

V1.0.0

Deductive Verification of Discrete-Time Quantum-Walk Simulators: Contracts, Cost Bounds, and Periodicity Targets

Rowan Brad Quni-Gudzin

QNFO-ULA-2.1

ABSTRACT

Discrete-time quantum walks are a universal model of quantum computation, a platform for topological phases, and a standard testbed for quantum algorithms, yet the simulators used to study them are rarely checked with the rigor applied to classical library code. We frame quantum-walk simulator verification as a contract-checking problem in the style of deductive verification of library functions, where formal contracts are extracted as preconditions and postconditions and discharged as proof obligations. We define four contracts — unitarity, probability conservation, exact support bounds, and integer-valued topological invariants — and derive their computational costs with explicit arithmetic. For a Hadamard walk evolved $t = 10^3$ steps, the state vector has dimension $d_t = 4002$; a dense unitarity check costs $d_t^3 = 64,096,048,008$ multiply-adds, while a sparse unitarity check exploiting two nonzeros per column costs $4d_t^2 = 64,064,016$ multiply-adds, a speedup factor of $d_t/4 = 1000.5$ over the dense check. [...]

CITE AS Quni-Gudzin, R. B. (n.d.). Deductive Verification of Discrete-Time Quantum-Walk Simulators: Contracts, Cost Bounds, and Periodicity Targets. QNFO Preprint v1.0.0.

Abstract

Discrete-time quantum walks are a universal model of quantum computation, a platform for topological phases, and a standard testbed for quantum algorithms, yet the simulators used to study them are rarely checked with the rigor applied to classical library code. We frame quantum-walk simulator verification as a contract-checking problem in the style of deductive verification of library functions, where formal contracts are extracted as preconditions and postconditions and discharged as proof obligations. We define four contracts — unitarity, probability conservation, exact support bounds, and integer-valued topological invariants — and derive their computational costs with explicit arithmetic. For a Hadamard walk evolved $t = 10^3$ steps, the state vector has dimension $d_t = 4002$; a dense unitarity check costs $d_t^3 = 64,096,048,008$ multiply-adds, while a sparse unitarity check exploiting two nonzeros per column costs $4d_t^2 = 64,064,016$ multiply-adds, a speedup factor of $d_t/4 = 1000.5$ over the dense check. Naive direct simulation costs $M_t = 8t(t + 2) = 8,016,000$ real multiplications, so normalization checking adds 50% multiplication overhead per step, independent of t , as derived in Section 4 (Derivation 6). We derive a double-precision drift bound of 8.88×10^{-10} after 10^6 steps (labeled projection) and an illustrative periodicity contract on the 3-cycle with return probability $1/3$ under an explicitly stated uniform-amplitude assumption. No empirical measurements are reported.

1. Introduction

A discrete-time quantum walk evolves a single quantum particle on a graph by alternating a local unitary "coin" operation with a conditional shift. Despite this simplicity, quantum walks display rich topological phenomena in one and two dimensions and are among the simplest systems in which to study topological phases [1]; they are the quantum counterparts of Markov chains with emphasis on algorithmic applications [8]; and they have been shown to constitute a universal model of quantum computation [3].

Two facts motivate this paper. First, quantum-walk simulators are software, and software that claims to implement a unitary evolution can be checked against explicit mathematical contracts. Deductive verification of unmodified library functions has been benchmarked on conventional code: a benchmark of 26 unmodified Linux kernel library functions implementing memory and string operations had formal contracts extracted from source code as preconditions and

postconditions, and the correctness of 23 of the 26 functions was completed [7]. The same discipline — extract the contract from the mathematical definition of the object, then discharge the proof obligations — applies to a quantum-walk evolution operator, where the contract is unitarity and the obligation is $U^\dagger U = I$.

Second, the most interesting correctness targets for a quantum walk are not generic matrix properties but discrete ones: integer-valued topological invariants that must be exactly integers if the simulator is faithful, and periodicity claims that are sharp, decidable properties of the implemented update rule. Prior work on cyclic quantum walks starts from the classical result that combining two chaotic random walks can yield an ordered (periodic) walk and seeks a quantum analog, focusing on a periodic quantum walk on a 3-cycle graph generated via a deterministic combination of two chaotic ones [4]. Such a periodicity statement can be expressed as a postcondition equating the simulator state at step t with the state at step $t + T$ for the claimed period T .

