

The Fundamental Limit of Quantum Computation (FLQC) v8

Revised version (v8) — corrected §2.2.1 arithmetic; explicit junction-vs-package scale distinction; competing local-mapping hypothesis; coherent-vs-stochastic distinction with a Josephson-metrology falsification route

(c) 2026 Sergej Materov

Revision Note (v8)

This revision **withdraws** the claim, present in earlier circulated versions of this preprint, that FLQC renders Shor's algorithm physically unrealisable for RSA-2048 and sets an obstruction threshold at $n^* \approx 100$ bits (formerly §3.1). That argument rested on an invalid premise: that a single physical rotation gate must directly realise an exponentially small ideal phase angle. This is contradicted by established results in quantum-circuit compilation (Solovay–Kitaev/Ross–Selinger synthesis into a fixed discrete gate set; Coppersmith's approximate quantum Fourier transform) and by an explicit, quantitative treatment of exactly this question already in the peer-reviewed literature (Fowler & Hollenberg, Phys. Rev. A 70, 032329, 2004), which shows that factoring integers of thousands of bits requires no controlled rotation finer than $\frac{\pi}{64}$ - roughly 28 orders of magnitude coarser than the $\Delta\theta_{\min}$ estimate used throughout this work. The corrected argument is given in §3.1 and Part V.

v8 makes explicit a second source of uncertainty in §2.2.1 that is distinct from, and larger than, the boundary-*shape* ambiguity already flagged in v7 (item (i)): the choice of length scale r_{dev} itself. Evaluating the same construction at the scale of the functional element that actually carries the physical phase - a single Josephson-junction loop, $r \sim 100$ nm - rather than at the scale of the surrounding chip or cryostat package ($r \sim 1$ mm – 1 m), changes $\Delta\theta_{\text{local}}$ by roughly two orders of magnitude. Both scales are now shown side by side in the table, explicitly labelled.

v8 adds a second, independent candidate resolution of the ontological-mapping problem of §2.2 (item (ii)): a statistical construction in which the device phase is modelled as an average over $N_{\text{dev}} \sim$ independent Planck-scale contributions, rather than derived from a holographic rescaling of area. This candidate predicts a *different functional form* - $\Delta\theta_{\text{local}} \propto r_{\text{dev}}^{-1}$, rather than the $r_{\text{dev}}^{-1/2}$ scaling of the Bousso construction - a disagreement in the shape of the prediction, not merely its normalisation. This is now flagged as an explicit open sub-problem, §2.2.1 item (iii), rather than left implicit.

Finally, v8 adds a candidate near-term empirical handle on item (iii). The AC Josephson voltage–frequency relation $V = \left(\frac{h}{2e}\right) f$ has been verified across independent national metrology institutes, junction materials and array geometries to a relative precision of order 10^{-10} , with no deviation from the relation reported (Niemeyer et al., 1986; NIST–BIPM and related bilateral comparisons, 2001–2009). A systematic (non-averaging) FLQC-type phase floor of the magnitude estimated in §2.2.1 for a real junction geometry, expressed as a fraction of a full cycle, is only ~ 5 orders of magnitude below this already-demonstrated precision - a far smaller safety margin than the ~ 40 -order margin found for Shor's algorithm in §3.1. This makes a *coherent* version of the local floor a considerably more attainable experimental target than the whole-device tests of Part IV -conditional on resolving item (iii): whether the predicted floor would in fact be coherent, or would average away over the $\sim 10^{13}$ – 10^{15} oscillation cycles of a typical measurement, the way the stochastic contributions discussed in the Appendix do. This revision does not resolve that question; it identifies it precisely enough to be worth resolving.

Abstract

This work explores the **Fundamental Limit of Quantum Computation (FLQC)**: a hypothesised precision limit on quantum phase operations, motivated by the conjectured discrete structure of spacetime at the Planck scale within the Quantumgraph v14 graph model. Part I derives a global quantity $\Delta\theta_{\min} = \frac{2\pi}{L}$ from a holographic bound on the total vertex count of the spacetime graph. Section 2.2 identifies an unresolved gap between this global quantity and the phase resolution of any specific local device, and constructs - via Bousso's covariant entropy bound applied to a device's own boundary - a candidate local rescaling, with an explicit falsifiable scaling prediction. Three problems are identified as remaining open: (i) the choice of boundary geometry, (ii) the choice of length scale (functional element versus surrounding package) at which the construction is evaluated, and (iii) whether a competing, statistically motivated construction - predicting a

different scaling exponent - or the Bousso construction is closer to correct, and whether either predicts a coherent or a stochastic local floor. Contrary to earlier versions of this manuscript, **no implication for the feasibility of Shor's algorithm or the security of RSA is claimed**: established gate-synthesis and approximate-QFT results show that the phase precision actually required to factor any cryptographically relevant modulus is many orders of magnitude coarser than any of the candidate local values of $\Delta\theta_{\text{local}}$, so FLQC - whatever its eventual local value and functional form - does not constrain the computational power of a quantum computer at any presently or foreseeably relevant scale. A separate, logically independent, and directly testable prediction is retained: a hardware-local critical temperature T_c (§2.2), derived from the graph's own noise Hamiltonian, together with an explicit experimental protocol (Part IV), a proposed near-term metrological test of item (iii) via existing high-precision Josephson-voltage-standard data, and a separate analysis of why chaotic (OTOC) amplification cannot rescue direct laboratory detection of $\Delta\theta_{\text{FLQC}}$ (Appendix).

Preface

Since Turing, computational limitations have been classified as algorithmic or engineering. Quantum computing added a third: complexity-theoretic. This work asks whether a fourth kind exists - one arising not from imperfect implementation, but from the possibility that spacetime itself is discrete at the Planck scale.

