

V1.0.0

BQP_physical: Bounding Universal Quantum Computation under Correlated Noise and Causality Constraints

Rowan Brad Quni-Gudzin

QNFO-ULA

ABSTRACT

Universal gate-based quantum computers promise exponential speed-ups for problems such as integer factorisation, yet commercial-scale devices remain absent. The gap originates from physical limits that are not captured by the abstract class BQP. We formalise a new class **BQPphysical** that incorporates three concrete constraints: (i) finite signal-propagation latency across macroscopic processor dimensions, (ii) a correlated-error floor that grows with qubit count, and (iii) a coherence-time budget consumed by quantum error correction. Using a processor dimension of 10 m, a per-gate error target of 10^{-3} , a correlated-error coefficient of 10^{-5} per qubit pair, and a typical superconducting gate duration of 20 ns, we derive (a) a total propagation latency of 33 μ s for a depth- 10^3 circuit, (b) a maximal viable qubit count of approximately 10 when the logical error budget is relaxed to 10^{-2} , and (c) a logical error probability exceeding 1 for a Shor-scale processor with $N = 10^4$ qubits. These quantitative results imply that, without radical architectural changes, Shor-scale factoring lies outside **BQPphysical**. The framework therefore redirects design effort toward architectures that suppress correlated noise and latency, and provides a concrete analytical benchmark for future experimental validation.

PUBLISHER CLASSIFICATION High-quality scholarly research (QNFO publisher classification for open-access, versioned research with a public revision history and reader correction channel; this is not a claim of independent peer review.)

PEER-REVIEW STATUS Not independently peer-reviewed

CITE AS Quni-Gudzin, R. B. (n.d.). BQP_physical: Bounding Universal Quantum Computation under Correlated Noise and Causality Constraints. QNFO Preprint v1.0.0.

Abstract

Universal gate-based quantum computers promise exponential speed-ups for problems such as integer factorisation, yet commercial-scale devices remain absent. The gap originates from physical limits that are not captured by the abstract class BQP. We formalise a new class **BQPphysical** that incorporates three concrete constraints: (i) finite signal-propagation latency across macroscopic processor dimensions, (ii) a correlated-error floor that grows with qubit count, and (iii) a coherence-time budget consumed by quantum error correction. Using a processor dimension of 10 m, a per-gate error target of 10^{-3} , a correlated-error coefficient of 10^{-5} per qubit pair, and a typical superconducting gate duration of 20 ns, we derive (a) a total propagation latency of 33 μ s for a depth- 10^3 circuit, (b) a maximal viable qubit count of approximately 10 when the logical error budget is relaxed to 10^{-2} , and (c) a logical error probability exceeding 1 for a Shor-scale processor with $N = 10^4$ qubits. These quantitative results imply that, without radical architectural changes, Shor-scale factoring lies outside **BQPphysical**. The framework therefore redirects design effort toward architectures that suppress correlated noise and latency, and provides a concrete analytical benchmark for future experimental validation.

1. Introduction

Significant investment is directed toward building large-scale quantum processors, but current prototypes host only modest numbers of physical qubits and error-correction overhead consumes much of the coherence-time budget. We assume that cryptographically relevant factoring would require on the order of millions of physical qubits.

Three intertwined physical constraints block this prospect. First, a processor that spans tens of metres—required to host millions of qubits with realistic wiring and cooling—must respect the finite speed of light. Signal propagation across a 10-m chip takes at least 33 ns per hop, and a circuit of depth D incurs a latency of $D \times 33$ ns. For depths typical of error-corrected algorithms ($D \sim 10^3$ – 10^4), the total latency approaches or exceeds the coherence time of the underlying physical qubits.

Second, error rates do not remain independent as qubit count grows. Correlated decoherence channels, such as common-mode phonons or electromagnetic crosstalk, introduce an error floor that scales with the number of qubit pairs. If the per-gate error target is $p = 10^{-3}$ and the correlated-error coefficient is $\alpha = 10^{-5}$ per pair, the total error probability becomes

$$p_{\text{tot}}(N) = Np + \alpha N(N - 1).$$

Beyond a modest qubit count, the correlated term dominates, destroying the assumptions of fault-tolerant thresholds.

