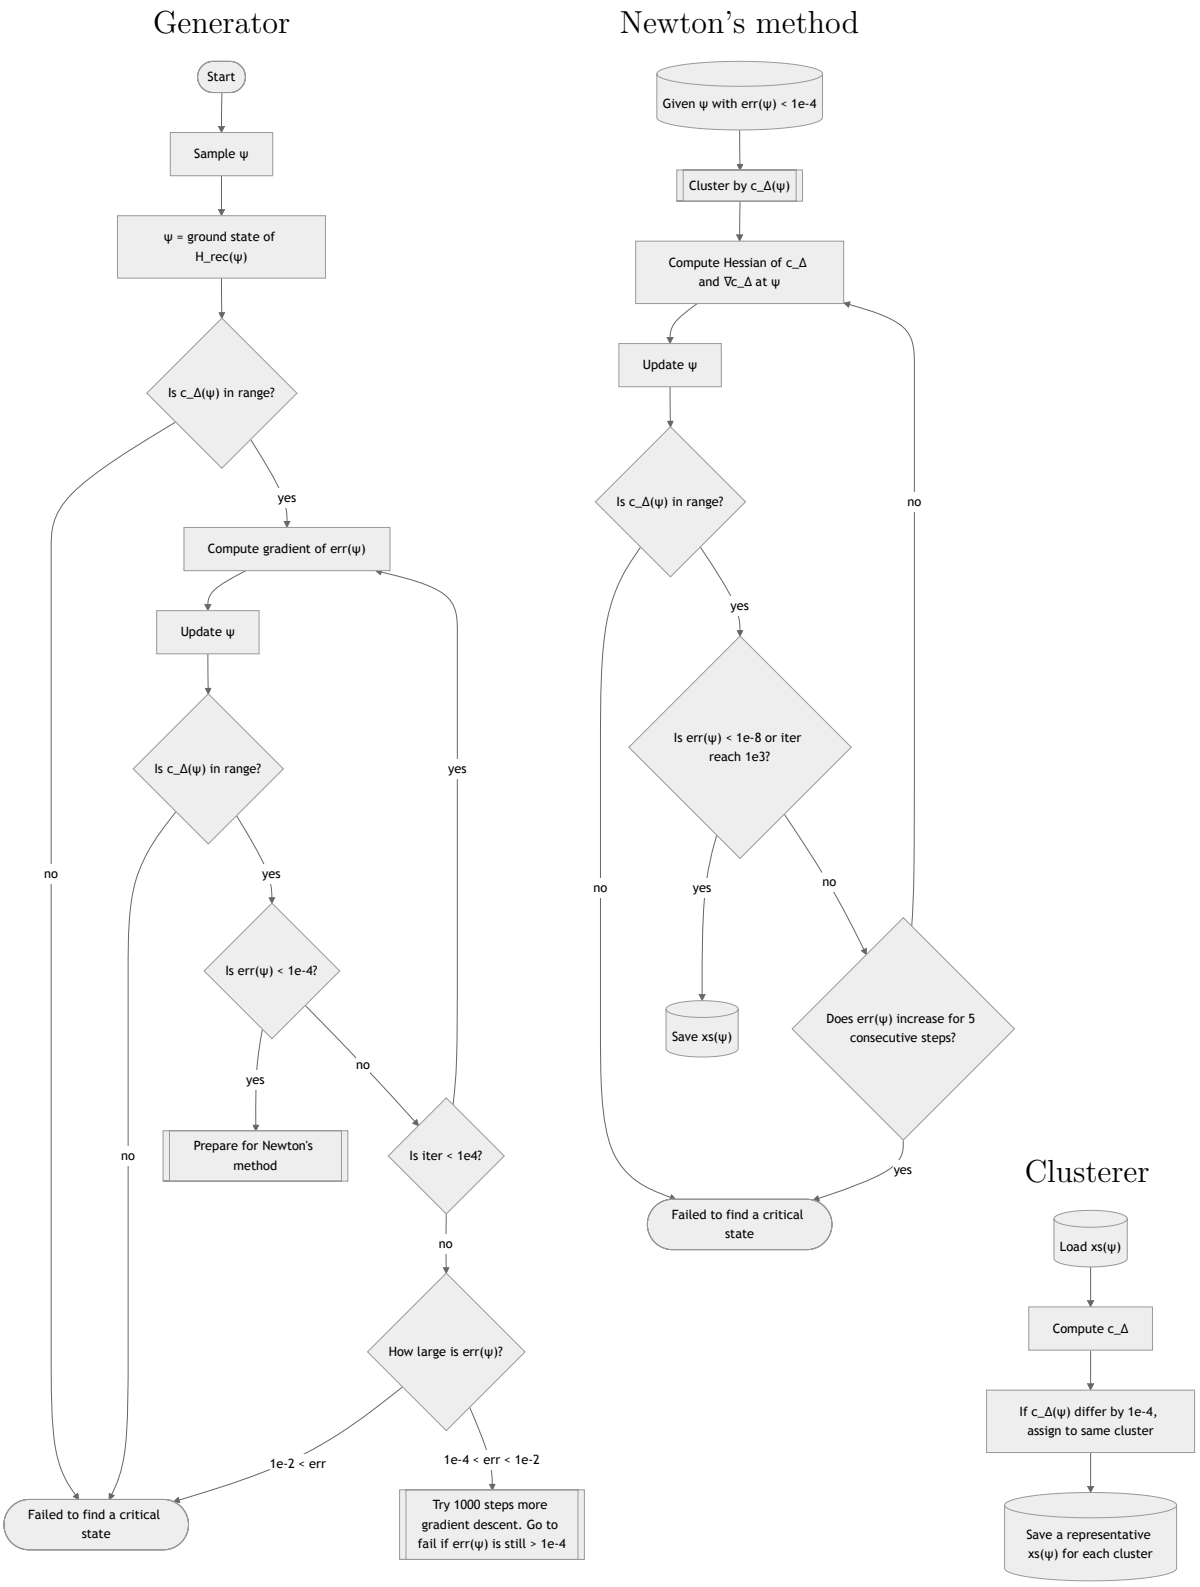

Figure 5.7: Flow diagrams of the generator and clusterer. The resulting output consists of a list of critical points.

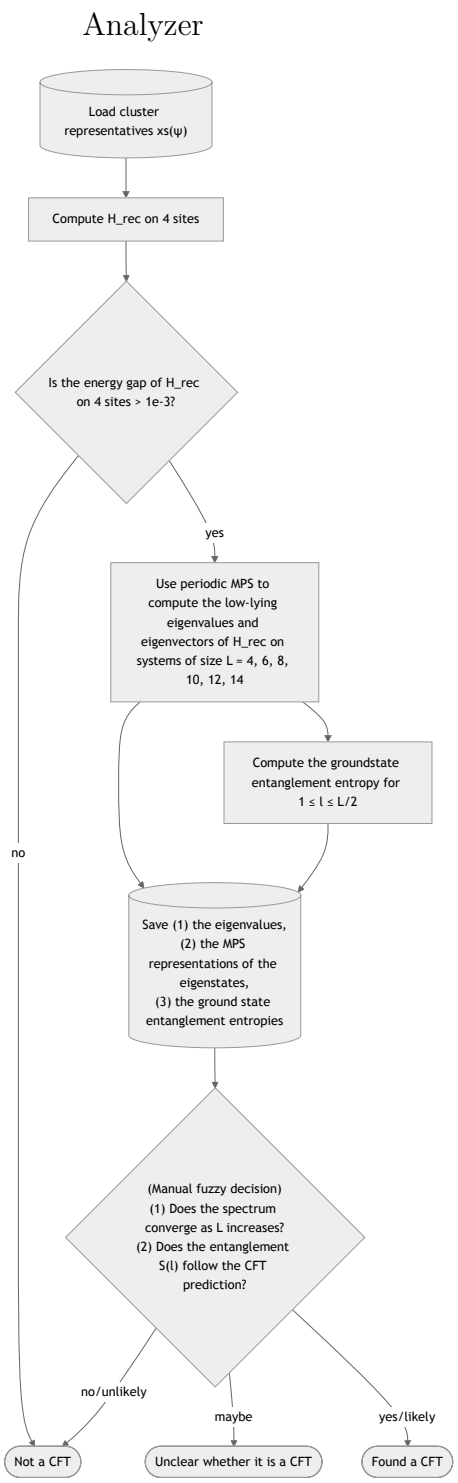

Figure 5.8: Flow diagram of the analyzer. The resulting output consists of a list of CFT candidates together with their associated properties, including the spectrum. The three final outcomes correspond to the three sets of tables at the end of the chapter.

show the entanglement entropy $S(\ell)$ versus $\log \sin \frac{\pi \ell}{L}$, whose slope should give $c/3$ for a CFT state. In an inset of the plot, we show the values of

$$c_{21} \equiv 3\left(S(\ell=2) - S(\ell=1)\right) / (\log\left(\frac{\sin(2\pi/L)}{\sin(\pi/L)}\right), \tag{5.16}$$

which gives a UV-insensitive estimate of c that is less sensitive to the effects of a potential relevant perturbation than the fits to the slope of $S(\ell)$ up to half the system size. From the spread of these numbers we can infer the uncertainty in our estimate of the actual central charge of the ideal CFT.

