quantum reports

MDPI

Article

From the First Triangle to an Error-Correcting Lattice: Simulations of a Growing Bell-Pair Network

Raghu Kulkarni

SSMTheory Group, IDrive Inc., Calabasas, CA 91302, USA; raghu@idrive.com

Abstract

In the Selection-Stitch Model (SSM), space is a lattice of Bell pairs that builds itself by two local rules: stitch adds a node in a triangle's plane, and lift adds one above or below it, both at exact close-packed sites, so that every bond has the same length. From one tetrahedron, stitch generates only the face-centered cubic (FCC) lattice spanned by its edges, so FCC order is a closure property of the rules; the simulations measure the defects that growth adds. In 688 simulations of up to 1.2×10^6 nodes, with the bare rule, about 98% of cluster nodes four or more layers below the surface are FCC. Slabs grown on one close-packed sheet form twin domains and are 84–92% FCC; on a three-layer ABC substrate they grow as one variant. Read as a quantum state, the grown network has an entanglement entropy that counts the bonds crossing a boundary. Node parity checks with local repair turn it deterministically into a state of the FCC stabilizer code. Its distance is at most three, so it protects the lattice locally, cell by cell; checks are intact around 96% of cluster nodes three or more layers deep. The error-correcting lattice emerges from the postulate "what exists is what can persist": only checked bonds of equal length persist, and the triangle is the smallest network whose bonds check one another.

Keywords: Bell pairs; GHZ state; foundations of quantum theory; entanglement area law; stabilizer codes; quantum error correction; face-centered cubic; crystal growth; discrete space; Selection-Stitch Model

1. Introduction

The idea that space is built from entanglement has a long history [1]. The Selection-Stitch Model (SSM) gives it a concrete form. It proposes that space is a discrete network that grows into the face-centered cubic (FCC) lattice [2]. Its nodes are joined by bonds of a single length L, which the SSM identifies with entangled (Bell) pairs. Two operators grow the network (Figure 1). Stitch adds a node in the plane of an existing triangle. Lift adds a node above or below a triangle, at the apex of a regular tetrahedron. FCC and hexagonal close packing (hcp) both reach the kissing number $K = 12$ and the Kepler packing bound [3]. Packing alone does not decide between them. For hard spheres the two differ in free energy by only about $10^{-3}\,kT$ per particle [4], and grown crystals often show random stacking [5]. This paper shows that FCC order follows from a closure property of stitch, and measures the defects that growth adds to it. It also shows that the grown network carries a stabilizer code that protects it locally, cell by cell, and traces both the rules and this protection to one postulate, "what exists is what can persist".

Check for updates

Academic Editor: Heliang Huang

Received: 14 October 2026
Revised: 8 November 2026
Accepted: 9 November 2026
Published: 10 November 2026

Copyright: © 2026 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

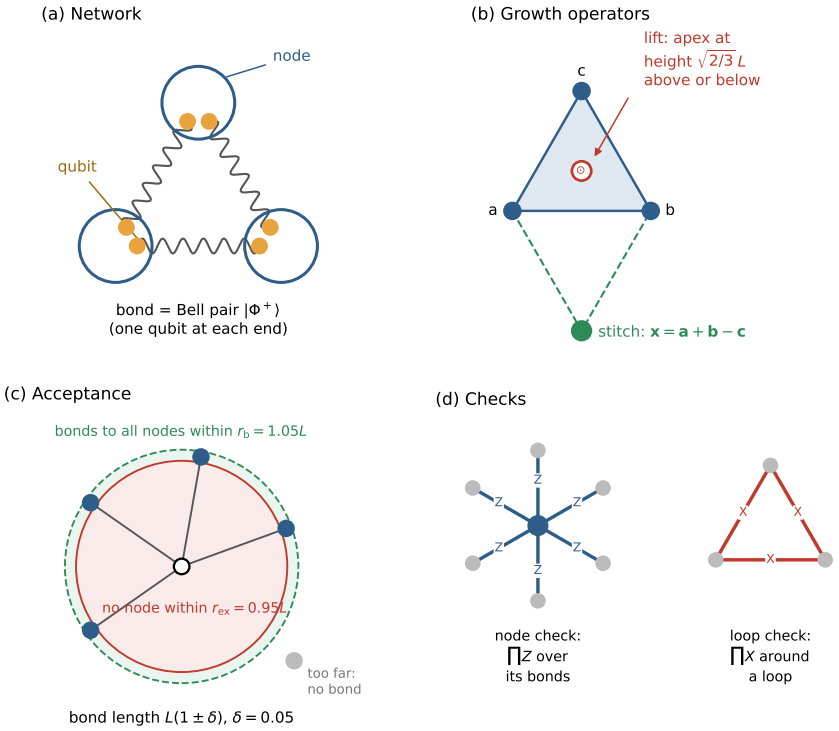


Figure 1. The ingredients of the model. (**a**) The network. Each node holds one qubit per bond, and each bond is a Bell pair shared by its two end nodes. (**b**) The growth operators. Stitch places a new node in the plane of a triangle, across one edge, at $\mathbf{x} = \mathbf{a} + \mathbf{b} - \mathbf{c}$. Lift places it at the apex of a regular tetrahedron above or below the triangle ($\odot$ marks the point above the centroid). (**c**) Acceptance. A candidate is rejected if any node lies closer than r_{ex}. An accepted node bonds to every node within r_b, so bonds are detected within the tolerance δ around L. (**d**) Checks. A node check is the product of Z over the qubits of its bonds. A loop check is the product of X around a closed loop of bonds; the FCC code uses the loops of octahedra. The rules are listed in Section 3.

Related experiments and systems. Two kinds of real systems bear on this work. On the quantum side, the checked network of Section 8.2 belongs to the family of toric-code states. Toric-code and related topological states have been prepared on quantum processors: by a unitary circuit on a superconducting processor [6], with entangled arrays of neutral atoms [7], and on a trapped-ion processor by measurement and feed-forward, where operations applied after mid-circuit measurements depend on their outcomes [8]. The deterministic preparation of Section 8.2, parity measurements followed by local repair, uses the same strategy. On the crystal side, close-packed growth with competing stackings is seen in colloidal crystals, which grow with random stacking [5] unless a patterned template fixes it [9], and in thin films, where grains whose fastest-growth direction points along the growth direction take over [10]. The twin domains and tilted grains of Section 7 are the model's analogs of these effects. The substrate test of Section 7.4 is a direct analog of template-directed growth: on one close-packed sheet the model patch forms twin domains, while a three-layer ABC substrate, like a template, fixes one stacking throughout. These comparisons are qualitative; a quantitative test, for example of the domain-growth exponent against colloidal or thin-film data, is left to future work.

Relation to previous work, and what is new. Table 1 sets out, topic by topic, what this paper takes from earlier work and what it adds or corrects. Ref. [2] proposed the SSM growth rules and simulated clusters of up to about a thousand nodes, judging structure by coordination number. Ref. [11] defined the FCC stabilizer code on the ideal lattice and studied its distance and decoding there. The main additions of this paper are these:

(i) the quantum structure of grown, imperfect networks: the area law and its deficit from defects; the local map from Bell pairs to the edge-qubit code state, a standard projected-entangled-pair construction [12,13], written out with a deterministic preparation; exact code parameters of grown patches; and the share of cells whose checks are intact, which the code protects locally; (ii) a local closure lemma, that from one tetrahedron stitch generates only the FCC lattice spanned by its edges; a test showing that the lattice and its code need equal bond lengths since nodes placed off site give neither; and a relation between the lift rate and the bond tolerance, under the assumption that weights are proportional to acceptance volumes; (iii) growth up to 10^6 nodes per cluster, with structure resolved by common neighbor analysis, and growth of patches, modeled as periodic slabs, which show twin domains that come from a single-sheet start, tilted grains and the trapping of remnant nodes; (iv) a first treatment of what precedes growth, which earlier work does not address: the model's postulate, with a concrete verification map, from which the growth rules and the error-correcting lattice follow, and a reason why the seed is a triangle. The paper also corrects several results of Ref. [2]: growth passes through no global $K = 4$ stage; rare lift does not leave disconnected sheets; the lift probability e^{-3} has no special status; the edge of the exclusion plateau is a property of the rules, not a phase transition; and the lower FCC fraction of the published simulator comes from two implementation choices (Appendix B).

Goals. The paper has three goals. The first is to test, by simulation at large scale, whether the two SSM growth rules build an FCC crystal, and to find out why they do. The second is to describe the quantum structure that the grown network carries if every bond is a Bell pair: how entangled it is, and which stabilizer code it supports. The third is to state the postulate behind the model, from which the growth rules and the error-correcting lattice follow, and to give a reason why growth starts from a triangle. The postulate is not derived, and parts of this third goal are proposals.

Far from the point where growth began, a region of the crystal grows in nearly one direction. We call such a region a patch. Its growth front is flat on a local scale and moves one way. In FCC each node sits between two kinds of tetrahedral void, whose tips point in opposite directions along a stacking axis. In a patch these two kinds differ: one points along the growth direction $\hat{g}$, and the other points against it. In the SSM, a remnant node left from an earlier disordered phase (a $K = 4$ phase in Ref. [2]) and trapped in a tetrahedral void is a baryon, and the two void orientations correspond to matter and antimatter [2]. Ref. [2] attributes the observed excess of matter to a local excess of one orientation. Suppose remnants that settle on the growing surface could be trapped in only one kind of void. Then a patch would hold matter but no antimatter, if matter is defined relative to the local growth direction (Section 10).

Within these goals we ask eight questions. Do the rules produce FCC, and not only 12-fold coordination? What sets the lift rate? How do new sheets form in a patch? Why is the stacking ABC? How does a patch grow? Is the void population biased along the growth direction? Where are remnant nodes trapped? And what quantum structure, and what error correction, does the grown network carry? The growth itself is classical geometry. Bonds are distance constraints, and no quantum state is evolved. But if every bond is a Bell pair, the grown graph fixes a quantum state. Its entanglement, and the FCC stabilizer code of Ref. [11] built on it, can then be measured directly on the simulated networks.

https://doi.org/10.3390/quantum8040121

Table 1. What this paper takes from earlier work and what it adds or corrects. Ref. [2] is the earlier simulation of the SSM, and Ref. [11] defines the FCC stabilizer code on the ideal lattice.

Topic	Earlier Work	This Paper
Before the seed	Not addressed; growth starts from a given seed [2]	New: what precedes growth, the postulate (P0) with a concrete verification map, from which the growth rules and the error-correcting lattice follow, and why the seed is a triangle; partly a proposal, with each claim's status in Section 2.6
Growth rules	Stitch, lift, exclusion and bonding proposed; stitch and lift counted as the only unit-length operators [2]	Same rules, listed in Section 3; the two operators traced to the loop and length conditions (Section 2.5); unit-bond condition added (Section 3.4)
Scale and structure test	About 10^3 nodes; coordination number [2]	Up to 10^6 nodes per cluster (Section 6) and 1.2×10^6 per slab (Section 7); common neighbor analysis, resolved by depth
Lift rate	Fixed at e^{-3} by a suppression argument [2]	If weights are proportional to acceptance volumes, $(2\sqrt{2}/\pi)\delta$ for both directions; volumes checked by Monte Carlo; e^{-3} has no special status (Section 4)
Bond tolerance δ	$\pm 5\%$, motivated by the Regge deficit [2]	Free parameter; the Regge estimate is a heuristic only (Section 4)
Exclusion radius	Plateau from $0.58L$ to $0.99L$; edge at $L/\sqrt{3}$ read as a phase transition [2]	Plateau confirmed ($0.58L$ to $0.98L$); the edge is a property of the rules, not a transition (Section 6.4)
FCC versus hcp	FCC is a Bravais lattice and hcp is not; ABC stacking reported [2]	From one tetrahedron, stitch generates only the FCC lattice spanned by its edges (a local closure lemma); every placement tested against the lattice; nodes placed off site give bonds of unequal length, no crystal and no code checks (Sections 7.3 and 9)
Growth sequence	Sheets, then a $K = 4$ tetrahedral phase, then FCC [2]	No global $K = 4$ stage; tetrahedra are local steps only; sheets grow from step edges (Sections 6.5 and 7.2)
Patches	–	New: twin domains from a single-sheet start, and their coarsening; a single variant on an ABC substrate; tilted grains (Section 7)
Matter and antimatter	Void orientations as matter and antimatter; a local excess assumed [2]	Void population balanced; trapping on a one-way front selects one class under the settling rule (Sections 7.6 and 7.7)
Entanglement	Bonds read as Bell pairs [2]	Exact area law on grown networks and its deficit from defects (Section 8.1)
Bell pairs and the code	Standard projected-entangled-pair construction [12,13]	Applied to grown networks; deterministic preparation written out and proved for these graphs (Section 8.2)
FCC stabilizer code	Defined on the ideal lattice; distance and decoder [11]	Built on grown networks: intact checks counted, exposed bonds located, exact parameters of grown patches; distance at most 3 at every size, so protection is local, cell by cell (Section 8.3)
Earlier implementation	Bulk $K = 12$ fraction from the published simulator, which bonds in batches [2]	The gap is traced to two implementation choices (Appendix B)

Outline. The paper has three parts that share one object, the grown network. The growth simulations produce the networks, and the quantum analysis is computed on them, so the defects that growth leaves are the defects that the code records. Section 2 comes first because it concerns the initial condition of every simulation, a triangle of Bell pairs. It states the postulate behind the model, with a verification map that is the measure-and-repair step proved in Section 8.2, and it proposes why growth starts from a triangle. The postulate is the source of the rules: the growth rules, and with them the error-correcting lattice,

follow the conditions it sets, and the simulations test their consequences. Parts of the section remain proposals, and Section 2.6 sorts its claims by status, from proved results to proposals. Section 3 lists the rules of the model. Section 4 derives the lift rate from the volumes of acceptable positions: if operator weights are proportional to these volumes, lift is rarer than stitch by a factor proportional to the bond tolerance δ. Section 5 gives the methods and the runs. Section 6 shows that clusters grown from a triangle are about 98% FCC in their interior, that this is insensitive to the lift rate when lift is rare, and that growth passes through no global $K = 4$ stage. Section 7 shows that new sheets start at step edges, that from one tetrahedron stitch generates only FCC sites, that a patch started on one close-packed sheet grows as twinned FCC whose domains coarsen with height while one started on an ABC substrate grows as a single variant, and that none of 2400 remnant nodes settling on a one-way front were trapped in a void pointing against growth. Section 8 shows that the grown network obeys an entanglement area law, that node parity checks turn it deterministically into a code state, and that the checks of the FCC stabilizer code, whose distance is at most 3 and which therefore protects the lattice locally, cell by cell, are intact around most interior nodes. Section 9 tests robustness, including placement off the close-packed sites, Section 10 discusses the results and their limits, and Section 11 concludes. Appendix A shows why the unit-bond condition is needed, and Appendix B compares the present simulator with that of Ref. [2]. Two interactive 3D views go with the paper. A growth viewer replays both kinds of growth node by node, with the same rules and parameters as this paper [14]. A second view shows how a new sheet starts at a step edge [15]. The simulator, scripts and data are in Ref. [16].

2. Before the Seed: Where the First Triangle Comes from

Every cluster in this paper grows from a three-node triangle; the patches of Section 7 stand for regions far from that seed. The growth rules say what happens after that triangle exists. They do not say what came before it, why it is a triangle, or why the rules are quantum. This section answers these questions. Its claims range from proved results to proposals about initial conditions, in the tradition of Refs. [17,18], and their status is set out claim by claim in Section 2.6 at its end. The growth rules that the simulations implement follow the conditions set out here (Section 2.5); the simulations of Sections 6–9 test the consequences of those rules, not the postulate itself.

2.1. What Existed Before the Triangle

Before the first bond there is no lattice. In the SSM, distance is a count of bonds, and a clock is a cycle of checks on bonds. So before the first bond there is no distance, no volume and no time. The questions usually asked of an initial state do not apply. Density is mass per volume, and there is no volume, so the state was neither dense nor dilute. "Point-like" and "infinite" are lattice words. The state had no size.

What did exist is a set of degrees of freedom with no position: two-level systems defined only by their algebra. We call them proto-qubits. A proto-qubit is not a particle since particles are defects of the finished lattice [2], and it has no location since location is a lattice property. It becomes a located qubit, one end of a Bell pair, once a bond gives it a place; each Bell pair then carries one qubit of the code [11] (Section 8.2).

The state also needs no energy budget. Energy conservation follows from time-translation symmetry, and there is no time before the first clock. Entanglement itself has no fixed energy cost: entangled and product states can have the same energy, and an isolated entangled state persists at no cost under unitary evolution. This is a discrete form of the zero-energy argument of Tryon [18]. It is an interpretation, and no result below uses it.

