

PARITY $\notin$ QAC⁰

(via Polarization and Direct Induction)

Shiva Kintali*

October 17, 2026

Abstract

We give a self-contained simplification of OpenAI’s proof of **parity is not in QAC⁰**. Our main simplification is to use polarization to express the terms in the reflection expansion as linear combinations of individual product projections. The uniform lower-depth induction hypothesis then applies directly to each projection, regardless of its support size, and its exponential decay absorbs the coefficient cost 18^k , where k is the original support size. This removes the intermediate low-count propagation argument while retaining OpenAI’s localization framework and inverse-circuit bootstrap.

1 Introduction

The parity of a binary string is 0 when the number of ones is even and 1 when it is odd. We study circuits built from arbitrary one-qubit unitaries and generalized Toffoli gates, which flip a target bit when all their control bits are 1. Gates applied simultaneously must act on disjoint sets of qubits. The depth is the number of such parallel layers. Extra qubits, called ancillas, start in the state $|0\rangle$; a measurement of one designated qubit supplies the answer. The other final qubits may be discarded in any state.

Main result. For every fixed depth and every fixed polynomial bound on the total number of qubits, these circuits cannot compute parity on all sufficiently large input lengths with success probability at least $1/2 + \epsilon$ on *every* input, for any fixed $\epsilon > 0$. This is the statement $\text{PARITY} \notin \text{QAC}^0$. Theorem 1 gives the precise result in a circuit normal form; Corollary 9 states it for the usual model. Intermediate measurements, postselection, and primitive unbounded-fanout gates are not allowed.

We simplify the proof in the OpenAI preprint [3]. Its organizing idea is *localization*: a product projection conjugated by a shallow circuit has very small matrix entries between suitable product states that differ on many qubits. The localization statement and inverse-circuit argument are retained here. We simplify the depth-reduction step by applying polarization to the source’s reflection expansion, allowing direct application of the lower-depth induction hypothesis. A different proof of the parity bound uses regular trajectories and pruning [4].

The main simplification. A product projection A fixes chosen one-qubit directions and leaves the other qubits unrestricted. Suppose its support has k qubits and R is one layer of disjoint product-state reflections. OpenAI’s proof expands RAR into ordered products $P_L A P_R$ of product

*Email: shiva.kintali@gmail.com Homepage: shivakintali.github.io

projections, then controls these products by a low-count propagation argument. Such a product need not itself be an orthogonal projection. We instead apply *polarization*, a standard identity expressing a rank-one operator as a linear combination of rank-one orthogonal projections, to the one-qubit factors on the support of A . This gives

$$RAR = \sum_j c_j A_j, \quad \sum_j |c_j| \leq 18^k,$$

where every A_j is again a product projection (Lemma 3). The bound depends on the original support size k , even when the gates and the resulting projections have large supports. The uniform lower-depth induction hypothesis applies to each A_j directly, and its stronger decay absorbs the coefficient cost. Applying induction term by term replaces the low-count propagation step; the polynomial approximation and inverse-circuit bootstrap remain part of the full proof.

How the proof fits together. First, a mismatch count records how many local directions differ from a chosen product condition. We prove that conjugated product projections have very small transitions from zero mismatches to many mismatches. The expansion above handles projections with small support; a polynomial approximation and the inverse circuit extend the estimate to arbitrary support. Finally, a change of orthonormal basis expresses parity correlation as an average of matrix entries between complementary product states. Localization makes this average tend to zero, whereas successful parity computation would keep it at least ϵ .

Section 2 introduces the linear-algebra notation and the localization statement. Section 3 proves the expansion and depth reduction. Section 4 establishes localization, and Section 5 derives the parity contradiction. Section 6 supplies the exact connection with standard circuit models and common gate-set variants. All proofs are included in the main text.

2 Linear-algebra setup and localization

2.1 Vectors, measurements, and product conditions

A qubit is a copy of $\mathbb{C}^2$, with orthonormal basis $|0\rangle = (1, 0)^\top$ and $|1\rangle = (0, 1)^\top$. The notation $|v\rangle$ denotes a column vector, $\langle v|$ its conjugate transpose, and $T^\dagger$ the adjoint of a matrix T . Thus $|u\rangle\langle v|$ is the operator $z \mapsto u(v^\dagger z)$; for a unit vector v , $|v\rangle\langle v|$ is an orthogonal projection. A pure state is represented by a unit vector, and a circuit acts by a unitary matrix. All input states in our circuit model are pure. For an orthogonal projection P , the associated measurement outcome has probability $\|P\psi\|^2$ in state ψ . These elementary rules are the quantum-mechanical input to the proof.

Operators act on $\mathcal{H}_N = (\mathbb{C}^2)^{\otimes N}$ for an integer $N \geq 1$. For a bit string x , the notation $|x\rangle = |x_1\rangle \otimes \cdots \otimes |x_N\rangle$ gives the *computational basis*. Operator norms are induced by the Euclidean vector norm; the same notation denotes the Euclidean norm of a vector. Logarithms are natural. For a self-adjoint operator Y and a set $\Omega \subseteq \mathbb{R}$, we write $[Y \in \Omega]$ for the spectral projection of Y onto the sum of its eigenspaces with eigenvalues in Ω ; for instance $[Y \geq r]$, $[Y < r]$, and $[Y = 0]$. A *product projection* is a tensor product whose one-qubit factors are either I or rank-one orthogonal projections. Its *support* is the set of qubits whose factor is not I ; the identity is the product projection with empty support. A *count* is

$$M = \sum_{i \in S} (I - |v_i\rangle\langle v_i|)_i,$$

where $S \subseteq \{1, \dots, N\}$ and the $v_i \in \mathbb{C}^2$ are unit vectors. A subscript i means tensoring with identities on the other qubits. We call the operators $(I - |v_i\rangle\langle v_i|)_i$ the *summands* of M . Each one-qubit operator $I - |v_i\rangle\langle v_i|$ is itself a rank-one projection, namely $|v_i^\perp\rangle\langle v_i^\perp|$ for a unit vector $v_i^\perp \perp v_i$. Summands on distinct qubits commute, so the spectrum of M is contained in $\{0, 1, \dots, |S|\}$ and $0 \leq M \leq NI$. Moreover $[M = 0]$ has factors $|v_i\rangle\langle v_i|$ on S and I elsewhere. Thus every product projection is $[M = 0]$ for an appropriate count; $S = \emptyset$ gives $M = 0$ and $[M = 0] = I$.