Unlike engineering limits, such a limit - if it exists, and if a coherent link between the global graph-theoretic quantity and a specific laboratory device can be established - could not be removed by better technology. Whether it would have any practical consequence for computation, even if it exists, is a separate question, addressed critically in Part III below. The two questions should not be conflated, and earlier versions of this manuscript did conflate them.

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, *in the idealised mathematical description of the circuit*. Whether the physical implementation of that circuit ever needs to realise that arbitrary precision directly is a separate question, taken up in Part III.

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

In Quantumograph, gauge phases on graph edges are treated as physical degrees of freedom of spacetime. A candidate 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 and $N \leq 10^{120}$ is the vertex count (a Bekenstein-Hawking-type bound applied to the graph as a whole). This gives an order-of-magnitude phase step $\Delta\theta_{\min} \approx 6.28 \times 10^{-30}$ rad. As §2.2 makes explicit, this quantity is defined for the graph as a whole and is not, by itself, a statement about any specific local device.

1.3 Orders of Magnitude - and a Caveat

Shor's algorithm's phase-estimation register nominally involves controlled-rotation angles as small as $2\pi/2^n$ for an n-qubit register; for RSA-2048 ($n \approx 4096$) this is $\theta \approx 2\pi/2^{4096} \approx 10^{-1233}$ rad, many orders of magnitude below $\Delta\theta_{\min} \approx 6.28 \times 10^{-30}$ rad. A naive numerical comparison of these two figures might suggest that such phases cannot be physically encoded and that the algorithm is therefore unrealisable. **Part III shows in detail why this comparison does not support that conclusion**: no physical implementation of Shor's algorithm

ever requires a single gate to realise an angle anywhere near $2\pi/2^n$. The gap between the two numbers does not, by itself, imply anything about feasibility, because the implicit premise behind the comparison — that one physical gate must reach that precision — is false. This is flagged here so the reader does not draw the incorrect conclusion before reaching Part III.

Part II. Formal Definition

2.1 Definition of FLQC (hypothesis)

Definition (hypothesis). If a coarse-graining map from the graph-global quantity $\Delta\theta_{\min}$ to the phase resolution of a specific physical device can be constructed (§2.2), it would set a hypothesised lower bound on that device's phase-operation error:

$$\varepsilon_{\text{FLQC}} \geq \Delta\theta_{\min} \approx 6.28 \times 10^{-30} \text{ rad}$$

Whether such a map exists, whether the resulting bound is device-independent (as assumed for concreteness in Parts I, III and IV below) or instead depends on a device's own physical scale and geometry (as the constructions in §2.2.1 suggest is more likely), and what its correct numerical value and functional form would then be, are not yet established. This is discussed in detail in §2.2.

2.2 Two Predictions, and an Open Problem

This theory makes two logically independent predictions, which differ in kind and not merely in magnitude: one is a global, currently unconnected-to-any-device quantity; the other is a local, hardware-dependent quantity measurable now.

Prediction 1: Minimal Phase Resolution (global; link to any device not yet established)

$\Delta\theta_{\min} = \frac{2\pi}{L}$, with $L = N^{1/4}$ and $N \leq 10^{120}$ the total vertex count of the graph (§1.2), giving $\Delta\theta_{\min} \approx 6.28 \times 10^{-30}$ rad. This is a property of the graph as a whole: it does not, by construction, refer to the couplings, topology, or temperature of any particular device. **No coarse-graining map from this global quantity to a local, per-qubit angle has been constructed.** It remains an in-principle, graph-level quantity rather than a laboratory target or a validated prediction about any specific piece of hardware.

2.2.1 Toward a Local Bound: Two Competing Constructions (worked, still provisional)

Two independents, and independently plausible, routes have now been worked through for turning the global quantity $\Delta\theta_{\min}$ into a statement about a specific local device. They disagree - not just in normalisation, but in the shape of the prediction - and that disagreement is itself informative.

Construction A: covariant-entropy-bound rescaling. Rather than asserting that the whole-universe quantity $N \leq 10^{120}$ applies unchanged to an arbitrary local subsystem, this construction applies the *same rule* Quantumograph already uses globally $N \leq \text{Area}/(4\ell_P^2)$, $L = N^{1/4}$, $\Delta\theta_{\min} = \frac{2\pi}{L}$ - to the boundary of a specific device's own spatial region, rather than to the cosmological horizon. This is not an ad hoc modification: it is precisely the generalisation the holographic principle is standardly taken to require, made explicit in Bousso's *covariant entropy bound* (Bousso, R., 1999, JHEP 9907:004; hep-th/9905177), which states that $S \leq \frac{A}{4\ell_P^2}$ holds for the light-sheet of *any* suitably defined bounded surface - not only a cosmological or black-hole horizon.

For a device occupying a spatial region of characteristic linear size r_{dev} , bounded for concreteness by a sphere of that radius:

$$N_{\text{dev}} \leq \frac{\text{Area}_{\text{dev}}}{4\ell_P^2} = \frac{\pi \cdot r_{\text{dev}}^2}{\ell_P^2}$$

$$l_{dev} = N_{dev}^{\frac{1}{4}} = \pi^{\frac{1}{4}} \cdot \left(\frac{r_{dev}}{\ell_p} \right)^{\frac{1}{2}}$$

$$\Delta\theta_{local}^A(r_{dev}) = \frac{2\pi}{l_{dev}} = 2\pi^{\frac{3}{4}} \cdot \sqrt{\frac{\ell_p}{r_{dev}}}$$

This expression introduces **no free parameter** beyond r_{dev} itself: it applies the model's own N–L– $\Delta\theta$ relation, unchanged, to a smaller boundary.

Construction B: statistical (central-limit) rescaling. An independent, and equally natural, hypothesis is that a device's macroscopic phase is not a single graph-edge holonomy but an emergent, coarse-grained *average* over the N_{dev} Planck-scale degrees of freedom within its causal patch. Modelling each microscopic contribution as an independent, order-unity phase, the central limit theorem gives a residual after averaging of

$$\Delta\theta_{local}^B(r_{dev}) \approx N_{dev}^{\frac{1}{2}} \approx \frac{\ell_p}{r_{dev}}$$

up to an order-unity prefactor - i.e. a *different scaling exponent* (r_{dev}^{-1} , not $r_{dev}^{-\frac{1}{2}}$) from Construction A.