2.2. What Can Exist There: The Persistence Postulate

Asking what chose the laws of physics seems to start a regress: every selecting law needs its own selector. The SSM starts instead from one postulate:

> *What exists is what can persist.* What cannot persist cannot become part of an experienced or observable universe.

The postulate can be read at two strengths. The weak, epistemic reading says that we can observe only structures that persist long enough to leave records. It is hard to dispute, but it says nothing about what those structures are. The strong, ontological reading says that only structures able to persist exist at all. It is metaphysical, and the weak reading does not establish it. The SSM adopts a definite reading between the two:

> (P0) Every structure that enters an observable record is built from relations that are fixed points of a verification map: one round of checking and repair returns them unchanged.

(P0) says more than the weak reading because it fixes the form of observable structure: checked relations. It says less than the strong reading because it makes no claim about configurations that never enter a record. It is a postulate, not a result.

The verification map. A postulate about fixed points is empty until the map is specified. In the SSM the map is concrete. On a network of Bell pairs, let Φ be one round of the measure-and-repair protocol of Section 8.2: every node measures the parity of the qubits it holds, and the nodes with outcome -1 are paired by a fixed pairing rule and repaired by strings of edge-qubit operators $\bar{X}_e$ (defined in Section 8.2). *The fixed points of* Φ *are exactly the states in which every node check has value* $+1$*, and* Φ *maps any state of the Bell-pair subspaces onto such a state in one round*, so $\Phi \circ \Phi = \Phi$. Proof: every output of Φ has all node checks $+1$ (Proposition (i)–(ii) of Section 8.2), so every fixed point lies in that subspace; and on that subspace every outcome is $+1$ and no repair is applied, so Φ acts as the identity. In this sense the checked code state of Section 8.2 is a persistent structure. Φ measures only node checks: it detects an X error on any bond but a Z error only through a loop check, which Φ does not measure (Section 8.3). This makes (P0) a statement with content. It does not make it a derivation: it specifies what persistence means in the SSM, and the results below are proved about this map.

Transient and metastable structures. Persistence is not all or nothing in a growing network, and the simulations show several grades of it. A rejected candidate leaves no trace. A vacancy is enclosed once no triangle can fill it, and then persists for good (about 0.3% of interior sites, Section 6). Twin domains and tilted grains are metastable. They persist locally, compete with their neighbors, and are overgrown or take over: in patches started on one sheet, twin imbalance grows with height and one variant wins (Section 7), and tilted grains died out in 9 of 15 runs and took over in 6. These structures shape the final crystal even when they vanish. Three measures quantify that influence: the twin imbalance against height, the takeover statistics of tilted grains, and the density of the code's X checks against height, which rises as twin boundaries disappear (Section 8.3). Because every node records the triangle it was built on, the full lineage of the final crystal could also be traced back through structures that vanished; we leave that analysis to future work. The postulate concerns what persists in the finished, checked structure, not what appeared on the way.

The argument then has a fixed order, and only part of it is postulated:

possible configurations
↓ (a)
persistence condition (P0)
↓ (b)
stable relational structures
↓ (c)
checked quantum network
↓ (d)
growing geometric lattice

Steps (a) and (b) are the postulate. Step (c) uses the conditional result of Section 2.3, which excludes classical theory only under two further premises and does not single out quantum theory, and the construction of Section 8.2; that the network is quantum remains a proposal. Step (d) is the particular SSM realization, stitch and lift, whose consequences Sections 6–8 test by simulation. The results of this paper concern steps (c) and (d). The postulate has known relatives. Quantum Darwinism selects the states whose records survive copying [19], and constructor theory founds physics on which transformations are possible [20].

What persists: checked bonds of one length. In the SSM a bond persists when it meets two conditions. It must be checked. In the Bell-pair reading (Section 2.4) a single pair already checks itself; in the edge-qubit reading of the code, a bond is checked by the others when it lies on a closed loop of bonds, so that an error on it is detected and can be repaired. And, once the network is placed in three-dimensional space, it must have the common length L; δ is the tolerance within which a bond is detected. The two conditions apply at different stages. Before any embedding a bond is only an adjacency, and "equal" bonds simply means that each joins two nodes; the length condition has no meaning yet. The length L and the tolerance δ belong to the embedding of the network in space, which begins with the first triangle. In this paper δ is a free parameter of that embedding, not a quantity derived from the state before it (Section 4). We propose that formation filtered by these two conditions is what builds the embedded network after the first triangle and, through the checks, maintains it. Before it, the selection of Section 2.4 uses adjacency alone. The simulations implement an idealized form of it, described in Section 2.5.

2.3. Why Quantum Mechanics

The postulate on its own does not select quantum theory. Classical systems also have fixed points, stable attractors, conserved quantities, lasting correlations and error-correcting codes. A selection result needs further conditions on the state space and on which correlations count as relations. We first give three motivations, and then a conditional result that states those conditions openly.

Relation without geometry. Before space exists, a relation between two degrees of freedom cannot be carried by anything in between. Entanglement is a relation that needs no medium, no distance and no history. Constructions of geometry from entanglement [1,21,22] then read as the promotion of that relation into space.

Variation with heredity. Building structure by selection needs variation, to try new configurations, and heredity, to keep the ones that work. A deterministic rulebook has heredity without variation, and pure noise has variation without heredity. Quantum theory has both: outcomes are random, but the law that governs them never changes, and unitary evolution keeps what exists.

Self-checking from relations alone. A structure that must check itself before any carrier exists can use only its relations. Quantum stabilizer states do this, and Section 2.4 shows the smallest one in the reading used for the code.

A conditional result. Two premises make the setting of Section 2.1 precise.

(P1) *Closed world.* Before the lattice there is nothing outside the system: no environment and no observer whose ignorance could define a probability mixture. The state of the whole is therefore pure.
(P2) *Relation is correlation.* Two parts are related, and so can become adjacent, only if their joint state is correlated (nonzero mutual information).

Lemma 1. *In classical probability theory the states of a composite system are probability distributions on the product of the parts' configurations. Its pure states are point masses, and every point mass is a product of point masses on the parts [23]. So every pure classical state has zero mutual information between any two parts.*

Consequence. If (P1) and (P2) hold, a world before geometry that has any relations at all needs pure states of composite systems that are not products. Classical probability theory has none, so it is excluded; quantum theory has them, the entangled states. This shows less than it may seem. Once (P1) and (P2) are granted, Lemma 1 is immediate, and the conclusion is close to a restatement of the premises. Persistence enters only through the requirement that relations exist. The content is in the premises, and each excludes a classical option by assumption. (P1) treats a probability mixture in a closed world as ignorance; a classical theory with objective chance could have a closed world in an objectively correlated mixed state, and (P1) rules that out. (P2) counts only correlations as relations; a classical pure state can encode a relation in the definite values of its parts, such as two bits that agree, and (P2) does not count that as a relation. Reasons for the premises are given above, but they are premises. Boxworld, the theory of Popescu–Rohrlich correlations, which are stronger than quantum ones, also has correlated pure states [23,24], so the result does not single out quantum theory; boxworld is excluded instead by its trivial reversible dynamics [25], in which nothing can be grown.

Lemma 2. (Bell pairs are exclusive while they form)**.** *If qubits A and B are in a pure maximally entangled state such as* $|\Phi^+\rangle$*, the joint state of A, B and any other system C is* $|\Phi^+\rangle\langle\Phi^+|_{AB} \otimes \rho_C$ *[26,27]. A qubit therefore belongs to one Bell pair at a time, and a node with K bonds needs K qubits, as in the SSM. The lemma constrains how bonds form. It does not by itself make the lattice sparse. After the node checks of Section 8.2, the entanglement becomes shared around loops, and pairs of qubits that do not share a bond can then be correlated too (in the triangle, all pairs are). So correlation in the checked state does not define adjacency; adjacency is the set of bonds that formed. The sparsity of the finished lattice, twelve neighbors per node, comes from the geometry of the growth rules: exclusion and the kissing number (Section 3).*

Lemma 3. (The triangle holds its information only in relations)**.** *The triangle state of Section 2.4 is pure, yet each of its three qubits, taken alone, is maximally mixed. All its information lies in the relations between parts, and its checks read those relations. By Lemma 1, no pure classical state has this property.*

We checked all three lemmas numerically (Section 5). These arguments connect to the reconstructions of quantum theory, which derive it from short lists of informational axioms [28–30]. Table 2 reads the six principles of Ref. [29] as conditions for persistence. Rival rulebooks are also known [31]. Quantum theory with real amplitudes fails when independent parts are composed, and experiments have ruled it out [32,33]; higher-order interference is tightly bounded by experiment [34]. Classical probability theory satisfies five of the six principles of Ref. [29] and fails only purification, the requirement that every mixed state be part of a pure state of a larger system. Lemma 1 is the reason: a correlated pure state is what a purification is. If persistence, made precise, implied purification

in full, then together with the other five principles the theorem of Ref. [29] would give finite-dimensional quantum theory. That step, and a justification of (P1) and (P2), are open.

Table 2. The six informational principles of Ref. [29] read as conditions for persistence. The readings are proposals (Section 2.6).

Principle	Content	Persistence Reading
Causality	No signaling from the future	Completed links stay completed; keeping structure needs a fixed past
Perfect distinguishability	Every state that is not completely mixed can be perfectly distinguished from some other state	Errors can be detected: orthogonal syndromes exist
Ideal compression	Lossless encoding exists	Structure can be encoded redundantly and recovered
Local distinguishability	States of a composite are distinguished by local measurements (equivalent to local tomography)	Local checks certify global structure, as in a stabilizer code on a lattice
Pure conditioning	Measuring part of a pure state leaves the rest pure	Reading a syndrome does not damage what it checks
Purification	Every mixed state has a pure extension, unique up to a reversible operation on the extension	Information is moved, never destroyed

2.4. How the First Triangle Forms, and Why It Lasts

Selection, not a process in time. Before the first clock there is no time, so the formation of the first triangle is not a sequence of events. Words such as "first" and "then" in this subsection describe logical order, not temporal order. Consider all configurations of entangled pairs among proto-qubits. A set of such pairs is a graph, and that graph is the first geometry: nodes are adjacent if a pair joins them. Under (P0), the configurations that can enter the observable world are those whose bonds are protected by checks. The question is which is the smallest.

Which reading. The answer depends on how bonds carry qubits, and we state the choice. In the Bell-pair reading (two qubits per bond), a single pair is already self-checking: it is stabilized by $X_{e,u}X_{e,v}$ and $Z_{e,u}Z_{e,v}$, and every single-qubit error flips one of them. So a single bond is already a persistent relation, and nothing smaller than one bond can carry one. What a single bond cannot do is check another bond. For that we use the reading of the code (Section 8.3): one qubit on each bond, a Z check on each node (the product of Z over its bonds), and as X-type checks all operators that commute with the Z checks. In this reading bonds check each other, and we ask for the smallest network in which every bond is protected. We assume simple graphs, in which two nodes share at most one bond. This is an assumption, not a result. It is consistent with the embedding that follows, in which a bond is a unit segment between two points, but before embedding it is imposed rather than derived; with parallel bonds allowed, two nodes joined twice would already form a checked loop.

The smallest checked network. An X-type operator commutes with every Z check exactly when each node meets an even number of its bonds, that is, when its bonds form closed loops. So the X-type checks available on a graph are its loops. An X error on a bond is always detected, by the Z checks at its two ends. A Z error is detected only if some loop passes through the bond. Every single-bond error is detected exactly when every bond lies on a loop, that is, when the graph has no bridge. The smallest simple graph without a bridge is the 3-cycle. This is an elementary fact of graph theory; the test suite confirms it (Section 5). (The FCC code of Section 8.3 keeps only the octahedral loops among all loops, which is why it can have exposed bonds and logical qubits.)

The triangle's checks are Z_1Z_2, Z_2Z_3 and $X_1X_2X_3$ on its three bond qubits (the third node check, Z_3Z_1, is the product of the first two). They fix a single state,

$$|\psi_\triangle\rangle = \tfrac{1}{\sqrt{2}}\big(|000\rangle + |111\rangle\big), \tag{1}$$

the three-qubit Greenberger–Horne–Zeilinger (GHZ) state. Every single-qubit error on it is detected since an X or Y error flips the node checks at both ends of its bond and a Z or Y error flips the loop check. In the Bell-pair reading the same state is obtained from three Bell pairs by one parity check at each node; it is written with two virtual qubits per bond (Section 8.2). We therefore claim less than "the first structure that can last": a single Bell pair also lasts. The triangle is the smallest network whose bonds check one another, and so the smallest seed from which a checked network can grow.

The triangle in the finished crystal. The triangle remains the unit of checking in the FCC code. The eight faces of an octahedron are triangles in two alternating classes, and every bond of the octahedron lies in exactly one face of each class. So each X check of the code, on the twelve bonds of an octahedron, is the sum of four triangular loops.

2.5. From the Triangle to the Lattice

The first clock. Once a checked structure exists, its cycle of checks defines an interval. We propose that the triangle's check cycle is the first clock; before it there is only logical order. A rule that changed from one application to the next could not keep any structure, so on this reading the uniformity of physical law is a condition for persistence. The Page–Wootters mechanism [35], since illustrated in the laboratory [36], is a known way for time to emerge from entanglement with a clock inside a state that is globally static. That the triangle plays this role is our proposal, not the theorem's.

Where stitch and lift come from. The two conditions of Section 2.2 fix which moves can add a node. A node joined by a single bond would hang from a bridge, a bond on no loop, and so would be unchecked (Section 2.4). A new node must therefore bond to at least two nodes that are already connected, and the shortest loop it can close is a triangle: it bonds to both ends of an existing bond. Once the network is embedded, each of these bonds must also have length L. The points at distance L from both ends of a bond form a ring around it; placing a node on this ring is stitch. The points at distance L from all three corners of a triangle are two apexes, one above and one below; placing a node there is lift. No point lies at distance L from all four corners of a regular tetrahedron since the only point equidistant from all four corners is the center, at $\sqrt{3/8}\,L \approx 0.61L$, so lift is the most constrained placement on mutually bonded nodes. Placements that bond to more nodes at once, such as the apex of an octahedron over a square, are special points on a stitch ring (as the lift apexes also are), and the rules reach them through stitch and the bonding rule. Stitch and lift are therefore the two elementary moves that add a node whose bonds all have length L and close the shortest loops, triangles, when they form: stitch bonds to two mutually bonded nodes, lift to three. Ref. [2] gave the ring, the apex pair and the empty four-sphere intersection from unit spheres alone; the loop condition adds why a new node needs at least two bonds, and the octahedral apex shows that four-sphere intersections are not always empty. The conditions do not fix where on the stitch ring the node goes. The rules take the in-plane point across the edge, a choice discussed next.

What the growth rules take from the postulate, and what they idealize. The growth rules of Section 3 take the two conditions of Section 2.2 from the postulate. A candidate is kept only if all its bonds have length L to within δ. And every bond the rules create lies on a loop when it forms: stitch bonds a new node to both ends of an existing bond, closing a triangle; lift bonds it to the three corners of a triangle, closing three; and any further bond joins two nodes that are already connected. The grown network therefore

never contains a bond outside every loop, which is why it has no bridge (Section 8.3). The FCC code uses only octahedral loops as checks, so in that code a bond on a triangle can still be exposed. The rules take the length condition exactly: a new node goes where its bonds have length L, and δ is only the window within which bonds are detected and checked. For lift this fixes the site, the apex at unit distance from all three corners. For stitch, unit distance from the two edge ends leaves a ring, and the rules choose its in-plane point, a site of the close-packed lattice. These choices are rules of the SSM, and they are motivated by the postulate: they are what gives the network stability and error correction. With equal bonds every interior node sits on complete octahedra, so its cell carries the checks of the code; off these sites the bonds have unequal length, the deviations grow from layer to layer, and the network loses both the crystal and the checks (Section 9). Off-site nodes would therefore not persist. The simulations show this direction, that without the rules there is no stability and no code; they do not derive the rules from a dynamics. The rules are also an idealization in one further way: the relative weights of stitch and lift are taken proportional to the volumes of their acceptance regions (Section 4), which assumes that all positions near the front are equally likely to be tried. The simulations model creation in this idealized form. They model maintenance only through the unit-bond condition applied at placement (Section 3.4); continued checking of the finished lattice is the role of the verification map above.

Growth from the seed. With the triangle in place, the rules of Section 3 take over. Stitch builds new nodes from triangles alone, by $\mathbf{x} = \mathbf{a} + \mathbf{b} - \mathbf{c}$. The first lift completes a tetrahedron, whose three edges at one corner are primitive vectors of the FCC lattice (Section 7.3). So the seed triangle and the first lift fix the lattice and its orientation: every node that stitch adds from triangles on that lattice lies on it. In the simulations this held for all but 17 of 485,073 such stitches, while off-lattice lifts start twin grains (Section 7.3). The simulations correct some expectations an origin story might have. Growth does not complete a flat sheet before it goes up. After a few lifts it spreads in three dimensions, by stitching on the triangles that join neighboring sheets (Section 7.2). It passes through no global $K = 4$ stage (Section 6). And the stacking order needs no separate explanation because stitch forces it (Section 7.3). This proposal does not supply the remnant nodes whose trapping Section 7.7 studies. In the SSM they come from an earlier disordered phase [2], which lies outside both this section and the simulations.