The terminology has a concrete interpretation. If all counted reference vectors are $|0\rangle$, then on a computational basis vector $|x\rangle$ the operator M has eigenvalue $\sum_{i \in S} x_i$, the number of mismatches on S . Choosing the local basis $(v_i, v_i^\perp)$ gives the same interpretation for general reference vectors. For example, with two counted qubits and reference vector $|0\rangle$ at both, the matrix of M in the ordered basis $(|00\rangle, |01\rangle, |10\rangle, |11\rangle)$ is $\text{diag}(0, 1, 1, 2)$. Thus $[M = 0]$ selects perfect agreement and $[M \geq r]$ selects at least r mismatches. These projections are tools of the analysis, not additional circuit operations; projected vectors are never renormalized in our estimates.

2.2 Reflection circuits and the output convention

Definition 1 (Reflection circuits). A *reflection gate* is a unitary $I - 2G$, where G is a product projection with nonempty support; the support of the gate is that of G . Equivalently, G is the rank-one projection onto a product vector on the qubits of its support, tensored with the identity elsewhere. A *reflection layer* is a product of reflection gates with pairwise disjoint supports; the empty product I is allowed. A *local layer* is a tensor product of one-qubit unitaries. For integers $N \geq 1$ and $d \geq 0$, let $\mathcal{U}_d(N)$ be the set of unitaries on $\mathcal{H}_N$ of the form

$$U = L_d R_d L_{d-1} R_{d-1} \cdots L_1 R_1 L_0, \quad (1)$$

where each L_i is a local layer and each R_i is a reflection layer.

Only reflection layers are counted in d ; local layers are free. Since I is both a local layer and a reflection layer, $\mathcal{U}_{d'}(N) \subseteq \mathcal{U}_d(N)$ for $d' \leq d$. Each reflection gate is a self-adjoint unitary, and gates with disjoint supports commute, so each R_i is a self-adjoint unitary. Therefore

$$U^\dagger = L_0^\dagger R_1^\dagger L_1^\dagger R_2^\dagger \cdots R_d^\dagger L_d^\dagger \in \mathcal{U}_d(N) \quad \text{whenever } U \in \mathcal{U}_d(N).$$

Reflections about product states of this kind also appear in [5].

Geometrically, $I - 2G$ changes the sign on the range of G and fixes its orthogonal complement. Conjugating a projection by U changes the subspace being tested: if P projects onto E , then $UPU^\dagger$ projects onto $U(E)$. This is why conjugated projections appear below. Generalized Toffoli gates have exactly this reflection form after choosing local bases. Proposition 8 proves the equivalence in Section 6; only the displayed reflection representation and its inverse closure are needed for the localization proof.

Definition 2 (Computing a Boolean function). Let $1 \leq n \leq N$, let $f: \{0, 1\}^n \rightarrow \{0, 1\}$, and let $0 < \epsilon \leq 1/2$. The first n qubits of $\mathcal{H}_N$ are the *input qubits*, and the other $N - n$ are *ancillas*. A unitary W on $\mathcal{H}_N$, together with a designated *output qubit* $o \in \{1, \dots, N\}$, *computes f with advantage ϵ* if, for every $x \in \{0, 1\}^n$, measuring qubit o of the state $W(|x\rangle \otimes |0\rangle^{\otimes(N-n)})$ in the computational basis gives the outcome $f(x)$ with probability at least $1/2 + \epsilon$.

No condition is imposed on the final state of the other qubits, and o may be an input qubit or an ancilla. There are no intermediate measurements and no postselection. Conventions that additionally require the answer to be exact, the ancillas to be returned to $|0\rangle$, or the input to be preserved are special cases of Definition 2, so lower bounds for Definition 2 apply to them as well.

2.3 Precise statements

Theorem 1. *Fix an integer $d \geq 0$, a real $c \geq 1$, and $0 < \epsilon \leq 1/2$. There exists $n_0 = n_0(d, c, \epsilon)$ such that, for all integers $n \geq n_0$ and N with $n \leq N \leq n^c$, no $W \in \mathcal{U}_d(N)$, with any choice of output qubit, computes $\text{PARITY}(x) = x_1 \oplus \cdots \oplus x_n$ with advantage ϵ .*

The estimate that implies the theorem. Theorem 1 is derived in Section 5 from the following uniform localization statement. For an integer $d \geq 0$ and real numbers $0 \leq s < t$, let $\mathcal{B}_d(s, t)$ denote the assertion that there exists $N_0(d, s, t)$ such that

$$\left\| [M \geq N^t] U[D = 0] U^\dagger [M = 0] \right\| \leq e^{-N^s} \quad (2)$$

for every integer $N \geq N_0(d, s, t)$, every $U \in \mathcal{U}_d(N)$, and every pair of counts M, D on $\mathcal{H}_N$. The threshold may depend only on d, s , and t : not on the circuit, its gate supports, or the subsets and reference vectors of the counts. If $\mathcal{B}_d(s, t)$ holds, the set of positive integers that serve as such a threshold is nonempty and closed upward; we write $N_0(d, s, t)$ for its least element.

Theorem 2 (Localization). *For every integer $d \geq 0$ and all real numbers $0 \leq s < t$, the assertion $\mathcal{B}_d(s, t)$ holds.*

Theorem 2 is proved by induction on d in Section 4, using the one-layer expansion of Section 3 and the polynomial of Subsection 4.2.