Numerical comparison - two length scales. A single “device size” is itself ambiguous (see item (ii) below): the functional element that actually carries the physical phase (a Josephson-junction loop, $r \approx 56\text{--}178\text{ nm}$ for typical transmon junction areas of $0.01\text{--}0.1\text{ }\mu\text{m}^2$) is four to five orders of magnitude smaller than the chip or package that is usually meant by “the device” ($r \sim 1\text{ mm} - 1\text{ m}$). The table below gives Construction A at both scales, together with the (previously mis-stated) comparison to $\Delta\theta_{min}$ and to current engineering gate-calibration precision ($\approx 10^{-3}\text{--}10^{-4}$ rad), recomputed directly from the formula above:

r_{dev}	N_{dev}	$\Delta\theta_{local}^A$	Orders coarser than $\Delta\theta_{min}$	Orders finer than eng. precision (1e-3 – 1e-4)
100 nm (junction)	$\sim 1.2 \times 10^{56}$	$\sim 6.0 \times 10^{-14}$ rad	~ 16.0	$\sim 9.2 - 10.2$
1 mm (package)	$\sim 1.2 \times 10^{64}$	$\sim 6.0 \times 10^{-16}$ rad	~ 14.0	$\sim 11.2 - 12.2$
1 cm	$\sim 1.2 \times 10^{66}$	$\sim 1.9 \times 10^{-16}$ rad	~ 13.5	$\sim 11.7 - 12.7$
10 cm	$\sim 1.2 \times 10^{68}$	$\sim 6.0 \times 10^{-17}$ rad	~ 13.0	$\sim 12.2 - 13.2$
1 m	$\sim 1.2 \times 10^{70}$	$\sim 1.9 \times 10^{-17}$ rad	~ 12.5	$\sim 12.7 - 13.7$

(The 100 nm row is representative of the 56–178 nm range typical of transmon junction areas of $0.01\text{--}0.1\text{ }\mu\text{m}^2$.)

Across this full range of length scales, $\Delta\theta_{local}^A$ is **$\sim 12\text{--}16$ orders of magnitude coarser** than the whole-universe estimate ($\Delta\theta_{min} \approx 6.28 \times 10^{-30}$ rad), and **$\sim 9\text{--}14$ orders of magnitude finer** than both current engineering gate-calibration precision and the Fowler–Hollenberg threshold for affecting Shor's algorithm ($\approx 5 \times 10^{-2}$ rad). (An earlier draft of this table stated a “13–16” and a “14–16” range for these two comparisons respectively; the corrected figures above are recomputed directly from the formula and supersede that statement. The qualitative conclusion - an enormous safety margin in both directions - is unaffected.) The correction of §3.1's conclusion is therefore robust to which length scale, and to either construction A or B, eventually turns out to be closer to correct.

The scaling law $\Delta\theta_{local}^A \propto r_{dev}^{-1/2}$ (Construction A) versus $\Delta\theta_{local}^B \propto r_{dev}^{-1}$ (Construction B) is itself a distinctive, falsifiable feature distinguishing the two constructions: a future cross-device or cross-scale comparison of any observed phase-error floor would in principle be able to tell them apart by the *exponent* of its size-dependence, not merely by its normalisation is a signature that would in principle also distinguish

either from an ordinary engineering-noise floor, which need not follow either scaling. This is proposed here as an additional candidate signature for the experimental programme of Part IV, alongside the n-independence test of §4.1.

Three problems remain open, and should not be understated.

(i) The choice of boundary. Construction A assumes a spherical region of radius equal to the device's characteristic size; a fully covariant treatment would instead specify a particular light-sheet associated with the device's causal history over some relevant timescale, which could plausibly differ from this choice by an order-1 geometric factor. Because $\Delta\theta_{\text{local}}^A$ scales only as $r_{\text{dev}}^{-1/2}$, this ambiguity affects the result weakly - an order-of-magnitude error in the assumed boundary size changes $\Delta\theta_{\text{local}}^A$ by only a factor of a few - but it is not eliminated. A more rigorous alternative would define the region via the modular Hamiltonian of the device's actual reduced quantum state, following Casini's entropy-positivity formulation of the Bekenstein bound (Casini, H., 2008, *Class. Quantum Grav.* 25, 205021; arXiv:0804.2182), rather than via an assumed geometric shape; this has not been attempted here.

(ii) The choice of length scale. Unlike (i), this is a large effect, not a weak one: evaluating Construction A at the scale of the physical functional element (a single junction, $r \approx 100$ nm) rather than the surrounding package ($r \sim 1$ mm – 1 m) changes $\Delta\theta_{\text{local}}^A$ by roughly two orders of magnitude (see table above), because it is not merely a different value of the same free parameter but a genuine ambiguity about *which physical region the construction should be applied to*. This manuscript does not resolve which scale is the physically correct one; the table above reports both explicitly rather than picking one silently, as earlier drafts did.

(iii) Which construction, and is the resulting floor coherent or stochastic? Constructions A and B are both physically motivated, agree on the order of magnitude of the effect (both give a floor many orders below current engineering precision - see the accumulated worked example above for A; B gives comparable or even smaller values for realistic r_{dev}), but disagree on the scaling exponent, and, more importantly, differ in kind: A is a geometric rescaling of a single global bound, while B is an explicitly statistical/averaging construction. A device that is well described by B would show the corresponding floor as *noise* - a small residual variance that shrinks further under any additional averaging - whereas a device well described by A, if the local bound is a genuine per-operation rounding (as posited in Part V), would show it as a *systematic, non-averaging bias*. Distinguishing these two possibilities matters for the experimental programme: only a coherent (systematic) floor is, in principle, within reach of near-term precision metrology (see the proposed test below); a stochastic floor of either construction's predicted magnitude is not, for the same statistical reasons given in §4.3 and the Appendix.

