

The Fundamental Limit of Quantum Computation (FLQC)

(c)2026 Sergej Materov

DOI: 10.5281/zenodo.21253441

Abstract.

This work introduces the concept of the **Fundamental Limit of Quantum Computation (FLQC)** - a precision limit of any quantum computing device arising from the discrete structure of spacetime at the Planck scale. The limit cannot be overcome by any technological progress. Consequences for Shor's algorithm, error correction, and an experimental programme on QPUs at 30–50 mK are discussed.

Preface

Since Turing, computational limitations have been classified as algorithmic or engineering. Quantum computing added a third complexity-theoretic. This work introduces a fourth.

The Fundamental Limit of Quantum Computation (FLQC) is a precision limit of any quantum computing device, arising not from imperfect implementation, but from the discrete structure of spacetime at the Planck scale.

This limit cannot be overcome by any technological progress - just as the speed of light cannot be overcome by any acceleration.

Part I. The Origin of the Limit

1.1 Phase Continuity - the Hidden Assumption

The entire architecture of modern quantum computers rests on one assumption so familiar it has ceased to be noticed: quantum phases are continuous.

The qubit state: $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, where α and β are complex numbers on the unit circle. Shor's algorithm, Grover's algorithm, the quantum Fourier transform - all exploit phase interference with arbitrary precision. But what if phases are **not continuous**?

1.2 Phase Discreteness from the $\mathbb{Z}^4$ Lattice

In Quantumograph v14, gauge phases on graph edges are physical degrees of freedom of spacetime. The minimum phase step is set by the torus structure:

$$\Delta\theta_{min} = \frac{2\pi}{L}$$

where $L = N^{1/4} = (10^{120})^{1/4} \approx 10^{30}$ is the linear torus size, $N \leq 10^{120}$ is the vertex count (Bekenstein–Hawking bound). Phase step: $\Delta\theta_{min} \sim 10^{-30}$ rad.

1.3 Why 10^{-30} Radians Matters

Shor's algorithm for breaking RSA-2048 requires phase precision:

$$\Delta\theta_{SHOR} \sim \frac{2\pi}{2^{2048}} \sim 10^{-617}$$

This is **587 orders of magnitude smaller** than $\Delta\theta_{min} \sim 10^{-30}$. Nature does not permit encoding a number with such precision into a qubit phase - not because hardware is imperfect, but because a phase space of the required granularity **does not exist**.

Part II. Formal Definition

2.1 Definition of FLQC

Definition.

The Fundamental (Topological) Limit of Quantum Computation (FLQC) is the lower bound on the phase operation error of any physical quantum computing device:

$$\epsilon_{FLQC} \geq \frac{2\pi}{L} \sim \frac{2\pi}{10^{30}} \sim 10^{-30} \text{ rad}$$

This error does not depend on device temperature, qubit material, error-correction codes, or number of qubits. It is an ontological constant.

2.2 Critical Temperature T_c

Phase space discreteness becomes observable only below the critical temperature:

$$T_c \sim \frac{\hbar \cdot \Delta\theta_{min}}{k_B} \sim 30\text{--}50 \text{ mK}$$

T_c falls exactly in the operating range of modern superconducting QPUs. This is a **prediction of the theory** testable right now.

2.3 Hierarchy of Quantum Computing Limitations

Type of Limitation	Source	Surmountable?
Engineering	Hardware imperfection, noise	Yes. By better technology
Thermodynamic	Decoherence at $T > T_c$	Yes. By cooling below T_c
Complexity-theoretic	Algorithmic complexity	Partially. By quantum speedup
Fundamental (FLQC)	Discreteness of spacetime	No/never

Part III. Consequences

3.1 Shor's Algorithm and Cryptography

The FLQC does not make RSA more secure: RSA-2048 as a mathematical object has not changed. The FLQC constrains the breaking tool, not the lock.

Analogy: opening the lock requires a key 10^{-617} mm thick, but nature does not permit manufacturing anything thinner than 10^{-30} mm. The lock is no stronger - **the required tool does not exist**.

The constraint depends on key length n :

- RSA-512: required precision $\sim 10^{-154}$ rad, gap 124 orders - threat remains
- RSA-2048: required precision $\sim 10^{-617}$ rad, gap 587 orders - Shor's algorithm physically unrealisable
- Boundary: $n \lesssim 100$ bits above this, FLQC physically blocks Shor's algorithm

3.2 Quantum Error Correction

The FLQC predicts a systematic, uncorrectable error component. As QPUs scale, error-correction codes will encounter a *floor* - an error level no correction can reduce below. Detecting this floor is direct experimental evidence of the FLQC.