Among the states that we believe are CFTs, we are able to identify several realizations of the Ising model, the tricritical Ising model (TCI), the Potts model, and many realization of CFTs with $c = 1$. Based on the spectra, most of the latter are either the XX point (where the first two excited states are degenerate at $\Delta = 0.25$) or the Heisenberg point (where the first three excited states are degenerate at $\Delta = 0.5$), though some realize radii in between these two values. We do not see a clear realization of the orbifold line of $c = 1$ theories in exact solutions of the VFPE on (qubits or) qutrits, though representatives will appear for $d = 4$ or if we allow finite error.

Farther down the list, we find 25 states that appear to correspond to CFTs with $c > 1$, but which we have not yet identified, so we label them 'Mystery CFTs'. Let us describe the criteria by which we selected these states. There are essentially two: (1) Does the spectrum behave like that of a CFT? In the leftmost plot of Figure 5.9, we show the finite-size behavior of the spectrum of the critical transverse-field Ising model Hamiltonian H_{TFIM}. The normalization is chosen so that at size L, the spectrum is that of aLH_{TFIM} where the constant a is chosen so that the state corresponding to ε is at 1 for $L = 4$. The spectrum is visibly converging to the correct CFT spectrum. In contrast, when we make a large relevant perturbation of the critical lattice model (as in the middle and right plots of Figure 5.9), the spacing between $E_i(L)$ for fixed level i becomes linear in L. We refer to

this as 'drift' of the spectrum, and infer that it indicates that the Hamiltonian is gapped. A direct comparison a state whose reconstructed spectra 'drift' (and therefore are not included) and one whose spectra do not can be seen in Figure 5.10.

(2) The second necessary condition for us to declare a state to be a CFT is that the entanglement entropy should behave like that of a CFT. Specifically, the fits of $S(\ell)$ to the expected CFT behavior should have high quality. For some of our states, the quality of these fits deteriorates at larger system sizes, which we take as a sign that the correlation length is in fact finite, and that H_{rec} is the critical Hamiltonian perturbed by a small relevant perturbation. This interpretation is corroborated by the fact that when we do infinite DMRG using the VUMPS method, the entanglement entropy of half the system plateaus as a function of bond dimension.

We emphasize that the input to our search has no connection to integrability or to rational CFT. So, although the CFTs we have managed to identify in our search all happen to be rational for some chiral algebra, we see no reason to expect that this will be the case for all the CFTs we have found. We have compared the reconstructed spectra with the list of known (rational) CFTs with $c \in (1, 2)$ (as helpfully listed in [8]), but did not yet find a good match. We note with some disappointment that the reconstructed spectrum of the state labelled Mystery CFT 1 with $c_\Delta = 1.29544$, whose c inferred from entanglement can be estimated as 1.15, does not match the spectrum found for the candidate irrational CFT of [8] with the same approximate c.

As was the case with the realizations of the Ising and $c = 1$ theories, not all of the solutions of the VFPE associated with CFTs appear to correspond to *distinct* CFTs. Thus, some of our Mystery CFTs are assigned the same number.

Both in the examples where we know the correct answer for c, and from our extractions of c from the entanglement, indeed we find that c_Δ always overestimates the correct value. We note that while some of the qutrit states identified as CFTs have $c_\Delta > 2$,

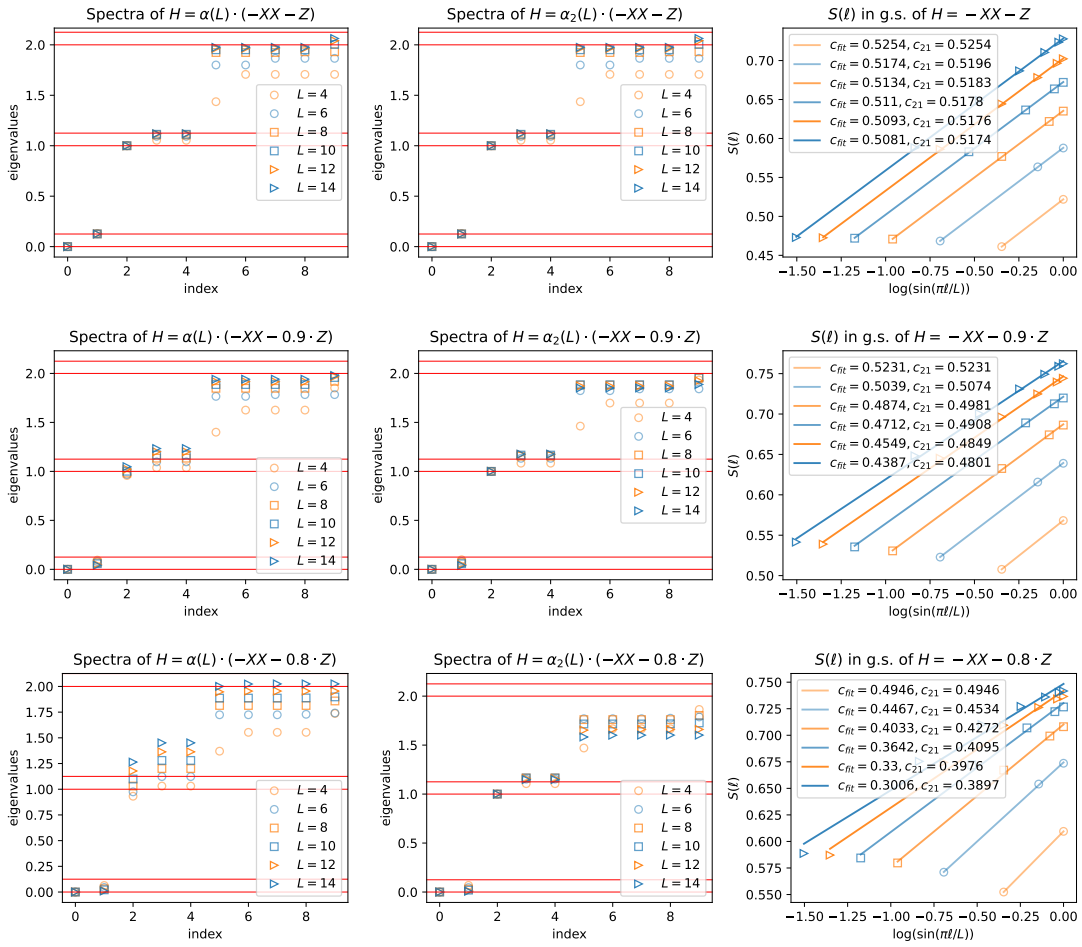

Figure 5.9: The effects of a relevant perturbation on the low-lying spectrum of a quantum critical lattice model. Top: the spectrum of the critical transverse-field Ising model $H = -\sum_i X_i X_{i+1} - h \sum_i Z_i$ at $h = 1$ at $L = 20$ with periodic boundary conditions, from DMRG. The normalization is chosen so that the ε state (third eigenstate) has energy 1.0 at $L = 4$. Middle: $h = 0.9$. Bottom: $h = 0.8$. Our purpose in showing this plot is for purposes of comparison with the spectra of H_{rec} from states that we identify as CFTs plus a relevant perturbation.

the inferred values of c from the entanglement are all less than 2, consistent with the c-d conjecture.

Amongst the states that we declare not to be CFTs, we do not distinguish between CFTs with a large relevant perturbation (where the short-distance behavior resembles CFT), and states which appear to have nothing to do with CFT at any scale (such as the W state).

In Figure 5.3 (bottom) and in Table 5.2, one can see a solution of the VFPE on qutrits that gives $c_\Delta \approx 0.63$, a place where no unitary CFT belongs. Surprisingly, the error can be made quite small, and the result of DMRG shows a state with a large correlation length! Closer inspection of the reconstructed spectrum shows that at system size L, H_{rec} has ℓ_L low-lying states below a large gap, where ℓ_L are again the Lucas numbers. Moreover, these low-lying states take the form of the lowest-energy states in the Ising spectrum. We infer that the reconstructed Hamiltonian is approximately $H_{\text{rec}} \approx \Gamma H_{\text{Fib}} + H_{\text{Ising}}$ where H_{Fib} is some realization of the Golden chain Hamiltonian (5.11) on qutrits, and Γ is a large number. As L grows, the spectrum rapidly becomes identical to the Ising CFT. Why c_Δ on 4 sites for this state is so far off we do not know.

We also studied a sampling of the many states with larger c_Δ. Among them, we did not find any states that look like good candidates for CFTs, consistent with the c-d conjecture.

In the tables of results, we define various normalizations of the Hamiltonian, $H = \alpha_i(L) H_{\text{rec}}$, with $\alpha_i(L)$ defined as follows: $\alpha_0(L) \equiv \frac{L}{2\pi}$. $\alpha_0(L)$ is the predicted theoretical prefactor, such that $D_{\text{rec}} \equiv \alpha_0(L) H_{\text{rec}}$ is the reconstructed dilatation operator, whose spectrum is the list of scaling dimensions of the CFT. $\alpha_i(L)$ is a prefactor such that the $i+1$-th level of the reconstructed Hamiltonian on L sites, $\alpha_i(L) H^L_{\text{rec}}$, is equal to the $i+1$-th level of D_{rec} on 4 sites. The other normalizations are an attempt to see if the spectra for different L collapse. In the table we show α_0 and the best collapse for $i = 1..3$.