(iv) The ontological mapping. Even granting a local $\Delta\theta_{\text{local}}(r_{\text{dev}})$ of either kind derived above, nothing in either construction explains *why* a specific measurable phase of a specific laboratory qubit - e.g. the Josephson phase of a transmon, or a spin's Larmor phase - should be identified with a graph-edge $U(1)$ holonomy of the causal patch described. This is a distinct problem from (i)–(iii), and is **not** resolved by either rescaling above: it requires an explicit correspondence between Quantumograph's plaquette-Hamiltonian formalism (already implemented for $U(1)/SU(2)/SU(3)$ in the accompanying software) and the specific control Hamiltonian of a real device. The most promising route toward this, not attempted in this manuscript, is not a further heuristic rescaling but a genuine lattice-continuum matching calculation: since the compact- $U(1)$ plaquette action already implemented is known to have Maxwell electrodynamics as its continuum limit, the physically measured Josephson phase is itself, to leading order, already a Wilson-loop phase in the continuum theory; the open task is to compute the leading residual lattice-spacing correction to that Wilson loop in Quantumograph's specific action, evaluated on the actual geometry of a real junction, and compare its scaling to the phenomenological $\Delta\theta_{\text{local}}$ constructions above. If they agree, the identification would no longer be assumed but derived. Constructing that correspondence - rather than assuming it, as the original Definition 2.1 implicitly did - is, in the author's assessment, the harder and more consequential of the four problems catalogued in this section, and is left as the central open question.

A near-term empirical handle on (iii). Independently of resolving (iv), a coherent version of the local floor already has a candidate empirical test using existing, published data, requiring no new experiment. The AC Josephson relation $V = \frac{h}{2e} f$ - the phase-evolution relation underlying every Josephson voltage standard - has been checked across independent national metrology institutes, junction materials, and array geometries to a relative precision of order 10^{-10} , with no deviation reported (Niemeyer et al., *Metrologia*, 1986; NIST–BIPM comparisons, 2001 and 2009; INMETRO–NIST comparison, 2009). Expressed as a fraction of a full 2π cycle, the Construction-A floor for a real junction ($r \approx 100$ nm, $\Delta\theta_{\text{local}}^A \approx 6 \times 10^{-14}$ rad) corresponds to a relative deviation of order 10^{-14} – 10^{-15} - only ~ 5 orders of magnitude below the precision already demonstrated, a far smaller safety margin than the ~ 40 -order margin found in §3.1 for Shor’s algorithm. If the floor is coherent (a fixed, same-sign bias per cycle, not averaged away by repetition), a few further orders of improvement in Josephson-array metrology - not a qualitatively new experiment - could in principle probe it directly. If instead the floor is stochastic, it averages over the $\sim 10^{13}$ – 10^{15} oscillation cycles accumulated during a typical voltage-standard measurement, and is suppressed far below reach, for the same statistical reason given in the Appendix (§3, coherent-noise-invariance argument). This manuscript does not determine which case applies; it identifies the alternative precisely enough, and ties it to data that already exists, to make resolving it a well-posed next step rather than an open-ended one.

For comparison, an earlier and cruder attempt at this same problem - translating the Karolyhazy length-uncertainty formula $\delta l \approx (\ell_p^2 \cdot r)^{1/3}$ (Karolyhazy, F., 1966, *Nuovo Cimento A* 42, 390; see also Ng, Y. J. & van Dam, H., 1999, gr-qc/9906003) into a phase via an assumed control wavelength - gives $\Delta\theta_{\text{local}}(1 \text{ cm}) \approx 10^{-22}$ rad, several orders away from either construction above, and with yet another scaling exponent, because it requires one further physical assumption (the choice of wavelength) that neither construction above requires. This three-way disagreement between independently reasonable local constructions confirms that the coarse-graining problem is genuinely underdetermined by the model as currently stated, not merely unfinished in its algebra. All three estimates agree, however, on the qualitative conclusion of §3.1: any physically reasonable local floor remains many orders of magnitude too fine to matter for Shor’s algorithm.

This section therefore leaves the numerical value and functional form of $\Delta\theta_{\text{min}}$ ’s local counterpart provisional, but adds, beyond the earlier sketch: (a) two principled, non-ad-hoc methods for rescaling the global quantity to a local one (Bousso’s bound, and a statistical central-limit alternative), in place of a single unmotivated construction; (b) an explicit, falsifiable scaling disagreement between them ($r_{\text{dev}}^{-1/2}$ vs. r_{dev}^{-1}) that a future cross-device or cross-scale comparison could test; (c) an explicit separation of the length-scale ambiguity (ii) from the boundary-shape ambiguity (i), since the former is quantitatively much larger; (d) an explicit coherent-vs-stochastic distinction (iii) with a concrete, already-existing metrological dataset that could probe the coherent case; and (e) an explicit statement of the ontological-mapping problem (iv) together with the specific lattice-matching calculation that would resolve it, rather than a generic call for “an explicit correspondence.”

Prediction 2: Critical Temperature T_c (local, hardware-dependent, testable now)

T_c is not derived from $\Delta\theta_{\text{min}}$. It is the mean-field critical temperature of the graph’s own microscopic noise Hamiltonian,

$$\hat{H}_{\text{noise}} = \sum_{i,j} J_{ij} \sigma_i^z \sigma_j^z + \sum_i h_i \sigma_i^x$$

with $T_c = \frac{\bar{J}\bar{z}}{k_B}$, where $\bar{J}$ is the mean measured qubit–qubit coupling and $\bar{z}$ the average coordination number of the actual device topology. Unlike $\Delta\theta_{\text{min}}$, T_c is local: it is a property of a specific QPU’s own graph and couplings, and different devices predict different T_c . It marks the point below which long-range correlations of $\hat{H}_{\text{noise}}$ persist and above which they vanish - a mean-field order/disorder transition, not a statement about phase quantisation.