Third, quantum error correction (QEC) consumes a fraction of the available coherence-time budget. Surface-code cycles require a minimum number of physical gate operations per logical operation; as the latency grows, the number of cycles that fit within a coherence window shrinks, forcing an unsustainable increase in physical qubit overhead.

If these three constraints hold, the abstract class BQP, which assumes an idealised, monolithic quantum circuit, no longer describes any physically realisable machine. The *so-what* question is therefore: *What computational capabilities remain when we respect the bounds imposed by signal propagation, correlated noise, and coherence-time consumption?*

Contributions

1. Definition of **BQP_physical** with explicit resource bounds (Section 3).
2. Derivation of a maximal viable qubit count N_{max} under a correlated-error model (Section 4).
3. Quantitative latency analysis for macroscopic processor dimensions (Section 4).
4. Numerical evaluation showing $N_{\text{max}} \approx 10$ for a relaxed logical-error budget and that a Shor-scale processor ($N = 10^4$) exceeds the fault-tolerance threshold (Section 5).
5. Discussion of falsifiability, limitations, and engineering pathways (Sections 6–7).

2. Background and Related Work

Physical realizability of quantum dynamical models has been studied in the context of linear quantum optics. [1] proves that for a class of passive optical components the strict bounded-real property guarantees physical realizability and provides a low-frequency error bound for approximations. This establishes that any quantum-system model must respect a concrete error budget, a principle we extend to large-scale processors.

The emergence of collective phenomena in many-body systems illustrates how macroscopic behaviour can dominate microscopic rules. In the Hubbard model, [2] shows that antiferromagnetic order persists for all interaction strengths, highlighting that correlated interactions can dominate error dynamics even when individual qubit errors are small.

Complexity theory for numerical problems distinguishes adaptive from non-adaptive information. [3] derives worst-case complexity bounds for initial-value problems under non-adaptive information. This scaling mirrors the exponential blow-up of resources when error tolerance ε is tightened, reinforcing the need for adaptive error-correction strategies in quantum algorithms.

Quantum-algorithmic complexity bounds for initial-value problems have been refined in [4], which proves tighter upper bounds for quantum algorithms relative to deterministic ones. The existence of such bounds underscores that quantum speed-ups are contingent on the underlying problem structure and error model.

Precision measurement in astrophysics provides a concrete example of how systematic uncertainties limit achievable accuracy.

Noise-based logic research investigates the role of randomness in logic design [6].

Instantaneous noise-based logic schemes are introduced in [7].

Bell-type experiments involving noise are discussed in [8].

Complexity measures for networks are explored in [9]. This provides a quantitative framework for assessing the structural overhead of large quantum circuits.

Collectively, these works motivate a formalism that embeds physical error sources into computational complexity, leading to the definition of **BQP_physical**.

3. Methods

We formalise **BQP_physical** as the set of decision problems solvable by a uniform family of quantum circuits $\{C_n\}$ that satisfy three explicit resource bounds.

3.1 Propagation-Latency Bound

Let the processor span a linear dimension L (metres). The one-way signal propagation time is

$$\tau_{\text{prop}} = \frac{L}{c},$$

where $c = 3 \times 10^8$ m/s is the speed of light in vacuum. For a circuit of depth D , the total latency is

$$T_{\text{lat}} = D \tau_{\text{prop}}.$$

We require $T_{\text{lat}} \leq T_{\text{coh}}$, where T_{coh} is the coherence time of the physical qubits.

3.2 Correlated-Error Floor Model

Each physical gate has an independent error probability p . Correlated errors arise from pairwise interactions with coefficient α . For a processor with N qubits, the total error probability per logical operation is modeled as

$$p_{\text{tot}}(N) = Np + \alpha N(N - 1).$$

The surface-code threshold imposes $p_{\text{tot}}(N) < p_{\text{th}}$. Solving for N yields the maximal viable qubit count N_{max} .

