Round-Robin CCZ Is All You Need

Zhiyang He,[1,†] Varun Menon,[2] and Theodore J. Yoder
[1] *Department of Mathematics, Massachusetts Institute of Technology*
[2] *Department of Physics, Harvard University*
(Dated: June 17, 2026)

Round-robin entangling gates are a primitive for implementing logical gates on stabilizer codes non-transversally. For qubit sets A, B, C, the *round-robin* gate $\mathrm{CCZ}(A, B, C)$ applies a physical CCZ to every triple $(a, b, c) \in A \times B \times C$. When A, B, C are the supports of three Z logical operators, $\mathrm{CCZ}(A, B, C)$ enacts a logical CCZ. When any one of them is the support of a Z stabilizer, it preserves the code space and acts trivially. In this note, we prove that on any CSS code, round-robin CCZ is all you need: every code-space-preserving CCZ circuit factors into a product of round-robin CCZ circuits, in each of which either all three sets are Z logical operators or at least one is a Z stabilizer. The proof is elementary and constructive: it yields a polynomial-time algorithm that returns the decomposition. The same argument carries over verbatim from CCZ to general $\mathrm{C^rZ}$. This characterization offers a combinatorial perspective on the design of codes with low-depth logical CCZ gates: existing frameworks, such as algebraic code families and topological properties, can be re-examined as conditions that enable effective cancellation within products of simple round-robin circuits. The proof of this observation was derived by Claude Opus 4.8 and verified in Lean by an auto-formalization model MerLean.

I. INTRODUCTION

Constructing quantum codes that admit low-depth logical CCZ gates is an important problem in quantum error correction (QEC). Substantial prior works have explored three broad approaches. Early studies on topological codes have found transversal T and CCZ gates on higher-dimensional topological codes [1, 2], such as 3D color codes [3] and 3D surface codes [4]. These results led to fault-tolerant schemes which perform universal fault-tolerant quantum computation (FTQC) by code-switching between 2D and 3D topological codes [4, 5]. Another line of work uses classical algebraic codes, such as Reed-Solomon codes, to construct quantum codes with good parameters and transversal CCZ. These codes often have high stabilizer check weights, but can be used for magic state distillation [6, 7]. Recent works constructed asymptotically good codes with transversal CCZ gates [8–10], which result in constant-overhead magic state distillation schemes [8]. The third and most recent direction combines the topological and algebraic approaches to construct high-rate LDPC codes with constant-depth CCZ gates, via a homological operation called cup-product [11–15]. These methods have led to discoveries of finite-length codes with interesting parameters and logical actions [16–19].

Relative to this rich literature on low-depth CCZ gates, round-robin CCZ circuits, which are natively high-depth, seem drastically different. Given qubit sets A, B, C, the round-robin gate $\mathrm{CCZ}(A, B, C)$ applies a physical CCZ to every triple $(a, b, c) \in A \times B \times C$, one qubit drawn from each set. When A, B, C are the supports of three Z logical operators, this acts as a logical CCZ on the corresponding logical qubits. Yoder, Takagi, and Chuang introduced these circuits [20] and showed that on certain stabilizer codes, such as the five-qubit code, they can implement logical CCZ gates fault-tolerantly through a method they call pieceable fault tolerance. A round-robin circuit is not transversal, so a single fault can spread to many qubits as the circuit runs. Pieceable fault tolerance cuts the circuit into a few pieces and runs a round of error correction between consecutive pieces, which keeps faults from spreading and makes the gate fault-tolerant. On the surface these round-robin gates sit far from the algebraic and topological conditions that have driven the constructions above.

In this note, we connect them with a structural observation: on any CSS code, round-robin CCZ is all you need. Precisely (Theorem 5), the diagonal CCZ circuits that act as the logical identity are exactly the products of round-robin gates anchored on a Z-stabilizer. This characterizes the trivial circuits, but it also reaches the nontrivial ones, including the transversal addressable CCZ gates sought in the constructions above. Given any diagonal CCZ circuit D with a nontrivial logical action, multiply D by a round-robin realization of the inverse action. The product acts as the logical identity, and decomposes into anchored round-robin gates by Theorem 5, and so D itself is a product of round-robin CCZ gates. Round-robin gates therefore generate every diagonal CCZ circuit. The same conclusion generalizes verbatim from CCZ to general $\mathrm{C^rZ}$.

