

Local Hilbert-Space Dimension as a Resolution Parameter: Information-Theoretic Bounds for Qudit Quantum Cellular Automata as Substrates of Emergent Holographic Geometry

Local Hilbert-Space Dimension as a Resolution Parameter: Information-Theoretic Bounds for Qudit Quantum Cellular Automata as Substrates of Emergent Holographic Geometry

Abstract

Discrete models of quantum gravity lack a computable criterion for when their emergent geometry becomes semiclassical, and the system sizes at which quantum cellular automata (QCA) show holographic entanglement structure remain out of reach of exact simulation. We ask whether enlarging the local Hilbert space from qubits to qudits lowers the system size N at which continuum-like behavior appears. We develop the analytical framework for qudit QCA with local dimension d : we derive the exact scaling of Hilbert-space dimension d^N , the maximum cut entanglement $S_{\max} = \min(|A|, |B|) \log_2 d$, the site-count equivalence $N_d = N_q / \log_2 d$ between qudit and qubit chains at matched information capacity, and the per-gate simulation overhead $d^2/4$ at matched Hilbert dimension. These bounds give concrete numbers: a $d = 3$ chain of 16 sites carries the entropy capacity of a 26-site qubit chain (a 36.9% site saving) at 2.25 $\times$ the per-gate simulation cost, and a $d = 4$ chain of 12 sites matches a 24-site qubit chain at 4 $\times$ cost. We propose a falsifiable protocol comparing Ryu–Takayanagi-style entropy curves across d at matched capacity. The central conjecture — that emergence thresholds scale as $1/\log_2 d$ — is stated as a projection with explicit assumptions, giving candidate thresholds $N^* \approx 20, 13, 10$ for $d = 2, 3, 4$, and the smallest simulation that would confirm or refute it is specified.

1. Introduction

Researchers building discrete models of quantum gravity face a concrete, unresolved problem: they cannot compute, for a given discrete substrate, at what system size its emergent geometry stops looking like a lattice and starts looking like a smooth manifold. In the holographic program, the bulk geometry of a region is encoded in the entanglement structure of its boundary degrees of freedom, so the question becomes operational: at what system size N does the entanglement entropy of a discrete quantum system follow the area-law–geodesic scaling that signals a semiclassical bulk? For qubit quantum cellular automata — lattices of finite-dimensional quantum systems evolving under local unitary dynamics — exact simulation is bounded by state vectors of dimension 2^N , and candidate holographic behavior has only been certified at sizes of order twelve qubits [9]. Larger qubit chains, where a genuine continuum limit could be observed, exceed classical memory: a 40-qubit state vector requires 2^{40} complex amplitudes. The field is therefore stuck between toy sizes too small to show geometry and continuum sizes too large to simulate.

This gap has a cost. Without a tunable parameter that moves the emergence threshold to smaller N , every investigation of emergent geometry in locally unitary dynamics is confined to the regime where finite-size effects dominate, and no statement about the approach to smooth-manifold behavior can be tested. If local Hilbert-space dimension d acts as a resolution parameter — so that a d -level chain reaches continuum-like entanglement scaling at a system size shrinking as $1/\log_2 d$ — then the accessible simulation window opens dramatically: the same classical computer that simulates 28 qubits exactly can simulate 14 four-level sites, and the entanglement structure at

those two sizes carries identical information capacity. The decision that changes is which automata a theorist screens for holographic behavior, and how small a prototype must be before a claim about emergent geometry is testable at all.

The physical motivation for treating d as a resolution knob is independent of simulation convenience. In the p-adic pre-geometric program, the continuum manifold of general relativity is replaced by a discrete hierarchical substrate precisely because the continuum description breaks down at the Planck scale [10]; the open question in that program is how a computable discrete substrate generates smooth effective geometry. In stochastic and classical cellular automata, the trade-off between alphabet size and neighborhood size is structural: the number of states can be decreased at the cost of a larger neighborhood, and conversely the number of inputs can be decreased by enlarging the alphabet [6]. Quantum mechanically, the analogous trade-off has never been quantified for geometric emergence: no prior work derives how the system size needed for semiclassical entanglement scaling depends on the local dimension d .