2.6. Status of This Section

Table 3 separates what this section proves from what it proposes. The section derives no rate for forming the first triangle and gives no dynamics for it. It makes no observational prediction about the state before the lattice. Two results would count against it: a derivation of quantum theory from a weaker principle, which would make persistence redundant, and a reason to reject premise (P1) or (P2), which would reopen the classical case.

Table 3. Status of the claims in Section 2.

Claim	Status
The fixed points of the measure-and-repair map Φ are exactly the states with every node check $+1$, reached in one round (Section 2.2)	Proved, from Section 8.2
With one qubit per bond and simple graphs, the smallest network in which every bond lies on a loop is the triangle (Section 2.4)	Elementary graph theory; the reading is a stated choice
The triangle's checks fix the three-qubit GHZ state, which detects every single-qubit error and whose single qubits are maximally mixed	Proved

Table 3. *Cont.*

Claim	Status
Bell pairs with one parity check per node give the edge-qubit code state, prepared deterministically for every outcome	Proved (Section 8.2); checked numerically
Each octahedral X check of the FCC code is the sum of four triangular loops	Proved
Stitch and lift are the elementary moves that add a node whose bonds have length L and close triangles when they form (Section 2.5)	Elementary geometry, given the two conditions; the point chosen on the stitch ring is a rule of the model
Equal bond lengths at close-packed sites give stability and the checks of the code; off-site placement loses both (Section 9)	Shown by simulation; motivates the rules, does not derive them
Every bond created by the growth rules lies on a loop when it forms (Section 2.5)	Proved from the rules; confirmed for every node of two grown clusters
If (P1) and (P2) hold, the theory needs non-product pure states, so classical probability theory is excluded (Section 2.3)	Conditional; close to a restatement of the premises
A qubit belongs to one Bell pair at a time (monogamy)	Standard result [26,27]; constrains formation, not the sparsity of the finished lattice
(P0): observable structure is built from relations that are fixed points of a verification map (Section 2.2)	Foundational postulate; not derived
Formation filtered by checking and, after embedding, by length within δ builds and maintains the network	Proposal; simulated with exact sites and in idealized form (volume weights); repair under continuing noise not modeled
δ and L belong to the embedding in space	$L = 1$ sets the unit; δ is a free parameter in this paper
Persistence selects quantum theory among all theories	Open; would follow if persistence implied purification, with the other five principles of Ref. [29]
The informational axioms read as conditions for persistence (Table 2)	Proposal; a research program
The triangle's checks define the first clock	Proposal; Page–Wootters is a known mechanism [35]

3. Model

The model rests on five rules, listed in Table 4: unit-length bonds, two growth operators, hard-shell exclusion, proximity bonding and a unit-bond condition. A kinetic rule says how sites are drawn. Lengths are in units of $L = 1$.

The rules are few, and four of them (rules 1 and 3–5) state one condition in three ways. Rule 1 says that bonds have length L. Rules 3 and 4 set the window $\delta = 0.05$ around it within which bonds are detected. Rule 5 restores it after placement. The free parameters are therefore δ and the lift weight, which Section 4 expresses through δ. In the sweeps δ enters only through w; r_{ex} and r_{b} stay at their defaults unless stated. The operators place nodes at close-packed sites. This is a rule of the model, motivated by the stability and error correction it gives (Sections 2.5 and 9); we do not derive it from a dynamics. The other choices, namely lift orientation, exclusion radius, bonding radius, kinetic scheme and the settings of rule 5, are varied in Sections 6 and 9. Within the tested ranges, and for exclusion radii on the plateau from $0.58L$ to $0.98L$, none of them changes the results.

Table 4. The rules of the model and their parameters.

Rule	Statement	Parameter
1. Bonds	Every bond is a Bell pair of length L	$L = 1$
2. Growth operators	Stitch: new node in the plane of a triangle, across one edge. Lift: new node at the apex of a regular tetrahedron above or below a triangle	Lift orientation (Section 3.5)
3. Exclusion	A candidate is rejected if any node lies closer than r_{ex}	$r_{ex} = 0.95L$
4. Bonding	An accepted node bonds to every node within r_b	$r_b = 1.05L$
5. Unit-bond condition	A placed node with three or more bonds moves to bring its bonds closest to L	Cap $0.02L$, 4 Gauss–Newton steps; used where stated
Kinetics	Sites are drawn at random, stitch with weight 1 and lift with weight w per direction	$w = (\sqrt{2}/\pi)\delta$ (Section 4)

3.1. Bonds and Initial Conditions

Every bond is a Bell-pair bond of length L. A cluster starts from three nodes joined by three bonds, which form an equilateral triangle of side L (Section 2.4). A slab starts from one complete close-packed sheet, periodic in x and y (Section 3.6).

3.2. Growth Operators

Section 2.5 shows why these two operators follow from the conditions of Section 2.2. With lift in both directions, every triangle offers five growth sites. Stitch offers three, one per edge. It places a node in the plane of the triangle, across that edge, on the side away from the third corner, at distance $(\sqrt{3}/2)L$ from the edge midpoint. Lift offers two. It places a node on the normal through the centroid, at the tetrahedral height

$$h = \sqrt{2/3}\, L \approx 0.8165\, L, \tag{2}$$

one above the triangle and one below (Figure 2a). A candidate is rejected if any node lies closer than the exclusion radius r_{ex} ($0.95L$ by default). An accepted node bonds to every node within r_b ($1.05L$ by default). Together these defaults set the bond tolerance, $\delta = 0.05$. New bonds make new triangles, and these add new sites.

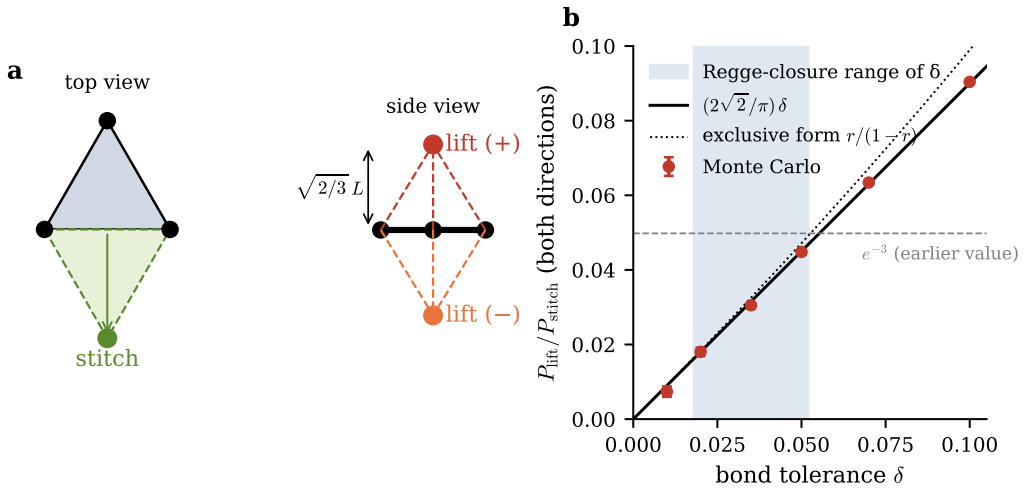


Figure 2. Growth operators and the lift rate. (**a**) Stitch places a node in the plane of a triangle across one edge (top view); lift places a node at height $\sqrt{2/3}\,L$ above or below the centroid (side view, triangle edge-on). (**b**) Ratio of lift to stitch acceptance volumes against bond tolerance δ, both lift directions: leading-order result $(2\sqrt{2}/\pi)\delta$ (solid), exclusive form $r/(1-r)$ (dotted) and Monte Carlo

integration (4×10^7 points per value; Table 5 gives higher-statistics values with error bars). The shaded band, $\delta = 1.8$–5.2%, is the range bracketed heuristically by the Regge deficit (Section 4); δ is a free parameter. The dashed line is the earlier assignment e^{-3}.

A triangle lies either within a sheet or between two neighboring sheets. A between-sheet triangle has two corners in one sheet and one in the next. Such triangles are faces of the tetrahedra and octahedra between sheets. They lie in the inclined close-packed planes of the crystal. Stitch and lift act on both kinds, and every sheet stays flat. The step-edge view [15] shows these triangles in 3D.

Table 5. Monte Carlo check of the acceptance volumes, with standard errors. Counting is non-exclusive: lift points also count for stitch. Theory is the leading order in δ; the exclusive form of the ratio is given for comparison.

	$\delta = 0.05$ (2×10^8 Points)		$\delta = 0.02$ (4×10^8 Points)	
Quantity	**Monte Carlo**	**Theory**	**Monte Carlo**	**Theory**
Stitch volume, one edge (L^3)	0.06281 ± 0.00009	0.06283	0.01011 ± 0.00003	0.01005
Lift volume, above ($10^{-3}L^3$)	1.395 ± 0.014	1.414	0.0959 ± 0.0026	0.0905
Lift volume, below ($10^{-3}L^3$)	1.403 ± 0.014	1.414	0.0903 ± 0.0025	0.0905
Lift volume, both ($10^{-3}L^3$)	2.798 ± 0.020	2.828	0.186 ± 0.004	0.181
Ratio, both directions	0.0445 ± 0.0003	0.0450	0.0184 ± 0.0004	0.0180
Exclusive form $r/(1-r)$	–	0.0471	–	0.0183

3.3. Kinetics

Sites are drawn at random by weight. Each stitch site has weight 1, and each lift site has weight w per direction. At each step, lift is chosen with probability $wN_\mathrm{L}/(wN_\mathrm{L} + N_\mathrm{S})$, where N_L and N_S count the open lift and stitch sites, and stitch otherwise. A site of the chosen kind is then picked uniformly. Section 4 derives $w = (\sqrt{2}/\pi)\delta \approx 0.0225$ for a bond tolerance $\delta = 0.05$. Apart from the unit-bond adjustment at placement (Section 3.4), nodes never move, and exclusion only grows. So a blocked site stays blocked, and a used site is filled. Each site is tried once and then removed. For comparison with earlier work, we also use a legacy rule, in which each step picks lift with a fixed probability p whenever a lift site is available. In both schemes a node's index equals its creation time. The first t nodes are thus the exact state of the cluster at time t. A cluster stops at N nodes, and a slab stops after a set number of layers' worth of nodes.

3.4. Unit-Bond Condition

After placement, a node with three or more bonds is moved to make its bonds as close to L as possible. It minimizes

$$E(\mathbf{x}) = \sum_j \left(|\mathbf{x} - \mathbf{p}_j| - L\right)^2 \tag{3}$$

over its bonded partners $\mathbf{p}_j$, using four Gauss–Newton steps. This is the direct reading of the rule that every bond has length L. The move is undone if it exceeds $0.02L$ or breaks exclusion or bonding. The cap and the number of steps were fixed once and were not tuned, and the results do not depend on them (Section 9). The condition is also needed for stability. Without it, position errors grow by about 30% per layer (Appendix A). All slab runs (Section 7), the lattice check of Section 7.3 and the quantum analysis of Section 8 use it. The cluster sweeps of Section 6 do not, and Section 9 shows that this changes them only slightly. We call growth without the condition the *bare rule*. Every caption of a table or figure with simulation results in Sections 5–9 says which form was used.

3.5. Lift Orientation Settings

We study three lift orientations. With *both*, every triangle offers a lift above and a lift below. With *axis*, a triangle offers only the lift on the side whose normal points toward $+z$ (both, if the triangle is vertical). This is one global lift orientation. With *radial*, a triangle offers only the lift that faces away from the seed.

3.6. Two Growth Geometries

A cluster models the start of growth. The crystal grows outward from a seed in every direction, and sheets form both above and below the seed plane.

A slab models a patch far from the seed. There the growth front is so large that it is flat on a local scale. It moves in one direction, with crystal behind it and empty space ahead. In the SSM, the observable universe would be such a patch. The slab's starting sheet, at $z = 0$, is the crystal already grown (Section 7.4 also uses a three-layer ABC substrate). No node may be placed below it, and it offers lifts only upward. The slab is periodic in x and y, with an $n_x \times n_y$ rectangular cell of a close-packed sheet (n_y even). Its periodic sides stand for a front much wider than the patch.

The growth viewer [14] shows both geometries node by node. The simulations test growth geometry in each case. They do not test whether our universe is such a patch.

4. The Lift Rate from Tolerance Volumes

Lift is rare because it must meet one more distance condition than stitch. Each extra condition costs one factor of the bond tolerance δ. Earlier work fixed a per-step lift probability $p = e^{-3}$ [2]. Here we derive the rate, under the assumption stated below.

4.1. Acceptance Volumes

A bond is accepted if its length lies within $L(1 \pm \delta)$. So a new node has a small volume of acceptable positions. We take each operator's weight to be proportional to that volume. This assumes that every position near the front is equally likely to be tried. It is an assumption of the model. The calculation below gives the volumes, but it does not show that weights should be proportional to them. The simulation does not try random positions either: it draws among the ideal sites with these weights (Section 2.5).

The scaling follows from the shapes. Points at distance L from one node form a thin shell. Two shells, as in stitch, meet in a thin ring. The ring is long, but thin in cross-section, so its volume scales as δ^2. Three shells, as in lift, meet in two small blobs. These are thin in every direction, so their volume scales as δ^3. The ratio is therefore linear in δ. The calculation below gives the prefactor.

A stitch needs unit distance to both ends of an edge. Two shells of radius L and thickness $2\delta L$, with centers L apart, meet in a thin torus. The torus runs around a circle of radius $(\sqrt{3}/2)L$. The shell normals meet at 60° there, so the cross-section is a parallelogram of area $(2\delta L)^2 / \sin 60^\circ$. To leading order,

$$V_{\text{stitch}} = 2\pi \cdot \tfrac{\sqrt{3}}{2} L \cdot \frac{(2\delta L)^2}{\sin 60^\circ} = 8\pi\,\delta^2 L^3. \tag{4}$$

A lift needs unit distance to all three corners of a triangle. Three shells meet in two small parallelepipeds, one above and one below. Each has volume $(2\delta L)^3$ divided by the determinant of the three unit normals. For a regular tetrahedron that determinant is $1/\sqrt{2}$. Counting both directions,

$$V_{\text{lift}} = 2 \times 8\sqrt{2}\,\delta^3 L^3 = 16\sqrt{2}\,\delta^3 L^3, \qquad \frac{P_{\text{lift}}}{P_{\text{stitch}}} = \frac{V_{\text{lift}}}{V_{\text{stitch}}} = \frac{2\sqrt{2}}{\pi}\,\delta \approx 0.90\,\delta. \tag{5}$$

https://doi.org/10.3390/quantum8040121

Per direction the ratio is $(\sqrt{2}/\pi)\delta \approx 0.45\delta$. The factor of two comes from the two directions. At $\delta = 0.05$ the ratio is 4.50% for both directions and 2.25% per direction. In the simulation this gives weight 1 per stitch site and $w = (\sqrt{2}/\pi)\delta$ per lift site. The stitch weight counts the whole ring, although stitch then places the node at the in-plane point on it. Each lift apex is at distance L from both ends of every edge, so the lift regions lie inside the stitch tori. If stitch means placement with exactly two bonds, the ratio becomes $r/(1-r)$, with $r = (2\sqrt{2}/\pi)\delta$. This changes 4.50% to 4.71% at $\delta = 0.05$, a change of order δ^2. The simulations use the first form.

4.2. Monte Carlo Check

We sampled points uniformly in the box $[-1.2, 2.2] \times [-1.2, 2.2] \times [-1.2, 1.2]$ around a unit triangle with corners A, B and C: 2×10^8 points at $\delta = 0.05$ and 4×10^8 at $\delta = 0.02$. The counting rule is as follows. A point counts for stitch on edge AB if its distances to A and B both lie within $L(1 \pm \delta)$. It counts for lift if its distances to all three corners do. The counts are not exclusive: a lift point also counts for stitch on each edge. So the Monte Carlo ratio estimates $V_{\text{lift}}/V_{\text{stitch}}$ as defined in Equation (5), not the exclusive form $r/(1-r)$. Errors are counting (Poisson) standard errors. Table 5 gives the results. At $\delta = 0.05$ the ratio is 0.0445 ± 0.0003. This is 1.5σ from the leading-order value 0.0450 and 8σ from the exclusive form 0.0471, so the data agree with the leading-order formula and are resolved from the exclusive form. All volumes agree with the leading-order formulas within about 2σ.