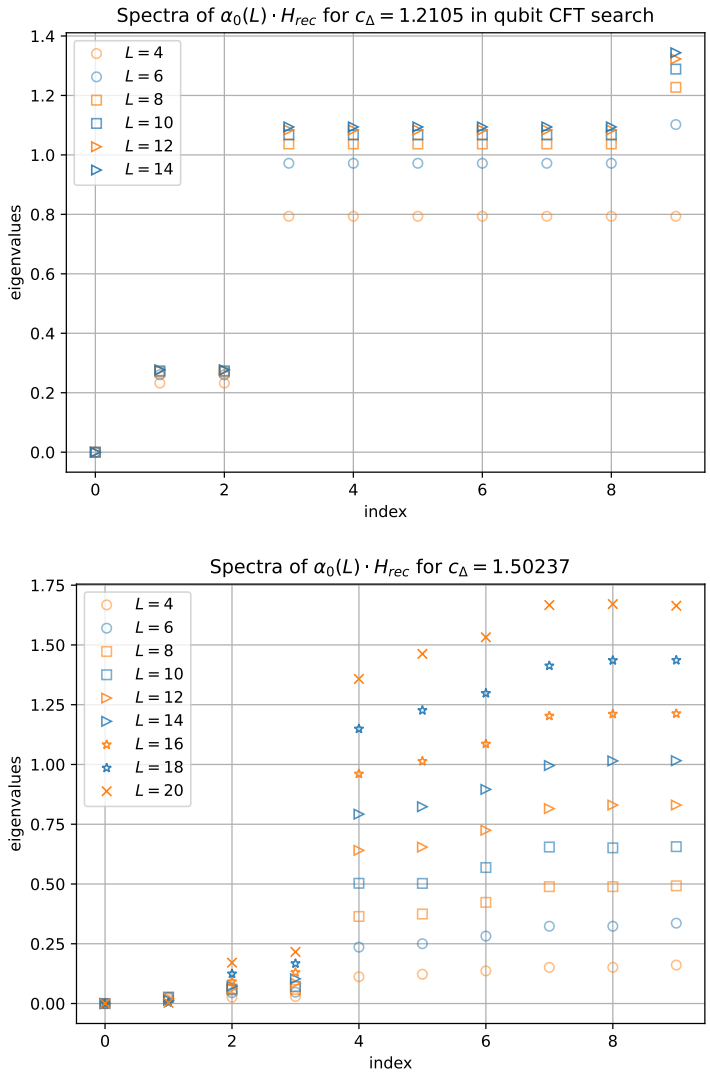

Figure 5.10: Left: The spectrum of $H_{\text{rec}} \cdot \frac{L}{2\pi}$ from the solution of the VFPE on 4 qubits associated with the XX model. The reconstructed Hamiltonian is that of the critical lattice model with infinite correlation length, and the spectrum is converging with system size. Right: The spectrum of $H_{\text{rec}} \cdot \frac{L}{2\pi}$ from the solution of the VFPE on 4 qutrits at $c_\Delta = 1.502373$ (labelled 'Gapped/GSD>1' in Figure 5.2). The spectrum (above the first four states) drifts up linearly with system size. States with this behavior are excluded from our list of CFT states Tables 5.2 to 5.10 and appear instead in Tables 5.11 to 5.18.

Table 5.2: States on four qutrits that we believe correspond to CFTs (part 1). $\star$ indicates that the spectrum depends on L mod 4.

c_{high}	name	spectra
0.51360	Ising	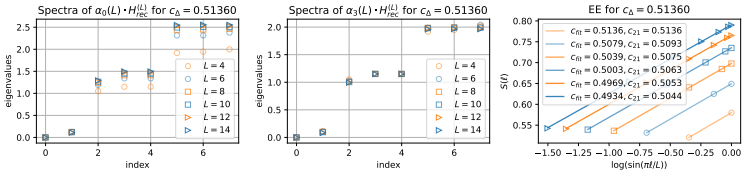
0.52445	$\approx$ qubit Ising	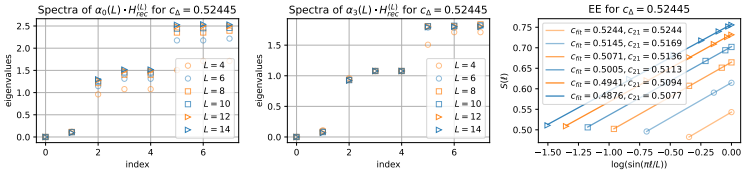
0.52639	qubit Ising	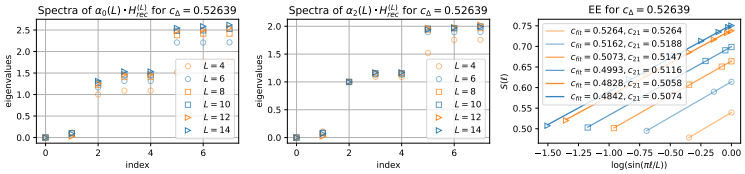
0.63776	Ising in Fib	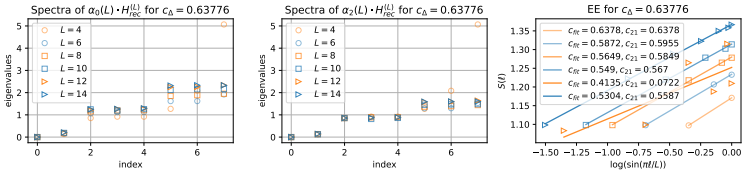
0.74138	TCI	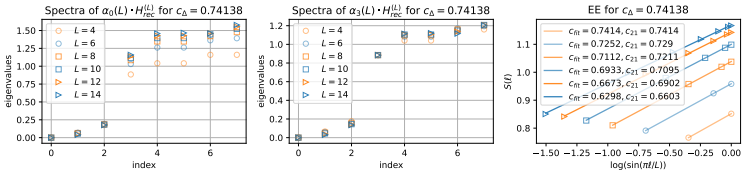
0.84475	Potts	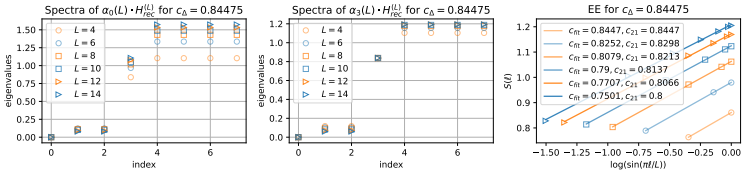

Table 5.3: States on four qutrits that we believe correspond to CFTs (part 2). $\star$ indicates that the spectrum depends on L mod 4.

c_{high}	name	spectra
1.09970	$c = 1$, XX	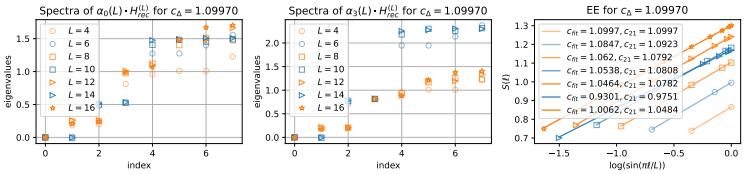 $\star$
1.14620	$c = 1$	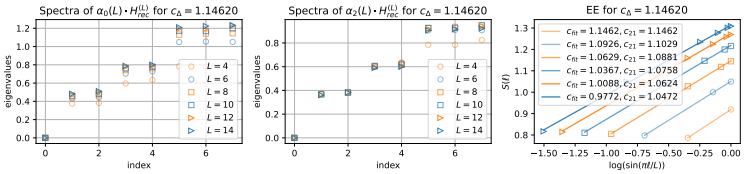 $R \in (R_{\text{XX}}, R_{\text{Heis}})$
1.14670	$c = 1$	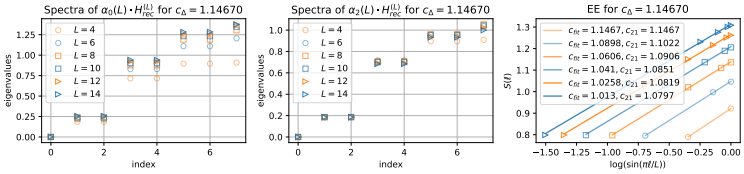 $R \in (R_{\text{XX}}, R_{\text{Heis}})$
1.15351	$c = 1$	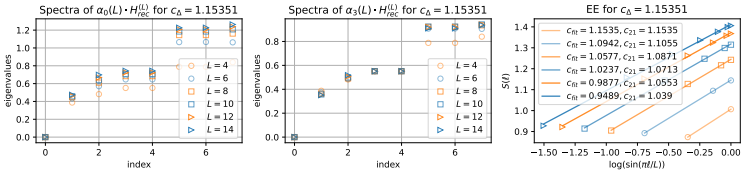 Heisenberg with relevant perturbation
1.15715	$c = 1$, XX	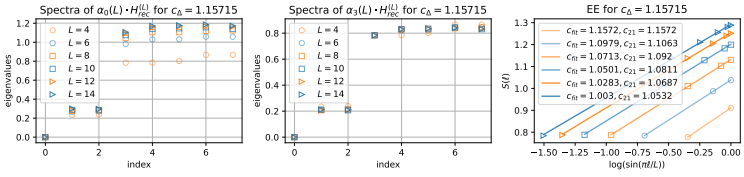
1.17090	$c = 1$, XX	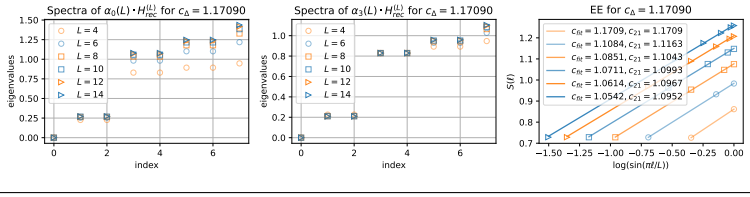