This characterization offers a combinatorial perspective on the design of codes with low-depth logical CCZ gates: existing frameworks, such as algebraic code families and topological properties, can be re-examined as conditions that enable effective cancellation within products of simple round-robin circuits. To endow a code with a desired logical CCZ, transversal and/or addressable, one may look for it among round-robin realizations

† szhe@mit.edu

and quotient by the stabilizer-anchored circuits. This strategy was used, for example, to find a transversal CCZ gate in the Bacon-Shor code [21]. Notably, round-robin CCZs are naturally addressable, while most aforementioned works design low-depth CCZ gates that act globally on all logical qubits. Recent works [22, 23] have constructed asymptotically good codes, via punctured classical algebraic codes, that admit transversal (depth-1) and arbitrarily addressable CCZ gates. [24] From the perspective of this note, we see that for these codes the round-robin CCZ circuit for any single logical CCZ gate can be sparsified into a transversal circuit.

A. The role of AI Models

Besides the characterization itself, we'd like to highlight another interesting facet of this note. The key observation of this note, namely that round-robin CCZ is all you need, was first hypothesized by the authors in a discussion. We first tried to prove it via induction on r, and reduced it down to a matrix decomposition problem. The notation became unwieldy and we left the thread open. After the announcement by OpenAI that their internal model disproved the longstanding unit distance conjecture [25], we asked Claude Opus 4.8 to prove this observation.

The model first produced a proof via tensor decomposition, in which the notation was heavy. We sent the proof to MerLean [26, 27], an auto-formalization model, which verified the proof in Lean 4 in six hours. Afterwards, we complained to Claude that the proof was hard to read and asked it to improve upon it. Claude then produced the proof in Section III, which is quite simple. Claude's response times to the prompts were each within 20 minutes. Arguably, the observation itself is perhaps at the level of a homework problem for a graduate-level course. However, upon talking to several experts in the field, none of them were aware of this equivalence. We therefore decided to write up this note.

Within this note, the abstract and introduction were almost entirely written by the authors. The remaining sections were first drafted by Claude, after which we provided comments and prompted Claude to revise. Once the writing started to take serious shape, we proceeded to revise the note ourselves. The authors have proofread all the contents and take responsibility for any mistakes. In our experience, we found that Claude is decent at writing proofs, but not very insightful at writing discussions, especially those around introductions. Despite Claude having a fairly thorough understanding of the literature, it needs (in our case) human input to interpret and narrate the result produced. The present note carries that purpose.

II. SETUP

A. Conventions

All linear algebra is over $\mathbb{F}_2 = \{0, 1\}$ with the standard symmetric bilinear form $\langle u, v\rangle = \sum_i u_i v_i$ and orthogonal complement $V^\perp = \{x : \langle x, v\rangle = 0 \ \forall v \in V\}$. For $w \in \mathbb{F}_2^n$ we write $\mathrm{supp}(w) = \{i : w_i = 1\}$ and identify a subset of qubits with its indicator vector. We use the Pauli operators $X(u) = \prod_i X_i^{u_i}$ and $Z(w) = \prod_i Z_i^{w_i}$, and write $\mathrm{C}^r\mathrm{Z}$ for the *r-qubit* controlled-Z gate, which applies the phase -1 exactly when all r of its qubits are in state 1. Thus $\mathrm{C}^1\mathrm{Z} = Z$, $\mathrm{C}^2\mathrm{Z} = \mathrm{CZ}$, and $\mathrm{C}^3\mathrm{Z} = \mathrm{CCZ}$. Section II introduces the objects studied, Section III proves the characterization, and Section IV gives the algorithm.

B. CSS codes and the codeword space

A CSS code on n qubits is specified by binary check matrices H_X, H_Z with

$$H_X H_Z^{\mathsf{T}} = 0. \tag{1}$$

Write $\mathcal{X} = \mathrm{rowsp}\, H_X$ and $\mathcal{Z} = \mathrm{rowsp}\, H_Z$ for the X- and Z-stabilizer support spaces in $\mathbb{F}_2^n$. The *code space* is the joint $+1$ eigenspace of the stabilizers $X(u)$, $u \in \mathcal{X}$, and $Z(w)$, $w \in \mathcal{Z}$. We work in the computational basis, where the relevant object is the *codeword space*

$$\mathcal{W} := \ker H_Z = \mathcal{Z}^\perp \subseteq \mathbb{F}_2^n.$$

The X-stabilizers partition $\mathcal{W}$ into $\mathcal{X}$-cosets: $\mathcal{W}/\mathcal{X}$ indexes the logical computational basis, and each code state is the uniform superposition over one coset,

$$|\bar{c}\rangle = |\mathcal{X}|^{-1/2} \sum_{s \in \mathcal{X}} |c + s\rangle\,, \qquad c + \mathcal{X} \subseteq \mathcal{W}. \tag{2}$$