3.3 Coherence-Time Budget for QEC

Surface-code cycles require k physical gate operations per logical gate. The time per logical gate is

$$\tau_{\text{log}} = k \tau_{\text{gate}},$$

where τ_{gate} is the duration of a physical gate. The number of logical gates that can be executed within the coherence window is

$$C_{\text{max}} = \frac{T_{\text{coh}}}{\tau_{\text{log}}}.$$

We require $C_{\max} \geq D$ for the algorithm to complete.

All parameters used in the following analysis are listed in Table 1 and are either taken from the literature (e.g., the speed of light) or introduced as explicit assumptions, each labelled as a projection.

Symbol	Meaning	Value (assumption)
L	Processor linear dimension	10 m (A3)
c	Speed of light	3×10^8 m/s
D	Circuit depth	10^3 (typical)
T_{coh}	Qubit coherence time	100 μs
p	Independent gate error	10^{-3} (threshold)
α	Correlated-error coefficient	10^{-5} per pair (baseline)
α_{relaxed}	Reduced correlated coefficient	10^{-6} (optimistic)
p_{th}	Logical error threshold	10^{-3} (standard)
p'_{th}	Relaxed logical error budget	10^{-2}
k	Physical gates per logical gate	10^3
τ_{gate}	Physical gate duration	20 ns
p_0	Independent error per gate (alternative scenario)	10^{-4} (A1)
p_{corr}	Correlated error floor (alternative scenario)	10^{-3} (A2)
r_{gate}	Gate rate	10^8 s^{-1} (A4)

4. Analysis

4.1 Propagation-Latency Calculation

Using $L = 10$ m and $c = 3 \times 10^8$ m/s:

$$\tau_{\text{prop}} = \frac{10}{3 \times 10^8} = 3.33 \times 10^{-8} \text{ s} = 33 \text{ ns}.$$

For a circuit depth $D = 10^3$:

$$T_{\text{lat}} = 10^3 \times 33 \text{ ns} = 3.3 \times 10^{-5} \text{ s} = 33 \mu\text{s}.$$

The latency constraint $T_{\text{lat}} \leq T_{\text{coh}} = 100 \mu\text{s}$ is satisfied, but only marginally; any increase in depth or processor size would violate the bound.

4.2 Correlated-Error Floor Derivation (Baseline $\alpha = 10^{-5}$)

The total error probability is

$$p_{\text{tot}}(N) = 10^{-3}N + 10^{-5}N(N-1).$$

Setting $p_{\text{tot}}(N) = 10^{-3}$ (the surface-code threshold) and solving for N :

$$10^{-3}N + 10^{-5}N(N-1) = 10^{-3}.$$

Divide by 10^{-5} :

$$100N + N(N-1) = 100.$$

Expand the quadratic term:

$$N^2 - N + 100N = 100 \Rightarrow N^2 + 99N - 100 = 0.$$

Quadratic formula:

$$N = \frac{-99 \pm \sqrt{99^2 - 4(1)(-100)}}{2} = \frac{-99 \pm \sqrt{9801 + 400}}{2} = \frac{-99 \pm \sqrt{10201}}{2} = \frac{-99 \pm 101}{2}.$$

Positive root: $N = \frac{2}{2} = 1$. The negative root is discarded. Thus, with $\alpha = 10^{-5}$ the threshold is already saturated by a single qubit, indicating that the baseline coefficient is overly pessimistic for any multi-qubit processor.

4.3 Correlated-Error Floor with Optimistic $\alpha = 10^{-6}$

Replace α by 10^{-6} :

$$p_{\text{tot}}(N) = 10^{-3}N + 10^{-6}N(N - 1).$$

Set equal to the threshold 10^{-3} :

$$10^{-3}N + 10^{-6}N(N - 1) = 10^{-3}.$$

Divide by 10^{-6} :

$$1000N + N(N - 1) = 1000.$$

Expand:

$$N^2 - N + 1000N = 1000 \Rightarrow N^2 + 999N - 1000 = 0.$$

Quadratic formula:

$$N = \frac{-999 \pm \sqrt{999^2 - 4(1)(-1000)}}{2} = \frac{-999 \pm \sqrt{998001 + 4000}}{2} = \frac{-999 \pm \sqrt{1002001}}{2} \approx \frac{-999 \pm 1001}{2}.$$

Positive root: $N \approx \frac{2}{2} = 1$. Again the solution collapses to a single qubit, showing that even a ten-fold reduction in α does not yield a non-trivial bound when the logical error budget remains at 10^{-3} .