This paper does not claim to have simulated qudit QCA at holographic scale. It derives the exact information-theoretic bounds that any such simulation must satisfy, quantifies the simulation cost of testing the conjecture, and specifies the smallest experiment that would confirm or refute it. The contributions are:

- Derivation of the exact entropy-capacity equivalence between qudit and qubit chains, $N_d = N_q / \log_2 d$, with worked numbers for $d = 3$ and $d = 4$ (Section 4.2).
- Derivation of the per-gate classical simulation overhead $d^2/4$ at matched Hilbert-space dimension, showing that qudit enlargement is polynomially, not exponentially, more expensive (Section 4.3).
- Derivation of the mutual-information and Ryu–Takayanagi-style scaling bounds $I(A : B) \leq 2 \min(|A|, |B|) \log_2 d$ and $S(L) \leq L \log_2 d$ that define the observable signatures of the conjecture (Section 4.4).
- A projection, with stated assumptions, that the emergence threshold scales as $N^*(d) \propto 1/\log_2 d$, giving candidate thresholds $N^* \approx 20$ ($d = 2$), 13 ($d = 3$), 10 ($d = 4$) at a 5% discretization criterion (Section 4.5, Section 5).
- Specification of the minimal falsification protocol: exact state-vector simulation at $d \in \{2, 3, 4\}$, $N \leq 14$, comparing entropy-curve convergence at matched Hilbert dimension (Sections 5 and 7).

2. Background and Related Work

Classical cellular automata research supplies the structural vocabulary for the quantum setting, and several of its results constrain what a qudit QCA program can assume.

The classification of one-dimensional automata is itself a difficult and partially undecidable problem. The survey of elementary cellular automata classifications [1] records that Culik II and Yu proved in 1988 that classifications of this kind are undecidable, while the same survey tabulates Wolfram's four behavioral classes — uniform, periodic, chaotic, and complex — over the 256 elementary rules, with 8 rules in class I. It also shows that composing a chaotic rule with a memory function can convert chaotic behavior into complex behavior. For the present program this is a warning and an opportunity: behavioral class in a qudit QCA cannot be inferred from d alone, but rule transformations that shift class suggest that structured rule families, not exhaustive search, are the way to navigate the d^{d^k} -sized rule spaces of higher-dimensional local rules.

Rule-space size is quantified explicitly in the iso-rule paradigm of [2], which gives the rule-table size of a generic CA with v states and k inputs as $S_r = v^k$ and shows that structured families (totalistic, outer-totalistic, iso-rules) shrink this to $S_k = (v + k - 1)! / (k! \times (v - 1)!)$ and smaller, with the general ordering

$S_r > S_i > S_k > S_t$. The same work argues that isotropy — equivalent dynamics in all directions — is the proper canvas for logical universality in CA. Both points transfer directly: a generic two-input qudit rule with $d = 3$ has $S_r = 3^2 = 9$ table entries, and imposing isotropy on a qudit lattice restricts the search in the same way the iso-rule paradigm restricts the binary case.

The trade-off between alphabet size and neighborhood size is stated in [6]: for one-dimensional CA the number of states and the number of inputs are not independent, since states can be traded against inputs in either direction. That work then proves that all 18 ternary CA rules induced by semigroups of order 3 are solvable, with explicit solution formulae $f^n(x_0, \dots, x_n) = x_0 \odot x_1 \odot \dots \odot x_n$ for idempotent semigroups, and shows these formulae compute expected values and block probabilities along orbits under Bernoulli initial conditions. This is the strongest existing analytical handle on $d = 3$ automata, and Section 4.5 uses the same idea — closed-form orbit statistics — as the analytical complement to simulation in the large- d limit.

Decidability results bound what can be proven about QCA phase spaces. The work on decidability of one-dimensional CA [3] establishes, building on Sutner's result that the first-order theory of the phase space $\mathcal{S}_{\mathcal{A}} = (Q^{\mathbb{Z}}, \longrightarrow)$ is decidable, that the graphs of one-dimensional CA are ω -automatic, and recalls that injectivity and surjectivity were shown decidable by Amoroso and Patt. For the qudit program this means that structural questions about the phase space — reversibility above all, which a unitary QCA requires — remain algorithmically decidable as d grows, so the enlargement of the local Hilbert space does not push the foundational questions past decidability.

Stochastic and asynchronous variants matter because any simulated qudit QCA will face noise and update-scheduling choices. The decomposition theory of stochastic CA [4] shows that a stochastic CA admits a cell-wise probability distribution representation and a continuous-CA (CCA) formulation based on a polynomial representation of the discrete rule, allowing the averaged behavior to be uncovered without sampling; it also notes that mean-field techniques apply to long-term behavior. The asynchronous survey [5] defines α -asynchronous updating, in which each cell applies the rule with probability α per step, and reports that asynchrony can break the complexity of a rule and render it simpler to study. Both results shape the noise model of Section 3: qudit QCA diagnostics should be computed on the averaged (density-matrix or dephased) dynamics, and update asynchrony is a control variable, not just a nuisance.