This paper makes three contributions: (i) a formal statement of verification contracts for discrete-time quantum-walk simulators, including a periodicity contract illustrated on the 3-cycle; (ii) explicit cost derivations for the unitarity, simulation, and contract-size obligations, with all arithmetic shown; (iii) an error-bound analysis for floating-point drift and for numerically extracted topological invariants, with explicit falsification conditions.

2. Background and Related Work

We discuss the supplied bibliography in its given order. Throughout, "coin" means the internal unitary acting on the walker's internal degree of freedom, and "support" means the set of graph positions with nonzero amplitude.

Topological phases in quantum walks [1]. The review of topological phenomena in quantum walks presents discrete quantum walks as dynamical protocols for controlling a single quantum particle and argues that, despite their simplicity, they display rich topological phenomena and constitute one of the simplest systems in which to study and understand topological phases, with elementary treatment of one and two dimensions. For verification this matters because topological invariants are exactly the quantity a buggy simulator can corrupt silently; our contracts include the support and normalization invariants that any downstream invariant computation presupposes.

Discrete-event simulation [2]. This work simulates two models of experimentally realizable quantum walks on a digital computer using models that comply with Einstein locality, in which particles follow well-defined trajectories and avoid concepts such as particle-wave duality and wave-function collapse, reproducing the quantum-theoretical results (the supplied summary is truncated at this point, so we attribute no further specific findings to it). This is a distinct verification target: a discrete-event simulator does not maintain a state vector, so the unitarity contract does not directly apply; the contract becomes statistical reproduction of quantum statistics rather than an algebraic identity.

Comprehensive review [3]. The comprehensive review describes quantum walks as the quantum mechanical counterpart of classical random walks, an advanced tool for building quantum algorithms, and states that quantum walks have been shown to constitute a universal model of quantum computation, with the field full of open problems for physicists, computer scientists, mathematicians, and engineers (the summary is truncated, so further specifics of its coverage are not attributed here). Universality raises the stakes of verification: a verified quantum-walk simulator is a verified simulator for a universal family of computations.

Order from chaos on cyclic graphs [4]. This study starts from the classical result that combining two chaotic random walks can yield an ordered (periodic) walk and seeks a quantum analog; it studies chaotic and periodic behavior of cyclic quantum walks and focuses on a periodic quantum walk on a 3-cycle graph generated via a deterministic combination of two chaotic ones (the summary is truncated beyond this). Periodicity is a canonical verification target, and [4] supplies a concrete family of such claims; we use it in Section 4 as an illustrative contract instance.

Limit distributions [5]. This work analytically shows the weak limit theorem — the asymptotic behavior — of the one-dimensional discrete-time quantum walk, and discusses the relation between the limit distribution of the discrete-time walk and the continuous-time walk (the summary is truncated at this point). Limit theorems supply asymptotic reference

distributions against which finite- t simulations can be statistically cross-checked, complementing the exact algebraic contracts.

Symmetries and noise [6]. This study examines discrete symmetries of unbiased (Hadamard) and biased quantum walks on a line and shows that these symmetries hold even when the walker is subjected to environmental effects, with noise modeled by phase flip, bit flip, and generalized amplitude damping channels, and numerical solutions obtained by evolving the density matrix (the summary is truncated beyond this). This is the natural bridge from verified unitary simulators to verified open-system simulators: the density-matrix loop has the same contract structure with trace conservation replacing probability conservation, and the named noise channels are the first operations whose contract clauses one would write.

Deductive verification of systems code [7]. The benchmark consists of 26 unmodified Linux kernel library functions implementing conventional memory and string operations; formal contracts were extracted from the functions' source code as preconditions and postconditions, and the correctness of 23 functions was completed (the summary is truncated at this point, so we attribute no further details of tooling or effort accounting). We use this as the methodological template — contracts extracted from the code as written, applied to unmodified functions, with a partial but high completion rate — and compute the verified fraction 23/26 explicitly in Section 4.

Algorithmic applications [8]. This overview characterizes quantum walks as quantum counterparts of Markov chains and gives a brief overview with emphasis on their algorithmic applications (the summary is short, and we attribute nothing further to it). It supplies the consumer-side motivation: every algorithmic application downstream of a walk simulation inherits the simulator's correctness.