Table 5.4: States on four qutrits that we believe correspond to CFTs (part 3). $\star$ indicates that the spectrum depends on L mod 4.

c_{high}	name	spectra
1.17120	$c = 1$, XX	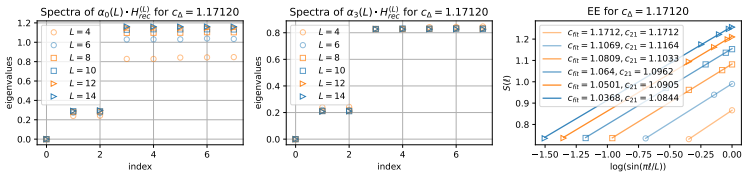 $R \in (R_{\text{XX}}, R_{\text{Heis}})$
1.17384	$c = 1$, XX	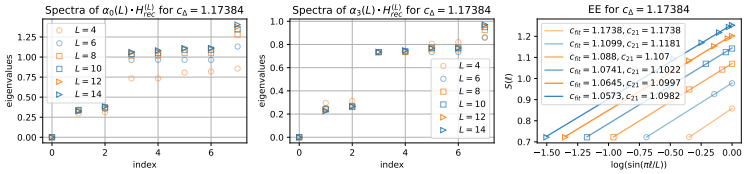
1.17828	Mystery CFT 0	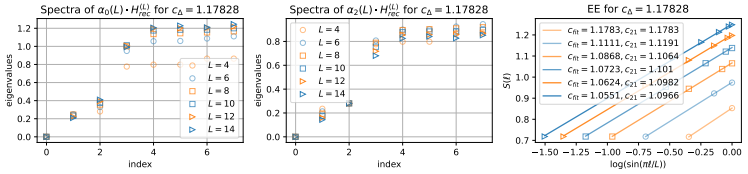
1.17936	$c = 1$, XX	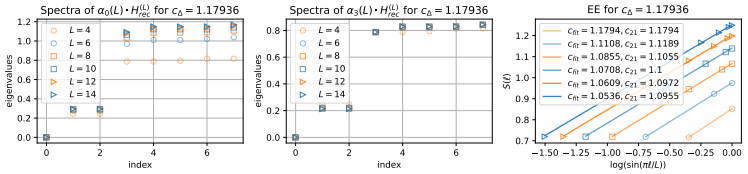
1.18098	$c = 1$, XX	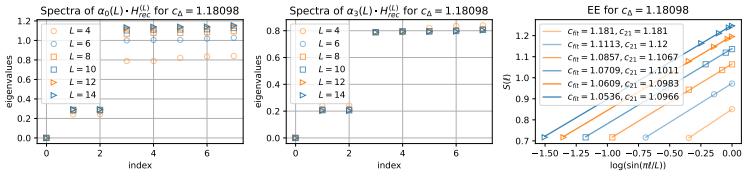
1.18233	$c = 1$, XX	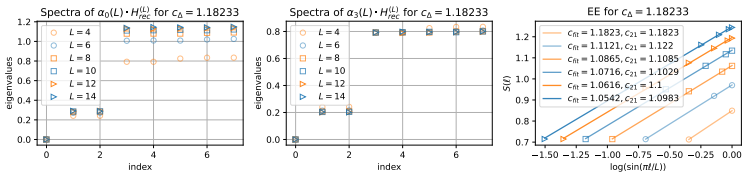

Table 5.5: States on four qutrits that we believe correspond to CFTs (part 4). $\star$ indicates that the spectrum depends on L mod 4.

c_{high}	name	spectra
1.18272	$c = 1$, XX	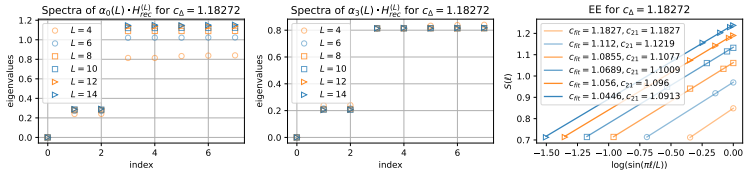
1.21030	$c = 1$, XX	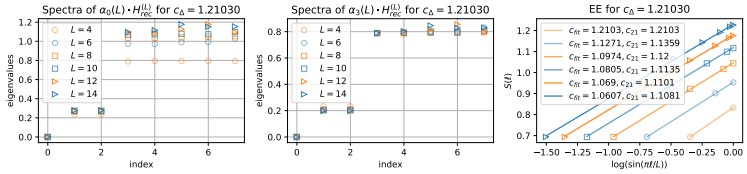
1.21051	$c = 1$, qubit XX	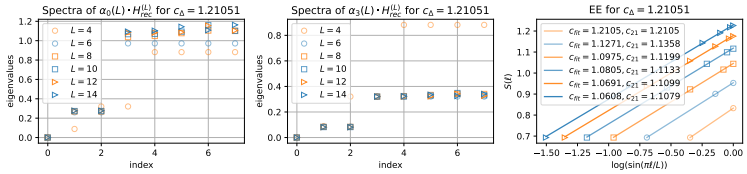
1.23873	$c = 1$, Heisenberg	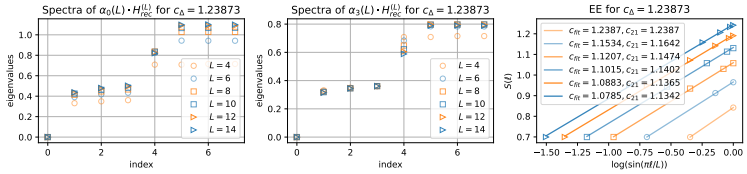
1.23906	$c = 1$, Heisenberg	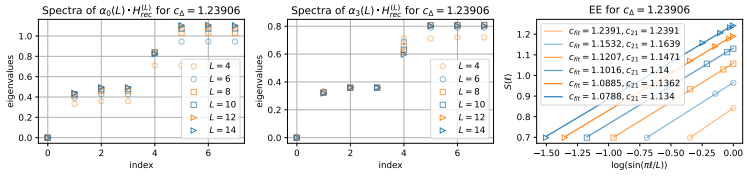
1.24511	$c = 1$, qubit Heisenberg	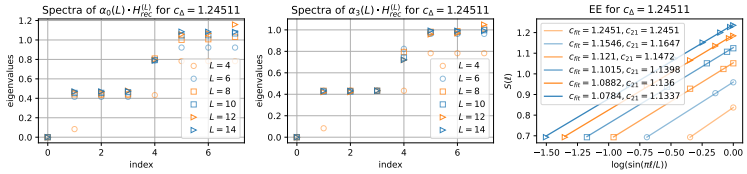

Table 5.6: States on four qutrits that we believe correspond to CFTs (part 5). $\star$ indicates that the spectrum depends on L mod 4.

c_{high}	name	spectra
1.29544	Mystery CFT 1	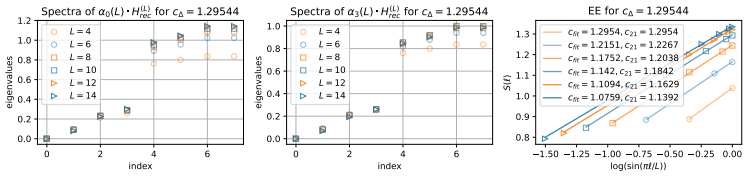
1.60459	Mystery CFT 2	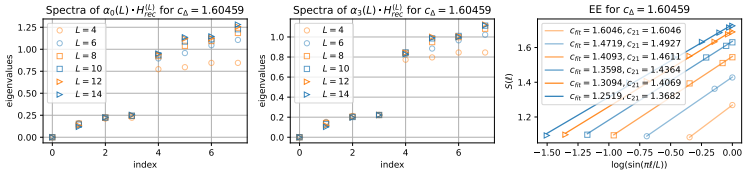
1.62156	Mystery CFT 3	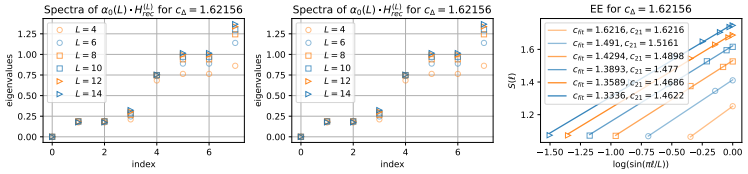
1.73220	Mystery CFT 4	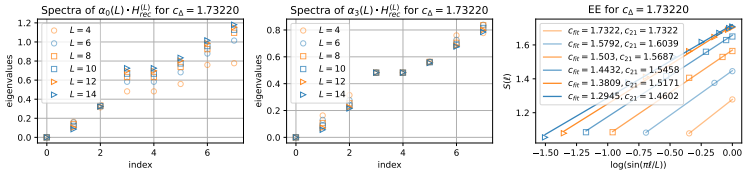
1.73405	Mystery CFT 4’	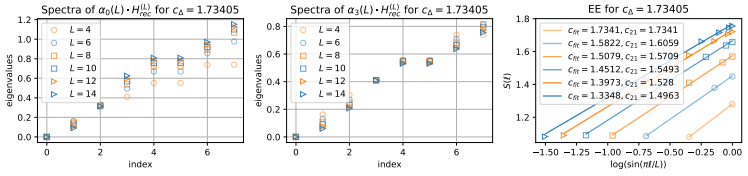 Related to previous
1.74111	Mystery CFT 5	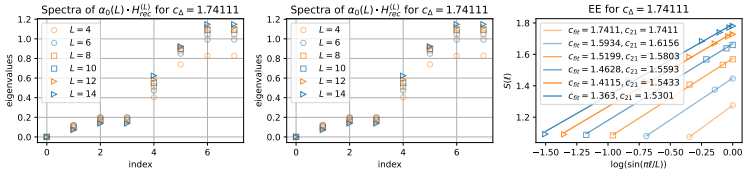

Dissertation Note

Chapter 5 is based on joint work with Xiang Li and John McGreevy, "A systematic search for conformal field theories in very small spaces," *arXiv:2509.04596*, 2025.