C. Diagonal circuits and phase polynomials

Diagonal gates in the computational basis can be written as Boolean polynomials. Such a gate multiplies each basis state $|x\rangle$ by a phase that depends on the bits x. For the controlled-Z gates we study this phase is ± 1, so it has the form $(-1)^{f(x)}$ for a function $f\colon \mathbb{F}_2^n \to \mathbb{F}_2$. That Boolean function is the polynomial we track. Because each coordinate satisfies $x_i^2 = x_i$ over $\mathbb{F}_2$, every function $f\colon \mathbb{F}_2^n \to \mathbb{F}_2$ agrees with a unique *multilinear* polynomial $f(x) = \sum_{T \subseteq \{1,\dots,n\}} c_T\, x^T$, where $x^T := \prod_{i \in T} x_i$ and $c_T \in \mathbb{F}_2$. We identify functions with their multilinear representatives and write $R := \mathbb{F}_2[x_1, \dots, x_n]/(x_i^2 - x_i)$ for the resulting ring. Here $\deg f$ denotes the multilinear degree.

We collect a few standard algebraic notions. The ring R is commutative: elements are added and multiplied as polynomials, then reduced by $x_i^2 \mapsto x_i$. An *ideal* $I \subseteq R$ is

a subset closed under addition and under multiplication by any ring element. The ideal $(\ell_1, \ldots, \ell_m)$ *generated* by $\ell_1, \ldots, \ell_m$ is the set of all combinations $\sum_i \ell_i h_i$ with $h_i \in R$. We say f *vanishes on* a set $V \subseteq \mathbb{F}_2^n$ if $f(a) = 0$ for every $a \in V$. Two homomorphisms of R are important for the arguments in Section III. First, *evaluation* at a point $a \in \mathbb{F}_2^n$, $f \mapsto f(a)$, is a ring homomorphism $R \to \mathbb{F}_2$. As evaluation respects sums and products, an element of $(\ell_1, \ldots, \ell_m)$ vanishes at a as soon as every ℓ_i does. Second, a *linear change of coordinates* (substituting for each x_i a linear form in new variables and reducing multilinearly) induces an $\mathbb{F}_2$-algebra automorphism of R, and this automorphism *preserves multilinear degree*: a degree-e monomial becomes a product of e linear forms, of degree $\le e$ after reduction, so the map cannot raise degree. The same applies to its inverse, so it cannot lower degree either. Finally, we use the *dual space* of $\mathbb{F}_2^n$, the linear maps (*functionals*) $\mathbb{F}_2^n \to \mathbb{F}_2$. By non-degeneracy of the form, every functional is $x \mapsto \langle v, x\rangle$ for a unique $v \in \mathbb{F}_2^n$. Hence any basis $g_1, \ldots, g_n$ of $\mathbb{F}_2^n$ yields a basis $\langle g_1, \cdot\rangle, \ldots, \langle g_n, \cdot\rangle$ of the dual space, and the n functionals $x \mapsto \langle g_i, x\rangle$ serve as a new coordinate system on $\mathbb{F}_2^n$.

For a set T of qubits, the gate $C^{|T|}Z(T)$ multiplies $|x\rangle$ by $(-1)^{x^T}$. Phases multiply, so a product of such gates is the diagonal operator $D_f\colon |x\rangle \mapsto (-1)^{f(x)}|x\rangle$ whose *phase polynomial* $f = \sum_T c_T x^T$ collects one monomial per gate (added mod 2), with $C^{|T|}Z(T)$ present iff $c_T = 1$. The map $f \leftrightarrow D_f$ is a bijection between multilinear polynomials and ± 1-valued diagonal operators, and $\deg f$ equals the largest gate arity. We call D_f a *diagonal* $C^{\le d}Z$ *circuit* when $\deg f \le d$. It is Clifford iff $\deg f \le 2$, and a CCZ *circuit* (our main case) when $\deg f \le 3$.

Definition 1 (Logical identity)**.** *D_f acts as logical identity* if it restricts to a scalar (a global phase) on the code space.

For Clifford D_f this coincides with "preserves the stabilizer group and induces the trivial logical gate." For $\deg f \ge 3$ the operator need not normalize the stabilizer group (conjugation can send a Pauli stabilizer to a non-Pauli operator), so we adopt the intrinsic Definition 1 throughout.