To read (2), start at the right: $[M = 0]$ imposes a product condition, $U[D = 0]U^\dagger$ performs a conjugated product test, and $[M \geq N^t]$ keeps only the part with many mismatches. The norm bound holds for every unit input vector satisfying the initial condition, including arbitrary states on uncounted qubits. The exponent t specifies the mismatch threshold, while s specifies the decay rate. In the parity application the two states have n mismatches; choosing $t < 1/c$ makes $n > N^t$ whenever $N \leq n^c$ and $n > 1$.

3 Direct depth reduction

The useful feature of the following expansion is its coefficient bound, not the number or support sizes of its terms. Disjointness ensures that only the original support of A incurs a decomposition cost.

Polarization in two dimensions. For $z \in \{1, -1, i, -i\}$, let P_z project onto the unit vector $(|0\rangle + z|1\rangle)/\sqrt{2}$. Direct multiplication gives

$$|0\rangle\langle 1| = \frac{1}{2}(P_1 - P_{-1} + iP_i - iP_{-i}).$$

Thus a non-self-adjoint operator is a linear combination of four rank-one orthogonal projections, with total absolute coefficient sum 2. The proof below uses the same identity for arbitrary local directions.

Lemma 3 (One-layer expansion). *Let A be a product projection whose support has k elements, and let R be a reflection layer. Then*

$$RAR = \sum_j c_j A_j, \quad \sum_j |c_j| \leq 18^k,$$

where the sum is finite, $c_j \in \mathbb{C}$, and each A_j is a product projection. The supports of the A_j need not be small.

Proof. Each gate $I - 2G$ of R satisfies $(I - 2G)^2 = I$, and gates of R commute with each other. If a gate's support is disjoint from $\text{supp}(A)$, then the gate also commutes with A , and its two occurrences cancel in RAR . In particular, $RAR = R^2 = I$ when $k = 0$. The remaining gates, say $I - 2G_1, \dots, I - 2G_\mu$, each meet $\text{supp}(A)$, and they meet it in nonempty pairwise disjoint sets; hence $\mu \leq k$, and $RAR = R'AR'$ with $R' = \prod_{\nu=1}^\mu (I - 2G_\nu)$. Expanding

$$R' = \sum_{T \subseteq \{1, \dots, \mu\}} (-2)^{|T|} G_T, \quad G_T = \prod_{\nu \in T} G_\nu,$$

on both sides of A writes RAR as a combination of the operators $G_T A G_{T'}$, with coefficients whose absolute values sum to $(\sum_T 2^{|T|})^2 = 3^{2\mu} \leq 9^k$. This cancellation and expansion are those in the proof of the small-support depth reduction [3, Lemma 6]. Each $G_T A G_{T'}$ is a tensor product of one-qubit operators. On a qubit outside $\text{supp}(A)$, its factor is I or a rank-one projection: at most one gate of R acts on that qubit, so if both G_T and $G_{T'}$ have a nonidentity factor there, the two factors are the same projection, and their product is that projection.

On a qubit $q \in \text{supp}(A)$, the factor has the form $Y_q = \Pi A_q \Pi'$, where A_q is the rank-one factor of A on q and each of Π, Π' is I or a rank-one projection. Thus Y_q has rank at most one and norm at most one. If $Y_q = 0$, the whole term vanishes and is discarded. Otherwise write $Y_q = \gamma |\phi\rangle\langle\psi|$ with unit vectors ϕ, ψ and $0 < |\gamma| \leq 1$. The polarization identity

$$|\phi\rangle\langle\psi| = \frac{1}{4} \sum_{z \in \{1, i, -1, -i\}} z |\phi + z\psi\rangle\langle\phi + z\psi|$$

holds because $\sum_z z = \sum_z z^2 = 0$ and $z\bar{z} = 1$. Normalizing the nonzero vectors $\phi + z\psi$ expresses Y_q as a linear combination of rank-one orthogonal projections, with coefficients whose absolute values sum to at most

$$\frac{|\gamma|}{4} \sum_z \|\phi + z\psi\|^2 = \frac{|\gamma|}{4} \sum_z (2 + 2 \text{Re}(z\langle\phi|\psi\rangle)) = 2|\gamma| \leq 2.$$

Substituting these expansions into the k factors on $\text{supp}(A)$ and multiplying out writes each nonzero $G_T A G_{T'}$ as a combination of tensor products whose factors are I or rank-one projections, that is, of product projections, with coefficients whose absolute values sum to at most 2^k . Altogether $\sum_j |c_j| \leq 9^k 2^k = 18^k$. $\square$

Lemma 3 gives the whole depth-reduction step, without any low-count lemma.

Corollary 4. *Fix an integer $d \geq 1$ and real numbers $0 \leq w < t$, and assume $\mathcal{B}_{d-1}(w, t)$. Let $N \geq N_0(d-1, w, t)$ and $K \geq 0$. Then*

$$\left\| [M \geq N^t] U A U^\dagger [M = 0] \right\| \leq 18^K e^{-N^w} \quad (3)$$

for every $U \in \mathcal{U}_d(N)$, every count M , and every product projection A whose support has at most K elements.

Proof. Write U with exactly d reflection layers and put $V = L_d R_d \cdots L_2 R_2 L_1$ (so $V = L_1$ if $d = 1$). Then $U = V R_1 L_0$ and $V \in \mathcal{U}_{d-1}(N)$. The operator $L_0 A L_0^\dagger$ is a product projection with the same support as A , since each rank-one factor is conjugated by a one-qubit unitary. Lemma 3 gives $R_1 L_0 A L_0^\dagger R_1 = \sum_j c_j A_j$ with $\sum_j |c_j| \leq 18^K$. Each A_j equals $[D_j = 0]$ for some count D_j , so $\mathcal{B}_{d-1}(w, t)$ applied to V and the counts M, D_j gives $\left\| [M \geq N^t] V A_j V^\dagger [M = 0] \right\| \leq e^{-N^w}$. Since $U A U^\dagger = \sum_j c_j V A_j V^\dagger$, the triangle inequality proves (3). $\square$

The assertion at depth $d - 1$ is uniform over all the resulting A_j , including those with very large supports. This uniformity is why the expansion can be used directly. For example, when $K \leq N^u$ and $u < w$, the bound is at most $\exp((\log 18)N^u - N^w)$, which decays rapidly. This elementary comparison is the reason the exponential expansion cost is affordable.