Finite-distinction quantum mechanics [9]. The QNFO program argues that a finite-entropy world is a finite-distinction world, that uncountable precision is unphysical while computable depth and p -adic valuation remain physically real, and that the state-space geometry of finite distinctions is combinatorial and ultrametric, with quantum mechanics read as thermodynamics (the supplied summary is truncated mid-sentence). We use only what the summary states: the finiteness thesis, the combinatorial-ultrametric geometry claim, and the thermodynamic reading. This motivates our emphasis on exactly checkable, finite, discrete contracts (integer invariants, exact support bounds) over asymptotic floating-point agreement.

Remaining entries. The entry on factoring, adelic complexity, and the silent-radix principle [10] has an empty supplied summary; we make no claim about its content and cite it only as part of the same corpus context as [9] and [11]. The Manifesto for Honest Computation [11] supplies, per its summary, five principles, a concrete portfolio, and a Falsification Pledge; we adopt its falsification stance methodologically (Section 6 states explicit falsification conditions) without attributing specific principles beyond what the summary names. The ZBW-Majorana hypothesis entry [12] establishes, per its summary and across four companion papers (P1–P4), that Zitterbewegung — the rapid trembling motion of Dirac fermions at the Compton scale — is a $\mathbb{Z}_2$ topological observable distinguishing Dirac from Majorana fermions at the hardware level (summary truncated). We cite it only as an analogy for the epistemic status of discrete-valued topological observables: like a $\mathbb{Z}_2$ invariant, a winding number is robust to small perturbations and vulnerable only to errors large enough to flip the integer.

3. Methods

3.1 Model

We consider the standard one-dimensional discrete-time Hadamard walk. The Hilbert space is $\mathcal{H}_p \otimes \mathcal{H}_c$, where $\mathcal{H}_p = \text{span}\{|x\rangle : x \in \mathbb{Z}\}$ and $\mathcal{H}_c = \text{span}\{|0\rangle, |1\rangle\}$ has dimension $d_c = 2$. The state at step t is

$$|\psi_t\rangle = \sum_{x \in \mathbb{Z}} \left(\alpha_{x,0}^{(t)} |x, 0\rangle + \alpha_{x,1}^{(t)} |x, 1\rangle \right),$$

with $\sum_x \left(|\alpha_{x,0}^{(t)}|^2 + |\alpha_{x,1}^{(t)}|^2 \right) = 1$. One step is

$$|\psi_{t+1}\rangle = \hat{U}|\psi_t\rangle, \quad \hat{U} = \hat{S}(\hat{I} \otimes \hat{C}),$$

where $\hat{C}$ is the Hadamard coin

$$\hat{C} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix},$$

and $\hat{S}|x, 0\rangle = |x - 1, 0\rangle$, $\hat{S}|x, 1\rangle = |x + 1, 1\rangle$. The walk starts at $x_0 = 0$ with $\alpha_{0,0}^{(0)} = 1$ and all other amplitudes zero. After t steps the support is contained in $\{-t, \dots, t\}$, of size $N_t = 2t + 1$, and the state-vector dimension is

$$d_t = d_c N_t = 2(2t + 1).$$

3.2 Contracts

In the extract-then-discharge style of [7], we attach to the step function $\text{step}(\psi_t) \rightarrow \psi_{t+1}$ the following contracts:

- **(K1) Unitarity:** $\hat{U}^\dagger \hat{U} = I_{d_t}$, i.e., $d_t(d_t + 1)$ real constraints.
- **(K2) Probability conservation:** $\sum_x (|\alpha_{x,0}^{(t)}|^2 + |\alpha_{x,1}^{(t)}|^2) = 1$ within tolerance ε , at every step.
- **(K3) Exact support bound:** $\text{supp}(\psi_t) \subseteq \{-t, \dots, t\}$, checkable with zero floating-point error.
- **(K4) Topological invariant:** the winding number $w \in \mathbb{Z}$ extracted from the walk's band structure equals its model value exactly.
- **(K5) Periodicity (illustrative, from [4]):** for a walk claimed periodic with period T , the state at step $t + T$ equals the state at step t .