D. Round-robin gates and anchored generators

Fix pairwise-disjoint qubit sets $A_1, \ldots, A_r$. The *round-robin* gate $\mathrm{C}^r\mathrm{Z}(A_1, \ldots, A_r)$ applies a physical $\mathrm{C}^r\mathrm{Z}$ to every transversal r-tuple, one qubit drawn from each set:

$$\mathrm{C}^r\mathrm{Z}(A_1, \ldots, A_r) := \prod_{\substack{(a_1,\ldots,a_r)\\ \in A_1\times\cdots\times A_r}} \mathrm{C}^r\mathrm{Z}(a_1, \ldots, a_r),$$

written $\mathrm{CZ}(A, B)$ and $\mathrm{CCZ}(A, B, C)$ for $r = 2, 3$. Its phase polynomial factors into one linear form per set, since expanding the product of sums gives

$$\sum_{(a_t)_t} \prod_{t=1}^{r} x_{a_t} = \prod_{t=1}^{r} \Big(\sum_{a\in A_t} x_a \Big) = \prod_{t=1}^{r} \langle A_t, x\rangle.$$

If the sets overlap, the factored form $\prod_t \langle A_t, x\rangle$ is still the correct phase polynomial (it is an algebraic identity), but it no longer corresponds to a product of full arity-r gates: under multilinear reduction ($x_a^2 = x_a$) any tuple $(a_1, \ldots, a_r)$ that repeats a qubit collapses to a gate of lower arity, e.g. $\mathrm{CZ}(a, a) = Z_a$.

The gate becomes the logical identity as soon as one of its sets is the support of a Z-stabilizer. To see this, take $A_1 = \mathrm{supp}(g)$ for a Z-stabilizer $g \in \mathcal{Z}$, and let the remaining sets be singletons indexed by a set J of qubits with $J \cap \mathrm{supp}(g) = \varnothing$ and $|J| = r$. The phase polynomial is then

$$\begin{aligned}\Psi_{J,g}(x) &:= \langle g, x\rangle \prod_{j\in J} x_j \\ &= \sum_{a\in\mathrm{supp}(g)} \prod_{i\in J\cup\{a\}} x_i,\end{aligned}$$

which we call the *anchored generator* with anchor g and legs J, and whose operator we write $D_{J,g}$. Explicitly,

$$D_{J,g} = \prod_{a\in\mathrm{supp}(g)} \mathrm{C}^r\mathrm{Z}\big(J \cup \{a\}\big). \tag{3}$$

Since $\langle g, x\rangle = 0$ for every codeword $x \in \mathcal{W} = \mathcal{Z}^\perp$, the polynomial $\Psi_{J,g}$ vanishes on $\mathcal{W}$, so $D_{J,g}$ acts as logical identity (Proposition 2). The empty leg set recovers the Z-stabilizers: $\Psi_{\varnothing,g} = \langle g, x\rangle$ is the phase of $Z(g)$, so the Z-stabilizers are exactly the degree-1 anchored generators. We write

$$\mathcal{G} := \mathrm{span}_{\mathbb{F}_2}\{\Psi_{J,g} : J \subseteq \{1, \ldots, n\}, |J| = r,\ g \in \mathcal{Z}\}$$

for the span of anchored generators. Every element of $\mathcal{G}$ acts as logical identity. The content of this note is the converse.

III. THE CHARACTERIZATION

We first specify the meaning of logical identity.

Proposition 2 (Logical-identity criterion)**.** *D_f acts as logical identity if and only if f is constant on the codeword space $\mathcal{W}$. Equivalently, after subtracting the global-phase constant $f(0)$, this holds if and only if $f|_{\mathcal{W}} = 0$.*

Proof. Use the codeword basis $\{|\overline{c}\rangle\}$ of (2), indexed by the $\mathcal{X}$-cosets $c + \mathcal{X} \subseteq \mathcal{W}$. Note $c + s \in \mathcal{W}$ since $\mathcal{X} \subseteq \mathcal{W}$. Then

$$D_f\,|\overline{c}\rangle = |\mathcal{X}|^{-1/2} \sum_{s\in\mathcal{X}} (-1)^{f(c+s)}\,|c + s\rangle\,.$$

If $f \equiv \kappa$ on $\mathcal{W}$, then $f(c+s) = \kappa$ for all c, s, so $D_f |\overline{c}\rangle = (-1)^\kappa |\overline{c}\rangle$ for every c, i.e. $D_f = (-1)^\kappa \mathrm{Id}$ on the code space. Conversely, if D_f acts as a scalar λ, then since $|\overline{c}\rangle$ has equal amplitudes on the distinct strings $\{c+s : s \in \mathcal{X}\}$, the equation $D_f |\overline{c}\rangle = \lambda |\overline{c}\rangle$ forces $(-1)^{f(c+s)} = \lambda$ for all $s \in \mathcal{X}$, with the same λ for all c. As $c+s$ ranges over all of $\mathcal{W}$, f is constant on $\mathcal{W}$. □

The criterion turns the problem into one about a polynomial vanishing on a subspace, governed by the following standard fact. We prove it with explicit degree control, since the algorithm of Section IV reads off its construction.