On the physics side, the nonequilibrium study of probabilistic CA [7] states plainly that there is no derivation of PCA dynamics from more microscopic physical rules, and entertains the reversal: that PCA are more fundamental and standard physical descriptions must be derived from them. It also shows that detailed balance is not naturally installed for parallel-updating PCA because simultaneous neighbor updates can prevent gradient flow, and constructs the entropy-production measure $J_{i,n}(\sigma) = \log \frac{p_i(\sigma_n(i)|\sigma_{n-1})}{p_i(\sigma_{n-1}(i)|\sigma_n)}$. A unitary QCA is the quantum analogue of the reversible, parallel-updating case, and the same tension — no energy function, no detailed balance — carries over, which is why Section 3 works with entanglement measures rather than Hamiltonian variational principles.

That CA dynamics can produce physically useful emergent behavior at all is demonstrated at applied scale in [8], which simulates granular flow in an hourglass on a 61×61 grid using a probabilistic CA with a modified Margolus neighborhood, reporting a few milliseconds per iteration and usable performance up to 4k resolution. The Margolus block-partition construction used there is also the standard way to build exactly unitary, reversible lattice dynamics, and its millisecond-scale cost on a classical machine is the practical precedent for the simulation-cost analysis of Section 4.3.

Within the quantum CA literature proper, the closest existing result is the twelve-qubit Fredkin-gate QCA study [9], which asks whether local unitary dynamics can generate the non-local information storage that holographic quantum error correction requires. It establishes a zero-versus-nonzero diagnostic contrast: Fredkin-class dynamics delocalize

information across the chain at every erasure scale tested, while transport-class dynamics collapse to zero segment correlations at the first erasure, and it packages this as a decoder-free erasure-screening pipeline runnable on a laptop. That work is bounded by twelve qubits and by the proxy status of mutual information — exactly the size ceiling this paper's resolution parameter is designed to lower.

Finally, the broader pre-geometric program supplies context. The p-adic spacetime toolkit [10] develops computable methods for a Bruhat–Tits-tree substrate, reports that a discrete Wheeler–DeWitt equation converged and that emergent time was mapped via thermodynamic RG flow, and frames the continuum manifold as mathematically exhausted at the Planck scale. The holographic optimization framework [11] proposes a structural duality between surface-area minimization in biological branching networks and spacetime-volume minimization in quantum circuits, reporting a biological fractal dimension $\mu \approx 2.51$ and a cross-domain correlation $r \approx 0.98$. The constructal-determinism synthesis [12] argues that deterministic, computationally dense substrates (with local computational index $LCI \approx 1.83$) on fractal invariant sets generate the observed stochasticity of quantum mechanics, with hidden variables in an ultrametric topology. These three works motivate the specific question studied here — how discrete substrates generate smooth effective geometry — but none of them addresses local Hilbert-space dimension as a control parameter; that gap is what this paper fills.

3. Methods

3.1 Model definition

We define a qudit quantum cellular automaton (QCA) as follows. The system is a one-dimensional chain of N sites, each carrying a local Hilbert space $\mathcal{H}_i \cong \mathbb{C}^d$ with $d \in \{2, 3, 4, \dots\}$; the total Hilbert space is

$$\mathcal{H}_N^{(d)} = \bigotimes_{i=1}^N \mathcal{H}_i, \quad \dim \mathcal{H}_N^{(d)} = d^N.$$

The dynamics is a translationally invariant local unitary: in one automaton time step, every site is updated by a two-site unitary U_d acting on $(\mathcal{H}_i, \mathcal{H}_{i+1})$, applied in a brickwork (Margolus-style) partition that alternates even and odd bonds between steps, following the block-partition construction used for reversible classical CA in [8]. A single time step is therefore

$$\mathcal{U}^{(d)} = (\mathbb{I}^{\otimes} U_d) (U_d \mathbb{I}^{\otimes}),$$

with the two factors acting on alternating bond partitions. Exact unitarity of the block partition guarantees reversibility, which the decidability results of [3] confirm is an algorithmically checkable property of the phase space at any d .

The two-site gate U_d is an element of $U(d^2)$. We do not fix a single gate; the protocol screens gate families, prioritizing conditional-swap (Fredkin-generalized) gates because the twelve-qubit study [9] certified that the Fredkin class, unlike transport class, delocalizes information across erasures — the redundancy signature that holographic error correction requires.

3.2 Observables

Three information-theoretic observables are computed on $\mathcal{H}_N^{(d)}$, all defined without reference to any decoder, following the decoder-free diagnostic philosophy of [9]:

1. **Cut entanglement entropy.** For a bipartition $A|B$ with $|A|$ contiguous sites, $S(A) = -\text{Tr } \rho_A \log_2 \rho_A$, bounded as derived in Section 4.1.

2. **Interval entropy curve.** $S(L)$ for intervals of length $L = 1, \dots, \lfloor N/2 \rfloor$; the Ryu–Takayanagi correspondence predicts, in a semiclassical bulk, a curve that is linear at small L and saturates, with the linear slope set by the central charge in CFT duals. Here the discrete analogue is the slope $\log_2 d$ derived in Section 4.4.
3. **Mutual information.** $I(A : B) = S(A) + S(B) - S(A \cup B)$, the erasure-robustness proxy used in [9], with the bound derived in Section 4.4.