Table 5.7: States on four qutrits that we believe correspond to CFTs (part 6). ⋆ indicates that the spectrum depends on L mod 4.

c_{high}	name	spectra
1.74382	Mystery CFT 6	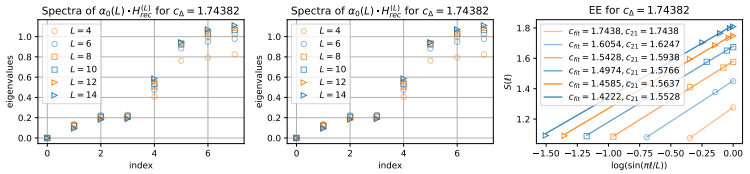 Similar to previous, spectrum is better
1.74805	Mystery CFT 6'	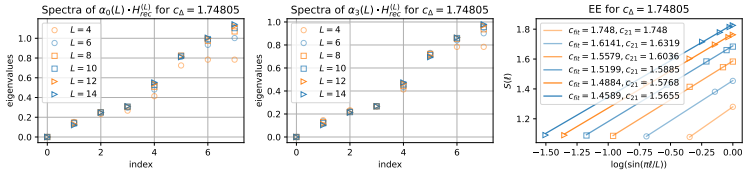 Related to previous
1.82871	Mystery CFT 7	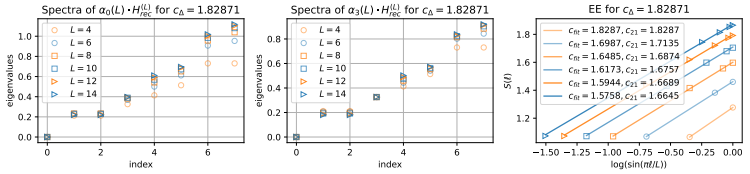
1.83075	Mystery CFT 8	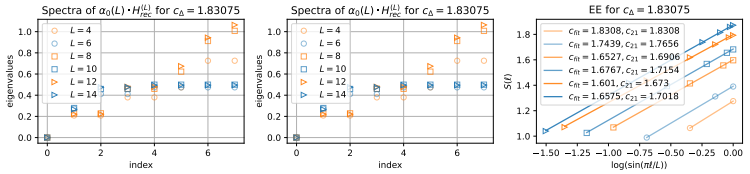 ⋆
1.84206	Mystery CFT 9	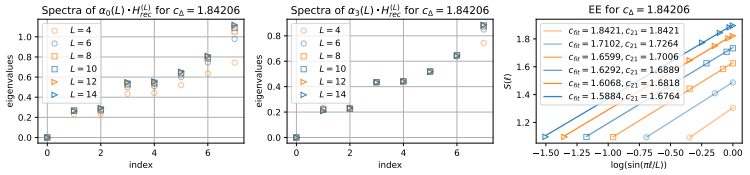 $E_2 = E_3, E_4 = E_5$, relevant perturbation visible
1.84526	Mystery CFT 10	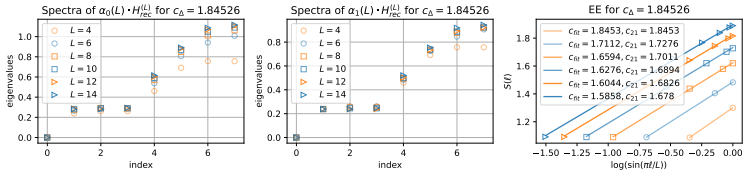 $E_2 = E_3 = E_4$

Table 5.8: States on four qutrits that we believe correspond to CFTs (part 7). $\star$ indicates that the spectrum depends on L mod 4.

c_{high}	name	spectra
1.84851	Mystery CFT 11	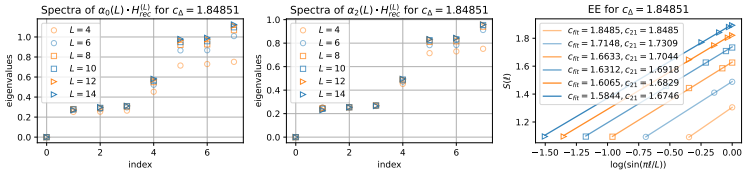 $E_2 = E_3 = E_4$, relevant perturbation visible
1.84926	Mystery CFT 11'	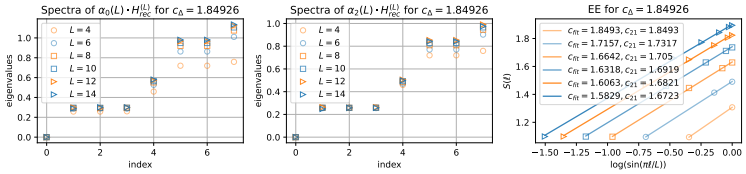 $E_2 = E_3 = E_4$, relevant perturbation visible
1.87447	Mystery CFT 11"	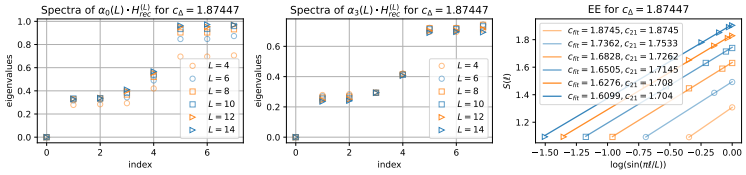 Relevant perturbation visible, same as previous two?
1.90104	Mystery CFT 12	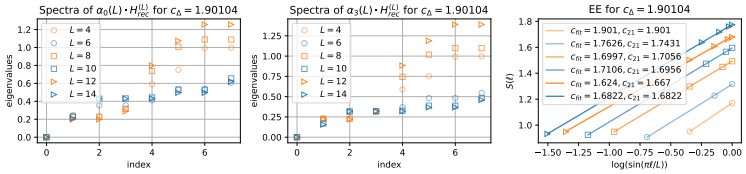 $\star$
1.93044	Mystery CFT 12'	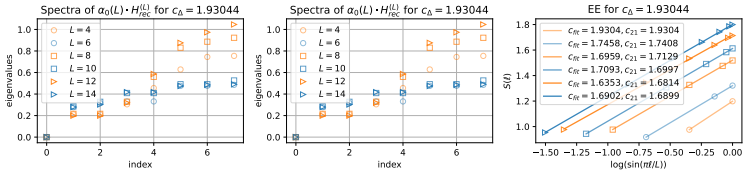 Related to previous $\star$
1.93838	Mystery CFT 13	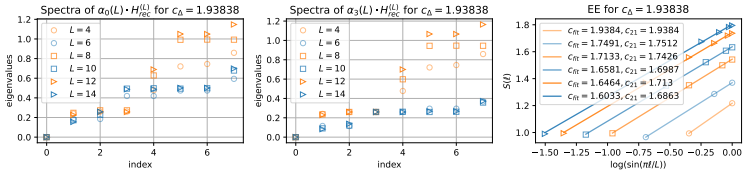 $\star$

Table 5.9: States on four qutrits that we believe correspond to CFTs (part 8). ⋆ indicates that the spectrum depends on L mod 4.

c_{high}	name	spectra
1.95423	Mystery CFT 14	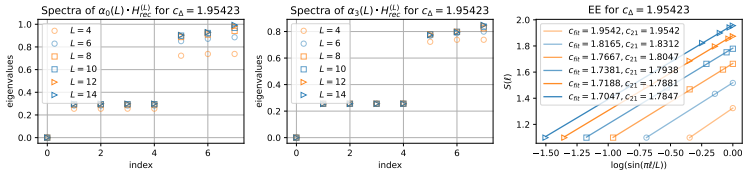
1.99222	Mystery CFT 15	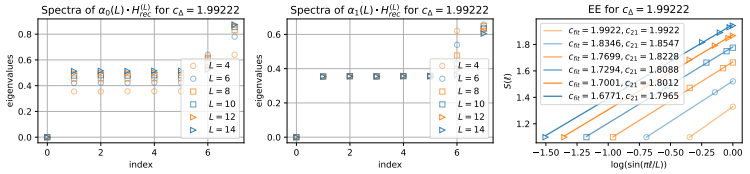 $E_2 = E_3 = E_4 = E_5 = E_6$
1.99325	Mystery CFT 16	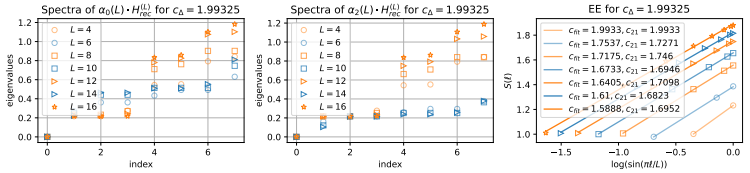 ⋆
1.99972	Mystery CFT 17	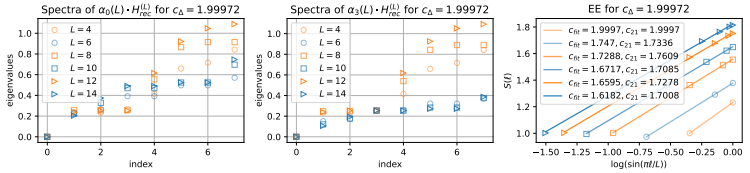 ⋆
2.06646	Mystery CFT 18	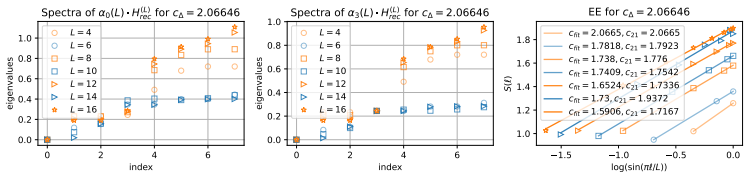 ⋆
2.08015	Mystery CFT 19	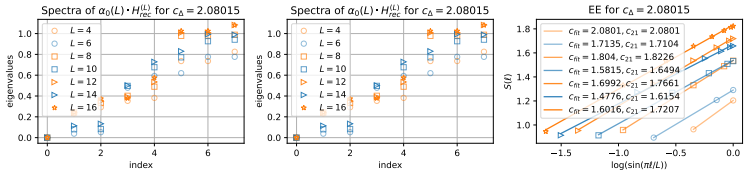 ⋆