Lemma 3 (Vanishing on a coordinate subspace)**.** *Let $V \subseteq \mathbb{F}_2^n$ be a subspace and $g_1, \ldots, g_m$ a basis of $V^\perp$, so $V = \{x : \langle g_i, x\rangle = 0,\ i \le m\}$. Write $\ell_i := \langle g_i, \cdot\rangle$. For $f \in R$ the following are equivalent:* (i) *$f|_V = 0$, and* (ii) *f lies in the ideal $(\ell_1, \ldots, \ell_m) \subseteq R$. Moreover, if* (i) *holds then $f = \sum_{i=1}^m \ell_i h_i$ with $h_i \in R$ and $\deg h_i \le \deg f - 1$.*

Proof. (ii)⇒(i). Evaluation at a point $a \in V$ is a ring homomorphism $R \to \mathbb{F}_2$, and it sends each generator to $\ell_i(a) = \langle g_i, a\rangle = 0$, because $a \in V = \{x : \langle g_i, x\rangle = 0\}$. A ring homomorphism that kills every generator of an ideal kills the whole ideal, so every element of $(\ell_1, \ldots, \ell_m)$ vanishes at a, i.e. on V.

(i)⇒(ii). The key idea is to rotate coordinates so that V becomes a coordinate subspace. In such coordinates, "vanishing on V" can be read off one monomial at a time. Extend $g_1, \ldots, g_m$ to a basis $g_1, \ldots, g_n$ of $\mathbb{F}_2^n$. The functionals $\ell_i = \langle g_i, \cdot\rangle$ are a basis of the dual space, so

$$\varphi\colon x \longmapsto y = \big(\ell_1(x), \ldots, \ell_n(x)\big)$$

is an invertible linear change of coordinates on $\mathbb{F}_2^n$. Let $F \in R$ be f rewritten in these coordinates, characterized by $f = F \circ \varphi$. Crucially, a linear change of coordinates *preserves* multilinear degree (§II C): each new variable y_i is a degree-1 form in x and each old variable x_i a degree-1 form in y, so neither direction of the substitution can raise degree, whence $\deg F = \deg f$.

Two features of the new coordinates now finish the proof.

- *V is a coordinate subspace.* By hypothesis $V = \{x : \ell_1(x) = \cdots = \ell_m(x) = 0\}$, which in the y-coordinates is exactly $\{y_1 = \cdots = y_m = 0\}$, with $y_{m+1}, \ldots, y_n$ free (as φ is onto and $\dim V = n - m$).

- *Vanishing is monomial-wise.* Restricting F to V sets $y_1 = \cdots = y_m = 0$, which retains exactly the monomials of F built from $y_{m+1}, \ldots, y_n$ alone and deletes the rest. Since the multilinear representative is unique, $f|_V = 0$ forces this retained part to vanish identically: *every* monomial of F is divisible by some y_i with $i \le m$.

Assign each monomial of F to the least index $i \le m$ of a variable y_i dividing it, and factor that variable out. Collecting terms by their assigned index gives

$$F = \sum_{i=1}^m y_i\, H_i(y)$$

with H_i multilinear and $\deg H_i \le \deg F - 1$. Finally pull back through φ: resubstituting $y_i = \ell_i(x)$ and setting $h_i := H_i \circ \varphi$ turns this into $f = \sum_{i=1}^m \ell_i\, h_i$, with $\deg h_i \le \deg F - 1 = \deg f - 1$. This places f in $(\ell_1, \ldots, \ell_m)$ with the stated degree control. □

Remark 4. For clarity, we (the human authors), note a stabilizer interpretation of the change of basis $\varphi : x \to y$. By definition $x_i = (1 - Z_i)/2$, a projector onto the -1 eigenspace of the single-qubit Pauli Z_i. Similarly, there are projectors $y_i = (1 - G_i)/2$ onto the -1 eigenspaces of the Z stabilizers $G_1, \ldots, G_m \in SU(2^n)$ and Z logical operators $G_{m+1}, \ldots G_n \in SU(2^n)$, which together $\{G_i : i = 1, \ldots, n\}$ is also a complete basis of the Z-type Pauli group. As such, it is possible to write each Z_i as a product $Z_i = \prod_j G_j^{\phi_{ij}}$ with all $\phi_{ij} \in \mathbb{F}_2$. Expand out