4.3. What Sets δ

The derivation does not fix δ. In this paper δ is a free parameter, with default 0.05. For orientation only, one scale is the strain needed to close the Regge deficit [37]. Five regular tetrahedra around a shared edge leave a gap of $2\pi - 5\arccos(1/3) \approx 0.128$ rad. Stretching the shared edge closes it at 5.2% strain. Stretching the five rim edges closes it at 1.8%, with less elastic energy. This arrangement is icosahedral, not FCC. In FCC each edge is shared by two tetrahedra and two octahedra, whose dihedral angles add to exactly 2π. So the Regge numbers say how much strain a tolerance must allow for growth to try a non-FCC local arrangement. They are not a property of the FCC crystal, and they do not derive δ. They suggest values of a few percent, between about 1.8% and 5.2%, for which the lift-to-stitch ratio lies between about 1.6% and 4.7%. The structure is not sensitive to this choice. FCC is best at the lowest lift rates (Section 6), and every lift rate in this range gives FCC.

The weight w is not the share of placements that are lifts. Most lift sites are blocked when drawn. At $w = 0.0225$ per direction, only 0.66% of placements are lifts. The legacy setting $p = 0.05$ gives about 1%. We report the realized lift fraction throughout.

5. Methods

Implementation. The simulator is written in Python 3 and compiled with Numba [38]. Neighbor search uses a uniform cell grid, periodic in x and y for the slab. Site lists use swap-removal, so each attempt costs $O(1)$ apart from the local search. A cluster of 10^6 nodes grows in about 20 s on one core. Runs are deterministic for a given seed, and each node records the triangle it was built on. A test suite checks the structure classifier on perfect FCC, hcp and icosahedral environments. It also checks that bond lengths lie in $[r_{\text{ex}}, r_{\text{b}})$, that exclusion holds and that runs reproduce exactly. Six further tests reproduce the published parameters of the FCC code [11] (Section 8.3), check the Bell-pair to edge-qubit construction, its entropy formula and its deterministic preparation (Section 8.2), check the exact triangle results of Section 2.4, check that every new node closes a triangle when it forms (Section 2.5), and check Lemmas 1–3 of Section 2.3.

Structure classification. We classify each node by common neighbor analysis (CNA) [39,40], with cutoff $(1+\sqrt{2})L/2 \approx 1.207L$. A node with 12 neighbors is FCC if all bonds have signature (4,2,1). It is hcp if six have (4,2,1) and six have (4,2,2), and icosahedral if all have (5,5,5). Any other node is labeled *other*. CNA uses geometric neighbors within the cutoff, while tetrahedra are counted on the bond graph. The median bond-length error is at most about $10^{-8}L$ (bare rule; Section 9), so a fixed cutoff is safe; results do not change for cutoffs from $1.15L$ to $1.25L$. We also count the bonded tetrahedra (4-cliques) that contain each node. A node in perfect FCC is in 8, an icosahedral node in 20, and a neighbor of a single vacancy in 6.

Bulk and depth. A node is bulk if at least 42 nodes lie within $2L$, the population of the first three FCC shells [2]. This test still admits the first layers below a rough surface, so we also compute depth. Non-bulk nodes have depth 0. A node at depth d lies within $1.2L$ of a node at depth $d-1$, with d the smallest such value. Depth is thus measured from the nearest non-bulk node, so the surroundings of an interior vacancy pocket also have small depth. We call nodes at depth 4 or more the interior; the code analysis of Section 8.3 uses depth 3 or more and says so.

Slab analysis. The stacking axis z is the substrate's [111] direction. A node is on-lattice if its height is within $0.05h$ of a layer height kh. Its lateral registry is the nearest of the close-packed positions A, B and C, with the substrate as A. In a layer with k not divisible by 3, an ABC domain has registry k mod 3, and its mirror twin ACB has $(-k)$ mod 3. The twin imbalance is $|f_{\mathrm{ABC}} - f_{\mathrm{ACB}}|/(f_{\mathrm{ABC}} + f_{\mathrm{ACB}})$, in bins of 10 layers. Twin fractions count only on-lattice nodes. A bonded tetrahedron is *up* if three corners share a layer and the fourth is one layer above. It is *down* if the fourth is one layer below. A perfect FCC or hcp crystal has equal numbers of each.

Runs. Table 6 lists the runs. With 2 sheet-formation slabs (Section 7.2), 6 lattice-check clusters (Section 7.3), 8 remnant-trapping slabs (Section 7.7), 24 robustness runs (Section 9) and 32 runs for the quantum analysis (Section 8), the study has 696 production runs. Eight of them, the exclusion-sweep runs at $r_{\mathrm{ex}} = 0.50L$ and $0.52L$, overflowed the triangle buffer and stopped early. They are left out, so 688 runs are analyzed. The four runs at $r_{\mathrm{ex}} = 1.02L$ are included because their result is that no growth is possible. Smaller diagnostic and test runs are not counted: the 12 cluster runs and 10 slab runs of the unit-bond sensitivity tests (Section 9), the 42 off-site placement runs (Section 9), the 32 substrate runs (Section 7.4), the six depth-bias slabs and the four layer-pair slabs of Section 7.6 and the comparison runs of Appendix B. Uncertainties, where shown, are standard errors over seeds; several tables in Sections 6–8 also give the range over seeds. Methods specific to Sections 7.2–7.7 and to Section 8 are given there.

Table 6. Simulation sets of Sections 6 and 7. The cluster sets use the bare rule, and the slab set uses the unit-bond condition. The exclusion sweep includes 8 runs, at $r_{\mathrm{ex}} = 0.50L$ and $0.52L$, that stopped early and are not analyzed.

Set	Geometry	Size	Varied	Seeds/Point	Runs
Size scaling	cluster, both	$N = 10^3$–10^6	N (7)	32–4	156
Lift sweep	cluster, both	$N = 10^5$	$w = 5\times10^{-4}$–1 (11)	8	88
Legacy rule	cluster, both	$N = 10^5$	$p = 0.01$–0.4 (6)	8	48
Exclusion sweep	cluster, both	$N = 3\times10^4$	$r_{\mathrm{ex}} = 0.50$–$1.02L$ (27)	4	108
Orientation, size	cluster, axis/radial	$N = 10^3$–10^6	N (4)	16–4	88
Orientation, lift	cluster, axis	$N = 10^5$	w (11)	8	88
Slab	periodic slab, axis/both	$32^2, 48^2, 64^2$	width, orientation	8	48
Total					624

6. Results I: Growth from a Seed

This section covers clusters grown from the three-node seed. Its sweeps use the bare rule, without the unit-bond condition (Section 3.4); Section 9 gives the small effect of adding it. The single run behind the lifted-node counts of Section 6.5 uses the unit-bond condition.

6.1. The Interior Is FCC

At $N = 10^6$, nodes at depth 4 or more are 98.5–98.7% FCC in every depth shell from 4 to 11. Their hcp fraction is below 10^{-4}, and no icosahedral node appears at any size (Figure 3a,b). Among bulk nodes, 94% belong to exactly 8 bonded tetrahedra and 3.4% to 6. Six is the count for a vacancy neighbor. Each vacancy has twelve neighbors, so this is consistent with about 0.3% vacant sites. The interior defects therefore look like isolated vacancies rather than stacking faults, and they show as isolated pockets in Figure 3a. A vacancy can form in a growth that only adds nodes: a site is enclosed after every triangle that could have filled it has been used.

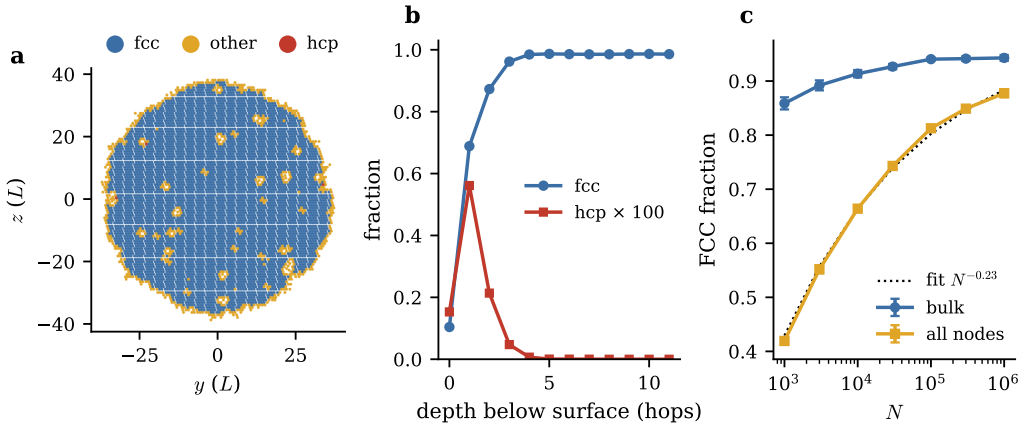


Figure 3. Growth from a seed. (**a**) Central cross-section ($|x - \bar{x}| < 0.45L$) of a 3×10^5-node cluster colored by common neighbor analysis: the interior is FCC, with isolated vacancy pockets and a disordered surface a few layers thick; hcp is essentially absent. (**b**) FCC and hcp ($\times 100$) fractions against depth below the surface, $N = 10^6$ (4 runs). (**c**) FCC fraction against N for bulk and all nodes, with the fit $1 - f \propto N^{-0.23}$ to all nodes. Bare rule.

6.2. Finite-Size Scaling

The FCC fraction rises with size (Table 7). The bulk fraction levels off near 94%, because the 42-neighbor test counts the first three layers below the surface. Those layers (depths 1–3) are 69–96% FCC. The depth-resolved value of 98.5–98.7% is the better measure of the interior. The non-FCC fraction over all nodes falls with an effective exponent of 0.23 over $N = 10^3$–10^6 (Figure 3c). A shell of fixed thickness w in a ball of radius R holds a fraction $1 - (1 - w/R)^3$ of the nodes, which falls as $N^{-1/3}$, the exponent used in Ref. [2], only when $w \ll R$. Here R grows from about $6L$ to $55L$, and a partly disordered shell 2.5–3L thick then gives an effective exponent of 0.26–0.27 over the same range. Interior defects, whose share rises slightly with N (from 0.6% to 0.9% of all nodes between 10^5 and 10^6), lower it further. The value 0.23 is therefore consistent with a disordered region a few layers thick.

Table 7. Finite-size scaling of clusters grown with the bare rule. Bulk and all-node FCC are the mean ± standard error over seeds; the range column gives the lowest and highest bulk FCC over seeds. Isotropy is the ratio of smallest to largest eigenvalue of the position covariance.

N	Seeds	Bulk FCC	Range	Bulk hcp	All-Node FCC	Lift Fraction	Isotropy
10^3	32	0.859 ± 0.011	0.72–0.92	0.0032	0.419 ± 0.007	0.0082	0.62
3×10^3	32	0.892 ± 0.009	0.80–0.95	0.0024	0.552 ± 0.007	0.0070	0.71
10^4	32	0.914 ± 0.007	0.80–0.96	0.0021	0.664 ± 0.006	0.0068	0.80
3×10^4	32	0.927 ± 0.005	0.85–0.95	0.0018	0.743 ± 0.005	0.0066	0.86
10^5	16	0.940 ± 0.003	0.90–0.95	0.0009	0.813 ± 0.004	0.0066	0.91
3×10^5	8	0.941 ± 0.004	0.92–0.95	0.0009	0.849 ± 0.006	0.0066	0.93
10^6	4	0.943 ± 0.005	0.93–0.95	0.0008	0.877 ± 0.007	0.0066	0.95

Why the cluster is round. The FCC interior of Figure 3a sits inside a roughly spherical cluster. The round shape is not imposed. No rule refers to a center or a radius, and a site is open wherever the front has a triangle. Stitch and lift act on triangles of every orientation, so the front advances at about the same speed in every direction, and the cluster rounds out as it grows. Its isotropy, the ratio of the smallest to the largest eigenvalue of the position covariance, rises from 0.62 at 10^3 nodes to 0.95 at 10^6 (Table 7); the lift share in rounding it is given in Section 6.3. A cluster is the seed region of a crystal, not a whole universe. Regions far from the seed are modeled separately, as patches (Section 7). The curved surface does bias quantities averaged over all nodes, because surface layers are disordered and make up a large share of a small cluster. We avoid this bias in two ways. Interior values are resolved by depth below the surface, and they are flat from depth 4 to 11. And the patches use periodic slabs, which have no curved surface.

6.3. Insensitivity to the Lift Rate

FCC is insensitive to the lift rate when lift is rare. The bulk FCC fraction is flat at 96–98% for realized lift fractions below about 0.2%, and it falls steadily above that (Table 8). The legacy rule gives the same curve against the realized lift fraction (Figure 4a). Rare lift does not leave disconnected sheets, contrary to the expectation in [2]. One lift per 10^4 placements still gives 97.5% FCC. The legacy setting $p = 0.05 \approx e^{-3}$ gives a realized lift fraction of 0.99% and bulk FCC 0.908.

The reason is that stitch acts on every triangle, including the between-sheet triangles that a lift creates. After a few lifts, growth spreads in three dimensions by stitching on these triangles (Section 7.2). Lift then has two remaining roles. It rounds the cluster: the isotropy is 0.73 at the smallest weight and 0.91 at the geometric one. It also creates stacking faults, and the hcp signal grows with the lift rate.

Table 8. Lift sweep at $N = 10^5$, 8 seeds per point. Bare rule.

w (per Direction)	Realized Lift Fraction	Bulk FCC	Bulk hcp	Isotropy
0.0005	0.0001	0.975 ± 0.006	0.0001	0.73
0.001	0.0003	0.964 ± 0.007	0.0004	0.79
0.002	0.0006	0.962 ± 0.007	0.0004	0.81
0.005	0.0015	0.971 ± 0.005	0.0002	0.83
0.0113	0.0033	0.940 ± 0.006	0.0017	0.86
0.0225 ($\delta_w = 0.05$)	0.0066	0.939 ± 0.004	0.0009	0.91
0.045	0.0132	0.893 ± 0.004	0.0020	0.90
0.1	0.0282	0.811 ± 0.006	0.0048	0.89
0.2	0.0554	0.691 ± 0.005	0.0107	0.86
0.5	0.129	0.505 ± 0.004	0.0286	0.82
1.0	0.225	0.370 ± 0.003	0.0576	0.82

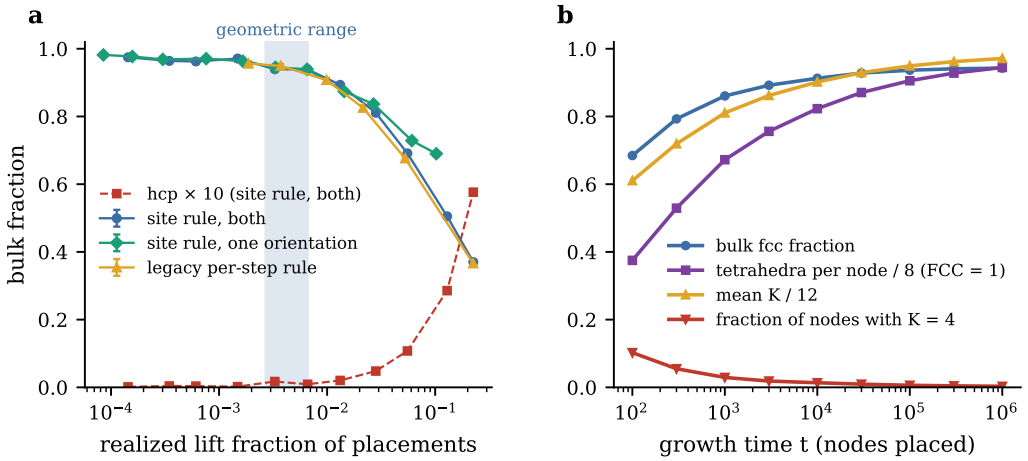


Figure 4. (**a**) Bulk FCC fraction against the realized lift fraction at $N = 10^5$ for the site-weighted rule with lift in both directions, with one global orientation, and for the legacy per-step rule; dashed: hcp ($\times 10$). The shaded band is the realized lift fraction for the geometric weight with δ between 1.8% and 5.2% (the Regge range of Figure 2, rounded in the plot). FCC is best when lift is rarest. (**b**) Growth history of the $N = 10^6$ runs: bulk FCC, tetrahedra per node relative to FCC, mean coordination relative to 12, and the fraction of nodes with $K = 4$, all against the number of nodes placed. No tetrahedra-rich stage occurs. Both panels use the bare rule.