3.3 Noise and update model

Two perturbations are part of the protocol because the classical literature shows they change behavior qualitatively. First, dephasing noise on individual sites converts the pure state to a mixture; the averaged dynamics is then analyzed with the cell-wise probability and continuous-CA machinery of [4], which computes averaged behavior without sampling. Second, asynchronous update is introduced as the parameter α , the per-step probability that each site participates in the update, following the α -asynchronous scheme defined in [5]; $\alpha = 1$ recovers the synchronous QCA. The survey [5] documents that asynchrony can break rule complexity, so α is swept rather than fixed.

3.4 Matched-capacity comparison design

The central methodological rule is that chains of different d are only compared at matched information capacity, i.e., matched total Hilbert-space dimension $d^N = 2^{N_q}$, because a $d = 4$ chain of N sites trivially carries more entropy than a qubit chain of the same N ; an unmatched comparison would confound local dimension with system size. The equivalence $N_q = N \log_2 d$ is derived in Section 4.2 and governs all cross- d comparisons in this paper.

4. Analysis

Every input number in this section is either a definition of the model of Section 3, a value cited from a specific source, or an explicitly stated assumption. All arithmetic is shown.

4.1 Hilbert-space dimension and entropy capacity

The dimension of the total Hilbert space is d^N by definition of the tensor product. For $N = 12$ sites (the largest size for which a qubit QCA has been exactly characterized [9]):

$$d = 2 : 2^{12} = 4096.$$

$$d = 3 : 3^{12} = 531441.$$

$$d = 4 : 4^{12} = (2^2)^{12} = 2^{24} = 16777216.$$

The ratio of qudit to qubit dimension at fixed N is

$$\frac{3^{12}}{2^{12}} = \frac{531441}{4096} = 129.7463\dots, \quad \frac{4^{12}}{2^{12}} = \frac{16777216}{4096} = 4096.$$

The maximum entanglement entropy of a bipartition $A|B$ is the logarithm of the smaller Schmidt rank, which is $\min(|A|, |B|)$ qudit levels:

$$S_{\max}(A) = \min(|A|, |B|) \log_2 d.$$

The per-site entropy capacity is $\log_2 d$ bits. Numerically, using $\log_2 3 = \ln 3 / \ln 2 = 1.0986123/0.6931472 = 1.5849625$:

$$\log_2 2 = 1, \quad \log_2 3 = 1.5849625, \quad \log_2 4 = 2.$$

So each $d = 3$ site carries 1.585 bits of entanglement capacity and each $d = 4$ site carries 2 bits, against 1 bit for a qubit.

4.2 Site-count equivalence at matched capacity

Matching total entropy capacity between a qubit chain of N_q sites and a d -level chain of N_d sites requires $N_q \log_2 2 = N_d \log_2 d$, hence

$$N_d = \frac{N_q}{\log_2 d}.$$

Worked examples, rounding up to whole sites:

- $N_q = 24, d = 3$: $N_d = 24/1.5849625 = 15.1424\dots$, so $N_d = 16$ sites. Site saving: $1 - 15.1424/24 = 1 - 0.630933 = 0.369067$, i.e. 36.9% fewer sites (16 vs 24 after rounding).
- $N_q = 24, d = 4$: $N_d = 24/2 = 12$ sites exactly, a 50% site saving.
- $N_q = 28, d = 4$: $N_d = 28/2 = 14$ sites.

Equivalently, a $d = 3$ chain of $N = 16$ sites has capacity $16 \times 1.5849625 = 25.3594$ qubit-equivalent bits, matching a 25–26 site qubit chain; a $d = 4$ chain of $N = 14$ has capacity $14 \times 2 = 28$ bits, matching 28 qubits exactly.

4.3 Classical simulation cost at matched dimension

Exact state-vector simulation stores d^N complex amplitudes. At double precision (16 bytes per amplitude), the memory is $16 d^N$ bytes. At matched Hilbert dimension $d^N = 2^{N_q}$, memory is identical across d by construction; the difference is the per-gate cost.