$$\begin{aligned} x_i = \frac{1}{2}(1 - Z_i) &= \frac{1}{2}(1 - \prod_j G_j^{\phi_{ij}}) \\ &= \frac{1}{2}(1 - \prod_j (1 - 2y_i)^{\phi_{ij}}) = \sum_j y_j^{\phi_{ij}} + 2P_i(y) \\ &= \sum_j y_j^{\phi_{ij}} \end{aligned}$$

where P_i is just some polynomial over the integers and the last equality uses $2P_i(y) = 0$ which results from our context, where x_i is a term in a phase polynomial over $\mathbb{F}_2$. We can replace each x_i in $f(x)$ by its expansion in y_j variables to create polynomial F such that $f(x) = F(y)$. Note that F has the same multilinear degree as f.

Theorem 5 (Round-robin generation)**.** *For a multilinear phase polynomial f, the following are equivalent:*

(1) *D_f acts as logical identity on the CSS code;*

(2) *f vanishes on the codeword space $\mathcal{W}$ (up to the constant $f(0)$);*

(3) *$f \in \mathcal{G}$, i.e. f is an $\mathbb{F}_2$-linear combination of anchored generators $\Psi_{J,g}$.*

The generators may be taken with g ranging over a fixed basis of $\mathcal{Z}$ and $|J| \le \deg f - 1$. Equivalently, D_f is a product of round-robin $C^{\le \deg f}Z$ gates each anchored on a Z-stabilizer, using no gate of arity exceeding $\deg f$.

Proof. (1)⇔(2) is Proposition 2. For (2)⇔(3), apply Lemma 3 with $V = \mathcal{W} = \mathcal{Z}^\perp$, whose orthogonal complement is $\mathcal{W}^\perp = \mathcal{Z}$ (the identity from §II), taking $g_1, \ldots, g_m$ a basis of $\mathcal{Z}$.

If (3) holds, each $\Psi_{J,g} = \langle g, \cdot\rangle\, x^J$ is a multiple of the ℓ_i (write $g = \sum_i a_i g_i$, so $\langle g, \cdot\rangle = \sum_i a_i \ell_i$), hence $f \in$

$(\ell_1, \ldots, \ell_m)$ and $f|_{\mathcal{W}} = 0$, which is (2). Conversely, if (2) holds, Lemma 3 gives $f = \sum_i \ell_i h_i$ with $\deg h_i \leq \deg f - 1$. Expanding each $h_i = \sum_J (h_i)_J x^J$,

$$\begin{aligned} f &= \sum_{i,J} (h_i)_J \, \langle g_i, x \rangle \, x^J \\ &= \sum_{i,J} (h_i)_J \, \Psi_{J,g_i}, \end{aligned}$$

an $\mathbb{F}_2$-combination of anchored generators with g_i in a basis of $\mathcal{Z}$ and $|J| \leq \deg f - 1$ (the constant terms $J = \varnothing$ contributing the Z-stabilizers). Through (3), D_f is a product of round-robin $C^{\leq \deg f}Z$ gates anchored on Z-stabilizers. □

IV. AN EFFICIENT DECOMPOSITION ALGORITHM

The proof of Theorem 5 is constructive: its only non-bookkeeping step is the change to coordinates in which the Z-stabilizers become coordinate functions (Lemma 3), and in those coordinates the decomposition is read off directly. Algorithm 6 carries this out.

Algorithm 6 (Decompose into anchored round-robin gates)**.** *Input:* the phase polynomial f (degree $\leq d$) and a basis $g_1, \ldots, g_m$ of $\mathcal{Z}$. *Output:* either a certificate that $f|_{\mathcal{W}} \neq 0$, or anchored generators with $f = \sum_i \Psi_{J_i,g_i}$.

1. Extend $g_1, \ldots, g_m$ to a basis $g_1, \ldots, g_n$ of $\mathbb{F}_2^n$. Let Φ be the matrix with rows g_i, so that $y = \Phi x$ are adapted coordinates with $y_i = \langle g_i, x \rangle$.
2. Compute $F(y) \leftarrow f(\Phi^{-1} y)$ by substituting $x = \Phi^{-1} y$ and reducing multilinearly.
3. Let ρ be the sum of the monomials of F containing no variable in $\{y_1, \ldots, y_m\}$. If $\rho \neq 0$, **return** "not logical identity" with witness ρ (this is $f|_{\mathcal{W}}$ in the y-coordinates).
4. Otherwise, for each monomial μ of F let $i = \min\{k \leq m : y_k \mid \mu\}$ and add μ / y_i to H_i, giving $F = \sum_{i \leq m} y_i H_i$.
5. **Return** the gates D_{J,g_i} for every monomial x^J of $H_i(\Phi x)$, $i = 1, \ldots, m$.