6.4. Exclusion Radius

Results are nearly identical for every r_{ex} from $0.58L$ to $0.98L$. Bulk FCC is 0.925–0.929 and mean coordination is 11.14–11.16 at $N = 3 \times 10^4$. At $r_{ex} = 0.54L$ and $0.56L$, no FCC forms and mean coordination exceeds 12. At $1.00L$ the FCC fraction falls to 0.77, and at $1.02L$ no growth is possible. The lower edge lies between $0.56L$ and $0.58L$. This matches $L/\sqrt{3} \approx 0.577L$, the distance from a triangle's centroid to its corners. The plateau is flat because candidates almost never fall at a distance in this range from an existing node. It is a check of the rules, not the phase transition reported in Ref. [2]. Runs at $0.50L$ and $0.52L$ overflowed the triangle buffer and are left out.

6.5. Growth History: Tetrahedral Units, but No $K = 4$ Stage

Node index equals creation time, so every earlier state of a run is available exactly. In the four $N = 10^6$ runs, all structural measures rise steadily toward their FCC values (Table 9, Figure 4b). The tetrahedron count never exceeds the FCC value of 8 per node, and no icosahedral node forms at any time. Growth thus goes from the sheet-like seed straight to FCC, with no tetrahedra-rich stage.

Locally, every lift makes a tetrahedral unit. The lifted node bonds to the three corners of its triangle. The four nodes are then all bonded to each other and form the complete graph K_4. These units carry growth into three dimensions. In this local sense the tetrahedron is the natural step between a sheet and the bulk.

This local unit differs from coordination $K = 4$, which is what K means in this paper. The new corner of a tetrahedron usually has coordination 3. In one run of $N = 10^5$ (seed 0, with the unit-bond condition), 173 of 665 lifted nodes (26%) had 3 bonds when placed, and 58 (9%) had 4. The rest land in partly built pockets with 5 to 12 bonds.

Nor is coordination 4 a stage that nodes stay in. Bonds are added one at a time, so a node placed with fewer than 4 bonds must pass through 4 on its way to 12. In the same run, only 33% of nodes do so, since most start with 5 or more bonds. Those that reach $K = 4$ leave it as fast as any other value (Table 10). Nodes stay longest at $K = 10$ and 11, just before they fill up.

Table 9. Growth history of the $N = 10^6$ runs (4 seeds). Bare rule.

t (Nodes Placed)	Bulk FCC	Mean K	Tetrahedra per Node	Fraction with $K = 4$
30	–	4.80	0.63	0.158
100	0.685	7.32	3.00	0.103
300	0.793	8.63	4.23	0.055
10^3	0.861	9.73	5.38	0.029
10^4	0.913	10.82	6.58	0.014
10^5	0.937	11.39	7.24	0.006
10^6	0.943	11.66	7.56	0.003

Table 10. Median number of placements a node spends at each coordination (nodes in the first half of an $N = 10^5$ run, with the unit-bond condition).

K	2	3	4	5	6	7	8	9	10	11
Placements at K	526	453	430	410	423	455	496	579	722	1095

6.6. Effect of Lift Orientation

The three lift orientations grow the same crystal (Table 11). The axis variant is slightly cleaner, and almost all of the difference comes from its halved lift rate. At equal realized lift fractions the two give the same structure. Lift in both directions gives 0.939 bulk FCC at 0.66%, and axis gives 0.939 at 0.65% (Figure 4a). One-sided lift leaves no mark on the cluster's shape. Under axis lift, 49.8% of nodes lie above the seed plane, and the cluster spans $z = -56L$ to $+56L$ at $N = 10^6$. Stitching on between-sheet triangles extends growth downward as fast as upward.

Table 11. Lift orientation at $N = 10^6$, 4 seeds each. Interior FCC is the mean over depth shells 4 to 8; ranges over seeds in parentheses. Bare rule.

Orientation	Interior FCC (Depth 4–8)	All-Node FCC	Lift Fraction	Nodes with 6 Tetrahedra
Both	0.986 (0.985–0.986)	0.877 (0.856–0.886)	0.0066	0.034
Axis ($+z$ only)	0.992 (0.991–0.992)	0.896 (0.883–0.910)	0.0033	0.023
Radial (outward)	0.985 (0.984–0.986)	0.884 (0.880–0.888)	0.0061	0.032

7. Results II: Growth Far from the Seed

A patch started on one close-packed sheet grows as twinned FCC. The two mirror-image stacking variants start evenly mixed. They coarsen with height until one dominates, and the winner is random. All slab runs use the unit-bond condition and $w = 0.0225$ per direction.

7.1. Overview

Table 12 summarizes the 48 production slabs. Their interior FCC fraction, 0.84–0.92, is below the cluster value, because twin boundaries leave layers near the base about 91% full (Section 7.4) and tilted grains lower it further in some runs (Section 7.5). All production slabs start from one close-packed sheet; on a three-layer ABC substrate the twins do not form (Section 7.4).

Table 12. Slab runs, 8 seeds per row. A tilted grain *appears* when the on-lattice fraction in any 10-layer bin drops below 0.98, and *takes over* when it averages below 0.5 in the top three bins. Interior FCC counts layers 2 to two below the top complete layer. For the last column see Section 7.6. All slabs use the unit-bond condition.

Width	Layers	Orient.	Grain Appears	Takes Over	Final Imbalance	ABC Wins	Int. FCC	(down − up)/Total
32^2	600	axis	2	0	1.00	4	0.922	6.5×10^{-4}
32^2	600	both	5	2	0.81	2	0.873	10.1×10^{-4}
48^2	400	axis	2	2	0.75	4	0.868	9.2×10^{-4}
48^2	400	both	3	0	0.85	7	0.884	7.9×10^{-4}
64^2	300	axis	2	1	0.37	6	0.843	9.5×10^{-4}
64^2	300	both	1	1	0.74	4	0.853	11.5×10^{-4}

7.2. How New Sheets Form

New sheets do not start from lifted nodes. They advance from step edges, mostly by stitching on the triangles that join neighboring flat sheets. This carries the stacking order from one sheet to the next.

We followed every new layer in one 32×32 and one 64×64 slab (99 layers each, axis lift; Table 13). The front is rough, so it always has step edges where an island of the top sheet ends. At such an edge, a triangle can have two corners on the island's edge and one in the sheet below. Stitching it across its top edge puts a node one layer higher, on the island (Figure 5). That node's own triangle is then stitched across its slanted edge. This places the next node one bond along, so a bonded row grows along the step. Triangles that bridge the row and the island add more rows. Flat triangles form in the new sheet once two rows meet. New layers thus advance from step edges, and their first in-sheet triangle appears only after most of the layer is placed. This is step-flow growth. An interactive version of Figure 5 is available [15].

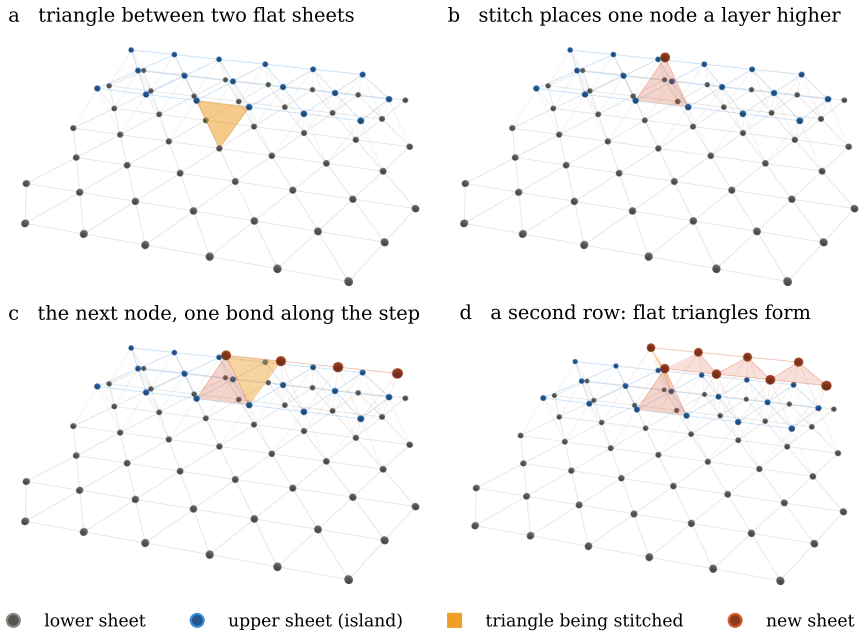


Figure 5. How a new sheet starts at a step edge. All sheets are flat. (**a**) At the edge of an island of the upper sheet, a triangle (amber) has two corners in the upper sheet and one in the lower sheet. (**b**) Stitching it across its top edge places a node one layer higher. (**c**) That node's own triangle is stitched across its slanted edge, placing the next node one bond along the step; repeating this lays a bonded row. (**d**) Triangles bridging the row and the island add a second row, and the two rows bond into flat triangles of the new sheet.

This mechanism explains the insensitivity of FCC to the lift rate below about 0.2% lift fraction (Section 6.3), because sheets are not made by lifts. It also explains why twin domains, once seeded (Section 7.4), stay separate: a layer advances from many step edges at once. Each edge grows on its own, and the regions meet at twin boundaries.

Table 13. How layers start in a growing patch (one slab per width, unit-bond condition).

	32 × 32	**64 × 64**
Layers whose first node was made by a lift	2%	1%
Lifted nodes per layer	≈6 of 1024	≈24 of 4096
Placements before the layer's first in-sheet triangle	≈650 (0.6 layer)	≈2900 (0.7 layer)

7.3. Why the Stacking Is ABC

The geometry of stitch forces the stacking. Stitch is a point inversion, and from one tetrahedron it generates only the FCC lattice spanned by the tetrahedron's edges. So every node that stitch adds from triangles on one FCC lattice lies on that lattice, and only lift can create another stacking. This is a closure property of exact placement, not a selection among competing structures; the simulations measure how lifts and placement errors depart from it.

Stitch is a point inversion. For an equilateral triangle with corners **a**, **b** and **c**, the stitch apex across the edge from **a** to **b** is

$$\mathbf{x} = \mathbf{a} + \mathbf{b} - \mathbf{c}. \tag{6}$$

This is the corner **c** reflected through the midpoint of the edge. The simulation places the apex at $(\sqrt{3}/2)L$ from the edge midpoint. For an equilateral triangle this is the same point, and the unit-bond condition keeps triangles equilateral.

Closure under stitch. The statement needed is local because stitch acts only on bonded equilateral triangles. Lemma: let a regular tetrahedron have a corner **o** and edge vectors $\mathbf{e}_1$, $\mathbf{e}_2$, $\mathbf{e}_3$ from that corner, and let $\Lambda = \mathbf{o} + \mathbb{Z}\mathbf{e}_1 + \mathbb{Z}\mathbf{e}_2 + \mathbb{Z}\mathbf{e}_3$, which is an FCC lattice. Every node obtained from the tetrahedron by repeated stitches, each on a triangle of nodes already obtained, lies in Λ. Proof: the four corners lie in Λ, and if $\mathbf{a}, \mathbf{b}, \mathbf{c} \in \Lambda$ then $\mathbf{a} + \mathbf{b} - \mathbf{c} \in \Lambda$; the claim follows by induction on the number of stitches. The lemma explains why the stacking cannot be hcp. Consider a step edge (Figure 5a,b), with $\mathbf{b}_1$ and $\mathbf{b}_2$ in sheet B and **a** in sheet A below. The new node $\mathbf{b}_1 + \mathbf{b}_2 - \mathbf{a}$ lies two layers above A. It sits over the hollow of sheet B that is not above a node of A. This is the C position. ABA would need the new node directly above **a**, and no stitch can put it there. FCC is a Bravais lattice and hcp is not, as Ref. [2] also notes, which is why only the ABC stacking is closed under this operation.

One tetrahedron fixes the lattice. The three edge vectors at a corner of a regular tetrahedron are primitive vectors of the FCC lattice. Once a lift makes one tetrahedron, every node that stitch adds from it lies on that FCC lattice. The stacking along each close-packed direction is then ABC. It is ACB, the mirror twin, if the first lift goes to the other side of its triangle.

Only lift leaves the lattice. Each triangle within a sheet has two lift apexes. One is a lattice site, and the other is the matching site of the mirror twin. A lift on the second side starts a twin or a stacking fault. A lift on a between-sheet triangle can start a grain with a different orientation. The two competing hollow positions above a sheet are $L/\sqrt{3}$ apart. With $r_{ex} > L/\sqrt{3}$, exclusion lets only one of them be filled. This is the lower edge of the exclusion plateau in Section 6.4.

Consequences. The rule explains the hcp trend of Section 6.3 and the twins and grains described below. The hcp fraction grows with the lift rate. A patch started on one sheet grows as twin domains (Section 7.4) because its flat starting sheet fits both twins; each first

lift picks one, and stitch then carries it forward. Finally, the exclusion radius must exceed $L/\sqrt{3}$.

Numerical check. This check uses clusters grown with the unit-bond condition, where the seed fixes one lattice. We recorded the triangle each node was built on. We then tested every placement against the FCC lattice set by the seed triangle and the first lift, in six clusters of $N = 10^5$ (Table 14). There were 485,073 stitches built on triangles lying fully on the lattice. All but 17 (0.004%) landed on the lattice. All 17 were in one high-lift run, and we have not traced them. Lifts from such triangles left the lattice 20–29% of the time. Raising the lift weight to 0.2 raised these off-lattice lifts about eightfold relative to the mean of the four $w = 0.0225$ seeds (six- to ninefold, seed by seed) and the bulk hcp fraction about sevenfold (six- to elevenfold). In seed 2, a large part of the cluster lies off the seed lattice. It belongs to a twin grain started by an early off-lattice lift.

Table 14. Placements tested against the seed's FCC lattice ($N = 10^5$, unit-bond condition).

w	Seed	On Seed Lattice	Stitches from On-Lattice Triangles	Of These Off	Lifts On/Off	Bulk hcp
0.0225	0	0.994	98,761	0	526/135	0.0009
0.0225	1	0.989	98,115	0	469/163	0.0011
0.0225	2	0.618	47,367	0	229/92	0.0018
0.0225	3	0.992	98,567	0	496/130	0.0009
0.2	0	0.861	78,678	17	3464/1163	0.0101
0.2	1	0.745	63,585	0	2831/934	0.0071

7.4. Twin Domains and Their Coarsening

The first layer above a single-sheet substrate places nodes at B and C positions in about equal numbers. These seed the two twin variants, ABC and ACB. Near the base, domain boundaries leave layers about 91% full; higher up, as domains coarsen, layers fill to 97–99%. Every third layer is the exception: it is A in both variants and fills completely. Over greater heights, one variant slowly takes over (Table 15; Figure 6a,c). A 32-wide slab becomes a single variant by about layer 300. Wider slabs take longer. An imbalance of 0.5 is reached near layers 65, 115 and 210 for widths 32, 48 and 64. A power-law fit to these three points gives a height that grows roughly as width$^{1.5}$–width$^{1.7}$, depending on how the crossing is interpolated, so domains widen roughly as height$^{0.6}$–height$^{0.7}$. This is a preliminary and poorly constrained estimate. The two successive pairs of widths give different exponents (about 0.9 and 0.5), and a bootstrap over runs gives a 68% interval of 0.5–0.9 for the domain exponent (95%: 0.4–1.4). The slab width also caps the domain size, so the box size enters the estimate directly; slabs wider than 64 would be needed to separate the two. An unbounded patch would never become a single variant at finite height. Instead it forms single-variant domains that keep growing.

Table 15. Mean twin imbalance against height, 16 runs per width, with the unit-bond condition. Bins whose on-lattice fraction is below 0.9 (tilted grains) are left out, so 13–16 runs enter each point.

Height (Layers)	10	50	80	120	200	300
Width 32	0.23	0.45	0.55	0.77	0.94	1.00
Width 48	0.16	0.30	0.40	0.51	0.62	0.80
Width 64	0.11	0.18	0.23	0.32	0.48	–

The role of the substrate. The production slabs start from one complete close-packed sheet. The first layer above it can fill both kinds of hollow, B and C, and so seeds both twin variants from the start. To test whether the twins come from this start, we grew slabs on a substrate of three complete layers stacked ABC, with the same settings: widths 32 and 48,

200 layers, axis lift, unit-bond condition, 8 seeds each, alongside the same runs on one sheet (Table 16). On the ABC substrate the first bin is already a single variant, the stacking stays ABC throughout (variant purity 0.997), and the interior FCC fraction is 0.94–0.95, against 0.82–0.92 on one sheet. One of the 16 ABC-substrate runs showed a tilted grain. The twin domains, their coarsening and the lower FCC fraction of the production slabs are therefore specific to a single-sheet start. They describe a patch whose first layers have no stacking memory, not a general consequence of the growth rules.

Table 16. Slabs on one close-packed sheet and on a three-layer ABC substrate (200 layers, axis lift, unit-bond condition, 8 seeds per row; ranges over seeds in parentheses). Purity: the larger of the two twin-variant fractions, averaged over layers. First-bin imbalance: twin imbalance in layers 0–9.