Applying a two-site gate U_d (a $d^2 \times d^2$ matrix) to the state vector requires, for each of the d^{N-2} configurations of the remaining sites, one $d^2 \times d^2$ matrix-vector product of cost $(d^2)^2 = d^4$ multiply-adds:

$$C_d(N) = d^{N-2} \cdot d^4 = d^{N+2}.$$

For a qubit chain, $C_2(N_q) = 2^{N_q+2}$. At matched dimension $d^N = 2^{N_q}$ the cost ratio is

$$\frac{C_d}{C_2} = \frac{d^{N+2}}{2^{N_q+2}} = \frac{d^N \cdot d^2}{d^N \cdot 4} = \frac{d^2}{4}.$$

Numerically:

$$d = 3 : \frac{9}{4} = 2.25, \quad d = 4 : \frac{16}{4} = 4.$$

This is the paper's key cost result: at equal information content, a qudit simulation costs a constant factor 2.25 ($d = 3$) or 4 ($d = 4$) more per gate than the equivalent qubit simulation — polynomial in d , not exponential in N . Concretely, the $d = 4, N = 14$ chain has $2^{28} = 268435456$ amplitudes, requiring $16 \times 268435456 = 4294967296$ bytes = 4.295 GB of memory — the same memory as a 28-qubit state vector, with 4× the per-gate arithmetic.

For orientation, the largest exactly characterized QCA study to date used twelve qubits [9]; its state vector is $2^{12} = 4096$ amplitudes, requiring $16 \times 4096 = 65536$ bytes = 64 KiB — laptop scale, which is why that study stopped at twelve qubits and why the memory wall, not the gate count, is the binding constraint this paper's

resolution parameter addresses.

4.4 Entropy and mutual-information bounds

Every observable of Section 3.2 obeys an exact bound set by d . For the interval entropy, the Schmidt rank of an interval of L sites against its complement is at most d^L , so

$$S(L) \leq \log_2 d^L = L \log_2 d.$$

For $d = 2, 3, 4$ the maximal slopes are 1, 1.5849625, and 2 bits per site. For the mutual information between two disjoint intervals A and B , each term is bounded by its own Schmidt rank, giving

$$I(A : B) = S(A) + S(B) - S(A \cup B) \leq 2 \min(|A|, |B|) \log_2 d.$$

These bounds define the observable signatures of the conjecture: if emergence thresholds shrink with d , then the entropy curve $S(L)$ should approach its saturating, continuum-like shape at smaller N for larger d at matched capacity, with the small- L slope fixed at $\log_2 d$.

4.5 Projection of the emergence threshold

We now state the central projection with its assumptions explicit. **Assumption A1 (baseline).** We assume a qubit emergence threshold $N^*(2) = 20$ sites, defined as the smallest qubit chain size at which the relative change of the interval entropy curve between successive N falls below 5% for the Fredkin-class gate family. This baseline is a stated assumption of this paper, adopted for definiteness; the twelve-qubit study [9] certified diagnostic contrast at $N = 12$ but did not establish a convergence threshold, so $N^*(2) = 20$ is a conservative round value chosen above the certified size. **Assumption A2 (scaling).** We assume the threshold scales inversely with per-site entropy capacity, $N^*(d) = N^*(2)/\log_2 d$, on the hypothesis that the number of sites needed to reach a fixed total entanglement capacity is the binding resource. **Assumption A3 (criterion).** The 5% relative-change criterion of Section 3.2 is applied at matched Hilbert dimension.

Under A1–A3 the projected thresholds are:

$$N^*(3) = \frac{20}{1.5849625} = 12.6186 \dots \Rightarrow 13 \text{ sites (rounding up),}$$

$$N^*(4) = \frac{20}{2} = 10 \text{ sites exactly.}$$

The corresponding matched qubit capacities are $13 \times 1.5849625 = 20.6045$ bits and $10 \times 2 = 20$ bits, both consistent with the 20-bit baseline capacity by construction. These are projections, not measurements: they follow from A1–A3 by arithmetic alone, and their uncertainty is dominated by the uncertainty in A1 and in the functional form of A2. If the true threshold instead scales with a power of $\log_2 d$ other than -1 , or with the rule-space size rather than the entropy capacity, the projected values shift; Section 6 quantifies this sensitivity.

5. Results

The quantitative results derived in Section 4 are summarized below. All values are exact derivations except the thresholds N^* , which are projections under Assumptions A1–A3 of Section 4.5.

Quantity	$d = 2$	$d = 3$	$d = 4$
Per-site entropy capacity $\log_2 d$ (bits)	1	1.5849625	2
Sites matching a 24-qubit chain	24	16	12

Quantity	$d = 2$	$d = 3$	$d = 4$
Site saving vs 24 qubits	—	36.9%	50%
Per-gate cost ratio $d^2/4$ at matched dimension	1	2.25	4
Rule-table entries $S_r = d^2$ (two-input rule)	4	9	16
Projected emergence threshold N^* (sites)	20 (assumed)	13	10

Three results follow. First, at matched information capacity a $d = 3$ chain of 16 sites carries the entropy capacity of a 25–26 site qubit chain (36.9% site saving) and a $d = 4$ chain of 12 sites matches 24 qubits exactly (50% saving). Second, the price of this saving is a constant per-gate factor of 2.25 ($d = 3$) or 4 ($d = 4$), independent of N — qudit enlargement is polynomially, not exponentially, more expensive. Third, under the stated assumptions the emergence threshold falls from an assumed 20 qubit sites to a projected 13 ($d = 3$) or 10 ($d = 4$) sites, bringing the threshold inside the range accessible to exact state-vector simulation (Section 4.3 shows $N = 14$ at $d = 4$ requires 4.295 GB). No empirical simulation data are introduced in this paper; every number in the table is either exact arithmetic or a labeled projection.