Clauses K1 and K2 follow from unitarity of $\hat{U}$: since $\hat{C}^\dagger \hat{C} = \hat{I}$ and $\hat{S}^\dagger \hat{S} = \hat{I}$, we have $\hat{U}^\dagger \hat{U} = \hat{I}$, so $\langle \psi_{t+1} | \psi_{t+1} \rangle = \langle \psi_t | \psi_t \rangle$. Clause K3 holds because one step moves each amplitude by at most one position. Each clause is a decidable predicate over the simulator's data structures, checkable by an SMT-backed verifier of the kind used in the benchmark of [7].

3.3 Cost model

We count multiply-add operations, not wall-clock time; where we convert to time we state the assumed machine rate explicitly as a projection. A dense unitarity check computes all d_t row inner products of $\hat{U}$ (each of length d_t), costing d_t^3 multiply-adds. A sparse unitarity check exploits that each column of $\hat{U}$ has exactly $k = 2$ nonzero entries, so each of the d_t^2 entries of $\hat{U}^\dagger \hat{U}$ costs $k^2 = 4$ multiply-adds, giving total cost $4d_t^2$ (Derivation 3). A cheaper $8d_t$ check computes only the d_t row norms and verifies normalization, not orthogonality, so it does not discharge K1. For the periodicity illustration (K5) we use the 3-cycle of [4]: vertices v_0, v_1, v_2 with edges $(v_0, v_1), (v_1, v_2), (v_2, v_0)$, evolved by a deterministic combination of two constituent walks in the spirit of [4]; the specific interleaving schedule is an illustrative assumption of this paper, not a detail reported by [4].

3.4 Finite-distinction alignment

The finite-distinction program [9] holds that uncountable precision is unphysical and that finite-distinction state spaces are combinatorial and ultrametric. We operationalize this as follows: the verified artifact is a finite map from integer pairs (x, c) with $|x| \leq t$ and $c \in \{0, 1\}$ to amplitude records — exactly $N_t = 2(2t + 1)$ records at step t . No continuum object appears in the verified artifact. We stress that this is a methodological alignment, not a derivation of quantum mechanics from [9]; the summary of [9] supports the finiteness and geometry claims we invoke and nothing more.

4. Analysis

Every input number below is either a model definition (Section 3), a structural fact derived from it, or an explicitly labeled projection. All arithmetic is shown.

Derivation 1 (state-space dimension). Inputs: $d_c = 2$ (Section 3.1); support size $N_t = 2t + 1$ (each step moves the walker by exactly one position, so the support after t steps lies in $\{-t, \dots, t\}$, which has $2t + 1$ integer points). Hence

$$d_t = d_c(2t + 1) = 2(2t + 1).$$

For $t = 10^3$: $2t + 1 = 2001$, so $d_{1000} = 2 \times 2001 = 4002$. For $t = 10^6$: $2t + 1 = 2,000,001$, so $d_{10^6} = 4,000,002$.

Derivation 2 (dense unitarity-check cost). Input: $d_t = 4002$ (Derivation 1).

$$d_t^2 = 4002 \times 4002 = 16,016,004,$$

$$d_t^3 = 16,016,004 \times 4002 = 16,016,004 \times 4000 + 16,016,004 \times 2 = 64,064,016,000 + 32,032,008 = 64,096,048,008.$$

So the dense check costs $64,096,048,008 \approx 6.41 \times 10^{10}$ multiply-adds. Converting to time under the stated assumption of $r = 10^9$ multiply-adds per second (projection; the operation count itself is exact):

$$T_{\text{dense}} = \frac{64,096,048,008}{10^9} = 64.096048008 \text{ s} \approx 64.096 \text{ s}.$$

Derivation 3 (sparse unitarity-check cost). Verifying $\hat{U}^\dagger \hat{U} = I$ requires evaluating all d_t^2 entries of $\hat{U}^\dagger \hat{U}$, each an inner product of two columns of $\hat{U}$. Each column has $k = 2$ nonzero entries, so each entry costs $k^2 = 4$ multiply-adds, giving total cost

$$4d_t^2 = 4 \times 16,016,004 = 64,064,016 \text{ multiply-adds},$$

using $d_t^2 = 16,016,004$ from Derivation 2. The speedup factor over the dense check is

$$\frac{d_t^3}{4d_t^2} = \frac{d_t}{4} = \frac{4002}{4} = 1000.5.$$

A cheaper check computing only the d_t row norms costs $8d_t = 8 \times 4002 = 32,016$ multiply-adds, but it verifies only normalization of rows, not orthogonality, and therefore discharges K2-style normalization clauses rather than K1.