Substrate	Width	Interior FCC	Purity	First-Bin Imbalance
One sheet	32	0.867 (0.820–0.918)	0.84 (0.69–0.96)	0.16 (0.05–0.36)
One sheet	48	0.841 (0.815–0.888)	0.67 (0.51–0.91)	0.13 (0.02–0.38)
ABC, 3 layers	32	0.949 (0.944–0.953)	0.997 (0.996–0.998)	1.00 (1.00–1.00)
ABC, 3 layers	48	0.947 (0.941–0.951)	0.997 (0.996–0.998)	1.00 (0.99–1.00)

The winning variant is random. In the six runs where a tilted grain takes over (defined in Table 12; Section 7.5), the winner is the variant that dominates the remaining on-lattice nodes at the top. ABC won 27 of 48 runs (56%), as expected for a fair coin, and lift orientation makes no difference. This follows from symmetry. A vertical mirror plane swaps ABC and ACB but leaves $+z$ unchanged. So no global lift orientation can favor either variant.

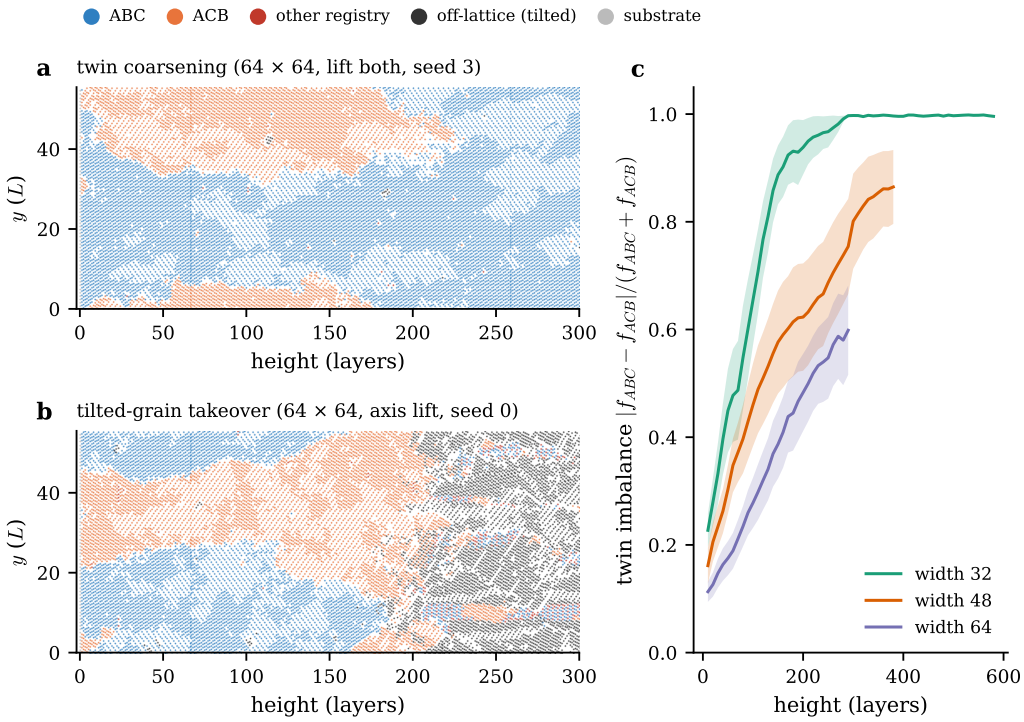


Figure 6. Growth of a patch. (**a**,**b**) Vertical sections ($0 \leq x < 0.9L$) through 64 × 64 slabs, colored by stacking variant (red: other registry). (**a**) The two twin variants ABC and ACB start mixed and coarsen until one dominates. (**b**) A tilted grain (black, off the substrate's layer heights) nucleates near layer 160 and takes over. (**c**) Mean twin imbalance against height for three slab widths, 16 runs each, leaving out bins taken over by tilted grains; shading is the standard error. All panels use the unit-bond condition.

7.5. Tilted Grains

In 15 of 48 slabs, a grain grew whose close-packed planes are inclined to the substrate. In the case we examined, it started when a lift on a between-sheet triangle placed a node off the substrate's layer heights. The grain then competed with the main crystal. It died out in 9 runs and took over in 6, which became polycrystalline (Figure 6b). At this sample size there is no clear effect of lift orientation. Axis lift gave 6 events and 3 takeovers, and both gave 9 events and 3 takeovers. In thin-film growth, grains whose fastest growth direction points along the growth direction tend to win [10]. This is a likely mechanism here too. The nucleation rate and the takeover criterion need more runs.

7.6. Tetrahedron Orientation

In the SSM the two orientations of tetrahedral voids correspond to matter and antimatter [2]. In a defect-free FCC crystal the two counts are equal by symmetry, so any asymmetry must sit in defects. A front that moves one way makes the two classes distinct, so the void population could carry an along-growth bias. We measured it directly in the slab, where depth below the front is clear. The bias is confined to the outermost layers (Table 17, Figure 7b). Below the top two layers there is no meaningful along-growth bias.

A small residual signal remains. Counted layer pair by layer pair, down tetrahedra exceed up by about 10^{-3} where twin boundaries are present, with a three-layer pattern that tracks the domains. In single-variant regions (layers 300–400 of 32-wide slabs), an excess of 4.2–5.3 $\times$ 10^{-4} remains in three of the four slabs checked (two per lift orientation); the fourth, taken over by a tilted grain, is left out. This is the same signal as the deep bins of Table 17. The excess is positive in every slab, with both lift orientations, so it is tied to the growth direction rather than the lift orientation. We have not found its cause. It may come from a defect that forms at the growth front, or from misclassified tetrahedra next to defects. It should not be read physically until this is settled.

Table 17. Along-growth void bias (up $-$ down)/total against depth below the outermost complete layer (6 slabs, 32 $\times$ 32 $\times$ 260, axis lift). Standard errors are 0.001–0.003 in the shallow bins and 0.0001 in the deep ones. Unit-bond condition.

Depth (Gaps)	0–1	2–3	4–5	6–7	8–9	10–11	12–19	20–49	50–99	100–200
Bias	$+0.017$	-0.001	-0.001	$+0.002$	$+0.002$	-0.002	0.000	-0.0003	-0.0004	-0.0006

7.7. Remnant Trapping and Matter–Antimatter Selection

None of the remnant nodes that we let settle on a one-way front was trapped in a tetrahedral void that points against growth. The void population itself is balanced, so any selection of one void class for matter must come from trapping. Geometry suggests why. On a flat close-packed surface, a remnant can rest only in a three-fold hollow, not on top of an atom. The hollows that the next layer covers become tetrahedral voids whose apex points along the growth direction (Figure 7a). We tested whether this holds on rough, realistic fronts.

Model. A remnant is a sphere that must stay at least $R = \sqrt{3/8}\,L \approx 0.612L$ from every lattice node. This is its closest approach inside a void. At a random time t, we drop a remnant at a random lateral position onto the surface formed by the first t nodes. It settles along $-z$, against the growth direction, in small downward steps that push it out of contact, until it comes to rest. The rest of the growth then runs unchanged, so the remnant acts as a passive tracer. We classify its final site by the nodes within $R + 0.04L$. Three below and one above is a tetrahedral void pointing along growth. One below and three above is a tetrahedral void pointing against growth. Three below and none above is the open hollow,

which becomes the octahedral side. Anything else is irregular. A later node closer than $R - 0.04L$ would mark a conflict, where the remnant would have blocked growth.

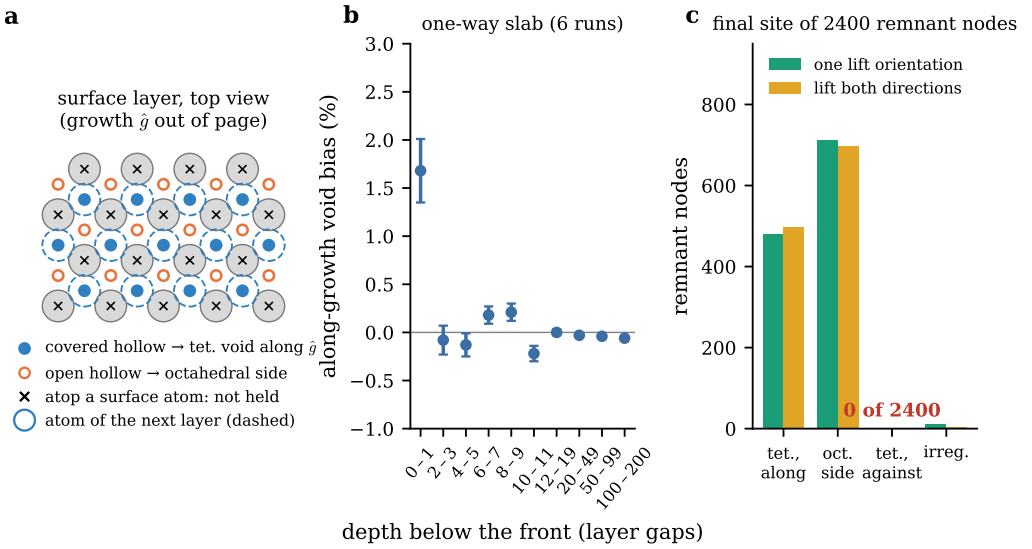


Figure 7. Void classes and remnant trapping. (**a**) Top view of a growing (111) surface. A remnant node can rest in either three-fold hollow but not atop an atom; the hollows that the next layer covers become tetrahedral voids pointing along the growth direction $\hat{g}$, and the others become the octahedral side. (**b**) Along-growth bias of the void population against depth below the front in one-way slabs: confined to the outermost layers. (**c**) Final sites of 2400 remnant nodes dropped onto one-way slabs: none ends in a tetrahedral void pointing against growth. Panels (**b**,**c**) use the unit-bond condition.

Validation. On a perfect ABC crystal, we dropped 40 remnants onto the third layer. Of these, 19 ended in voids along growth and 21 on the octahedral side. None ended against growth, as the geometry requires.

Results. Each run used a 32×32 slab of 150 layers with 300 remnants, with four seeds per lift orientation (Table 18, Figure 7c). Of 2400 remnants, 2394 came to rest on three contacts and 6 on two. Of the 2394, 977 ended in voids along growth, 1409 on the octahedral side, 7 with one node below and two above, and 1 with two below and one above; Table 18 counts these 8 and the 6 two-contact sites as irregular (14). No conflict occurred, so the passive-tracer approach is consistent. The two lift orientations give the same result. We examined the seven one-below-two-above sites one by one. In each, one node lies $0.52L$ below the remnant and the other two lie level with it, only $0.025L$ higher; all three were placed before the remnant arrived, and no node later covered it. These are ledge sites against a step, not voids pointing against growth. The split between along-growth voids (41%) and the octahedral side (59%) differs from the 50/50 of a perfect surface; we have not analyzed why.

Table 18. Final sites of remnant nodes dropped onto one-way slabs grown with the unit-bond condition.

Lift Orientation	Remnants	Tet., Along	Octahedral Side	Tet., Against	Irregular
Axis	1200	479	711	0	10
Both	1200	498	698	0	4
Total	2400	977 (41%)	1409 (59%)	0	14 (0.6%)

The result is largely built into the settling rule. A remnant that arrives from the growth side and moves against it cannot reach a void beneath a layer that is already in place, so a

void pointing against growth is out of reach. What the simulation adds is that rough fronts, steps and twin boundaries create no overhang that would let a remnant reach one. It does not test the premise behind the rule. The selection enters through the settling direction: the assumption that remnants arrive from the growth side and rest on the front. Remnants might instead be overgrown from the side, move by thermal motion or be captured from below. Whether they settle as assumed is a physical assumption of the model.

8. Results III: Quantum Structure of the Grown Network

The SSM reads every bond as a Bell pair. The simulation evolves no quantum state, but two quantum objects are defined by the bond graph alone. One is the network of Bell pairs itself. The other is the FCC stabilizer code of Ref. [11], which puts one qubit on every bond. Section 8.2 shows that they are related by one local operation: Bell pairs checked by one parity measurement per node form a code state on one qubit per bond. Both are stabilizer states or codes [41–43]. They depend only on which nodes are bonded, so the grown graph fixes them. Their defects are the crystal's defects. This section measures both on the grown networks. Everything here is an exact property of states defined on the graph. Nothing is time-evolved, and we do not simulate noise.

8.1. Bell Pairs and the Area Law

Give each bond two qubits, one at each end, in the Bell state $|\Phi^+\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$. The network state is then a product of Bell pairs. Take any region A of nodes. A bond with both ends in A, or both outside, adds nothing to the entanglement between A and the rest. A bond that crosses the boundary leaves one qubit in A in a maximally mixed state. So the entanglement entropy is exactly

$$S(A) = |\partial A| \ln 2, \tag{7}$$

where $|\partial A|$ is the number of bonds that cross the boundary. This is an area law [44], and its coefficient is the number of bonds crossing unit area. An ideal FCC has $6\sqrt{2}$ bonds per L^3. A bond of length L crosses a randomly oriented plane with probability proportional to $|\cos\theta|$, whose mean is $1/2$. A randomly oriented surface is therefore crossed by $3\sqrt{2} \approx 4.243$ bonds per L^2. A (111) plane is crossed by three bonds per node and $2/\sqrt{3}$ nodes per L^2, giving $2\sqrt{3} \approx 3.464$ bonds per L^2.

We counted the bonds crossing spheres about the center of four clusters of 10^5 nodes, for radii from $4L$ to $16L$ (Figure 8b). Above $R = 8L$ the count per unit area is 4.14 ± 0.03, with a spread over seeds from 4.06 to 4.18. The ideal value is $3\sqrt{2} = 4.243$, and ideal FCC balls cut in the same way give 4.23–4.24. The grown network thus has an area law with a coefficient about 2.5% below the ideal one. In patches we counted bonds crossing horizontal planes between layers (four 48×48 slabs, 80 planes each). The mean is 3.268 ± 0.003 per L^2, 5.7% below $2\sqrt{3}$ (Figure 8c). The deficit has a simple source. Each missing node removes the bonds it would have had. In the analyzed patch region (defined in Table 19) about 2.6% of lattice sites are empty, mostly at twin boundaries (Section 7). A missing node also leaves its neighbors under-coordinated, so the bond deficit is about twice the vacancy fraction. In clusters the deficit varies from seed to seed: 1.1%, 2.4%, 4.0% and 1.3% for $R > 8L$. It tracks the fraction of empty lattice sites in the analyzed region of each cluster (Table 19), which is 0.5%, 1.4%, 2.3% and 0.5%; seed 2 holds a large low-density region. The bond deficit is about twice the empty-site fraction, as in patches. The vacancy estimate of about 0.3% in Section 6 is smaller because it counts only bulk nodes, and the 42-neighbor test excludes the surroundings of larger empty pockets by construction. The entanglement density of the network is therefore a direct measure of its defects. These counts are for the Bell pairs alone; the node checks of Section 8.2 lower the entropy of each such region by

exactly ln 2 when the region and its complement are both connected, and leave the per-area coefficient unchanged.

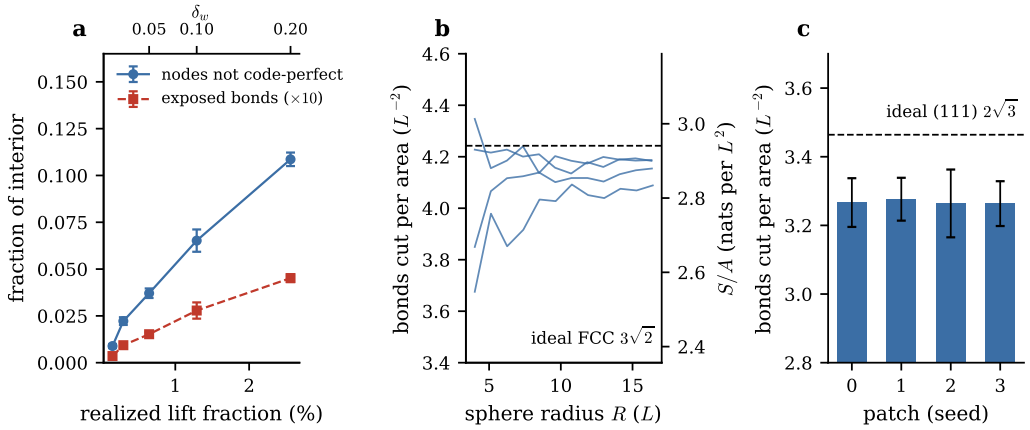


Figure 8. Quantum structure of the grown network. (**a**) Fraction of interior nodes that are not code-perfect and fraction of exposed interior bonds (×10; a Z error there goes undetected) against the realized lift fraction, $N = 10^5$, mean over 4 seeds; the top axis gives δ_w, the tolerance assigned to each lift weight by the volume relation (Section 8.4). (**b**) Bonds crossing a sphere of radius R per unit area in four clusters; with one Bell pair per bond the entanglement entropy per unit area is this number times ln 2 (right axis). Dashed: ideal FCC, $3\sqrt{2}$. (**c**) Bonds crossing (111) planes per unit area in four 48×48 patches (mean of 80 planes, bars: standard deviation); dashed: ideal, $2\sqrt{3}$. All networks use the unit-bond condition.