6. Discussion

6.1 What the results do and do not establish

The exact results of this paper are three information-theoretic identities: the entropy capacity $\log_2 d$ per site, the site-count equivalence $N_d = N_q / \log_2 d$, and the per-gate cost ratio $d^2/4$ at matched Hilbert dimension. These hold for any qudit chain regardless of the gate U_d , because they follow from the Schmidt-rank structure of the tensor product alone. They establish that qudit enlargement is a feasible simulation strategy, not that it accelerates geometric emergence. The acceleration claim — that emergence thresholds scale as $N^*(d) = N^*(2) / \log_2 d$ — is a projection under Assumptions A1–A3, and its status is hypothesis, not result.

6.2 Limitations

The projection rests on the following assumptions, each with its sensitivity:

- **A1 (baseline threshold $N^*(2) = 20$).** This is a stated assumption, not a measurement. The twelve-qubit study [9] certified a diagnostic contrast at $N = 12$ but established no convergence threshold, so the true qubit threshold could lie anywhere above 12. The projected thresholds scale linearly in A1: if $N^*(2) = 30$, then $N^*(3) = 30/1.5849625 = 18.927 \dots \Rightarrow 19$ sites and $N^*(4) = 15$ sites; if $N^*(2) = 16$, then $N^*(3) = 10.096 \dots \Rightarrow 11$ and $N^*(4) = 8$. The $d = 4$ projection is therefore between 8 and 15 sites for any baseline in $[16, 30]$ — still inside the exact-simulation window derived in Section 4.3, which is the robust part of the claim.
- **A2 (scaling exponent -1).** The hypothesis that threshold scales as $1/\log_2 d$ is the paper's central conjecture and its most fragile assumption. If the threshold instead scales with a power $\log_2 d$ other than -1 , say $N^*(d) = N^*(2)/(\log_2 d)^\gamma$ with $\gamma = 2$, then $N^*(3) = 20/1.5849625^2 = 7.96 \dots \Rightarrow 8$ and $N^*(4) = 20/4 = 5$; if $\gamma = 1/2$, then $N^*(3) = 20/\sqrt{1.5849625} = 15.89 \dots \Rightarrow 16$ and $N^*(4) = 20/\sqrt{2} = 14.14 \dots \Rightarrow 15$. The conjecture is thus distinguishable from its alternatives only by simulation, which is why Section 5 specifies the falsification protocol rather than asserting the exponent.
- **A3 (matched-dimension criterion).** The 5% relative-change criterion is a convention. A stricter criterion (1%) moves all thresholds upward proportionally; the cross- d ordering is unaffected, the absolute numbers are not.
- **Gate-family dependence.** All entropy bounds are gate-independent maxima, but actual entropy curves depend strongly on U_d . The Fredkin-class prioritization follows the delocalization contrast certified in [9]; a gate

family that fails to delocalize information would show no convergence of the entropy curve at any d , and the projection would be vacuous for it.

- **No empirical data.** This paper contains no simulation of any qudit QCA. Every number is exact arithmetic on the model definition or a labeled projection. The comparison of entropy curves across d at matched capacity has not been performed.
- **Bibliography scope.** The cited literature is drawn from classical cellular automata theory, one twelve-qubit QCA study, and pre-geometric and holographic-optimization programs; no prior work quantifies emergence thresholds versus local dimension, so there is no external benchmark against which A2 can be checked before simulation.

6.3 Failure modes and falsification

The main result — the projected threshold shrinkage — is falsified by a specific observation: at matched Hilbert dimension, simulate the Fredkin-generalized gate family at $d \in \{2, 3, 4\}$ and compute the interval entropy curve $S(L)$ for each N up to 14. If the $d = 4$ chain reaches the 5% relative-change criterion at a site count N_4^* that is not smaller than the $d = 2$ threshold measured at the matched capacity $N_2^* \log_2 2$ — i.e., if $N_4^* \log_2 4 > N_2^*$ — then Assumption A2 is false and local dimension is not a resolution parameter in this sense. A weaker falsification of the practical motivation occurs if the $d^2/4$ per-gate overhead is swamped by a gate-dependent prefactor: if applying a generic U_d requires $O(d^4)$ primitive operations per two-site gate rather than the single $d^2 \times d^2$ matrix-vector product assumed in Section 4.3, the cost ratio becomes $(d^2/4) \times c_d$ with c_d the gate-decomposition overhead, and the advantage can vanish for $c_4 > 1$ at $d = 4$.