Derivation 4 (unitarity constraint count). The contract K1 imposes $d_t(d_t + 1)$ real constraints. For $d_t = 4002$:

$$d_t(d_t + 1) = 4002 \times 4003 = 4002 \times 4000 + 4002 \times 3 = 16,008,000 + 12,006 = 16,020,006.$$

Derivation 5 (simulation multiplication count). Input: the coin entries have magnitude $2^{-1/2}$ (Section 3.1). Applying $\hat{C}$ at one site computes two output amplitudes,

$$\beta_0 = 2^{-1/2}(\alpha_0 + \alpha_1), \quad \beta_1 = 2^{-1/2}(\alpha_0 - \alpha_1).$$

Each product of a complex amplitude by the real scalar $2^{-1/2}$ costs 2 real multiplications (real and imaginary parts). There are 2 output amplitudes $\times$ 2 input terms each = 4 such products, hence $4 \times 2 = 8$ real multiplications per site per step; the shift is data movement and costs none. At step s the occupied support has at most $2s + 1$ sites (worst case), so

$$M_t = \sum_{s=1}^t 8(2s + 1) = 16 \sum_{s=1}^t s + 8t = 16 \cdot \frac{t(t+1)}{2} + 8t = 8t^2 + 16t = 8t(t+2).$$

Check of the closed form: $\sum_{s=1}^t (2s + 1) = 2 \cdot \frac{t(t+1)}{2} + t = t^2 + 2t = t(t+2)$; multiplying by 8 gives $M_t = 8t(t+2)$. For $t = 1000$: $1000 \times 1002 = 1,002,000$, and $8 \times 1,002,000 = 8,016,000$. For $t = 100$: $8 \times 100 \times 102 = 800 \times 102 = 81,600$.

Derivation 6 (verification overhead ratio). Evaluating clause K2 requires computing $\sum_x (|\alpha_{x,0}|^2 + |\alpha_{x,1}|^2)$ over d_t records; each term costs 2 real multiplications, so one check costs $2d_t$ real multiplications. Two checks per step (input K2 and output K2) cost $4d_t$. The simulation itself costs $8d_t$ per step (Derivation 5, per-step rate). The overhead ratio is

$$\frac{4d_t}{8d_t} = \frac{1}{2},$$

i.e., normalization checking adds exactly 50% multiplication overhead per step, independent of t . At $t = 1000$: $4 \times 4002 = 16,008$ verification multiplications per step versus $8 \times 4002 = 32,016$ simulation multiplications per step.

Derivation 7 (contract clause count). The minimal per-step contract of Section 3.2 has, in the clause-counting convention of Derivation 6's framework, 2 preconditions (input normalization, input support) and 3 postconditions (output normalization, output support, coin correctness), i.e., 5 clauses per step-function instance. For a t -step run verified step by step, the aggregate count is

$$K_t = 5t; \quad K_{1000} = 5 \times 1000 = 5,000, \quad K_{100} = 500.$$

If instead the verifier discharges the step lemma once and applies it t times by induction, the aggregate obligation drops to one 5-clause lemma plus t applications; we treat $K_t = 5t$ as the upper bound for the naive scheme and the inductive scheme as a projected lower bound with constant lemma cost.

Derivation 8 (floating-point drift bound, projection). Assumptions: IEEE double precision with unit roundoff $\varepsilon_{\text{fp}} = 2^{-52} \approx 2.2204 \times 10^{-16}$; each rounding contributes at most ε_{fp} relative error; errors accumulate linearly (no favorable cancellation assumed). Input: the coin update at one site forms each output amplitude from 2 products (Derivation 5), each product costing 2 real multiplications, i.e., 4 roundings per output amplitude per step. The worst-case relative drift after t steps is therefore

$$\delta_t \leq 4t \varepsilon_{\text{fp}} = 4t \cdot 2^{-52}.$$

For $t = 10^6$:

$$\delta_{10^6} \leq 4 \times 10^6 \times 2.2204 \times 10^{-16} = 8.8816 \times 10^{-10} \approx 8.88 \times 10^{-10}.$$

This is a projection: it assumes the linear-accumulation model and no error growth beyond per-step roundoff; it is not an empirical measurement.