4 Proof of localization

There are two inductions: first on circuit depth, then backwards along a finite chain of mismatch thresholds at the same depth:

all supports at depth $d - 1$ $\xrightarrow{\text{polarization}}$ small supports at depth d ,

$$\underbrace{\mathcal{B}_d(s_J, t_J)}_{t_J > 1: \text{zero tail}} \implies \cdots \implies \underbrace{\mathcal{B}_d(s_0, t_0)}_{\text{desired estimate}} .$$

Each backward step combines the small-support bound with a polynomial approximation and the inverse circuit. Subsection 4.4 constructs the fixed exponent pairs; starting from the zero tail supplies every same-depth hypothesis before it is used.

4.1 The depth-zero estimate

With no reflection layers, conjugation preserves product projections. A tensor-product calculation already gives a strong mismatch bound.

Lemma 5 (Depth zero). *Let M be a count, $P = [M = 0]$, and Q a product projection. Then, for every real $\eta \geq 0$,*

$$\|2^{M/2}QP\| \leq 1, \quad \|[M \geq \eta]QP\| \leq 2^{-\eta/2}. \quad (4)$$

Consequently $\mathcal{B}_0(s, t)$ holds for all real $0 \leq s < t$.

Proof. Write $M = \sum_{i \in S} (I - |v_i\rangle\langle v_i|)_i$. Since the summands act on distinct qubits, $2^{M/2}$ is the tensor product of the one-qubit operators $\Delta_i = |v_i\rangle\langle v_i| + \sqrt{2}(I - |v_i\rangle\langle v_i|)$ for $i \in S$ and I elsewhere. Hence $2^{M/2}QP$ is a tensor product of one-qubit operators, and its norm is the product of their norms. On a qubit outside S , the factor is the factor of Q , a contraction. On $i \in S$ where Q has factor I , the factor is $\Delta_i|v_i\rangle\langle v_i| = |v_i\rangle\langle v_i|$, of norm one. On $i \in S$ where Q has factor $|\phi_i\rangle\langle\phi_i|$, the factor is the rank-one operator $\langle\phi_i|v_i\rangle|\Delta_i\phi_i\rangle\langle v_i|$. With $\alpha_i = |\langle\phi_i|v_i\rangle|^2 \in [0, 1]$, its squared norm is

$$\alpha_i \|\Delta_i\phi_i\|^2 = \alpha_i(\alpha_i + 2(1 - \alpha_i)) = \alpha_i(2 - \alpha_i) \leq 1.$$

This proves the first bound. For the second, $[M \geq \eta]$ commutes with $2^{-M/2}$, so $[M \geq \eta]QP = ([M \geq \eta]2^{-M/2})(2^{M/2}QP)$, and $\|[M \geq \eta]2^{-M/2}\| \leq 2^{-\eta/2}$ by the spectral theorem.

If $U \in \mathcal{U}_0(N)$, then U is a product of one-qubit unitaries, so $Q = U[D = 0]U^\dagger$ is a product projection for every count D . Taking $\eta = N^t$ gives the bound $2^{-N^t/2} = \exp(-(\log 2)N^t/2)$, which is at most e^{-N^s} as soon as $N^{t-s} \geq 2/\log 2$. This threshold depends only on s and t , proving $\mathcal{B}_0(s, t)$. $\square$

4.2 A polynomial approximation to a projection

For a count D , the projection $[D = 0]$ is one at eigenvalue zero and zero at every positive integer eigenvalue. Its polynomial approximation must meet three requirements: equal one at zero to preserve the desired subspace; be small at positive integers below a chosen threshold to suppress mismatches; and have manageable degree and coefficients so that the small-support estimate applies. Indeed, D^k expands into product projections on at most k qubits, whose estimates we must sum. The inverse circuit will control the error beyond the approximation threshold.

For a real polynomial $p(x) = \sum_k c_k x^k$ and an integer $N \geq 1$, put $\text{wt}_N(p) = \sum_k |c_k| N^k$.

Lemma 6. *Fix real numbers $0 < a < b$ and $u > (a + b)/2$. There is $N_1(a, b, u)$ such that for every integer $N \geq N_1(a, b, u)$ there exists a real polynomial p with*

$$p(0) = 1, \quad \deg p \leq N^u, \quad \text{wt}_N(p) \leq e^{N^u}, \quad |p(j)| \leq e^{-N^a} \text{ for every integer } 1 \leq j \leq \lceil N^b \rceil. \quad (5)$$

The construction, a Chebyshev polynomial multiplied by a product vanishing at the first few integers, follows a classical approximation-and-interpolation pattern [6, Section 3.1]; it is the filter of the source's polynomial approximation lemma [3, Lemma 3]. Relative to the source, only the coefficient estimate differs.

Proof. Set $r = \lceil N^a \rceil$, $m = \lceil N^b \rceil$, and $\ell = \lceil 2\sqrt{mr} \log m \rceil$. Since $a < b$, for all sufficiently large N we have $m \geq 2r$ and $m \geq 3$; assume this from now on, so that $\ell \geq 1$. Define the Chebyshev polynomials by $T_0(x) = 1$, $T_1(x) = x$, and $T_{j+1}(x) = 2xT_j(x) - T_{j-1}(x)$. By the recurrence, for $j \geq 1$ the polynomial T_j has degree j and leading coefficient 2^{j-1} . The cosine and hyperbolic-cosine addition identities give $T_j(\cos \theta) = \cos(j\theta)$ and $T_j(\cosh y) = \cosh(jy)$ by induction. Hence T_j has magnitude at most one on $[-1, 1]$, and for $j \geq 1$ its j roots are the distinct numbers $\cos((2\nu - 1)\pi/(2j))$, $1 \leq \nu \leq j$, all in $(-1, 1)$. Set

$$p(x) = \prod_{i=1}^r \left(1 - \frac{x}{i}\right) \cdot \frac{T_\ell((m+r-2x)/(m-r))}{T_\ell((m+r)/(m-r))}.$$