A second failure mode is diagnostic rather than quantitative: the entropy curve $S(L)$ may converge in shape without the convergence tracking any geometric content. The twelve-qubit study [9] already showed that mutual information is a proxy — its conclusion explicitly bounds the result by "the proxy status of mutual information" — and the same caveat applies with full force here. Convergence of $S(L)$ to a saturating curve is necessary but not sufficient for a semiclassical bulk; a falsification protocol that passes could still reflect finite-entanglement physics with no geometric interpretation.

A third failure mode concerns the noise model. The dephased, averaged dynamics analyzed with the continuous-CA machinery of [4] and the α -asynchronous sweeps of [5] may destroy the very entanglement structure whose convergence is being measured: asynchrony can break rule complexity [5], and if it breaks the Fredkin-class delocalization as well, the $\alpha < 1$ branches of the protocol measure nothing about emergence. The protocol treats α as a control variable for this reason, but the interaction between asynchrony and entanglement delocalization is unquantified.

6.4 Open questions

Three questions are open and material. First, does the large- d limit admit the closed-form orbit statistics that [6] provides for ternary semigroup rules, and if so, does the limiting entropy curve approach a continuum form? Second, is there a rule-family construction for qudit QCA that plays the role the iso-rule paradigm [2] plays for isotropic binary CA, shrinking the $d^{d^{2k}}$ -sized two-input rule space to a navigable structured subset? Third, does the undecidability of behavioral classification [1] imply that no finite-size diagnostic can certify semiclassicality in general, so that the 5% criterion is at best a per-family heuristic? These questions bound the reach of the program and are not resolved by the bounds derived here.

7. Practical Implications and Engineering Pathways

7.1 Who would use the result and what changes

The users are researchers who screen discrete substrates for holographic behavior by exact simulation — the practitioners addressed directly by the erasure-screening pipeline of [9], which is runnable "in minutes on a laptop" and produces a pass/fail contrast between Fredkin-class and transport-class dynamics. For this group, the result changes two decisions. First, the screening budget: at matched information capacity, a $d = 4$ chain of 12 sites replaces a 24-qubit chain at $4\times$ per-gate cost and identical memory (4.295 GB at the 28-qubit-equivalent scale derived in Section 4.3), so a fixed compute budget buys roughly $2.25\times$ ($d = 3$) or $4\times$ ($d = 4$) more entropy capacity per site count. Second, the threshold claim: if the projection holds, a candidate automaton can be tested for continuum-like entropy scaling at 10–13 sites instead of 20, moving the test inside the regime where exact state-vector simulation is routine.

No commercial or industrial user exists for this result; the work is foundational. Its value is that it converts an untestable question (does d accelerate emergence?) into a bounded, falsifiable simulation protocol with stated costs.

7.2 Readiness level

Against the NASA Technology Readiness Level definitions, this work is TRL 1–2: it provides analytical predictions and basic principles (the exact bounds of Section 4) with an application concept formulated (the protocol of Section 5), but no simulation has been run, no code has been validated, and the paper has not been peer reviewed. Analytical work of this kind cannot exceed TRL 3, and the honest current level is TRL 2, since the concept and its application have not yet been supported by even a proof-of-concept calculation on a real gate family.

7.3 The smallest next experiment

The minimal de-risking calculation is: implement exact state-vector simulation of the brickwork QCA of Section 3.1 for $d \in \{2, 3, 4\}$, $N \leq 14$, with a single Fredkin-generalized two-site gate U_d per d ; compute $S(L)$ for all $L \leq \lfloor N/2 \rfloor$ at each N ; and evaluate the 5% relative-change criterion at matched Hilbert dimension. The measured quantities are the entropy curves $S(L; N, d)$ and the thresholds N_d^* . The outcome that refutes the central conjecture is $N_4^* \log_2 4 > N_2^*$ (no shrinkage in capacity units); the outcome that refutes the practical motivation is a measured per-gate cost ratio exceeding $d^2/4$ by a gate-dependent factor that erases the site saving. Memory for the largest case is 4.295 GB (Section 4.3), so the entire experiment fits a single workstation.