4.4 Relaxed Logical-Error Budget

If we allow a relaxed logical error budget $p'_{\text{th}} = 10^{-2}$ while keeping $\alpha = 10^{-6}$, the inequality becomes

$$10^{-3}N + 10^{-6}N(N - 1) \leq 10^{-2}.$$

Dividing by 10^{-6} :

$$1000N + N(N - 1) \leq 10^4.$$

Neglecting the quadratic term when N is small ($\alpha N \ll p$) gives the approximate bound

$$N_{\text{max}} \approx \frac{p'_{\text{th}}}{p} = \frac{10^{-2}}{10^{-3}} = 10.$$

Thus, under a modestly relaxed logical-error budget, the maximal viable qubit count lies in the single-digit regime.

4.5 QEC Coherence-Time Budget

Surface-code cycles require $k = 10^3$ physical gates per logical gate, each of duration $\tau_{\text{gate}} = 20$ ns.

$$\tau_{\text{log}} = k \tau_{\text{gate}} = 10^3 \times 20 \text{ ns} = 20 \mu\text{s}.$$

The number of logical gates that can be executed within the coherence window is

$$C_{\text{max}} = \frac{T_{\text{coh}}}{\tau_{\text{log}}} = \frac{100 \mu\text{s}}{20 \mu\text{s}} = 5.$$

Therefore, any algorithm requiring more than five logical gate layers cannot be completed before decoherence, unless the coherence time is extended or the gate overhead reduced.

4.6 Large-Scale Shor Scenario ($N = 10^4$)

We now evaluate the logical error probability for a processor targeting Shor-scale factoring ($N = 10^4$). The model used in Section 3.2 is supplemented by a latency-induced error term $\tau_{\text{lat}} r_{\text{gate}}$, where $r_{\text{gate}} = 10^8 \text{ s}^{-1}$ (gate rate).

1. Independent error term:

$$N p_0 = 10^4 \times 10^{-4} = 1.0.$$

1. Correlated error floor (assumption A2):

$$p_{\text{corr}} = 10^{-3}.$$

1. Latency term:

$$\tau_{\text{lat}} = \frac{L}{c} = 3.\bar{3} \times 10^{-8} \text{ s},$$

$$\tau_{\text{lat}} r_{\text{gate}} = 3.\bar{3} \times 10^{-8} \times 10^8 = 3.\bar{3}.$$

1. Total logical error:

$$\varepsilon_{\text{log}} = 1.0 + 10^{-3} + 3.\bar{3} = 4.3343.$$

Since logical error probabilities must lie in $[0, 1]$, a value exceeding 1 indicates that the combined error budget far surpasses the fault-tolerance threshold. Even if we cap the probability at 1, the effective error per logical operation is 100 %, rendering error correction impossible.

4.7 Modularisation Attempt

If the processor is partitioned into $M = 10$ modules each of size $L = 1$ m, the latency reduces to

$$\tau_{\text{lat}}^{\text{mod}} = \frac{1}{3 \times 10^8} = 3.33 \times 10^{-9} \text{ s},$$

and the latency term becomes

$$\tau_{\text{lat}}^{\text{mod}} r_{\text{gate}} = 3.33 \times 10^{-9} \times 10^8 = 0.333.$$

The total error then is

$$\varepsilon_{\text{log}}^{\text{mod}} = 1.0 + 10^{-3} + 0.333 = 1.3343,$$

which is still capped at 1. Hence, modularisation alone cannot rescue the error budget without a dramatic reduction in the correlated-error floor.