The denominator. Put $\kappa = \sqrt{r/m} \in (0, 1)$ and $y = \log \frac{1+\kappa}{1-\kappa}$. A direct computation gives $\cosh y = \frac{1+\kappa^2}{1-\kappa^2} = \frac{m+r}{m-r}$, and $y = \int_0^\kappa \frac{2d\vartheta}{1-\vartheta^2} \geq 2\kappa$. Hence the denominator is

$$\cosh(\ell y) \geq \frac{1}{2}e^{\ell y} \geq \frac{1}{2}e^{2\ell\kappa} \geq \frac{1}{2}m^{4r},$$

where the last step uses $\ell\kappa \geq 2\sqrt{mr}\kappa \log m = 2r \log m$. In particular the denominator is positive, and $p(0) = 1$.

Values at integers. The integers $1, \dots, r$ are roots of p . For real $x \in [r, m]$, the Chebyshev argument $(m+r-2x)/(m-r)$ lies in $[-1, 1]$, so the numerator has magnitude at most one; and $0 \leq x/i - 1 < m$ for $1 \leq i \leq r$, so the product has magnitude at most m^r . Hence $|p(x)| \leq 2m^{-3r}$ on $[r, m]$. Since $\log m \geq 1$ and $r \geq N^a$, we get $2m^{-3r} \leq 2e^{-3r} \leq e^{-r} \leq e^{-N^a}$. This proves the last condition in (5).

Coefficients. Besides $1, \dots, r$, the roots of p are the ℓ distinct points x with $(m+r-2x)/(m-r) = \cos((2\nu-1)\pi/(2\ell))$, and these lie in (r, m) . Thus p has degree $r+\ell$ and $r+\ell$ real roots $\rho_1, \dots, \rho_{r+\ell} \geq 1$. Since $p(0) = 1$, this gives $p(x) = \prod_{i=1}^{r+\ell} (1 - x/\rho_i)$. The coefficient of x^k is therefore $(-1)^k$ times the k th elementary symmetric polynomial of the positive numbers $1/\rho_i$. So the coefficients alternate in sign, and

$$\text{wt}_N(p) = p(-N) = \prod_{i=1}^{r+\ell} \left(1 + \frac{N}{\rho_i}\right) \leq (1 + N)^{r+\ell}.$$

Finally, $r \leq 2N^a$, $m \leq 2N^b$, and $a < (a+b)/2$ give $r + \ell = O(N^{(a+b)/2} \log N)$, so $\deg p = r + \ell$ and $\log \text{wt}_N(p) \leq (r + \ell) \log(1 + N)$ are both $o(N^u)$. All the implied constants and thresholds depend only on a , b , and u . $\square$

4.3 The inverse-circuit step

We now combine the expansion with the polynomial. The target is $\mathcal{B}_d(s, t)$. We assume a stronger decay estimate at a larger mismatch threshold, $\mathcal{B}_d(t, b)$, and use it to control the polynomial's high-spectrum error. This is a conditional implication; the next subsection supplies its hypothesis through a finite chain starting from a zero tail.

A direct polynomial substitution leaves an error on vectors whose high-count part is uncontrolled. Squaring the desired norm inserts a second copy of the projected test on the right of the error, producing exactly the transition bounded by the inverse-circuit hypothesis.

The parameters have the following roles:

- s, t target decay exponent and mismatch exponent,
- b larger mismatch exponent used for the inverse circuit,
- a accuracy exponent of the polynomial approximation,
- u exponent controlling polynomial degree and coefficient growth,
- w decay exponent supplied by the preceding circuit depth.

The strict inequalities in the proof ensure that decay dominates every coefficient cost. Their precise numerical values are otherwise immaterial.

Lemma 7 (One backward step). *Fix an integer $d \geq 1$, and assume $\mathcal{B}_{d-1}(\sigma, \tau)$ for all real $0 \leq \sigma < \tau$. Let s, t, b be real numbers with $0 \leq s < t < b < 2t - s$. If $\mathcal{B}_d(t, b)$ holds, then $\mathcal{B}_d(s, t)$ holds.*

Proof. Exponents. Put $a = (s + 2t - b)/2$ and $\beta = (a + b)/2$, and set $u = (2\beta + t)/3$ and $w = (\beta + 2t)/3$. Since $b < 2t - s$, we have $s < a < 2t - b$; since $t < b$, also $2t - b < b$, so $0 < a < b$. Moreover $a + b < 2t$ gives $\beta < t$, hence $\beta < u < w < t$. Also $\beta > s$ because $a > s$ and $b > s$. In summary,

$$0 < a < b, \quad \frac{a+b}{2} < u < w < t, \quad a > s, \quad u > s. \quad (6)$$

These exponents depend only on s, t, b . Let N be at least the threshold $N_1(a, b, u)$ of Lemma 6, the threshold $N_0(d - 1, w, t)$ of the hypothesis $\mathcal{B}_{d-1}(w, t)$, and the threshold $N_0(d, t, b)$ of $\mathcal{B}_d(t, b)$. A final threshold depending only on s, t, b is imposed at the end.

Setup. Fix $U \in \mathcal{U}_d(N)$ and counts M, D . Put

$$P = [M = 0], \quad H = [M \geq N^t], \quad C = UDU^\dagger, \quad Q = [C = 0] = U[D = 0]U^\dagger.$$

The operator C is unitarily equivalent to D , so its spectrum is contained in $\{0, 1, \dots, N\}$. Let $p(x) = \sum_k c_k x^k$ be the polynomial of Lemma 6 for a, b, u and this N .