3.3 Quantum Simulation as a Measurement Tool

A QPU operating below T^c is itself a physical realization of a $\mathbb{Z}^4$ -graph fragment. Its systematic errors are a **direct measurement of spacetime structure**. What hinders computation helps physics.

3.4 Relation to Existing Finite-Information Frameworks

The claim that physical systems carry only a finite amount of information is not new. The Bekenstein bound limits the entropy — and hence the information — that can be enclosed within a finite spatial region to a quantity proportional to the boundary area rather than the volume. The holographic principle, developed by 't Hooft and Susskind, generalizes this to the claim that the number of independent degrees of freedom in any spatial volume is finite and scales with its bounding area. Loop Quantum Gravity independently finds that geometric operators (area, volume) possess discrete spectra. Seth Lloyd's bound on the computational capacity of physical systems and the Margolus–Levitin theorem further constrain, respectively, the total number of operations available to a system and the rate at which it can perform them, from energy and entropy alone. FLQC does not claim priority over this line of work: the finiteness of the number of vertices on the $\mathbb{Z}^4$ graph (§1.2) is itself a holographic-type bound in the spirit of Bekenstein and 't Hooft.

What distinguishes FLQC from this literature is the object being constrained. Existing entropy and holographic bounds limit the number of distinguishable physical states — the cardinality of the accessible state space. The present hypothesis is different: it concerns the minimum physically realizable phase increment, and therefore the precision of physical transformations rather than the cardinality of the state space. The logical chain is: finite spacetime resolution $\Rightarrow$ finite phase resolution $\Rightarrow$ finite gate precision $\Rightarrow$ computational limitation. This last inference — from a bound on state count to a bound on the resolution of the transformation group acting on that space — is the step that remains comparatively unexplored in the finite-information literature, and is the specific claim advanced here.

Among these related but formally distinct proposals, the closest is Palmer's Rational Quantum Mechanics (RaQM, arXiv:2510.02877), which also predicts a fundamental breakdown of ideal quantum computation and likewise targets RSA-2048 as a falsification benchmark. The two proposals are formally distinct and constrain different physical quantities. RaQM discretizes the

Hilbert-space basis itself: qubit amplitudes are required to satisfy a rationality condition derived from Diósi–Penrose gravitational decoherence, and the constrained quantity is the number of maximally entangled qubits N , bounded by a qubit information capacity $N_{\max} \approx 200\text{--}1000$ depending on qubit mass/energy. FLQC instead discretizes the gauge phase field on the $\mathbb{Z}^4$ spacetime graph: the constrained quantity is the angular resolution $\Delta\theta$ of individual phase operations, set by the graph’s holographic vertex bound and independent of N . Where RaQM predicts that entanglement capacity saturates before enough qubits can be brought to bear on Shor’s algorithm, FLQC predicts that phase-rotation precision saturates regardless of qubit count. The two mechanisms are logically independent and, in principle, separately falsifiable: an obstruction to factoring RSA-2048 that appears as a hard wall in qubit count would support RaQM, while one that appears as a phase-error floor growing smoothly with circuit depth would support FLQC. Neither framework’s status here depends on the other.

Part IV. Experimental Programme

4.1 What to Measure

1. Implement QFT on n qubits at $T < T_c$. Measure systematic phase deviation as a function of n .
2. Verify: does the deviation grow with n (engineering noise) or remain constant (FLQC)? Constancy is the FLQC signature.
3. Implement random walks on graph directly on QPU. Measure $d_s(\sigma)$: passage through $d_s = 4$ at σ^* is evidence of $\mathbb{Z}^4$ structure.
4. Compare phase errors with $\Delta\theta_{\min} \sim 10^{-30}$ rad. Order-of-magnitude agreement confirms FLQC.