Theorem 7 (Efficient decomposition)**.** *Given the phase polynomial f of a diagonal $C^{\leq d}Z$ circuit and a basis of $\mathcal{Z}$, Algorithm 6 runs in time $O(n^{2d})$ (in particular $O(n^6)$ for* CCZ *circuits) and either*

(i) *certifies that D_f is* not *the logical identity, returning the nonzero restriction $f|_{\mathcal{W}}$ as a witness; or*

(ii) *returns an explicit list of anchored generators Ψ_{J_i,g_i} with $|J_i| \leq d - 1$ and $f = \sum_i \Psi_{J_i,g_i}$, a decomposition of D_f into Z-stabilizer-anchored round-robin $C^{\leq d}Z$ gates.*

Proof. Correctness. Since $y = \Phi x$ gives $y_i = \langle g_i, x \rangle$, the codeword space is $\mathcal{W} = \{y_1 = \cdots = y_m = 0\}$ and ρ is exactly $f|_{\mathcal{W}}$ in the y-coordinates. Thus $\rho \neq 0$ iff f does not vanish on $\mathcal{W}$ iff, by Proposition 2, D_f is not the logical identity. This is output (i). If $\rho = 0$, every monomial of F contains an anchor variable, so step 4 produces $F = \sum_{i \leq m} y_i H_i$. Substituting $y = \Phi x$ gives

$$f(x) = \sum_i \langle g_i, x \rangle \, H_i(\Phi x) = \sum_{i,J} \Psi_{J,g_i}$$

with $|J| \leq \deg f - 1$, a valid decomposition by Theorem 5, which is output (ii).

Complexity. Extending the basis and inverting Γ are $O(n^3)$. The substitutions in steps 2 and 5 act on multilinear polynomials of degree $\leq d$, of which there are $O(n^d)$ monomials. Expanding each through a linear substitution costs $O(n^d)$, for $O(n^{2d})$ in total, and the peeling loop is $O(n^d)$. The dominant cost is $O(n^{2d})$, i.e. $O(n^6)$ for CCZ ($d = 3$). □

Remark 8 (Non-uniqueness and minimality). The anchored generators are linearly dependent, so the decomposition is not unique. The least-anchor-index rule in step 4 selects one canonical representative. Finding a decomposition with the *fewest* gates is the minimum-weight solution of an $\mathbb{F}_2$ linear system, a separate optimization that is hard in general and that we do not address.

Remark 9 (Worked example: round-robin to transversal CZ on the Steane code). Take two copies of the Steane code, the smallest two-dimensional color code, on qubits 1 to 7, with the three Z-stabilizer generators equal to the X-stabilizers (the code is self-dual):

$$S_1 = \{1, 3, 5, 7\}, \quad S_2 = \{2, 3, 6, 7\}, \quad S_3 = \{4, 5, 6, 7\}.$$

See [28] for more details on the Steane code. Write t, u for the computational labels of the two copies, placed on common qubits, and take the logical Z of each copy to be $\{1, 2, 3\}$, so the logical bits are $t_1 + t_2 + t_3$ and $u_1 + u_2 + u_3$. Two diagonal CZ circuits realize the same logical CZ between the encoded qubits:

$$\text{round-robin:} \quad \mathrm{CZ}(\{1,2,3\}, \{1,2,3\}) = \sum_{i,j \in \{1,2,3\}} t_i u_j,$$

$$\text{transversal:} \quad \prod_{i=1}^{7} \mathrm{CZ}(\{i\}, \{i\}) = \sum_{i=1}^{7} t_i u_i.$$

The round-robin gate uses nine physical CZs and places qubit 1 of each copy in three of them (degree 3). The transversal gate uses seven, one per qubit (degree 1). They differ by a product of anchored gates, each a CZ between a Z-stabilizer of one copy and a single qubit of the other. Explicitly, the transversal gate is the round-robin gate composed with the seven free gates

$$\begin{aligned} &t_1 \langle S_2, u \rangle, && t_2 \langle S_1, u \rangle, \\ &t_3 \langle S_1 \triangle S_2, u \rangle, && t_4 \langle S_3, u \rangle, \\ &\langle S_2 \triangle S_3, t \rangle \, u_5, && \langle S_1 \triangle S_3, t \rangle \, u_6, \\ &\langle S_1 \triangle S_2 \triangle S_3, t \rangle \, u_7, \end{aligned}$$

that is, CZs pairing t_1 with u-qubits $\{2,3,6,7\}$, t_2 with $\{1,3,5,7\}$, t_3 with $\{1,2,5,6\}$, t_4 with $\{4,5,6,7\}$, and pairing t-stabilizers $\{2,3,4,5\},\{1,3,4,6\},\{1,2,4,7\}$ with u_5,u_6,u_7. The sum can be verified directly. This sparsifies the round-robin CZ into a transversal one, and equivalently writes the transversal CZ as a product of eight round-robin CZ gates: one anchored on the logical operators and seven anchored on Z-stabilizers, the latter acting as the logical identity.