7.4 Binding constraints

The constraints that bind first, in order: (i) **memory**, which grows as $16 d^N$ bytes and is identical across d at matched dimension — the wall at $N_q \approx 30$ –34 (16–64 GB) bounds the whole program; (ii) **gate decomposition**, the open question of whether a physically motivated U_d (e.g., a qudit Fredkin) decomposes into primitives with overhead c_d that stays $O(1)$; (iii) **rule-family selection**, since the two-input qudit rule space has d^{d^2} elements ($3^9 = 19683$ for $d = 3$, $4^{16} \approx 4.29 \times 10^9$ for $d = 4$) and structured families in the spirit of the iso-rule paradigm [2] have not been constructed for the quantum case; (iv) **criterion choice**, since the 5% threshold is a convention whose sensitivity is unquantified until the first simulation. Regulation, manufacturing and energy constraints do not bind: the experiment is classical computation on commodity hardware.

7.5 Cost and performance bounds (projections)

As a projection with stated assumptions: if the twelve-qubit study's laptop runtime ("in minutes on a laptop" [9]) is taken as a baseline for 2^{12} -dimensional state vectors, and runtime scales linearly in the number of gates and in state dimension, then a $d = 4$, $N = 14$ run (2^{28} amplitudes, $2^{28}/2^{12} = 65536$ times the state dimension, with $4\times$ per-gate cost) would take on the order of $65536 \times 4 \approx 2.6 \times 10^5$ times the per-iteration cost of the twelve-qubit run. This is a scaling projection only; it assumes linear scaling in dimension, no memory-bandwidth saturation, and the same number of time steps, and it is offered solely to size the compute commitment, not as a measurement.

7.6 Risks and misuse paths

There are no physical, safety or dual-use risks: the work concerns classical simulation of abstract quantum systems. The relevant misuse is rhetorical — presenting the projected thresholds $N^* \approx 13, 10$ as established results rather than projections under A1–A3, which would overstate the evidence for dimension-as-resolution by the entire gap between arithmetic and simulation.

8. Conclusion

This paper derives the exact information-theoretic bounds governing qudit quantum cellular automata as candidate substrates of emergent holographic geometry, and quantifies what it would cost to test whether local Hilbert-space dimension acts as a resolution parameter. The exact results are: per-site entropy capacity $\log_2 d$; the site-count equivalence $N_d = N_q / \log_2 d$, giving a 36.9% site saving at $d = 3$ and 50% at $d = 4$ against a 24-qubit chain; and the per-gate simulation overhead $d^2/4$ at matched Hilbert dimension — 2.25 for $d = 3$, 4 for $d = 4$ — which establishes that qudit enlargement is polynomially, not exponentially, more expensive. On these foundations the paper states, as an explicitly labeled projection under three assumptions, that emergence thresholds scale as $N^*(d) \propto 1/\log_2 d$, yielding candidate thresholds of 13 sites ($d = 3$) and 10 sites ($d = 4$) against an assumed 20-site qubit baseline, all inside the exact-simulation window whose memory cost the paper derives (4.295 GB at the 28-qubit-equivalent scale). The contribution is not a demonstration that dimension accelerates emergence; it is the conversion of that conjecture into a falsifiable protocol with a specified smallest experiment, a stated refutation condition ($N_4^* \log_2 4 > N_2^*$), and quantified sensitivity to each assumption. Running that experiment is the single step that separates the conjecture from either a result or a refutation.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.18.3-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 zai-org/glm-5.3-flash, zai-org/glm-5.3, openai/gpt-oss-120b, 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] A Note on Elementary Cellular Automata Classification. arXiv:1306.5577v2. <https://arxiv.org/abs/1306.5577v2>
- [2] Isotropic Cellular Automata: the DDLab iso-rule paradigm. arXiv:2008.11279v2. <https://arxiv.org/abs/2008.11279v2>
- [3] On Decidability Properties of One-Dimensional Cellular Automata. arXiv:0903.4615v2. <https://arxiv.org/abs/0903.4615v2>
- [4] On the decomposition of stochastic cellular automata. arXiv:1503.03318v2. <https://arxiv.org/abs/1503.03318v2>
- [5] A guided tour of asynchronous cellular automata. arXiv:1406.0792v2. <https://arxiv.org/abs/1406.0792v2>
- [6] Ternary cellular automata induced by semigroups of order 3 are solvable. arXiv:2601.00486v1. <https://arxiv.org/abs/2601.00486v1>
- [7] Nonequilibrium physics aspects of probabilistic cellular automata. arXiv:1605.02876v1. <https://arxiv.org/abs/1605.02876v1>
- [8] Probabilistic Cellular Automata for Granular Media in Video Games. arXiv:2008.06341v1. <https://arxiv.org/abs/2008.06341v1>
- [9] QNFO: Computational Toy Model of Non-Local Information Storage in a Quantum Cellular Automaton
- [10] QNFO: Computational Toolkit for p-Adic Spacetime
- [11] QNFO: Spacetime Volume and Fractal Surfaces
- [12] QNFO: Topological and Computational Unification of Emergent Agency