The simulation for the hypothesis is located in the repository:
<https://github.com/SergejMaterov/FLQC-Protocol>

4.2 What This Gives Science

- First direct evidence of spacetime discreteness
- A new fundamental technological limit
- Reassessment of quantum computer threats to cryptography
- A new quantum metrology tool

4.3 Present Feasibility of Direct Laboratory Verification

A direct measurement of $\Delta\theta_{\text{FLQC}} \sim 10^{-30}$ rad is not achievable with current or near-term qubit technology. The best demonstrated superconducting-qubit coherence times are of order 1 ms, and the best achievable measurement-cycle repetition rates are of order 10^6 Hz; even under an idealized Heisenberg-limited phase-estimation scheme with $N \sim 10^3$ entangled qubits, resolving a phase at the 10^{-30} rad level would require on the order of 10^{54} independent measurements, i.e. roughly 10^{48} seconds of continuous operation — far beyond the age of the universe. Independently, ordinary decoherence-induced phase noise at currently achievable coherence times already exceeds the predicted floor by a comparable margin, so the FLQC signal would remain masked by conventional technical noise even before any statistical averaging could recover it. This gap is not an artefact of the estimate in §1.2–1.3: it reflects the same ~ 587 -order-of-magnitude separation that makes RSA-2048 unreachable by Shor’s algorithm under FLQC. The experimental programme in §4.1 should accordingly be read as identifying the qualitative signature to search for — a non-vanishing, n -independent phase floor distinguishable in principle from engineering noise by its scaling behaviour

- rather than as a near-term quantitative target. Closing this gap would require either a qualitatively different metrological principle or a physical channel with more favourable scaling than direct phase-noise averaging; this question is left open for future work. The implementation attempt “Theorem on the absence of asymptotic chaos” is described in the Appendix.

Part V. Phase Resolution, Asymptotic Quantum Algorithms, and the FLQC Constraint

The Fundamental Limit of Quantum Computation (FLQC) has a direct consequence for algorithms whose computational advantage depends on arbitrarily fine phase resolution. The Quantum Fourier Transform (QFT), which forms the basis of Shor's order-finding procedure, is the most direct example.

For an input size M , Shor's algorithm uses a phase-estimation register with approximately $n \approx 2 \log_2(M)$ qubits. The controlled phase rotations in the QFT are represented by

$$R_k = \begin{pmatrix} 1 & 0 \\ 0 & e^{2\pi i/2^k} \end{pmatrix}$$

The standard formulation assumes that the phase angle can be implemented with arbitrary precision. Under the FLQC hypothesis this assumption is replaced by a physical resolution limit:

$$\Delta\theta \geq \Delta\theta_{\text{FLQC}} \approx 10^{-30} \text{ rad}$$

The experimentally implemented operation is therefore not the ideal rotation, but a discretized operation:

$$R_k^{\text{FLQC}} = \begin{pmatrix} 1 & 0 \\ 0 & \exp\left(i \left\lfloor \frac{\theta_{\text{ideal}}}{\Delta\theta_{\text{FLQC}}} \right\rfloor \Delta\theta_{\text{FLQC}}\right) \end{pmatrix}$$

This does not represent ordinary calibration error. It represents a finite resolution of the physical phase space itself.

For RSA-2048, the required register size is approximately $n \approx 4096$ qubits.

The smallest QFT rotations have: $\theta_{\text{ideal}} = \frac{2\pi}{2^{4096}} \approx 10^{-1233} \text{ rad}$.

This value is many orders of magnitude below the assumed FLQC resolution. Therefore, rotations with sufficiently large k cannot encode additional physical information in the phase degree of freedom. The mathematical circuit continues to exist as an abstract operator sequence, but the corresponding physical implementation may become indistinguishable from an identity operation. The accumulated effect can be estimated through the fidelity between the ideal and physical QFT states:

$$F = |\langle \psi_{\text{ideal}} | \psi_{\text{act}} \rangle|^2.$$

Approximating the accumulated phase deviations by $\delta\theta_j$,

$$F \approx \prod_j \cos^2(\delta\theta_j)$$

For a systematic phase floor,

$$\sigma^2 \approx n \left(\frac{\Delta\theta_{\text{FLQC}}}{2} \right)^2$$

The resulting interference degradation behaves as:

$$P_{\text{success}} \propto e^{-\sigma^2}$$

The significance of this expression is not the exact numerical coefficient,

which depends on the architecture, but the existence of a non-zero asymptotic floor that does not vanish with improved engineering.

This leads to a distinction between two classes of errors. Engineering errors arise from noise sources, control imperfections, thermal fluctuations, and material defects. They can in principle be reduced by improved design. The FLQC contribution would instead remain after these effects are minimized because it originates from the assumed discrete structure of the physical phase space. The consequence for quantum error correction is a change in interpretation. Standard fault-tolerant schemes assume that arbitrarily small coherent rotations can be represented and projected into correctable error channels. If the physical Hilbert-space realization has a minimum phase granularity, error correction may encounter a lower bound below which additional correction cycles cannot improve the result.