The polynomial term. Write $D = \sum_{i \in S} \delta_i$, where the δ_i are the summands of D (Section 2): rank-one projections on distinct single qubits, tensored with identities. Then D^k is the sum of the $|S|^k \leq N^k$ products $\delta_{i_1} \cdots \delta_{i_k}$ with $i_1, \dots, i_k \in S$. Since the δ_i commute and are idempotent, each such product is the product of the distinct δ_i occurring in it: a product projection whose support has at most k elements. For $k = 0$ the single term is I . As $p(C) = Up(D)U^\dagger$ and $k \leq \deg p \leq N^u$ for every nonzero coefficient, Corollary 4 with exponents (w, t) and $K = N^u$ bounds every resulting term. By the triangle inequality,

$$\|Hp(C)P\| \leq \sum_k |c_k| N^k 18^{N^u} e^{-N^w} = \text{wt}_N(p) 18^{N^u} e^{-N^w} \leq \exp((1 + \log 18)N^u - N^w). \quad (7)$$

Since the spectrum of C lies in $[0, N]$, also

$$\|p(C)\| \leq \max_{0 \leq \lambda \leq N} |p(\lambda)| \leq \text{wt}_N(p) \leq e^{N^u}. \quad (8)$$

Squaring. Since P is a projection, $(HQP)(HQP)^\dagger = HQPQH$, so $\|HQP\|^2 = \|HQPQH\|$. Replacing only the leftmost Q by $p(C) + (Q - p(C))$ gives $HQPQH = Hp(C)PQH + H(Q - p(C))PQH$. Since $\|H\|, \|QH\| \leq 1$,

$$\|HQP\|^2 \leq \|Hp(C)P\| + \|(Q - p(C))PQ\|. \quad (9)$$

No commutation of P with C or Q is used. Squaring is essential here: it places a second Q on the right of the error. The resulting factor $[C \geq N^b]PQ$ will match the inverse-circuit hypothesis exactly.

The error term. Let $E_- = [C < N^b]$, $E_+ = [C \geq N^b]$, and $F = Q - p(C)$. These are functions of C , so they commute, and $E_- + E_+ = I$. Each eigenvalue λ of C with $\lambda < N^b$ is an integer with $0 \leq \lambda \leq \lceil N^b \rceil$. On the corresponding eigenspace, F acts as $1 - p(0) = 0$ if $\lambda = 0$ and as $-p(\lambda)$ if $\lambda \geq 1$. By (5), $\|FE_-\| \leq e^{-N^a}$. Since $N^b > 0$, we have $QE_+ = 0$, so $FE_+ = -p(C)E_+$. Hence $FPQ = FE_-PQ - p(C)E_+PQ$, and (8) gives

$$\|(Q - p(C))PQ\| \leq e^{-N^a} + e^{N^u} \|E_+PQ\|. \quad (10)$$

The inverse circuit. Since $E_+ = U[D \geq N^b]U^\dagger$ and $Q = U[D = 0]U^\dagger$, unitary invariance of the norm gives

$$\|E_+PQ\| = \|[D \geq N^b]U^\dagger[M = 0]U[D = 0]\| \leq e^{-N^t}. \quad (11)$$

The inequality is $\mathcal{B}_d(t, b)$ applied to the circuit $U^\dagger \in \mathcal{U}_d(N)$ (Section 2), with the roles of the counts M and D exchanged.

Conclusion. Combining (7)–(11),

$$\|HQP\|^2 \leq e^{(1+\log 18)N^u - N^w} + e^{-N^a} + e^{N^u - N^t}.$$

Multiplied by e^{2N^s} , each term on the right tends to zero as $N \rightarrow \infty$, because $w > u$ and $w > s$, $a > s$, and $t > u$ and $t > s$ by (6). This includes $s = 0$, where $e^{2N^s} = e^2$. Hence there is a threshold, depending only on s, t, b , beyond which the right side is at most e^{-2N^s} . Taking square roots proves (2) for the exponents (s, t) , uniformly in U , M , and D . $\square$

4.4 Closing the induction with finitely many thresholds

The previous lemma transfers an estimate from a larger mismatch threshold to a smaller one. Above N mismatches the relevant subspace is zero. We therefore construct a finite chain of thresholds leading above N and then apply the lemma backwards. The gaps decrease but stay bounded below, so the chain cannot stall before reaching this trivial case.

Proof of Theorem 2. We use induction on d . The case $d = 0$ is Lemma 5. Let $d \geq 1$ and assume $\mathcal{B}_{d-1}(\sigma, \tau)$ for all $0 \leq \sigma < \tau$. Fix $0 \leq s < t$.

If $t > 1$, then $N^t > N \geq \|M\|$ for every $N \geq 2$ and every count M , so $[M \geq N^t] = 0$, and $\mathcal{B}_d(s, t)$ holds with threshold 2.

Suppose $t \leq 1$ and put $g = t - s > 0$. For $i \geq 0$ define

$$g_i = (1 + 2^{-i}) \frac{g}{2}, \quad t_0 = t, \quad t_{i+1} = t_i + g_{i+1}, \quad s_i = t_i - g_i.$$

Then $g_0 = g$, so $s_0 = s$; and $s_{i+1} = t_i$, so every $s_i \geq 0$. Since $0 < g_{i+1} < g_i$,

$$t_i < t_{i+1} < t_i + g_i = 2t_i - s_i.$$

Moreover $g_i > g/2$, so $t_i > ig/2$, and therefore $t_i > 1$ for $i \geq 2/g$. Let J be the least index with $t_J > 1$; it depends only on s and t , and $J \leq \lceil 2/g \rceil$. By the previous case, $\mathcal{B}_d(s_J, t_J)$ holds. For $i = J-1, J-2, \dots, 0$ in turn, apply Lemma 7 with $(s, t, b) = (s_i, t_i, t_{i+1})$: its hypothesis $\mathcal{B}_d(t_i, t_{i+1}) = \mathcal{B}_d(s_{i+1}, t_{i+1})$ was established at the previous step, and its conclusion is $\mathcal{B}_d(s_i, t_i)$. The case $i = 0$ is $\mathcal{B}_d(s, t)$.