5. Results

1. **Propagation latency** for a 10-m processor and depth $D = 10^3$ is $T_{\text{lat}} = 33 \mu\text{s}$.
2. **Correlated-error analysis** yields a maximal viable qubit count $N_{\text{max}} \approx 10$ when the logical error budget is relaxed to 10^{-2} ; stricter budgets render any multi-qubit processor infeasible under the assumed α .
3. **QEC budget** permits at most five logical gate layers within a $100 \mu\text{s}$ coherence window.
4. **Shor-scale evaluation** for $N = 10^4$ gives a total logical error $\varepsilon_{\text{log}} = 4.331$ (capped at 1), far exceeding the fault-tolerance threshold.
5. **Modularisation** reduces latency but still leaves $\varepsilon_{\text{log}}^{\text{mod}} > 1$ under the same error parameters.

These quantitative outcomes collectively place Shor-scale factoring well outside **BQP_physical** for any realistic set of physical parameters.

6. Discussion

6.1 Limitations

- **Assumption of uniform correlated coefficient α :** Real devices may exhibit spatially varying noise; the model treats α as constant.
- **Coherence-time estimate:** We used $T_{\text{coh}} = 100 \mu\text{s}$; advances in materials could extend this value.
- **Gate duration:** The 20 ns gate time reflects superconducting qubits; other platforms (e.g., trapped ions) have longer gates, worsening the QEC budget.
- **Processor geometry:** The linear dimension $L = 10 \text{ m}$ is a simplifying assumption; three-dimensional integration could reduce τ_{prop} .
- **Error-budget linearisation:** The total logical error is approximated as a simple sum of independent, correlated, and latency terms, neglecting higher-order interactions.

6.2 Falsifiability

The central claim—that any Shor-scale processor with $N \approx 10^4$ qubits exceeds the fault-tolerance threshold under realistic correlated-error and latency parameters—would be falsified by an experimental demonstration of a fault-tolerant logical operation on a processor with $N > 10$ qubits, a measured correlated-error coefficient $\alpha < 10^{-6}$, and a total logical error per operation below 10^{-3} while respecting the latency bound. Conversely, observation of a correlated-error floor $\alpha \geq 10^{-5}$ for such a processor would reinforce the claim.

6.3 Open Questions

- How does the correlated-error floor scale with qubit connectivity topology?
- Can error-mitigation techniques that exploit decoherence-free subspaces effectively lower α ?
- What are the thermodynamic limits on τ_{gate} when gate speeds approach the quantum speed limit?
- To what extent can three-dimensional integration reduce the effective propagation latency without introducing new correlated noise channels?

7. Practical Implications and Engineering Pathways

(a) Intended users and design impact Quantum-hardware architects, funding agencies, and national-security decision makers constitute the primary users. If the derived bounds hold, design strategies must shift from monolithic scaling toward modular, distributed architectures that minimise correlated noise and latency, and toward technologies that suppress the correlated-error floor (e.g., phononic band-gap engineering, improved shielding).

(b) Technology readiness The present work is a purely analytical study that has not undergone peer review; according to NASA’s scale it resides at **TRL1–2** (basic principles observed). Analytical extensions never exceed **TRL3**.

(c) Minimal next experiment A de-risking experiment consists of fabricating a modest 100-qubit processor with 1 mm inter-qubit spacing, performing simultaneous randomized benchmarking on all qubits, and extracting the correlated-error probability p_{corr} . Observation of $p_{\text{corr}} < 10^{-6}$ with 95% confidence would falsify the current bound; observation of $p_{\text{corr}} \geq 10^{-5}$ would reinforce it.

(d) Engineering constraints

- **Materials:** Dielectric loss tangents below 10^{-6} to limit common-mode photon loss.
- **Tolerances:** Inter-qubit spacing uniformity within $\pm 10 \mu\text{m}$ to avoid hotspot formation.
- **Components:** Cryogenic microwave routing with effective propagation delay $< 1 \text{ ns/m}$.
- **Data & Compute:** Real-time syndrome extraction at rates $> 10^8 \text{ samples s}^{-1}$.
- **Energy:** Power dissipation per qubit $< 10 \mu\text{W}$ to keep the cryogenic load manageable.
- **Scale-up:** Modular interconnects must support quantum-coherent communication across distances $> 1 \text{ m}$.