Remark 10 (Cross-block circuits). The round-robin literature usually wires gates between distinct code blocks. The characterization above applies to these circuits directly, as multiple blocks of CSS code are simply one larger CSS code.

ACKNOWLEDGMENTS

We thank Adam Wills, Louis Golowich, Katie Chang, Rohan Mehta, Chris Pattison and Alex Kubica for helpful discussions. Z.H. thanks Peter Shor and the MIT Department of Mathematics for funding of research activities.

[1] H. Bombín, New Journal of Physics **17**, 083002 (2015).
[2] A. Kubica, B. Yoshida, and F. Pastawski, New Journal of Physics **17**, 083026 (2015).
[3] A. Kubica and M. E. Beverland, Physical Review A **91**, 032330 (2015).
[4] M. Vasmer and D. E. Browne, Physical Review A **100**, 012312 (2019).
[5] H. Bombín, New Journal of Physics **18**, 043038 (2016).
[6] S. Bravyi and A. Kitaev, Physical Review A **71**, 022316 (2005).
[7] S. Bravyi and J. Haah, Physical Review A **86**, 052329 (2012).
[8] A. Wills, M.-H. Hsieh, and H. Yamasaki, Nature Physics **21**, 1842–1846 (2025).
[9] L. Golowich and V. Guruswami, in *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, STOC '25 (Association for Computing Machinery, New York, NY, USA, 2025) p. 707–717.
[10] Q. T. Nguyen, in *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, STOC '25 (Association for Computing Machinery, New York, NY, USA, 2025) p. 697–706.
[11] G. Zhu, S. Sikander, E. Portnoy, A. W. Cross, and B. J. Brown, PRX Quantum **6**, 040361 (2025).
[12] L. Golowich and T.-C. Lin, in *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, STOC '25 (Association for Computing Machinery, New York, NY, USA, 2025) p. 689–696.
[13] T.-C. Lin, arXiv preprint arXiv:2410.14631 (2024).
[14] N. P. Breuckmann, M. Davydova, J. N. Eberhardt, and N. Tantivasadakarn, Communications in Mathematical Physics **407**, 86 (2026).
[15] G. Zhu, arXiv preprint arXiv:2501.19375 (2025).
[16] V. Menon, J. P. Bonilla Ataides, R. Mehta, A. Gu, D. B. Tan, and M. D. Lukin, Physical Review X **16**, 10.1103/ghhp-cytl (2026).
[17] A. Jacob, C. McLauchlan, and D. E. Browne, arXiv preprint arXiv:2508.08191 (2025).
[18] C. Li, J. Preskill, and Q. Xu, arXiv preprint arXiv:2510.07269 (2025).
[19] R. Tiew and N. P. Breuckmann, arXiv preprint arXiv:2602.23307 (2026).
[20] T. J. Yoder, R. Takagi, and I. L. Chuang, Physical Review X **6**, 031039 (2016).
[21] T. J. Yoder, arXiv preprint arXiv:1705.01686 (2017).
[22] Z. He, V. Vaikuntanathan, A. Wills, and R. Y. Zhang, arXiv preprint arxiv:2502.01864 (2025).
[23] Z. He, V. Vaikuntanathan, A. Wills, and R. Y. Zhang, arXiv preprint arXiv:2507.05392 (2025).
[24] Here by arbitrary addressability we mean that the logical CCZ gate can target any triple of logical qubits inside one code block or across multiple blocks.
[25] OpenAI, Planar point sets with many unit distances, `https://cdn.openai.com/pdf/74c24085-19b0-4534-9c90-465b8e29ad73/unit-distance-proof.pdf` (2026), technical report.
[26] Y. Ren, J. Li, and Y. Qi, arXiv preprint arXiv:2602.16554 (2026).
[27] J. Li, Z. Zhu, and Y. Ren, arXiv preprint arXiv:2605.26959 (2026).
[28] [[7, 1, 3]] steane code, in *The Error Correction Zoo*, edited by V. V. Albert and P. Faist (2026), arxiv:2606.11484.