Derivation 9 (illustrative 3-cycle return probability). Assumption (this paper, illustrative): the periodic orbit of the 3-cycle contract K5 places equal-amplitude magnitude $3^{-1/2}$ on each of the three vertices at the return step. The probability of finding the walker at the start vertex v_0 is then

$$P_{\text{ret}} = \left| 3^{-1/2} \right|^2 = \frac{1}{3} \approx 0.3333.$$

This follows from the stated uniform-amplitude assumption only; bibliography entry [4] does not report this number.

5. Results

All numbers below are computed in Section 4 with shown arithmetic; none is an empirical measurement.

- State-vector dimension at $t = 10^3$: $d_{1000} = 4002$ (Derivation 1); at $t = 10^6$: $d_{10^6} = 4,000,002$.
- Dense unitarity-check cost at $t = 10^3$: $d_t^3 = 64,096,048,008$ multiply-adds $\approx 6.41 \times 10^{10}$; projected time ≈ 64.1 s at the stated rate $r = 10^9 \text{ s}^{-1}$ (Derivation 2).
- Sparse unitarity-check cost: $4d_t^2 = 64,064,016$ multiply-adds; speedup factor $d_t/4 = 1000.5$ (Derivation 3). The cheaper row-norm check $8d_t = 32,016$ verifies normalization only, not unitarity.
- Unitarity constraint count: $d_t(d_t + 1) = 16,020,006$ at $t = 10^3$ (Derivation 4).
- Simulation cost: $M_t = 8t(t + 2)$; $M_{1000} = 8,016,000$ and $M_{100} = 81,600$ real multiplications (Derivation 5).
- Verification overhead: verification $4d_t = 16,008$ versus simulation $8d_t = 32,016$ multiplications per step at $t = 10^3$, ratio $4d_t/8d_t = 1/2$, i.e., 50% multiplication overhead per step, independent of t (Derivation 6).

- Contract clause count: $K_t = 5t$; $K_{1000} = 5,000$, $K_{100} = 500$ for the naive scheme (Derivation 7).
- Projected double-precision drift bound: $\delta_{10^6} \approx 8.88 \times 10^{-10}$ (Derivation 8, projection).
- Illustrative 3-cycle return probability: $P_{\text{ret}} = 1/3$ under the stated uniform-amplitude assumption (Derivation 9).

6. Discussion

Limitations. The cost model counts multiply-adds, not wall-clock time; the 64.1 s figure is a projection at an assumed rate $r = 10^9 \text{ s}^{-1}$. The 50% overhead figure is a worst-case bound resting on the assumption that the support at step s occupies all $2s + 1$ sites; the true support of the Hadamard walk is smaller in aggregate, so the realized overhead on actual trajectories may be lower, though the per-step bound holds. The drift bound of Derivation 8 assumes linear accumulation of roundoff and could be violated by ill-conditioned amplitude cancellation or by error growth in long runs; it is a projection, not a measurement. The 3-cycle return probability $1/3$ is an illustrative consequence of a uniform-amplitude assumption made in this paper; entry [4]’s supplied summary does not state it, and the interleaving schedule of the constituent walks is likewise assumed here, not sourced from [4]. The winding-number contract K4 is stated but not instantiated against a specific band structure in this paper, so no numerical margin is claimed.

Failure modes and falsification. The contract framework is falsified if (i) a faithful simulator violates K1–K3 under exact arithmetic, which would contradict the algebraic derivation of Section 3.2; (ii) the measured drift of a concrete implementation exceeds the projected bound $4t \varepsilon_{\text{fp}}$ by a large factor, refuting the linear-accumulation model; or (iii) the clause-count convention of Derivation 7 misrepresents real verifier obligation counts. The methodological alignment with the finite-distinction program [9] is explicitly not a derivation of quantum mechanics and would be misleading if read as one.

Open questions. Whether the inductive verification scheme of Derivation 7 achieves the projected constant lemma cost in practice; how the statistical contracts appropriate to discrete-event simulators [2] compose with the algebraic contracts used here; and how trace-conservation contracts for open-system evolution [6] relate to the unitarity contract K1.