Table 5.10: States on four qutrits that we believe correspond to CFTs (part 9). $\star$ indicates that the spectrum depends on L mod 4.

c_{high}	name	spectra
2.49527	Mystery CFT 20	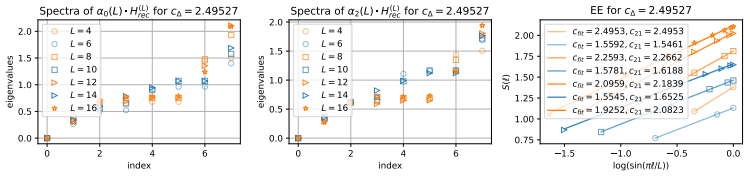 $\star$

Table 5.11: States on four qutrits where we cannot tell whether they correspond to CFTs (part 1). $\star$ indicates that the spectrum depends on L mod 4.

c_{high}	spectra	comments
1.15139	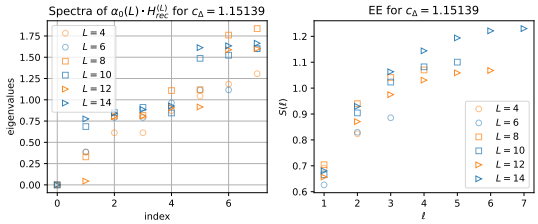	
1.18670	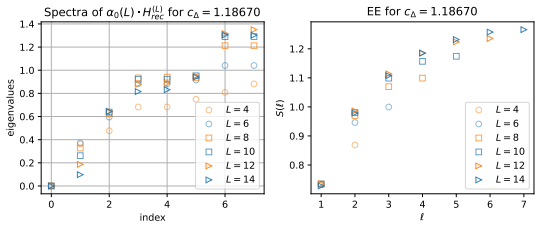	GSD $= 2$
1.26197	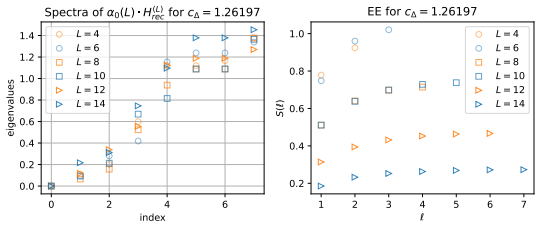	Perhaps a big relevant perturbation of a CFT. Larger L has smaller S.
1.50674	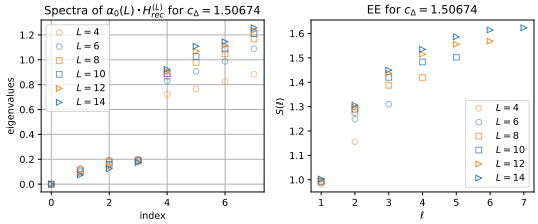	Spectrum is drifting?
1.54380	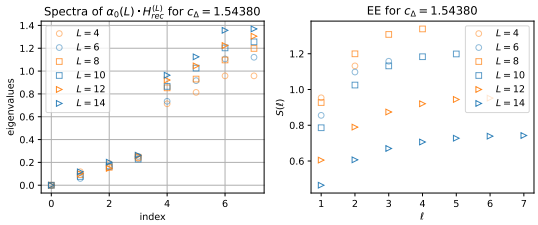	Larger L has smaller S
1.74272	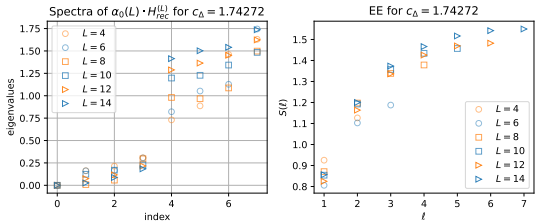	Maybe the spectrum is converging?

Table 5.12: States on four qutrits where we cannot tell whether they correspond to CFTs (part 2). $\star$ indicates that the spectrum depends on L mod 4.

c_{high}	spectra	comments
1.88979	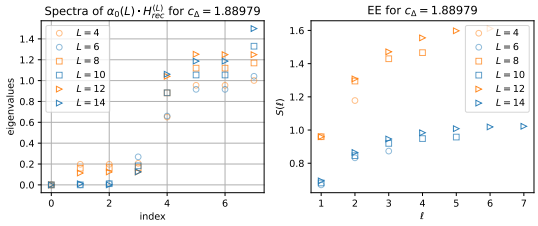	
1.99557	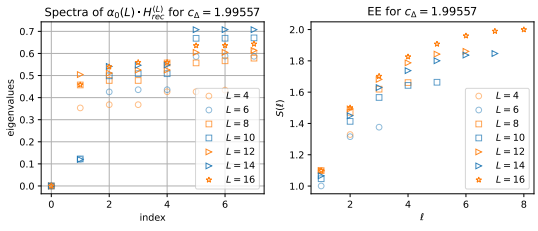	$E_2 = E_3 = E_4 = E_5$, $E_6 = E_7 = E_8$
1.99790	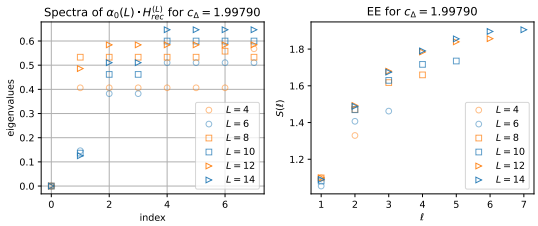	Related to previous
1.99834	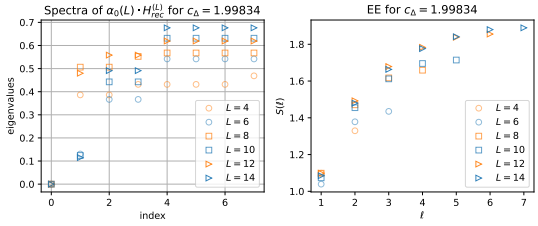	Related to previous
1.99851	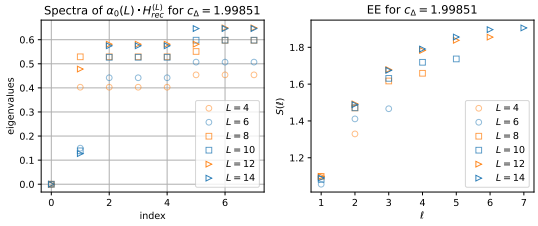	Related to previous
2.00000	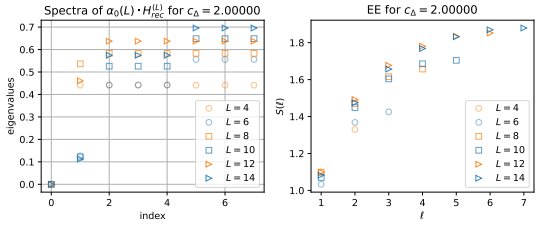	

Table 5.13: States on four qutrits where we cannot tell whether they correspond to CFTs (part 3). $\star$ indicates that the spectrum depends on L mod 4.