For a representative coupling $\bar{J} \sim 20$ MHz ($2\pi f$, with f in Hz - using $E = hf$, not $h \cdot 2\pi f$) and $\bar{z} \sim 5$, $T_c \approx 4.8$ mK; for denser topologies ($\bar{z} \sim 15$), T_c scales up to ≈ 14 mK. These are testable predictions: T_c is extracted from independently measured $\bar{J}$ and $\bar{z}$ on a given processor’s own coupling graph, with no free

parameters left to tune after that measurement. A full experimental protocol - coupling spectroscopy, temperature sweep, correlation and error-rate crossover detection, and an explicit falsification criterion - is given in a companion supplement.

The two predictions should not be conflated: a null result for T_c (no crossover at the predicted temperature) would falsify the existence or form of $\hat{H}_{\text{noise}}$, but would say nothing about $\Delta\theta_{\text{min}}$ phase quantisation, which remains untestable on any timescale discussed in this work for a stochastic floor, and only marginally so - per the new argument above - for a coherent one (§4.3).

2.3 Hierarchy of Quantum Computing Limitations

Type of Limitation	Source	Surmountable?	Notes
Engineering	Hardware imperfection, noise	Yes: better technology	—
Thermodynamic	Decoherence at $T > T_c$	Yes: cooling below T_c	Testable now, §2.2/Part IV
Complexity-theoretic	Algorithmic complexity	Partially: quantum speedup	—
Fundamental (FLQC)	Discreteness of spacetime	No: in principle - if it exists	Conditional on §2.2; per §3.1, not binding for any known algorithm

Note: the “Fundamental (FLQC)” row is conditional on the coarse-graining problem of §2.2 being resolved - items (i)–(iv). Even conditional on that, §3.1 shows it would not be surmountable by better engineering for the reason stated - but also that it would not, at the scale estimated here, be a binding constraint on any known quantum algorithm.

Part III. Consequences

3.1 Shor’s Algorithm and Cryptography - Corrected

An earlier version of this manuscript argued that FLQC renders Shor’s algorithm physically unrealisable for RSA-2048, with an obstruction threshold at $n^* \approx 100$ bits. That argument assumed that the controlled rotations $R_k = \text{diag}(1, e^{2\pi i/2^k})$ in the phase-estimation register must each be physically realised as a single native gate reaching the full ideal precision $\frac{2\pi}{2^k}$. This assumption fails for two independent, well-established reasons.

(i) Gate synthesis. No physical qubit implements arbitrary continuous rotations directly. Fault-tolerant architectures compile every rotation into a sequence of gates drawn from a fixed, discrete universal set; for example Clifford+T, where $T = \text{diag}(1, e^{i\pi/4})$. By the Solovay–Kitaev theorem (Solovay, 1995, unpublished; Kitaev, Russ. Math. Surv. 52, 1191, 1997) and its refinements (Ross, N. J. & Selinger, P., Quantum Inf. Comput. 16, 901, 2016), any target rotation is approximated to precision ϵ using $O(\log(\frac{1}{\epsilon}))$ to $O(\log^c(\frac{1}{\epsilon}))$ fixed, ordinary-precision gates - never one arbitrarily fine gate.

(ii) The algorithm does not need that precision. Coppersmith’s approximate quantum Fourier transform (IBM Research Report RC19642, 1994; arXiv:quant-ph/0201067; see also Griffiths, R. B. & Niu, C.-S., Phys. Rev. Lett. 76, 3228, 1996) shows that controlled rotations R_k for k beyond $O(\log n)$ can be dropped from the circuit entirely, at the cost of only exponentially small fidelity loss. Fowler, A. G. & Hollenberg, L. C. L. (Phys. Rev. A 70, 032329, 2004; arXiv:quant-ph/0306018) quantify this directly for Shor’s algorithm: factoring integers *thousands of bits long* — matching or exceeding the RSA-2048 register size discussed here — requires no controlled rotation finer than $\pi/64 \approx 4.9 \times 10^{-2}$ rad. This is roughly **28 orders of magnitude coarser** than $\Delta\theta_{\text{min}} \approx 6.28 \times 10^{-30}$ rad, and, per the extended table in §2.2.1, still 9–14 orders of magnitude coarser than either local construction’s estimate of $\Delta\theta_{\text{local}}$.

Because the precision actually required of any physically realised gate in Shor's algorithm is bounded below by $\sim 10^{-2}$ rad - independent of key length, and already coarser than current best engineering gate-calibration precision ($\sim 10^{-3}$ – 10^{-4} rad) - FLQC cannot constrain the feasibility of Shor's algorithm at RSA-2048 or any larger key length, **regardless of whether the coarse-graining problem of §2.2 is ever resolved, regardless of which of the two constructions of §2.2.1 is correct, and regardless of the ultimate numerical value.**

This can also be checked directly against the accumulated-error model of Part V: with $\Delta\theta_{\min} \approx 6.28 \times 10^{-30}$ rad and any physically realistic gate count (n up to $\sim 10^9$, including fault-tolerant overhead), the resulting fidelity loss is negligible to more than forty orders of magnitude - see Part V for the explicit calculation, which corrects the conclusion of earlier versions of this manuscript.

The FLQC hypothesis, if eventually supported by a completed coarse-graining derivation, should therefore be read purely as a statement about a residual gate-calibration floor several to many orders below current engineering noise (§3.2) - not as a statement with consequences for computational complexity, cryptography, or the security of RSA-class cryptosystems at any key length in current or foreseeable use. RSA-2048's security, in particular, should be assessed exactly as it would be in the absence of this hypothesis.