The chain has finitely many steps, determined by s and t , and each step invokes the depth- $(d-1)$ hypothesis at a single fixed exponent pair. Hence $\mathcal{B}_d(s, t)$ holds with a threshold that is the maximum of finitely many thresholds, each depending only on the depth and fixed exponents; no threshold uniform over infinitely many exponent pairs is required. This completes the induction. $\square$

5 From localization to parity

It remains to connect the operator estimate with a Boolean computation. Half the difference between the average output-zero probabilities on even and odd inputs is the parity correlation used below. A change from the computational basis to the product basis $|+\rangle, |-\rangle$ expresses this correlation through matrix entries between complementary product vectors. They differ on every input qubit, placing them exactly in the regime controlled by localization.

Proof of Theorem 1. Fix d, c , and ϵ , and put $s = 1/(3c)$ and $t = 2/(3c)$, so that $0 < s < t < 1/c$. Let $N_0 = N_0(d, s, t)$ be the threshold of Theorem 2, and choose $n_0 \geq \max\{2, N_0\}$ such that $e^{-n^s} < \epsilon$ for all $n \geq n_0$. This n_0 depends only on d, c , and ϵ . Let $n \geq n_0$ and $n \leq N \leq n^c$, and suppose, for contradiction, that some $W \in \mathcal{U}_d(N)$, with output qubit o , computes PARITY with advantage ϵ .

Let O be the projection $|0\rangle\langle 0|$ on qubit o , tensored with the identity on the other qubits. Put

$$\hat{O} = W^\dagger O W, \quad \Phi = (I \otimes \langle 0^{N-n} |) \hat{O} (I \otimes | 0^{N-n} \rangle),$$

so Φ acts on the n input qubits. By the Born rule, the probability that qubit o is measured as 0 on input x is $p_0(x) = \|OW|x, 0^{N-n}\rangle\|^2 = \langle x | \Phi | x \rangle$, whatever the final state of the other qubits.

A trace identity. Let $Z = \text{diag}(1, -1)$ and $|\pm\rangle = (|0\rangle \pm |1\rangle)/\sqrt{2}$. For $h \in \{0, 1\}^n$, let $|\chi_h\rangle$ be the tensor product whose i th factor is $|+\rangle$ if $h_i = 0$ and $|-\rangle$ if $h_i = 1$, and let $\bar{h}$ be the complementary string. Since $Z|\pm\rangle = |\mp\rangle$, we have $Z^{\otimes n}|\chi_h\rangle = |\chi_{\bar{h}}\rangle$. Computing $\text{Tr}(\Phi Z^{\otimes n})$ in the computational basis and in the basis $\{|\chi_h\rangle\}$, and then substituting $\bar{h}$ for h , gives

$$2^{-n} \sum_x (-1)^{|x|} p_0(x) = 2^{-n} \text{Tr}(\Phi Z^{\otimes n}) = 2^{-n} \sum_h \langle \chi_h | \Phi | \chi_h \rangle = 2^{-n} \sum_h \langle \chi_{\bar{h}} | \Phi | \chi_h \rangle, \quad (12)$$

where $|x| = x_1 + \dots + x_n$.

Upper bound. Fix h , and let M_h be the count on the n input qubits whose reference vectors are the factors of $|\chi_h\rangle$. The vector $|\chi_h\rangle|0^{N-n}\rangle$ lies in the range of $[M_h = 0]$. Each input factor of $|\chi_{\bar{h}}\rangle$ is orthogonal to the corresponding factor of $|\chi_h\rangle$, so $|\chi_{\bar{h}}\rangle|0^{N-n}\rangle$ is an eigenvector of M_h with eigenvalue n . As $n \geq 2$, we have $N^t \leq n^{ct} = n^{2/3} < n$, so this vector lies in the range of $[M_h \geq N^t]$. Let D be the count $(I - |0\rangle\langle 0|)$ on qubit o , so that $[D = 0] = O$. Since $W^\dagger \in \mathcal{U}_d(N)$ and $N \geq n \geq N_0$, Theorem 2 applied to $U = W^\dagger$ and the counts M_h, D gives

$$|\langle \chi_{\bar{h}} | \Phi | \chi_h \rangle| = |\langle \chi_{\bar{h}}, 0^{N-n} | [M_h \geq N^t] \hat{O} [M_h = 0] | \chi_h, 0^{N-n} \rangle| \leq e^{-N^s} \leq e^{-n^s} < \epsilon.$$

The same threshold serves every h , by the uniformity in the counts. The right side of (12) is an average of 2^n such entries, so the parity correlation on its left side has absolute value less than ϵ .

Lower bound. If $|x|$ is even, the correct answer is 0, so $p_0(x) \geq 1/2 + \epsilon$; if $|x|$ is odd, the correct answer is 1, so $p_0(x) \leq 1/2 - \epsilon$. Exactly 2^{n-1} strings have each parity, so the left side of (12) is at least $\frac{1}{2}(\frac{1}{2} + \epsilon) - \frac{1}{2}(\frac{1}{2} - \epsilon) = \epsilon$. This contradicts the upper bound. $\square$

6 Standard circuit models and gate-set variants

This section verifies the connection used in the introduction between reflection circuits and the usual QAC model. It also records two common gate-set variants. These facts identify the scope of the theorem; they are not additional ingredients in the localization argument.

6.1 The standard model

Definition 3 (QAC⁰). A *generalized Toffoli gate* with $m \geq 1$ controls is the unitary on $m + 1$ qubits that applies $X = |0\rangle\langle 1| + |1\rangle\langle 0|$ to the target qubit if all m controls are $|1\rangle$, and acts as the identity otherwise. A *standard circuit* on N qubits is a finite sequence of *layers*; each layer is a set of one-qubit gates and generalized Toffoli gates acting on pairwise disjoint sets of qubits, and it implements the product of its gates. The *depth* is the number of layers. A sequence of functions $f_n: \{0, 1\}^n \rightarrow \{0, 1\}$ belongs to QAC⁰ [2] if there are an integer $d \geq 0$, a polynomial q , and, for every n , a standard circuit C_n on $N(n)$ qubits with $n \leq N(n) \leq q(n)$ and depth at most d , which together with some output qubit computes f_n with advantage $1/6$, that is, with success probability at least $2/3$.