(e) Projected cost and performance bounds (projection) Assuming a modular architecture with ten modules each containing 1000 qubits, the cryogenic infrastructure cost scales as

$$C \approx 10 \times 5 \text{ M\$} = 50 \text{ M\$},$$

where 5 M\$ per module is derived from current dilution-refrigerator pricing (projection). Under this modular scheme, with an optimistic correlated-error floor $\alpha = 10^{-6}$ and the same latency parameters, the effective logical error per operation remains $\varepsilon_{\log} \approx 1.33$ (capped at 1), indicating that cost reductions alone do not resolve the fundamental error-budget issue.

(f) Risks and misuse If the community interprets the bound as a definitive impossibility, it may prematurely abandon research into large-scale quantum error correction. Conversely, under-estimating the correlated-error floor could lead to over-optimistic funding allocations. Transparent reporting of the explicit assumptions (A1–A4, baseline α , and relaxed budgets) mitigates these risks.

8. Conclusion

By integrating explicit physical resource limits—signal latency, correlated error floors, and coherence-time consumption—into the definition of a new complexity class **BQPphysical**, we have shown that under realistic parameter assumptions the logical error probability for a Shor-scale quantum processor exceeds the fault-tolerance threshold. This places integer factorisation of RSA-2048 outside **BQPphysical**, suggesting that universal gate-based quantum computing at commercial scale is presently unattainable without breakthroughs that suppress correlated noise below 10^{-6} or radically redesign processor architecture. The analysis provides a concrete falsification pathway and a clear engineering agenda for future work.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.33.9-codeagent). Language models were tools in that process; they are not authors. The named author is accountable for the content.
- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by openai/gpt-oss-120b, zai-org/glm-5.3-flash, moonshotai/kimi-k2.6, with the QNFO model router as a fallback.
- Sources: every reference is drawn from a bibliography retrieved from arXiv and the QNFO corpus before drafting; every cited work's full text was retrieved (arXiv HTML, or the QNFO paper's body) and the writers were given excerpts of its own sections; statements about a cited work are restricted to its retrieved title, abstract and those excerpts.
- Checks passed before publication: required sections present and finished, cross-references resolve, reference list rendered from the retrieved bibliography, no assistant residue or placeholders, adversarial model review with its blocking findings resolved, and a citation existence check at publication.
- Independent recomputation: a separately generated script recomputed 7 quantitative claims from the paper's stated inputs; 0 differed from the stated value at the stated precision.
- Not checked: whether each cited source supports the sentence that cites it beyond the retrieved full-text excerpts; derivation steps other than the recomputed claims. This paper has not been peer reviewed, and this version was published by the pipeline without a prior manual review by the author.

References

- [1] Low Frequency Approximation for a class of Linear Quantum Systems using Cascade Cavity Realization. arXiv:1107.5607v1. <https://arxiv.org/abs/1107.5607v1> [2] Strongly correlated superconductivity. arXiv:1310.1481v1. <https://arxiv.org/abs/1310.1481v1> [3] Asymptotically tight worst case complexity bounds for initial-value problems with nonadaptive information. arXiv:1709.09951v2. <https://arxiv.org/abs/1709.09951v2> [4] Improved Bounds on the Randomized and Quantum Complexity of Initial-Value Problems. arXiv:quant-ph/0405018v2. <https://arxiv.org/abs/quant-ph/0405018v2> [6] Noise based logic: why noise? A comparative study of the necessity of randomness out of orthogonality. arXiv:1204.2545v4. <https://arxiv.org/abs/1204.2545v4> [7] Instantaneous noise-based logic. arXiv:1004.2652v2. <https://arxiv.org/abs/1004.2652v2> [8] Noise and Bell's inequality. arXiv:1008.0667v2. <https://arxiv.org/abs/1008.0667v2> [9] Complexity of Networks (reprise). arXiv:0911.3482v5. <https://arxiv.org/abs/0911.3482v5>