3.2 Quantum Error Correction

If a coarse-graining map (§2.2) eventually establishes a genuine local phase-resolution floor, standard error-correction codes would in principle be unable to reduce coherent phase error below that floor, however far engineering noise is reduced. Given the estimates of §2.2 and §3.1, this floor - if it exists - lies many orders of magnitude below current engineering noise ($\sim 10^{-3}$ – 10^{-4} rad) and is **not** expected to be the binding constraint on any near-term or foreseeable error-correction scheme. Its interest, at present, is conceptual - a genuinely irreducible error channel of geometric or statistical origin - rather than practical.

3.3 Quantum Simulation as a Measurement Tool

A QPU operating below T_c is itself a physical realisation of a $\mathbb{Z}^4$ -graph fragment. Its systematic errors are, on this hypothesis, a direct probe of graph structure. What hinders computation, on this view, helps physics - independent of whether §3.1's conclusion about Shor's algorithm holds.

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, generalises 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 instead concerns the minimum physically realisable phase increment, and therefore the precision of physical transformations rather than the cardinality of the state space. The logical chain proposed is: finite spacetime resolution $\Rightarrow$ finite phase resolution $\Rightarrow$ finite gate precision. As §3.1 shows, the last link in this chain - from finite gate precision to any consequence for a specific algorithm's feasibility - does not follow without the further, quantitative step of checking what precision that algorithm actually requires; for Shor's algorithm, that check shows no consequence follows at any realistic scale.

Among 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 references RSA-2048 as a benchmark. RaQM discretises the Hilbert-space basis itself: qubit amplitudes are required to satisfy a rationality condition derived from Diósi–Penrose gravitational decoherence, with the constrained quantity being the number of maximally entangled qubits N , bounded by a qubit information capacity $N_{\text{max}} \approx 200\text{--}1000$ depending on qubit mass/energy. FLQC instead concerns the angular resolution of individual phase operations. Whether either mechanism has any consequence for a specific algorithm is a separate, quantitative question in each case, to be checked explicitly as in §3.1 - not assumed from the existence of a bound alone.

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 signature)? Constancy is the FLQC signature. If a constant floor is found, compare its scaling with device or junction size against the two competing exponents of §2.2.1 ($r_{\text{dev}}^{-1/2}$ for Construction A, r_{dev}^{-1} for Construction B) to discriminate between them.
3. Implement random walks on the graph directly on the QPU. Measure the spectral dimension $ds(\sigma)$: passage through $ds = 4$ at some σ^* is evidence of $\mathbb{Z}^4$ structure.
4. Independently of 1–3, re-examine existing high-precision Josephson-voltage-standard comparison data (or commission a dedicated new comparison at the highest attainable precision) for any *coherent, junction-geometry-dependent* deviation from $V = \frac{h}{2e} f$ at the $10^{-14}\text{--}10^{-15}$ relative level predicted in §2.2.1. A null result at improved precision would constrain the coherent-floor hypothesis of item (iii) directly, using existing metrological infrastructure, without requiring the QPU-based tests above.

Compare any observed phase-error floor with the local estimates of §2.2.1 ($\approx 10^{-14}\text{--}10^{-17}$ rad depending on scale and construction) rather than with the whole-universe $\Delta\theta_{\text{min}} \approx 10^{-30}$ rad, and note that this numerical value is provisional pending resolution of items (i)–(iv).

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, if observed.
- A candidate new fundamental technological limit, conditional on §2.2.
- A worked example of how to correctly assess whether a Planck-scale bound has algorithmic consequences (§3.1) - of independent methodological interest given how easily this kind of comparison can be made incorrectly, as this manuscript’s own earlier version illustrates.
- A new quantum metrology tool, via §3.3.
- A demonstration that an existing, high-precision metrological dataset (Josephson voltage standards) can be repurposed as a Planck-scale probe, without new hardware.

4.3 Present Feasibility of Direct Laboratory Verification

A direct measurement of the whole-universe $\Delta\theta_{\text{min}} \approx 10^{-30}$ 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 idealised 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 is far beyond the age of the universe. Independently, ordinary decoherence-induced phase noise at currently achievable coherence times already exceeds this whole-universe floor by a comparable margin, so any such signal would remain masked by conventional technical noise even before statistical averaging could recover it.

The *local* estimates of §2.2.1 (10^{-14} – 10^{-17} rad, depending on scale and construction) sit in a materially different position, as the metrological comparison introduced there shows: while still far below what a QPU-based direct phase measurement can resolve, a coherent version of the junction-scale estimate is only ~ 5 orders of magnitude below already-demonstrated Josephson-metrology precision - a gap closable, in principle, by continued improvement of an existing measurement technique rather than by a qualitatively new one. A stochastic version of either local estimate remains as inaccessible as the whole-universe case, for the same statistical reasons.

This gap is a distinct comparison from the one in §3.1: here the concern is whether $\Delta\theta_{\text{FLQC}}$ can be measured *at all*; in §3.1 the concern is whether it would affect Shor's algorithm *if it exists* - these are independent questions with independent (negative for the whole-universe case and for any stochastic local floor; cautiously more promising for a coherent local floor; negative for Shor's algorithm regardless) answers. The experimental programme in §4.1 should 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 for the whole-universe quantity specifically. The implementation attempt "Theorem on the absence of asymptotic chaos" is described in the Appendix.

Part V. Phase Discretisation and the Accumulated-Error Model

If FLQC's coarse-graining problem (§2.2) is eventually resolved, physical implementations of phase-rotation gates would round the ideal target angle to the nearest point on a phase lattice:

$$R_k^{\text{FLQC}} = \text{diag}(1, \exp(i \cdot [\theta_{\text{ideal}} / \Delta\theta_{\text{FLQC}}] \cdot \Delta\theta_{\text{FLQC}}))$$

As established in §3.1, the relevant physical question is **not** whether a single gate can reach θ_{ideal} for large k - no implementation of Shor's algorithm requires this - but how a per-gate rounding error of order $\frac{\Delta\theta_{\text{FLQC}}}{2}$, accumulated over the actual number of gates n used in a compiled circuit, affects overall fidelity.