8.2. From Bell Pairs to Edge Qubits

The Bell-pair network of Section 8.1 has two qubits per bond, one held by each end node. The code of Section 8.3 has one qubit per bond. We now show that the second is obtained from the first by one local operation at each node: a parity check on the node's own qubits. This is the standard projected-entangled-pair construction of $\mathbb{Z}_2$ gauge states and of the toric code [12,13]; we write it out for the grown networks and add a deterministic preparation.

Each Bell pair is one edge qubit. Label the two virtual qubits of bond $e = (u, v)$ as $a_{e,u}$ and $a_{e,v}$. The Bell state $|\Phi^+\rangle$ lies in the two-dimensional subspace with $Z_{a_{e,u}} Z_{a_{e,v}} = +1$, spanned by $|00\rangle$ and $|11\rangle$. Define the isometry $V_e : |0\rangle \mapsto |00\rangle$, $|1\rangle \mapsto |11\rangle$. On this subspace an edge qubit has the operators

$$\bar{Z}_e = Z_{a_{e,u}} = Z_{a_{e,v}}, \qquad \bar{X}_e = X_{a_{e,u}} X_{a_{e,v}}, \tag{8}$$

and $|\Phi^+\rangle = V_e|+\rangle$. The Bell-pair network is therefore $\bigotimes_e V_e|+\rangle$: a product state of edge qubits, each in $|+\rangle$.

Each node checks its own bonds. Let node v apply the projector

$$P_v = \tfrac{1}{2}\Big(1 + \prod_{e \ni v} Z_{a_{e,v}}\Big), \tag{9}$$

which acts only on the qubits that v holds. Under the isometry, $\prod_{e \ni v} Z_{a_{e,v}} = \prod_{e \ni v} \bar{Z}_e = A_v$, the Z check of node v in Section 8.3. So

$$|\Psi\rangle \propto \prod_v P_v \bigotimes_e |\Phi^+\rangle = \Big(\bigotimes_e V_e\Big) \prod_v \tfrac{1}{2}(1 + A_v)\, |+\rangle^{\otimes N_E}, \tag{10}$$

where N_E is the number of bonds.

Result. $|\Psi\rangle = (\bigotimes_e V_e)|\psi_G\rangle$, where $|\psi_G\rangle$ is the edge-qubit stabilizer state with checks A_v on every node and $\prod_{e\in C}\bar{X}_e$ on every closed loop C. In the Z basis it is the equal superposition of all sets of bonds that meet every node an even number of times. *Proof.* $|+\rangle^{\otimes N_E}$ is stabilized by every $\bar{X}$ product. Such a product commutes with all A_v exactly when it meets every node an even number of times, that is, when it is a product over loops; these survive the projection, and the A_v are added. With N_V nodes and N_E bonds on a connected graph this gives $N_V - 1$ independent node checks and $N_E - N_V + 1$ independent loops, N_E in all, so $|\psi_G\rangle$ is a single pure state. It is nonzero since the empty set of bonds has amplitude $2^{-N_E/2}$ in $P_+|+\rangle^{\otimes N_E}$. For the triangle, $|\psi_G\rangle = (|000\rangle + |111\rangle)/\sqrt{2}$, the GHZ state of Section 2.4. The two readings used in this paper are thus related by the node projections: the checked Bell-pair network is the code state, written with two virtual qubits per bond. The Bell pairs are the resource. After the checks each bond keeps $Z_{a_{e,u}}Z_{a_{e,v}} = +1$, so its two qubits are perfectly correlated, but they are no longer a Bell pair: for the triangle, each bond's two qubits are an equal classical mixture of $|00\rangle$ and $|11\rangle$. The entanglement has become shared around loops. Across any cut it still amounts to one ebit (one maximally entangled pair) per crossing bond, less $n_A + n_B - 1$ (see below). We confirmed this directly for the triangle, and checked purity on random graphs and on a grown cluster.

Preparation as a physical process. The projectors P_v are not unitary. Applying them by post-selection, keeping only runs in which every node finds parity $+1$, succeeds with probability $2^{-(N_V-1)}$, which is exponentially small. A deterministic protocol avoids this. Every node measures its parity, and the outcomes are then repaired by $\bar{X}$ strings. Write the outcome of node v as $s_v = \pm 1$ and the set of nodes with $s_v = -1$ as D.

Proposition (deterministic preparation). On a connected graph:

(i) Every outcome pattern that can occur has $|D|$ even, and every such pattern can be repaired: there is a set of bonds T whose odd-degree nodes are exactly D, and applying $\bar{X}_T = \prod_{e\in T}\bar{X}_e$ removes every -1.

(ii) Every step preserves the encoded Bell-pair subspaces: the state keeps $Z_{a_{e,u}}Z_{a_{e,v}} = +1$ on every bond throughout.

(iii) For every outcome pattern, the repaired state is exactly $(\bigotimes_e V_e)|\psi_G\rangle$, with no phase, and it does not depend on which repair set T is chosen.

Proof. (i) On the Bell-pair subspaces, $\prod_v A_v = \prod_e \bar{Z}_e^2 = 1$, so the outcomes multiply to $+1$ and $|D|$ is even. Pair the nodes of D arbitrarily and join each pair by a path of bonds, which exists because the graph is connected. Let T be the set of bonds used an odd number of times. A path has odd degree only at its two ends, and degrees add modulo 2, so the odd-degree nodes of T are exactly D. Since $\bar{X}_e$ anticommutes with A_v exactly when v is an end of e, $\bar{X}_T$ anticommutes with A_v exactly when $v \in D$, and so turns every $s_v = -1$ into $+1$. (ii) Each parity measurement is a product of Z operators, and each $\bar{X}_e = X_{a_{e,u}}X_{a_{e,v}}$; both commute with every $Z_{a_{e,u}}Z_{a_{e,v}}$. (iii) Write $P_s = \prod_v \frac{1}{2}(1 + s_v A_v)$. The state after measurement is $P_s|+\rangle^{\otimes N_E}/\sqrt{p_s}$. Because $\bar{X}_T$ flips exactly the A_v with $v \in D$, $\bar{X}_T P_s = P_+ \bar{X}_T$, where $P_+ = \prod_v \frac{1}{2}(1 + A_v)$ is the all-$+1$ projector. And $\bar{X}_T|+\rangle^{\otimes N_E} = |+\rangle^{\otimes N_E}$. Hence $\bar{X}_T P_s|+\rangle^{\otimes N_E} = P_+|+\rangle^{\otimes N_E}$ for every allowed s. Taking norms, all allowed patterns have the same probability; there are 2^{N_V-1} of them and their probabilities sum to 1, so $p_s = 2^{-(N_V-1)}$, and the repaired state is $P_+|+\rangle^{\otimes N_E}/\sqrt{p_+} = |\psi_G\rangle$ exactly. Two repair sets with the same odd-degree nodes differ by a set of bonds of even degree everywhere, which is a product of loops and fixes $|\psi_G\rangle$ with eigenvalue $+1$. □

Physically, each node measures the parity of the qubits it holds, and each correction $\bar{X}_e$ is one single-qubit X gate at each end of a bond. The only non-local resource is classical: the outcomes must be communicated so that defects can be paired. Pairing nearby defects, for example by minimum-weight matching, keeps the strings short, but any pairing gives the

same state by (iii). The corrections can also be left unapplied and tracked as a known Pauli frame. We checked all three properties, with one repair set per pattern, with explicit state vectors on the two virtual qubits of every bond, for every allowed outcome pattern on three small graphs (the triangle, K_4 and a five-node graph with a chord). This is the repair step of a $\mathbb{Z}_2$ gauge theory, and it is the node-level checking of Section 2.2 written as an operation: random outcomes, repaired locally, always give the same state. Preparation is not yet persistence. The proposition shows that one round of checking and repair reaches the target state from any outcome. Persistence also needs the state to be kept under continuing disturbance, which requires specified dynamics of repeated verification and repair together with a noise model. Those dynamics, and the threshold below which they keep the lattice, are not given here.

Entanglement after the checks. For a region A of nodes, with n_A and n_B connected components of A and of its complement, the checked state has

$$S(A) = \big(|\partial A| - n_A - n_B + 1\big) \ln 2. \tag{11}$$

Each connected component's product of node checks fixes the parity of the bonds it shares with the cut, and these parities give $n_A + n_B - 1$ independent constraints on the $|\partial A|$ crossing pairs. For a connected region with a connected complement, such as the spheres and slabs of Section 8.1, $S(A) = (|\partial A| - 1) \ln 2$. The area-law coefficient of Section 8.1 is unchanged, and the checks lower the entropy by one constant $\ln 2$, the topological entanglement entropy of a $\mathbb{Z}_2$ gauge theory [45,46]. We confirmed the formula exactly, by stabilizer ranks over GF(2), for 75 regions on random graphs, the octahedron and a grown cluster of 300 nodes.

8.3. The FCC Stabilizer Code on the Grown Graph

The code. Ref. [11] defines a Calderbank–Shor–Steane (CSS) code on an ideal FCC. It is built like the toric code [47], with qubits on bonds and checks on nodes and voids. A qubit sits on every bond. Each node carries a Z check, the product of Z over its 12 bonds. Each octahedral void carries an X check, the product of X over the 12 bonds among its six corner nodes. On a periodic cube of side ℓ (in units of $L/\sqrt{2}$, so with $\ell^3/2$ nodes and $3\ell^3$ bonds) the code has $k = 2\ell^3 + 2$ logical qubits, $[[192, 130, 3]]$ at $\ell = 4$. We build the same checks on the grown graph: a Z check on every node, over the bonds it actually has, and an X check on every *complete* octahedron, six nodes joined by all 12 octahedral bonds. We find octahedra as pairs of nodes $\sqrt{2}L$ apart whose four common neighbors form a bonded square. Our rank code reproduces $k = 130$ and $k = 434$ for the ideal codes at $\ell = 4$ and 6.

Relation to the FCC code. $|\psi_G\rangle$ satisfies every node check and every octahedral loop check, so it is a code state of the FCC code. It is the particular code state in which the remaining loops, such as single triangles, also take the value $+1$. Those loops are logical X operators of the code, so the node projections prepare one definite codeword, with every logical qubit fixed. A single triangle has weight 3, so the X distance of the code is at most 3 at every lattice size, as in Ref. [11], where $d = 3$ at $\ell = 4$; where a bond is exposed the Z distance is 1.

Local protection. Distance 3 does not grow with size, but it is enough to protect each cell. In an ideal FCC an X error on a bond flips the two node checks at its ends, and a Z error flips the X checks of the two octahedra that contain it. Both patterns belong to that bond alone and lie within one cell, so any single error is found and repaired from the checks around it. Errors in different cells are repaired independently, and no global decoding is needed for them, much as in edge computing, where each node handles its own data. A cell fails only if two errors arrive in it before its checks are read, which for independent errors at rate p has probability of order p^2 at every lattice size. Ref. [11] gives

a minimum-weight matching decoder for the ideal code and finds a tenfold coding gain at $p = 10^{-3}$. The lattice is the union of its cells, so it is protected cell by cell. This is local protection, not the protection of a topological code whose distance grows with size: with a fixed per-cell failure rate, the chance that some cell fails grows with the number of cells, so the checks must be read round after round. The verification map of Section 2.2 does this for the node checks; a full round would also read the octahedral checks. On a grown network, protection holds wherever the checks of the cell are intact. The simulations below therefore measure how many cells are protected.

The checks commute. An X check and a Z check overlap on the bonds of the octahedron that meet the node. Each corner of a complete octahedron meets exactly four of its twelve bonds, and every other node meets none. The overlap is always even, so $H_X H_Z^{\mathsf{T}} = 0$ over GF(2) on every graph that contains complete octahedra, grown or not; this is automatic, and our numerical check of it served only to test the code. Defects remove checks, but they cannot make the code invalid.

What a defect does. In an ideal FCC every node has 12 bonds and is a corner of 6 complete octahedra, and every bond lies in 2 of them. We call a node *code-perfect* when it has 12 bonds, 6 complete octahedra and 8 bonded tetrahedra as in an ideal FCC (in our data the first two conditions imply the third). A bond is *exposed* when it lies in no complete octahedron. A Z error on an exposed bond commutes with every X check and goes undetected. It is not a stabilizer either. A product of Z checks over a set of nodes is Z on the bonds leaving that set, so a single bond is a stabilizer only if it is a bridge of the graph. The grown graph has no bridge, because the seed triangle (or the starting sheet of a patch) has none, every new node closes a triangle when it joins (Section 2.5), and bonds are never removed. Where exposed bonds occur, the code as grown has distance 1, and the cells around them are not protected. Every exposed bond still lies on some loop, so adding such loops as X checks would detect these errors, at the cost of logical qubits; with every loop as a check the code would encode none (Section 8.2). The positions of exposed bonds are known from the graph, so they could also be removed or treated as gauge qubits. We do not analyze these options here.

Clusters. Table 19 gives the results for two regions. The first is a structural interior, nodes at depth 3 or more (one layer shallower than the depth 4 used in Section 6). It leaves out the surface and the surroundings of vacancy pockets, so it gives the favorable figure. The second is purely geometric: the ball within 0.8 times the median node radius, which keeps every defect. At the geometric lift weight, 96.3% of interior nodes are code-perfect and 0.15% of interior bonds are exposed. In the geometric region these become 89% and 1.1%, with a large spread over seeds. The spread comes from twin grains started by early off-lattice lifts (Section 7.3). Seed 2, which has a large twin grain, has the lowest values. Raising the lift rate degrades the interior code steadily (Figure 8a). The region values are noisier and flat between $\delta_w = 0.022$ and 0.05.

Patches. The code in a young patch started on one sheet is dominated by twin boundaries. In 48×48 slabs of 100 layers, 76% of nodes in the geometric region (layers 3 to about 85) are code-perfect, the region has 86% of the ideal number of X checks, and 3.3% of bonds are exposed. The code heals as the patch grows. In two 32×32 slabs of 300 layers, the number of X checks per lattice site rises from 0.68 in the first ten layers to 0.85–0.90 near layer 70 and to 0.93–1.00 above layer 200, as one twin variant takes over (Section 7).

Exact code parameters. For a direct count we cut four 16×16 patches flat after 12 layers and computed the ranks of H_Z and H_X over GF(2). Bonds are taken as all node pairs within r_{b}, which is the simulator's bonding rule. Table 20 gives the results. Each grown patch is compared with a perfect ABC patch of the same size and shape. Defects, mostly twin boundaries, remove 4–8% of nodes and 22–39% of X checks. All remaining X checks are

independent, as in the perfect patch. The number of logical qubits per physical qubit is higher, 0.679–0.690 against 0.667. This follows directly from $k = n - \operatorname{rank} H_Z - \operatorname{rank} H_X$ with fewer X checks, and it is not a gain or a quality measure. The number of logical qubits itself falls, and wherever a bond is exposed the code has a weight-1 logical operator. We did not compute code distances of these patches. At $w = 0.005$, under a quarter of the geometric weight, the patch still loses 27% of its X checks. Single seeds at other weights thus suggest that the loss in young patches is set mainly by twin boundaries rather than by the lift rate.

Table 19. The FCC stabilizer code on grown networks ($N = 10^5$ clusters with the unit-bond condition; patches $48 \times 48 \times 100$). Interior: depth ≥ 3. Region: the ball within $0.8\times$ the median node radius (clusters), or from layer 3 up to 8 layers below the 90th height percentile (patches). X checks are per lattice site of an ideal FCC, which has one. Values are means over the number of seeds in parentheses, with the range over seeds below each mean.

	Lift Fraction	**Code-Perfect Nodes**		**Exposed Bonds**		***X* Checks**
		Interior	**Region**	**Interior**	**Region**	**per Site (Region)**
Clusters, $\delta_w = 0.011$ (4)	0.0015	0.991 (0.989–0.994)	0.978 (0.968–0.988)	0.0004	0.0015	0.989 (0.984–0.994)
Clusters, $\delta_w = 0.022$ (4)	0.0030	0.978 (0.972–0.981)	0.893 (0.844–0.969)	0.0009	0.014	0.937 (0.906–0.986)
Clusters, $\delta_w = 0.05$ (4)	0.0065	0.963 (0.957–0.968)	0.891 (0.812–0.940)	0.0015	0.011	0.940 (0.890–0.970)
Clusters, $\delta_w = 0.10$ (4)	0.0129	0.935 (0.918–0.945)	0.820 (0.770–0.862)	0.0028	0.018	0.903 (0.872–0.928)
Clusters, $\delta_w = 0.20$ (4)	0.0256	0.891 (0.884–0.901)	0.750 (0.702–0.781)	0.0045	0.024	0.867 (0.837–0.884)
Patches, $\delta_w = 0.05$ (4)	0.0067	0.957 (0.954–0.962)	0.764 (0.761–0.773)	0.0016	0.033	0.858 (0.856–0.864)

Table 20. Exact parameters of the FCC code on 16×16 patches cut after 12 layers, against a perfect ABC patch of the same shape (periodic in x and y, open at top and bottom). n: physical qubits (bonds); $k = n - \operatorname{rank} H_Z - \operatorname{rank} H_X$. Unit-bond condition.