c_{high}	spectra	comments
2.03145	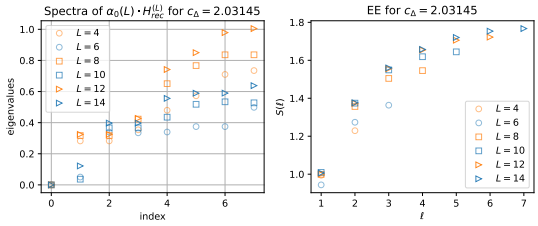	
2.05983	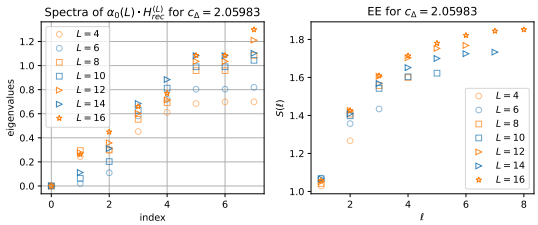	
2.06179	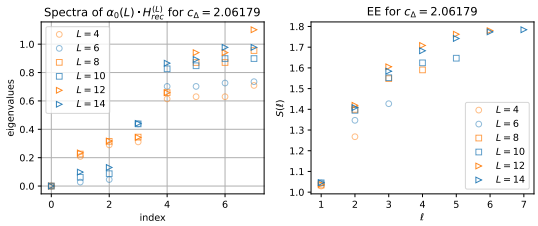	Related to previous
2.06193	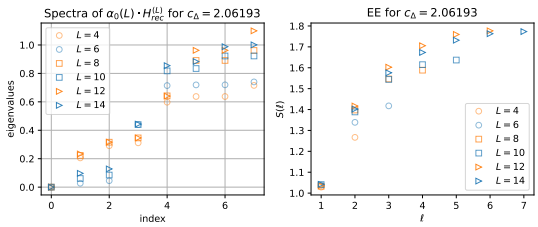	Related to previous
2.06802	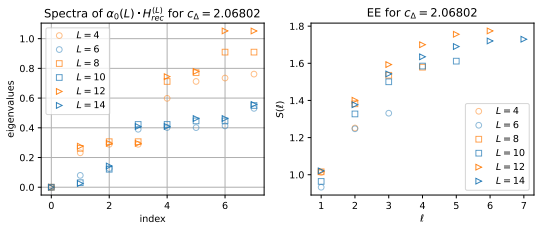	Related to previous
2.19797	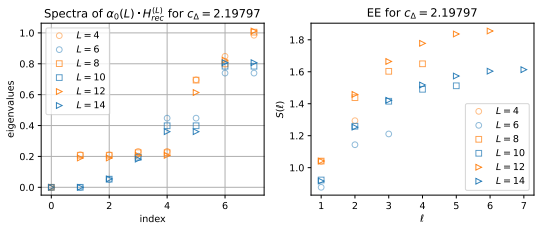	Same as previous

Table 5.14: States on four qutrits where we cannot tell whether they correspond to CFTs (part 4). $\star$ indicates that the spectrum depends on L mod 4.

c_{high}	spectra	comments
2.49499	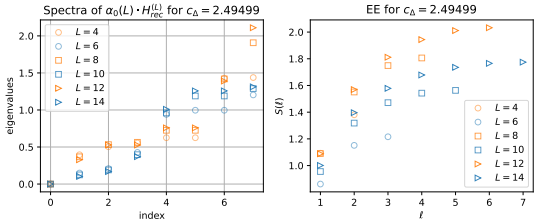	
2.50322	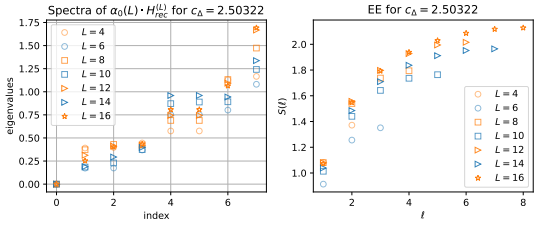	

Table 5.15: States on four qutrits that we believe do not correspond to CFTs (part 1). $\star$ indicates that the spectrum depends on L mod 4.

c_{high}	spectra	comments
1.13233	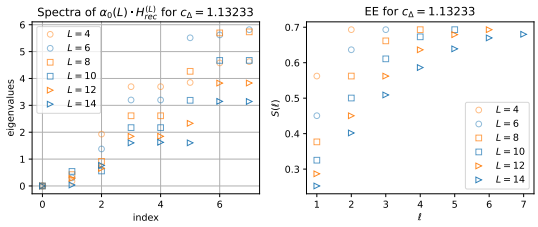	qubit W state
1.50237	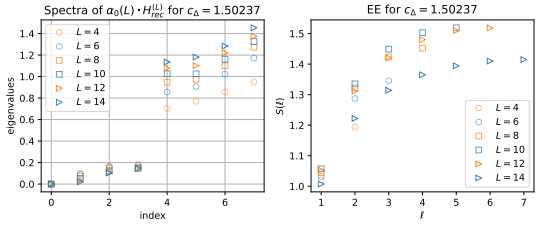	Spectrum is drifting. The number of low-lying states (below a gap) is ℓ_L.
1.58353	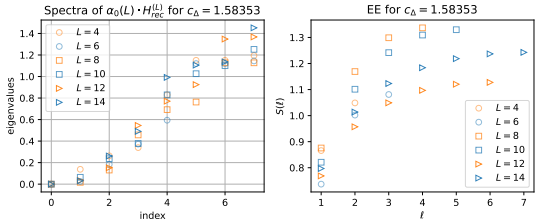	missing
1.59520	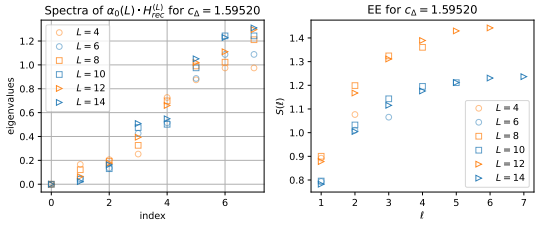	Perturbation of a CFT into a gapped phase with GSD
1.63246	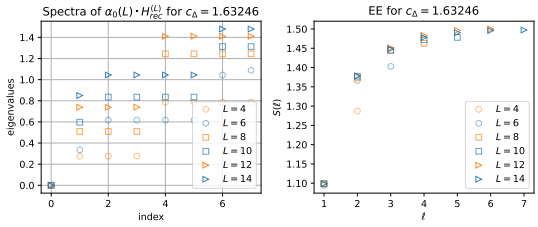	Gapped
1.63298	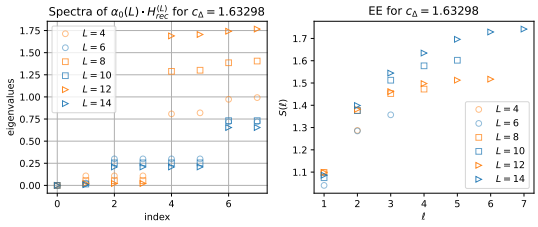	Gapped, GSD=4 for L=0 mod 4, GSD=2 for L=2 mod 4.

Table 5.16: States on four qutrits that we believe do not correspond to CFTs (part 2). $\star$ indicates that the spectrum depends on L mod 4.

c_{high}	spectra	comments
1.63357	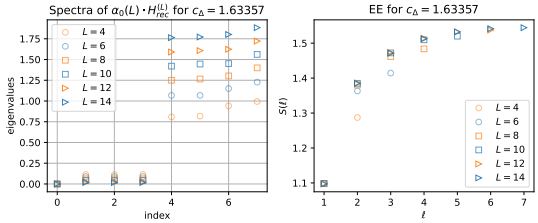	Gapped, GSD=4
1.71387	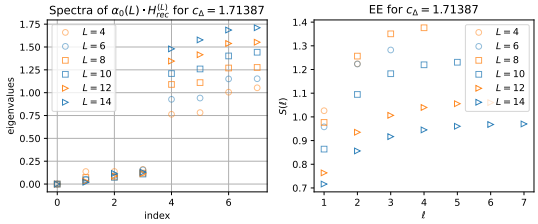	Gapped, related to previous
1.73428	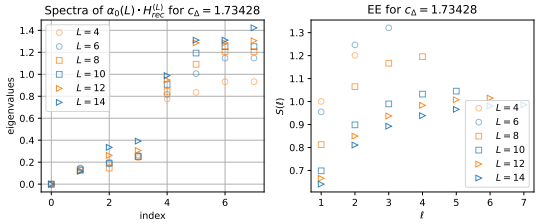	Spectrum is drifting
2.04119	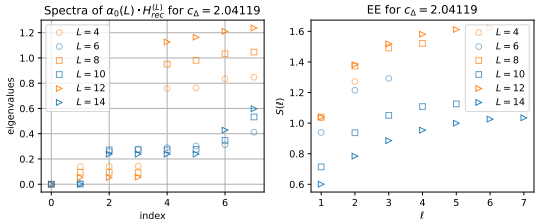	GSD
2.15965	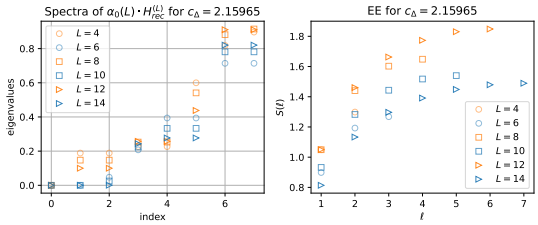	GSD
2.16231	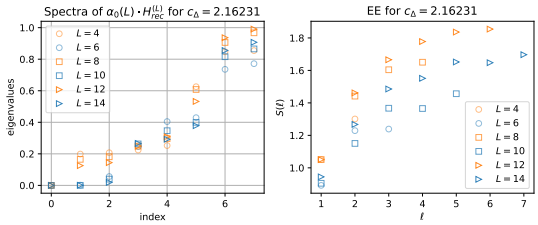	Same as previous

Table 5.17: States on four qutrits that we believe do not correspond to CFTs (part 3). $\star$ indicates that the spectrum depends on L mod 4.