The accumulated effect can be estimated through the fidelity between the ideal and physical states,

$$F = |\langle \psi_{\text{ideal}} | \psi_{\text{act}} \rangle|^2 \approx \prod_j \cos^2(\delta\theta_j)$$

For a systematic phase floor,

$$\sigma^2 \approx n \cdot (\Delta\theta_{\text{FLQC}}/2)^2, P_{\text{success}} \propto e^{-\sigma^2}$$

Worked example (corrected). For $n \sim 4096$ (the naive gate count) up to $n \sim 10^9$ (a generous fault-tolerant overhead), and $\Delta\theta_{\text{FLQC}} \sim 10^{-30}$ rad (the whole-universe value; the local values of §2.2.1 make the effect even smaller), σ^2 ranges from $\sim 10^{-57}$ to $\sim 10^{-42}$. In every case $P_{\text{success}} \approx 1$: the effect is entirely negligible for any circuit currently conceivable. A non-negligible degradation ($\sigma^2 \sim 1$) would require $n \sim \left(\frac{2}{\Delta\theta_{\text{FLQC}}}\right)^2 \sim 10^{59}$ gates - far beyond any circuit that could ever physically be built. This is the corrected conclusion of the accumulated-error model; it does not support the claim, made in earlier versions of this manuscript, of a Shor-algorithm-blocking floor at RSA-2048 or at any bit-length.

This does not remove interest in the underlying distinction between two classes of error. Engineering errors arise from noise sources, control imperfections, thermal fluctuations, and material defects, and can in principle be reduced by improved design. An FLQC contribution, if it exists, would instead remain after these effects are minimised, because it would originate from the assumed discrete structure of the physical phase

space. Its magnitude, however, is - on the estimates above - too small to be the binding constraint on any algorithm's success probability, and its relevance to error correction is addressed in §3.2.

A direct experimental test follows the programme of Part IV: a QPU operated below T_c should be used to implement QFT circuits of increasing depth, with the measured systematic phase deviation separated into a component that scales with hardware noise and a component that remains constant. A persistent, n -independent component approaching some fundamental order of magnitude would constitute evidence for a phase-resolution limit of geometric or statistical origin - whatever consequences, or lack of them, that limit turns out to have for specific algorithms.

Note on Interpretation

The FLQC hypothesis, even if eventually supported by a completed coarse-graining derivation, does not imply that quantum computation is impossible, nor - as §3.1 shows - that any specific currently envisioned algorithm is blocked. It proposes that the ultimate computational capacity of a physical quantum system may, in principle, differ from that of an abstract mathematical quantum computer, and offers two candidate mechanisms and a testable local signature (T_c , Part IV) for investigating whether that is so. Whether the difference has any operational consequence is a question that must be checked explicitly, algorithm by algorithm and precision requirement by precision requirement, rather than assumed from the existence of a bound - the error made in earlier versions of this manuscript.

Conclusion

The Fundamental Limit of Quantum Computation is a candidate physical consequence of the conjectured finiteness and discreteness of spacetime within the Quantumograph framework. Its status remains that of a hypothesis with four open, logically related but distinct problems: the choice of boundary geometry (i), the choice of length scale (ii), the choice between at least two competing coarse-graining constructions with different scaling exponents and different coherent/stochastic character (iii), and - separately and most consequentially - the demonstration of an explicit ontological correspondence between graph-edge holonomy and a measured device phase (iv). Independently of all four, §3.1 shows that even a genuine local floor at any of the scales estimated here would have no bearing on the feasibility of Shor's algorithm or on the security of RSA-class cryptosystems. What survives as concrete, falsifiable, near-term predictions are the critical temperature T_c of §2.2, the experimental programme of Part IV, and the newly proposed metrological test of a coherent local floor against existing Josephson-voltage-standard precision.

This can be verified, at least in its T_c and Josephson-metrology forms. The equipment and, in the latter case, the data largely already exist. The experiments remain 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 the whole-universe $\Delta\theta_{FLQC} \sim 10^{-30}$ 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. This is a question about *measurability*, independent of §3.1's question about *algorithmic consequence*, and its

coherent/stochastic distinction (§2.2.1, item (iii)) is the same one invoked for the local floor and the Josephson-metrology test above.

1. The Optimistic Ceiling

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

To obtain an optimistic ceiling on the achievable chaos rate without importing any thermal or continuum-field-theoretic bound, we use the quantum-mechanical Margolus–Levitin speed limit in place of the Maldacena–Shenker–Stanford chaos bound: for a system with mean operating energy E above its ground state, the minimum time to reach an orthogonal state is $\Delta t_{\min} = \frac{\pi\hbar}{2E}$, giving a maximum admissible rate $\lambda_{ML} = \frac{2E}{\pi\hbar}$. This bound is purely quantum-mechanical - no chaos, no thermal ensemble, and no external temperature scale enter it. For a typical superconducting-qubit transition frequency $f \sim 5$ GHz ($E = hf \approx 3.3 \times 10^{-24}$ J), this gives $\lambda_{ML} \approx 2.0 \times 10^{10}$ rad/s.

If a system saturated this bound, the required evolution time would be $t = \frac{N_{\text{efold}}}{\lambda_{ML}} \approx 3.5$ 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 realisation of a rotation gate under FLQC rounds the target angle to the nearest point on the phase lattice: $\varphi_{\text{phys}} = \varphi + \delta(\varphi)$, with $\delta(\varphi) = \varphi - \text{round}(\varphi/\Delta\theta_{\min}) \cdot \Delta\theta_{\min}$. Because the rounding lattice is symmetric about zero, $\text{round}(-x) = -\text{round}(x)$, so $\delta(-\varphi) = -\delta(\varphi)$ (δ is odd).