A direct experimental test follows. A QPU operated below the critical temperature range predicted by Quantumograph should be used to implement QFT circuits of increasing depth. The measured systematic phase deviation should be separated into a component that scales with hardware noise and a component that remains constant. A persistent component approaching the predicted order of magnitude would constitute evidence for a fundamental phase-resolution limit.

Note on Interpretation

The FLQC hypothesis does not imply that quantum computation is impossible. It proposes that the ultimate computational capacity of a physical quantum system may differ from the capacity of an abstract mathematical quantum computer. The question is therefore not only whether algorithms exist, but whether the physical universe permits the exact operations assumed by those algorithms.

Conclusion

The Fundamental Limit of Quantum Computation is a physical consequence of the finiteness of the world. Any device existing in discrete spacetime inevitably inherits its granularity.

The precision of computation is bounded by the precision of reality itself.

This can be verified. The equipment exists. The experiment remains to be done.

Funding Acknowledgement: This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors. It was conducted entirely as an independent research initiative.

Competing Interests: The author declares no competing financial or non-financial interests that could inappropriately influence or bias the content of this manuscript.

Appendix Chaotic (OTOC) Amplification Cannot Rescue Direct Detection of $\Delta\theta_{FLQC}$

Motivation

Given that a direct measurement of $\Delta\theta_{FLQC} \sim 10^{-30}$ rad is ruled out on statistical and decoherence grounds (§4.3), it is natural to ask whether exponential sensitivity to initial conditions in chaotic quantum circuits (quantified by the out-of-time-order correlator, OTOC, and its associated quantum Lyapunov exponent λ) could amplify the tiny floor to an observable level. This note works out that possibility and shows it fails for two independent, structural reasons - not for lack of engineering effort.

1. The optimistic ceiling

The number of e-folds needed to amplify $\Delta\theta_{FLQC} \sim 10^{-30}$ to $O(1)$ is $N_{efold} = \ln(10^{30}) \approx 69$.

The quantum bound on chaos (Maldacena–Shenker–Stanford) states $\lambda \leq \frac{2\pi k_B T}{\hbar}$. At the FLQC operating point $T_c \sim 30$ mK this gives $\lambda_{max} 2.5 \times 10^{10}$ rad/s.

If a system saturated this bound (as only maximally chaotic systems such as the SYK model or black-hole duals are believed to do), the required

evolution time would be $t = N_{efold} / \lambda_m \approx 2.8$ ns, comfortably inside any realistic T2 budget. Unlike the direct-measurement case (§4.3), timescale is *not* the obstruction here. Two other obstructions are fatal instead.

2. Structural cancellation in echo/OTOC protocols

Standard OTOC and Loschmidt-echo protocols require implementing the time-reversed evolution, which for a phase-rotation gate means substituting the target angle $\varphi \rightarrow -\varphi$ (since $R(\theta)^{-1} = R(-\theta)$).

The physical realization of a rotation gate under FLQC rounds the target angle to the nearest point on the phase lattice:

$\varphi_{phys} = \varphi + \delta(\varphi)$, $\delta(\varphi) = \varphi - \text{round}\left(\frac{\varphi}{\Delta\theta_{min}}\right) \cdot \Delta\theta_{min}$ because the rounding lattice $\{k \cdot \Delta\theta_{min}\}$ is symmetric about zero, $\text{round}(-x) = -\text{round}(x)$, and therefore $\delta(-\varphi) = -\delta(\varphi)$ (δ is an odd function of φ).

Consider a forward gate with target angle θ_i and its reversal with target angle $-\theta_i$, both physically realized on the same hardware:

forward: $\theta_i + \delta(\theta_i)$

backward: $-\theta_i + \delta(-\theta_i) = -\theta_i - \delta(\theta_i)$

sum = $[\theta_i + \delta(\theta_i)] + [-\theta_i - \delta(\theta_i)] = 0$.

The FLQC contribution cancels *exactly*, to leading order, in any protocol that implements time-reversal as angle negation. This is not a practical limitation but a direct consequence of (i) δ being odd and (ii) reversal being negation - both structural features of the discretization and of standard

echo protocols. Any OTOC/Loschmidt-echo measurement of this type is blind to $\Delta\theta_{FLQC}$ by construction.

Implication: detection would require comparing the real QPU output against an independent reference (classical simulation for small N, or a statistical estimator such as cross-entropy benchmarking, XEB) rather than against a hardware-implemented time-reversal of itself. This removes the cancellation, but reintroduces the scaling problem below.