The gate set and depth follow [2]; the output convention is the permissive one of Definition 2. The gates of C_n may depend arbitrarily on n ; no uniformity or gate-precision condition is imposed. A layer on N qubits contains at most N gates, so C_n has at most $dN(n)$ gates; the polynomial bound on the number of qubits therefore also bounds the size. Primitive fanout gates are not allowed. The choice $2/3$ is immaterial, since Theorem 1 excludes every fixed advantage $\epsilon > 0$.

6.2 Equivalence with reflection circuits

The next proposition shows that $\mathcal{U}_d(N)$ is a normal form for standard circuits, not an extension of them.

Proposition 8. *For all $N \geq 1$ and $d \geq 0$, $\mathcal{U}_d(N)$ is exactly the set of unitaries implemented by standard circuits on N qubits in which at most d layers contain a generalized Toffoli gate. In particular, every standard circuit on N qubits of depth at most d implements an element of $\mathcal{U}_d(N)$.*

Proof. Let $|\pm\rangle = (|0\rangle \pm |1\rangle)/\sqrt{2}$. Since $X = I - 2|-\rangle\langle -|$, a generalized Toffoli gate with target q_0 and controls $T \setminus \{q_0\}$ equals $I - 2\Pi_T$, where

$$\Pi_T = |-\rangle\langle -|_{q_0} \otimes \bigotimes_{q \in T \setminus \{q_0\}} |1\rangle\langle 1|_q;$$

this is a reflection gate with support T . The gates of a layer act on disjoint qubits and commute, so a layer implements the product of a local layer (its one-qubit gates) and a reflection layer (its generalized Toffoli gates); a layer without generalized Toffoli gates implements a local layer. Products of local layers are local layers, and I is a local layer. Hence a standard circuit in which at most d layers contain a generalized Toffoli gate implements a unitary of the form (1).

Conversely, let $I - 2G$ be a reflection gate with support T and $G = \bigotimes_{q \in T} |g_q\rangle\langle g_q|$. If $|T| = 1$, it is a one-qubit unitary. If $|T| \geq 2$, fix $q_0 \in T$ and one-qubit unitaries V_q with $V_{q_0}|-\rangle = |g_{q_0}\rangle$ and $V_q|1\rangle = |g_q\rangle$ for $q \neq q_0$. With $V_T = \bigotimes_{q \in T} V_q$ we get $V_T \Pi_T V_T^\dagger = G$, so $I - 2G = V_T(I - 2\Pi_T)V_T^\dagger$ is a generalized Toffoli gate conjugated by one-qubit unitaries. Since the gates of a reflection layer have disjoint supports, the layer equals $L' \Theta L'^\dagger L''$, where L' is the tensor product of the V_T with $|T| \geq 2$, Θ is the corresponding layer of generalized Toffoli gates, and L'' is the product of the gates with $|T| = 1$. Substituting this into (1) and merging adjacent local layers gives a standard circuit in which at most d layers contain a generalized Toffoli gate. $\square$

Corollary 9. $\text{PARITY} \notin \text{QAC}^0$. *More precisely, for every integer $d \geq 0$ and every polynomial q , if n is sufficiently large, then no standard circuit of depth at most d on N qubits, with $n \leq N \leq q(n)$, computes PARITY on n bits with advantage $1/6$.*

Proof. Choose a real $c \geq 1$ and an integer n_1 with $q(n) \leq n^c$ for all $n \geq n_1$. Let $n \geq \max\{n_1, n_0(d, c, 1/6)\}$, and let C be a standard circuit of depth at most d on N qubits, where $n \leq N \leq q(n) \leq n^c$. By Proposition 8, C implements some $W \in \mathcal{U}_d(N)$. By Theorem 1, no choice of output qubit makes W compute PARITY with advantage $1/6$. $\square$

6.3 Gate-set variants

A NOT gate controlled on an arbitrary pattern $\omega \in \{0, 1\}^m$ equals $I - 2(\bigotimes_i |\omega_i\rangle\langle \omega_i| \otimes |-\rangle\langle -|)$, again a reflection gate, so Proposition 8 and Corollary 9 remain valid if such gates are allowed. Some definitions of QAC^0 also allow arbitrary gates on at most κ qubits, for a constant κ . Each such gate is a product of one-qubit gates and controlled-NOT gates whose number depends only on κ [1]. Replacing the gates of each layer by these products, in parallel, multiplies the depth by a factor depending only on κ , so Corollary 9 holds for such circuits as well.

Use of AI: The author's main contribution is to *use polarization to express the terms in the reflection expansion as linear combinations of individual product projections*. The author prompted Claude and ChatGPT to apply this idea to simplify OpenAI's proof and work out details interactively and verified them manually.

References

- [1] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, *Elementary gates for quantum computation*, Phys. Rev. A 52 (1995), 3457–3467.
- [2] C. Moore, *Quantum circuits: fanout, parity, and counting*, preprint, 1999. [arXiv:quant-ph/9903046](#).
- [3] OpenAI, *Product-projection localization and the QAC^0 parity lower bound*, preprint, September 24, 2026. Source of the theorem, localization strategy, filter construction, and inverse bootstrap. [Repository manuscript \(21 pages\)](#).
- [4] OpenAI, *Regular trajectories, pruning and quantum parity*, preprint, September 24, 2026. [Repository manuscript \(27 pages\)](#).
- [5] G. Rosenthal, *Bounds on the QAC^0 complexity of approximating parity*, in 12th Innovations in Theoretical Computer Science Conference (ITCS 2021), LIPIcs 185, 32:1–32:20, 2021. [doi:10.4230/LIPIcs.ITCS.2021.32](#).
- [6] A. A. Sherstov, *Approximate inclusion-exclusion for arbitrary symmetric functions*, Comput. Complexity 18(2) (2009), 219–247.