7. Conclusion

We framed discrete-time quantum-walk simulator verification as contract extraction and discharge, following the template of deductive verification of unmodified library functions [7]. We stated five contracts (K1–K5), derived their costs with explicit arithmetic — a sparse-over-dense unitarity-check speedup of $d_t/4 = 1000.5$ at $t = 10^3$, a worst-case 100% normalization-check overhead, and a naive clause count $K_t = 5t$ — and gave a projected double-precision drift bound $\delta_{10^6} \approx 8.88 \times 10^{-10}$ plus an illustrative 3-cycle periodicity contract with return probability $1/3$ under stated assumptions. The framework favors exactly checkable, finite, discrete contracts, in line with the finite-distinction emphasis of [9], and every quantitative claim is either computed here with shown arithmetic or explicitly labeled a projection.

References

- [1] Topological phenomena in quantum walks; elementary introduction to the physics of topological phases. arXiv:1112.1882v1. <https://arxiv.org/abs/1112.1882v1> [2] Discrete-event simulation of quantum walks. arXiv:2005.03401v1. <https://arxiv.org/abs/2005.03401v1> [3] Quantum walks: a comprehensive review. arXiv:1201.4780v2. <https://arxiv.org/abs/1201.4780v2> [4] Order from chaos in quantum walks on cyclic graphs. arXiv:2008.00316v3. <https://arxiv.org/abs/2008.00316v3> [5] From Discrete Time Quantum Walk to Continuous Time Quantum Walk in Limit Distribution. arXiv:1307.3384v1. <https://arxiv.org/abs/1307.3384v1> [6] Symmetries and noise in quantum walk. arXiv:quant-ph/0607188v3. <https://arxiv.org/abs/quant-ph/0607188v3> [7] Deductive Verification of Unmodified Linux Kernel Library Functions. arXiv:1809.00626v1. <https://arxiv.org/abs/1809.00626v1> [8] Quantum walks and their algorithmic applications. arXiv:quant-ph/0403120v3. <https://arxiv.org/abs/quant-ph/0403120v3> [9] QNFO: Finite-Distinction Quantum Mechanics: Unitary Evolution and Superposition as the Large-Distinction Limit of Stochastic

Thermodynamics [10] QNFO: FACTORING, Adelic Complexity, and the Silent-Radix Principle [11] QNFO: Manifesto for Honest Computation [12] QNFO: Vanishing ZBW Signal: The ZBW-Majorana Hypothesis as a Unified Framework for Topological Fermion Distinction

Appendix A. Divergence report

One quantitative divergence arose across drafts: the value of $M_{1000} = 8t(t+2)$ for $t = 1000$. One draft reported 8,024,000; another computed 8,016,000. The disagreement stems from an arithmetic slip in evaluating $8 \times 1000 \times 1002$: the correct product is $1000 \times 1002 = 1,002,000$, giving 8,016,000. The main text adopts 8,016,000 (Derivation 5, shown arithmetic) and the Abstract has been corrected accordingly. No other divergences between drafts were substantive.

Appendix B. Claim attribution

Claim	Source drafts	Status
C1: Contract framework K1–K5 modeled on [7]	A, B, C	CONVERGENT
C2: $d_{1000} = 4002$; dense cost $d_t^3 = 64,096,048,008$	A, B, C	CONVERGENT
C3: Sparse unitarity-check cost $4d_t^2 = 64,064,016$; speedup $d_t/4 = 1000.5$ (row-norm check $8d_t = 32,016$ verifies normalization only)	A, B, C	CONVERGENT
C4: $M_{1000} = 8,016,000$ (one draft: 8,024,000)	A, B (C divergent)	CONVERGENT after correction
C5: Worst-case 50% normalization overhead	A, B	CONVERGENT
C6: Clause count $K_t = 5t$	A	SINGLE
C7: Drift bound $\delta_{10^6} \approx 8.88 \times 10^{-10}$ (projection)	A, B	CONVERGENT
C8: 3-cycle return probability $1/3$ (illustrative assumption)	B	SINGLE
C9: Finite-distinction alignment via [9]	A, C	CONVERGENT