3. Signal-to-noise ratio is invariant under chaotic growth

Write the real Hamiltonian as

$$H_{real} = H_{ideal} + \varepsilon_{tech} \cdot V_{tech} + \varepsilon_{FLQC} \cdot V_{FLQC},$$

with $\varepsilon_{tech} \sim 10^{-3}-10^{-4}$ (typical coherent two-qubit gate miscalibration on current best hardware) and $\varepsilon_{FLQC} \sim \Delta\theta_{min} \sim 10^{-30}$.

In OTOC theory, the Lyapunov exponent λ characterizes the chaotic *background* dynamics H_{ideal} , not the perturbation itself: any small, generic perturbation grows at the same rate,

$$d_{tech}(D) \approx \varepsilon_{tech} \cdot e^{\lambda D},$$

$$d_{FLQC}(D) \approx \varepsilon_{FLQC} \cdot e^{\lambda D}.$$

Hence

$$\frac{d_{FLQC}(D)}{d_{tech}(D)} = \frac{\varepsilon_{FLQC}}{\varepsilon_{tech}} = \text{const}, \text{ independent of depth } D.$$

Chaotic amplification boosts signal and background by the *same factor* at every step, so their relative visibility never improves, no matter how many e-folds are accumulated (up to OTOC saturation). Resolving a relative deviation of $\frac{\varepsilon_{FLQC}}{\varepsilon_{tech}} \sim 10^{-27}$ statistically requires

$$M \sim \left(\frac{\varepsilon_{tech}}{\varepsilon_{FLQC}} \right)^2 \sim 10^{54} \text{ measurements, the same astronomical figure as the direct-measurement case}$$

in §4.3. Amplification does not remove the statistical burden; it merely relocates it.

If instead the technical noise is stochastic (decoherence) rather than coherent, the invariance argument above does not apply directly – but decoherence then bounds the usable circuit depth to $D \sim \frac{1}{\gamma}$ (γ = per-gate decoherence rate), which is reached long before $N_{efold} \approx 69$ can accumulate, reproducing the T2-budget obstruction already established in §4.3 by a different route.

Conclusion

Both branches - coherent technical noise and stochastic decoherence - independently close this path, for structurally different reasons:

- **Coherent noise:** chaotic growth amplifies signal and background proportionally; their ratio is a depth-independent invariant, so no amount of scrambling improves distinguishability.
- **Decoherence:** the available circuit depth is capped well below the depth needed to accumulate the required 69 e-folds.

Additionally, the most natural experimental implementation (echo/OTOC via Hamiltonian reversal) is “structurally blind” to $\Delta\theta_{FLQC}$ because the rounding error is an odd function of the target angle and reversal is angle negation - the two effects cancel exactly, independent of the two points above.

This is a stronger and more specific statement than the earlier general observation that $\Delta\theta_{FLQC}$ lies far below current noise floors (§4.3): here, the negative result follows from the mathematical structure of chaotic amplification and of the discretization itself, not merely from present-day hardware limits. It is unlikely to be overturned by future improvements in coherence time, qubit count, or control precision, since none of these change the invariance argument in §3 or the parity argument in §2.

References

- Shor, P. W. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer / P. W. Shor // SIAM Journal on Computing. - 1997. - Vol. 26, no. 5. - P. 1484–1509.
- Grover, L. K. Quantum mechanics helps in searching for a needle in a haystack / L. K. Grover // Physical Review Letters. - 1997. - Vol. 78, no. 2. - P. 325–328.
- Bombelli, L., Lee, J., Meyer, D., & Sorkin, R. D. (1987). Space-time as a causal set. Physical Review Letters, 59(5), 521–524.
- Dawid, R. (2013). String Theory and the Scientific Method. Cambridge University Press.
- Fredkin, E. (1990). Digital mechanics: An informal introduction. Physica D, 45(1–3), 254–270.
- Wolfram, S. (2002). A New Kind of Science. Wolfram Media.
- Materov, S. (2025 - 2026). Quantumograph v13. A Testable Quantum Graph Theory of Spacetime. <https://doi.org/10.5281/zenodo.18410313>
- Materov, S. (2026). Theorem on the absence of asymptotic chaos. <https://doi.org/10.5281/zenodo.19343624>
- Materov, S. (2026). Computational Reducibility theorem. <https://doi.org/10.5281/zenodo.19325404>
- Materov, S. (2026). Simulation FLQC-Protocol. <https://github.com/SergejMaterov/FLQC-Protocol>