Consider a forward gate with target angle θ_i and its reversal with target angle $-\theta_i$, both physically realised on the same hardware: forward $\theta_i + \delta(\theta_i)$; backward $-\theta_i + \delta(-\theta_i) = -\theta_i - \delta(\theta_i)$; sum = 0. The FLQC contribution cancels exactly, to leading order, in any protocol that implements time-reversal as angle negation - a direct consequence of δ being odd and of reversal being negation. Any OTOC/Loschmidt-echo measurement of this type is blind to $\Delta\theta_{FLQC}$ by construction for either a coherent or a stochastic floor.

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) 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 $H_{\text{real}} = H_{\text{ideal}} + \varepsilon_{\text{tech}} \cdot V_{\text{tech}} + \varepsilon_{FLQC} \cdot V_{FLQC}$, with $\varepsilon_{\text{tech}} \sim 10^{-3}$ – 10^{-4} (typical coherent two-qubit gate miscalibration on current best hardware) and $\varepsilon_{FLQC} \sim \Delta\theta_{\min} \sim 10^{-30}$ (for the whole-universe quantity; substituting either local estimate of §2.2.1 only strengthens the conclusion below).

In OTOC theory, the Lyapunov exponent λ characterises the chaotic background dynamics, not the perturbation itself: any small, generic *coherent* perturbation grows at the same rate, so

$\frac{d_{FLQC(D)}}{d_{\text{tech}(D)}} = \frac{\varepsilon_{FLQC}}{\varepsilon_{\text{tech}}} = \text{const}$, independent of circuit depth D . Chaotic amplification boosts signal and background by the same factor at every step, so their relative visibility never improves. Resolving a relative deviation of

$\sim 10^{-27}$ (whole-universe case) statistically requires $M \sim \left(\frac{\varepsilon_{\text{tech}}}{\varepsilon_{FLQC}} \right)^2 \sim 10^{54}$ measurements is 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 does not apply directly - but decoherence then bounds the usable circuit depth to $D \sim \frac{1}{\gamma}$ (γ = per-gate decoherence rate), reached long before $N_{efold} \approx 69$ can accumulate, reproducing the T2-budget obstruction of §4.3 by a different route.

Conclusion (Appendix)

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, so their ratio is a depth-independent invariant. Decoherence: the available circuit depth is capped well below what is 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 odd and reversal is negation - the two effects cancel exactly, independent of the two points above. This is unlikely to be overturned by future improvements in coherence time, qubit count, or control precision, since none of these change the invariance argument or the parity argument above. It is precisely because OTOC-based amplification is closed off that the direct, non-amplified metrological route proposed in §2.2.1 and §4.1 - for the coherent, junction-scale floor specifically - is worth pursuing instead.

References

- Shor, P. W. (1997). Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM Journal on Computing*, 26(5), 1484–1509.
- Grover, L. K. (1997). Quantum mechanics helps in searching for a needle in a haystack. *Physical Review Letters*, 78(2), 325–328.
- Solovay, R. (1995). Unpublished manuscript / conference remarks; see Kitaev (1997) below for the standard citation of the Solovay–Kitaev theorem.
- Kitaev, A. Yu. (1997). Quantum computations: algorithms and error correction. *Russian Mathematical Surveys*, 52(6), 1191–1249.
- Ross, N. J., & Selinger, P. (2016). Optimal ancilla-free Clifford+T approximation of z-rotations. *Quantum Information & Computation*, 16(11–12), 901–953.
- Coppersmith, D. (1994). An approximate Fourier transform useful in quantum factoring. IBM Research Report RC19642; arXiv:quant-ph/0201067 (2002).
- Griffiths, R. B., & Niu, C.-S. (1996). Semiclassical Fourier Transform for Quantum Computation. *Physical Review Letters*, 76(17), 3228–3231.
- Fowler, A. G., & Hollenberg, L. C. L. (2004). Scalability of Shor’s algorithm with a limited set of rotation gates. *Physical Review A*, 70, 032329; arXiv:quant-ph/0306018.
- Karolyhazy, F. (1966). Gravitation and quantum mechanics of macroscopic objects. *Il Nuovo Cimento A*, 42(2), 390–402.
- Ng, Y. J., & van Dam, H. (1999/2000). Measuring the foaminess of space-time with gravity-wave interferometers. *Foundations of Physics*; arXiv:gr-qc/9906003.
- Bousso, R. (1999). A Covariant Entropy Bound. *Journal of High Energy Physics*, 9907:004; arXiv:hep-th/9905177.
- Casini, H. (2008). Relative entropy and the Bekenstein bound. *Classical and Quantum Gravity*, 25, 205021; arXiv:0804.2182.

- 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.
- Niemeyer, J., Hinken, J. H., Vollmer, E., Grimm, L., & Meier, W. (1986). A new Josephson voltage standard with tunnel-junction series arrays. *Metrologia*.
- Solve, S., Chayramy, R., Stock, M., Tang, Y., & Sims, J. E. (2009). Comparison of the Josephson Voltage Standards of the NIST and the BIPM. *Metrologia / BIPM KCDB*.
- Tang, Y., Landim, R., Edson, A., & Ferreira, V. (2009). Comparison of the INMETRO and NIST Josephson Voltage Standards. *Metrologia*.
- Materov, S. (2025–2026). *Quantumograph. A Testable Quantum Graph Theory of Spacetime*.
- Materov, S. (2026). Theorem on the absence of asymptotic chaos. <https://doi.org/10.5281/zenodo.21865133>
- Materov, S. (2026). Computational Reducibility theorem. <https://doi.org/10.5281/zenodo.21864926>
- Materov, S. (2026). Simulation FLQC-Protocol. <https://github.com/SergejMaterov/FLQC-Protocol>