c_{high}	spectra	comments
2.45903	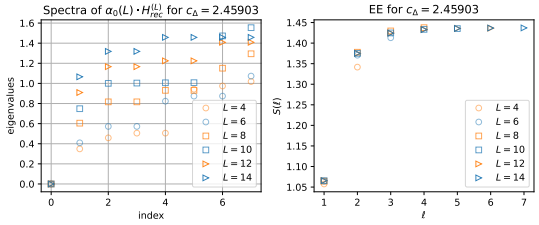	Gapped: spectrum collapses, but drifts away. Entropy is area law.
2.46425	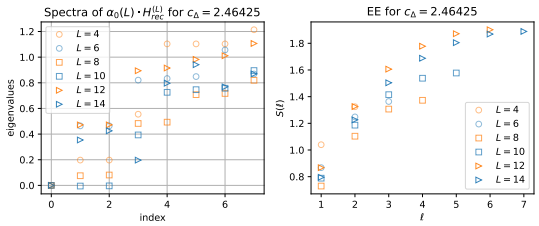	The DMRG routine failed to find the correct eigenstates at L=14.
2.48588	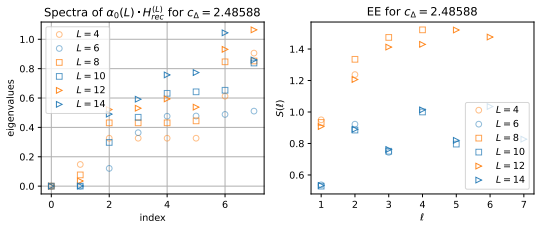	GSD = 2
2.49803	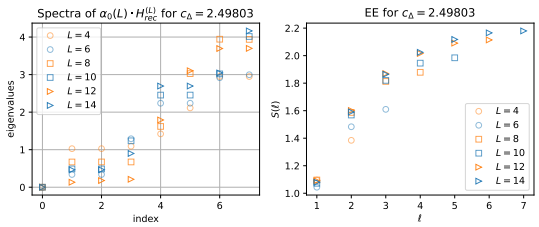	GSD = 4
2.49947	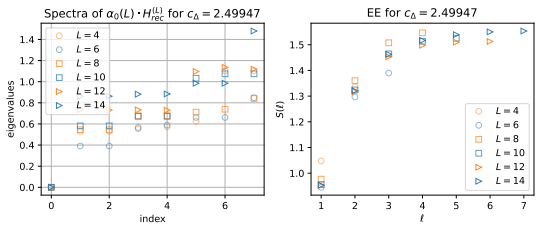	Gapped
2.50198	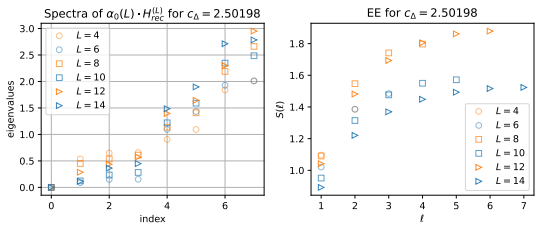	GSD = 2

Table 5.18: States on four qutrits that we believe do not correspond to CFTs (part 4). $\star$ indicates that the spectrum depends on L mod 4.

c_{high}	spectra	comments
2.51738	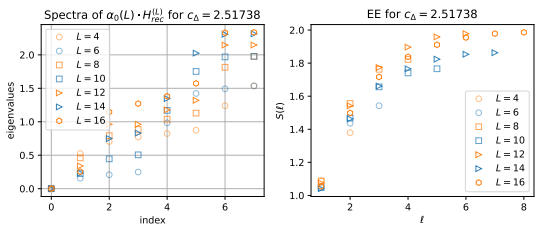	Spectrum is drifting.

Bibliography

[1] A. A. Belavin, Alexander M. Polyakov, and A. B. Zamolodchikov. Infinite Conformal Symmetry in Two-Dimensional Quantum Field Theory. *Nucl. Phys. B*, 241:333–380, 1984.

[2] Daniel Friedan, Zong-An Qiu, and Stephen H. Shenker. Conformal Invariance, Unitarity and Two-Dimensional Critical Exponents. *Phys. Rev. Lett.*, 52:1575–1578, 1984.

[3] P. Di Francesco, P. Mathieu, and D. Senechal. *Conformal Field Theory*. Graduate Texts in Contemporary Physics. Springer-Verlag, New York, 1997.

[4] Vladimir Dotsenko, Jesper Lykke Jacobsen, Marc-André Lewis, and Marco Picco. Coupled Potts models: Self-duality and fixed point structure. *Nucl. Phys. B*, 546:505–557, 1999.

[5] Vladimir S. Dotsenko, Jesper Lykke Jacobsen, Xuan Son Nguyen, and Raoul Santachiara. Universality of coupled Potts models. *Nucl. Phys. B*, 631:426–446, 2002.

[6] António Antunes and Connor Behan. Coupled Minimal Conformal Field Theory Models Revisited. *Phys. Rev. Lett.*, 130(7):071602, 2023.

[7] António Antunes and Connor Behan. Coupled minimal models revisited II: Constraints from permutation symmetry. *SciPost Phys.*, 18(4):132, 2025.

[8] António Antunes and Junchen Rong. Irrational CFTs from coupled anyon chains with non-invertible symmetries? 7 2025.

[9] Bowen Shi and Yuan-Ming Lu. Characterizing topological order by the information convex. *Phys. Rev. B*, 99(3):035112, 2019.

[10] Bowen Shi. Seeing topological entanglement through the information convex. *Phys. Rev. Research.*, 1:033048, 2019.

[11] Bowen Shi, Kohtaro Kato, and Isaac H. Kim. Fusion rules from entanglement. *Annals Phys.*, 418:168164, 2020.

[12] Bowen Shi. Verlinde formula from entanglement. *Phys. Rev. Res.*, 2(2):023132, 2020.

[13] Bowen Shi and Isaac H. Kim. Domain Wall Topological Entanglement Entropy. *Phys. Rev. Lett.*, 126(14):141602, 2021.

[14] Bowen Shi and Isaac H. Kim. Entanglement bootstrap approach for gapped domain walls. *Phys. Rev. B*, 103(11):115150, 2021.

[15] Jin-Long Huang, John McGreevy, and Bowen Shi. Knots and entanglement. *arXiv:2112.08398*, 2021.

[16] Bowen Shi, Jin-Long Huang, and John McGreevy. Remote detectability from entanglement bootstrap I: Kirby's torus trick. 1 2023.

[17] Ting-Chun Lin and John McGreevy. Conformal Field Theory Ground States as Critical Points of an Entropy Function. *Phys. Rev. Lett.*, 131(25):251602, 2023.

[18] Horacio Casini, Marina Huerta, and Robert C. Myers. Towards a derivation of holographic entanglement entropy. *JHEP*, 05:036, 2011.

[19] John Cardy and Erik Tonni. Entanglement hamiltonians in two-dimensional conformal field theory. *J. Stat. Mech.*, 1612(12):123103, 2016.

[20] Xiang Li, Ting-Chun Lin, and John McGreevy. On the strength of the vector fixed-point equation. *to appear*, 2025.

[21] Isaac H. Kim, Xiang Li, Ting-Chun Lin, John McGreevy, and Bowen Shi. Conformal geometry from entanglement. *SciPost Physics*, 18(3):102, March 2025.

[22] Isaac H. Kim, Xiang Li, Ting-Chun Lin, John McGreevy, and Bowen Shi. Chiral Virasoro algebra from a single wavefunction. *Annals Phys.*, 471:169849, 2024.

[23] A. B. Zamolodchikov. Irreversibility of the Flux of the Renormalization Group in a 2D Field Theory. *JETP Lett.*, 43:730–732, 1986.

[24] Ting-Chun Lin, Isaac H. Kim, and Min-Hsiu Hsieh. A new operator extension of strong subadditivity of quantum entropy. *Lett. Math. Phys.*, 113(3):68, 2023.

[25] Zane Ozzello and Yannick Meurice. Multipartite entanglement from ditstrings for 1+1D systems. 7 2025.

[26] David Poland, Slava Rychkov, and Alessandro Vichi. The conformal bootstrap: Theory, numerical techniques, and applications. *Reviews of Modern Physics*, 91(1):015002, 2019.

[27] Lei Gioia and Ryan Thorngren. W state is not the unique ground state of any local Hamiltonian. 10 2023.

[28] Adrian Feiguin, Simon Trebst, Andreas W. W. Ludwig, Matthias Troyer, Alexei Kitaev, Zhenghan Wang, and Michael H. Freedman. Interacting anyons in topological quantum liquids: The golden chain. *Phys. Rev. Lett.*, 98(16):160409, 2007.

[29] David Aasen, Roger S. K. Mong, and Paul Fendley. Topological Defects on the Lattice I: The Ising model. *J. Phys. A*, 49(35):354001, 2016.

[30] David Aasen, Paul Fendley, and Roger S. K. Mong. Topological Defects on the Lattice: Dualities and Degeneracies. 8 2020.

[31] José I. Latorre and Germán Sierra. The c − d conjecture. *J. Stat. Mech.*, 2024(11):113103, 2024.

[32] Paul H. Ginsparg. Curiosities at c = 1. *Nucl. Phys. B*, 295:153–170, 1988.

[33] Paul H. Ginsparg. APPLIED CONFORMAL FIELD THEORY. 1988.

[34] Anran Jin and Ling-Yan Hung. Note on searching for critical lattice models as entropy critical points from strange correlator. *to appear*, 2025.

[35] Matthew Fishman, Steven R. White, and E. Miles Stoudenmire. The ITensor Software Library for Tensor Network Calculations. *SciPost Phys. Codebases*, page 4, 2022.

[36] Matthew Fishman, Steven R. White, and E. Miles Stoudenmire. Codebase release 0.3 for ITensor. *SciPost Phys. Codebases*, pages 4–r0.3, 2022.

[37] Yijian Zou. Universal information of critical quantum spin chains from wavefunction overlap. *Phys. Rev. B*, 105(16):165420, 2022.