Patch	**Nodes**	***n***	***X* Checks**	**Rank H_X**	***k***	***k*/*n***
Perfect ABC	**3072**	**17,664**	**2816**	**2816**	**11,777**	**0.667**
Grown, seed 0	2834	14,659	1716	1716	10,110	0.690
Grown, seed 1	2889	15,363	1985	1985	10,490	0.683
Grown, seed 2	2894	15,428	2006	2006	10,529	0.682
Grown, seed 3	2940	15,997	2202	2202	10,856	0.679
Grown, seed 0, $w = 0.005$	2904	15,552	2051	2051	10,598	0.681
Grown, seed 0, $w = 0.09$	2802	14,279	1619	1619	9859	0.690

8.4. Lift Weight and Code Quality

In these runs the acceptance window stays at its default, $r_{\mathrm{ex}} = 0.95L$ and $r_{\mathrm{b}} = 1.05L$; only the lift weight w is varied. We label each weight by the tolerance δ_w that the volume relation $w = (\sqrt{2}/\pi)\delta$ assigns to it (Section 4), so the labels inherit the assumption that weights are proportional to acceptance volumes. Lifts on the twin side of a triangle start twins and grains, and these remove checks and expose bonds. Table 19 and Figure 8a give the result. In clusters, over $\delta_w = 0.011$ to 0.20, the fraction of interior nodes that are not code-perfect rises from 0.9% to 10.9%, so a lower lift weight gives a cleaner code; if lift is weighted by acceptance volume, a tighter bond tolerance would do the same. This is a statement about the lift rate only. A lift weight of zero is trivially best, so code quality gives

no independent argument for any particular δ. With four seeds per row, Table 19 gives the range over seeds; the region values vary most, because one seed per row can hold a large twin grain. In young patches the twin boundaries dominate instead (Section 8.3). We do not convert δ into a Bell-pair fidelity. That needs a physical model of how a stretched bond loses entanglement, which the SSM does not yet supply.

9. Robustness

The main results hold across kinetic schemes (below), lift orientations (Sections 6.6, 7.4 and 7.5), the unit-bond condition, the structure classifier and the bonding radius. Appendix B compares the present simulator with that of [2].

Kinetic scheme. The site-weighted and legacy rules give the same FCC fraction at the same realized lift fraction, over the full range (Figure 4a).

Unit-bond condition. We regrow the four $N = 10^6$ clusters of Section 6 with the condition. In every seed it lowers the interior (depth ≥ 4) FCC fraction from 98.6% to 98.0%. The mean bulk FCC fraction changes from 0.943 to 0.939, and bulk hcp from 0.0008 to 0.0011. The effect is small and systematic, and no conclusion changes. Without the condition, the median error in bond length grows from 10^{-15} at radius $10L$ to 10^{-8} at $60L$, far below the 5% tolerance. Clusters several times larger would need the condition.

Settings of the unit-bond condition. We regrow two $N = 10^5$ clusters with the cap and the number of Gauss–Newton steps varied (Table 21). The cap is never reached: caps of $0.01L$, $0.02L$ and $0.04L$ give identical clusters. With two or eight steps instead of four, the interior FCC fraction changes by at most 0.4 percentage points, about the difference between the two seeds. Any change in a placement alters the later random draws, so the clusters differ in detail, but not in structure. In every setting of the condition the median bond error is below 10^{-15}. We repeated the test on slabs, where the condition matters most (Appendix A): two 32×32 slabs of 300 layers with axis lift. Caps of $0.01L$, $0.02L$ and $0.04L$ again give identical slabs, with median bond error about 10^{-14}. With eight steps the slabs stay ordered as with four (on-lattice fraction 0.997–0.999). With two steps one slab keeps residual errors (median 4×10^{-8}, largest $0.036L$, still within the tolerance) and its least ordered 10-layer bin falls to 0.95, so four steps leave a margin that two do not. Interior FCC ranges over 0.85–0.91 across all settings and seeds, within the spread between seeds in Table 12.

Table 21. Sensitivity to the settings of the unit-bond condition ($N = 10^5$, seeds 0 and 1). Interior: depth ≥ 4. Lift fraction and bulk FCC are means, and bond error the median, over both seeds.

Setting	Lift Fraction	Interior FCC		Bulk FCC	Median Bond Error
		Seed 0	Seed 1	Mean	
Bare rule	0.0064	0.988	0.986	0.948	4×10^{-14}
Cap $0.01L$, $0.02L$ or $0.04L$; 4 steps	0.0065	0.980	0.981	0.936	6×10^{-16}
Cap $0.02L$; 2 steps	0.0064	0.984	0.983	0.945	5×10^{-16}
Cap $0.02L$; 8 steps	0.0064	0.980	0.983	0.945	6×10^{-16}

Structure classifier. Reclassifying one $N = 10^6$ cluster with CNA cutoffs of $1.15L$, $1.207L$ and $1.25L$ gives identical fractions. At $1.30L$ the bulk FCC fraction changes only in the sixth decimal place.

Bonding radius. At $N = 3 \times 10^4$ (4 seeds each, with the unit-bond condition), bulk FCC is 0.914, 0.915, 0.925 and 0.917 for $r_b = 1.02L$, $1.05L$, $1.08L$ and $1.10L$. All agree within one standard error (0.010–0.016).

Placement off the close-packed sites. Rule 1 requires every bond to have length L, and the rules meet it by placing each node at an exact close-packed site. To test what this

exactness does, we place each new node at a random point near its site instead: uniform in a cube of half-width ε around the site, kept only if its distances to the anchor nodes (the two edge ends for stitch, the three corners for lift) lie within $L(1 \pm \varepsilon)$. We also let stitch place the node at a random angle on the full ring of unit distance from both edge ends. Each case was run with the bare rule, with the unit-bond condition at its default settings, and with a stronger relaxation (cap $0.1L$, eight steps), for two clusters of 5×10^4 nodes each (Table 22). Because the 42-neighbor test is not robust to position noise, we measure in a geometric core, the nodes within half the cluster radius of its center. Off-site nodes break the equal bond lengths. With the bare rule, a placement error of $10^{-3}L$ gives a median bond deviation of $4.5 \times 10^{-3}L$, because errors grow from layer to layer, as Appendix A shows for the exact rule. With unequal bonds the network loses both its order and its code. At $\varepsilon = 10^{-3}L$ the core FCC fraction falls from 0.94 to 0.76, and the share of core nodes on all six complete octahedra, the X checks of the code (Section 8.3), falls from 0.91–0.93 to 0.48–0.49. At $\varepsilon = 0.05L$ there is almost no FCC (3%), the mean coordination is about 6, and no core node keeps its six checks. On the full stitch ring there is no FCC and no complete octahedron in the core. The unit-bond condition holds the core FCC fraction at 0.80–0.87 up to $\varepsilon = 5 \times 10^{-3}L$, and the stronger relaxation at 0.5–0.7 up to $0.02L$. So the bond tolerance δ is a window for detecting bonds, not a freedom in where nodes go: the lattice and its code exist only when bonds have equal length to much better than $10^{-3}L$, which is what the rules provide. This is the reason for the rules: equal lengths and close-packed sites are what make the network stable and error-correcting.

Table 22. Placement off the close-packed sites ($N = 5 \times 10^4$, 2 seeds each). FCC: core FCC fraction; K: mean coordination in the core (12 in ideal FCC); $|\ell - L|$: median bond-length deviation in the core; Checks: share of core nodes on all six complete octahedra. Relaxed: unit-bond condition with cap $0.1L$ and eight steps.

	Bare Rule				Unit-Bond	Relaxed
Placement	**FCC**	K	$\lvert\ell - L\rvert/L$	**Checks**	**FCC**	**FCC**
Exact site	0.94	11.9	10^{-15}	0.92	0.93	0.94
Local, $\varepsilon = 0.001L$	0.76	11.2	0.005	0.49	0.87	0.70
Local, $\varepsilon = 0.005L$	0.42	9.7	0.005	0.13	0.80	0.61
Local, $\varepsilon = 0.01L$	0.34	9.2	0.007	0.07	0.39	0.72
Local, $\varepsilon = 0.02L$	0.18	8.1	0.011	0.01	0.19	0.53
Local, $\varepsilon = 0.05L$	0.03	6.1	0.021	0.00	0.03	0.09
Full stitch ring	0.00	5.0	0.020	0.00	0.00	0.00

Not yet tested. This includes a different seed shape, such as a tetrahedron, and a finer r_{ex} sweep near $L/\sqrt{3}$. Section 2 argues that the triangle is the natural seed, but the simulations do not test whether the bulk results depend on it.

10. Discussion

What the results show. With nodes placed at exact close-packed sites, the FCC order is a closure property of stitch: from one tetrahedron, stitch generates only the FCC lattice spanned by its edges (Section 7.3). The simulations therefore do not show an independent selection of FCC over other structures. They measure the defects that growth adds to this lattice, and how those depend on the settings. Every node that stitch adds from triangles on one FCC lattice stays on it in exact geometry (in simulation, all but 17 of 485,073 stitches), so only lift creates hcp, twins and misoriented grains. Bulk FCC accordingly exceeds 90% for realized lift fractions below about 1%, across the exclusion plateau from $0.58L$ to $0.98L$, for three lift orientations and two kinetic schemes. Equal bond lengths are essential. Nodes placed off their sites have bonds of unequal length, the deviations grow from layer to

layer, and the network loses both the crystal and the octahedral checks of the code; order already degrades at placement errors of $10^{-3}L$ (Section 9). The bond tolerance is therefore a window for detecting bonds, not a freedom in placement. Sequential deposition of hard spheres, which has no equal-length rule, gives random close packing with polytetrahedral order [48]. The two operator types follow from the loop and length conditions (Section 2.5). With exact lengths, lift's site is fixed; stitch's in-plane site is the model's choice on its ring. Both choices are motivated by what they give: stable bonds and a cell structure that corrects errors, which is the sense of persistence in Section 2.2. If operator weights are proportional to acceptance volumes, geometry gives the lift rate in terms of the bond tolerance: lift is suppressed by one power of δ. The tolerance itself is a free parameter, and every lift fraction below about 1% gives FCC.

Corrections and limits. The network as a whole does not pass through a $K = 4$ phase. Locally, each lift makes a tetrahedral K_4 unit, and these units carry growth into three dimensions. The tetrahedron is the natural local step between a sheet and the bulk. But coordination 4 is neither common nor lasting. Only a third of nodes ever have it, they leave it as fast as any other value, and the tetrahedral content never exceeds that of FCC. The sequence $K = 6 \to K = 4 \to K = 12$ of Ref. [2] (sheets, then tetrahedra, then FCC) holds as a local docking step, not as a global stage. If the SSM needs a global $K = 4$ phase, that phase is an assumption about the initial state. A rule set that produces it would be a separate test, for example lift-first growth or a high early lift rate that falls off. The value e^{-3} has no special status. As a per-step probability it gives about 1% realized lift and 91% FCC, slightly below the values at rarer lift. The bare rule also cannot keep long-range order. The unit-bond condition is therefore part of the model. It is the most direct form of the Bell-pair rule in the simulation. More generally, the model is pure geometry. It has no energetics, no thermal motion and no motion of remnant nodes.

Growth far from the seed. A global lift orientation leaves no trace on the shape of a cluster or on which twin variant wins in a patch. The twin domains themselves come from the starting surface. One close-packed sheet fits both twin variants, and the first layer above it fills both kinds of hollow. On a three-layer ABC substrate a patch grows as a single variant with interior FCC 0.94–0.95 (Section 7.4). The domain coarsening and the slab FCC fraction of 0.84–0.92 are therefore specific to a single-sheet start. Within that setting, domains widen roughly as height$^{0.6}$–height$^{0.7}$, a poorly constrained estimate from three widths. In [2], the two orientations of tetrahedral voids stand for matter and antimatter, and the observed excess of matter is attributed to a local excess of one orientation. They are balanced in any FCC crystal by symmetry, and in the simulation they are balanced to within about 2×10^{-3} below the outermost layers. So selection of one class for matter must come from trapping. On a one-way front, none of 2400 remnants that settled on the surface was trapped in a tetrahedral void pointing against growth (Section 7.7), a result largely built into the settling rule. The two classes need a definition. They are related by inversion, and "along growth" is defined relative to the local growth direction. Two patches growing in opposite directions on the same lattice would therefore select opposite global classes. If matter and antimatter are the two global classes of Ref. [2], trapping does not give matter everywhere. It does so only if matter is defined as the class pointing along the local growth direction, which is then an assumption of the SSM. With that definition, a front that grows outward everywhere would select matter everywhere; this is a conjecture. The simulation confirms the geometry of the selection, but not the settling premise. The small excess of down tetrahedra points against growth. At about 10^{-3} or less (Table 12), it is too small to affect this.

Quantum structure. Read as Bell pairs, the grown network is a quantum state whose entanglement follows an exact area law. Its coefficient is the density of bonds crossing a

surface. Every missing node lowers it, so the entanglement density records the defects that growth leaves. One parity check per node, acting only on qubits that node holds, turns the network into a code state with one qubit per bond, so the Bell-pair and code readings are related by local operations. The preparation is deterministic: random parity outcomes are repaired by local X gates along paths, and every outcome gives the same state. The checks lower the entropy of a connected region with a connected complement by exactly $\ln 2$ and leave the coefficient unchanged. The map from Bell pairs to the code state is the standard projected-entangled-pair construction [12,13], and the area-law coefficient is the density of bonds crossing a surface; what the paper adds is their measurement on grown networks. The checks of the FCC stabilizer code [11] commute on any graph that contains complete octahedra, so defects only remove checks. In the interior (depth ≥ 3, away from the surface and vacancy pockets), about 96% of nodes carry the full structure of the ideal code; in a region that keeps every defect, 89%. A single triangle is a weight-3 logical operator, so the distance is at most 3 at every size. Protection is therefore local: each cell with intact checks corrects any single error, and the lattice is protected cell by cell (Section 8.3). Exposed bonds, at 0.15% of interior bonds and 1.1% in a region that keeps every defect, reduce the distance to 1 where they occur. And twin boundaries in young patches remove a large share of the X checks. Both are set by the growth rules. Exposed bonds grow with the lift rate (and, under the volume relation, with the bond tolerance δ), and twin boundaries become sparser as domains coarsen. In the SSM picture, the density of intact checks would thus be lowest where space is young and would rise as it grows, if a patch starts without stacking memory; on an ABC substrate no twins form. Because protection is local, this is also the density of protected cells. These results are exact for the stabilizer states defined on the graph. We do not simulate how the state evolves under physical noise; the per-cell failure rate of Section 8.3 is an estimate from the distance.

Before the seed. Section 2 is the foundation of the model and also its least secure part. Its exact content is small. The fixed points of the measure-and-repair map are exactly the checked states. With one qubit per bond, the triangle is the smallest simple network whose node and loop checks detect every error on its bonds (elementary graph theory), and its checks fix the GHZ state. And the octahedral checks of the FCC code are built from triangles. The growth rules take two conditions from the postulate, bonds of one length and bonds created on loops, and these are what make the grown lattice error-correcting. They place nodes at exact close-packed sites, a rule motivated by the stability and error correction it gives (Section 2.5), and they idealize formation by taking operator weights proportional to acceptance volumes. They model maintenance only through the unit-bond condition at placement, not continued checking of the finished lattice. If a closed world is in a pure state and two parts can be related only if they are correlated, classical probability theory is excluded. This is close to a restatement of the two premises, and it does not single out quantum theory. The persistence postulate itself is stated as a postulate and is not derived. The rest, from the postulate to the state before the lattice, is a proposal.

Open questions. These include: (1) whether remnants settle onto the front as assumed: this is the premise behind the trapping result, and a model of remnant motion near the front would test it; (2) the cause of the tetrahedron excess: classifying the local setting of every unbalanced tetrahedron, and testing the classifier's tolerance, would settle it; (3) the nucleation rate and takeover criterion of tilted grains; (4) the domain-growth exponent, from slabs wider than 64, and the twin structure of patches started on other substrates; (5) the off-lattice lift rate: Section 7.3 traces stacking disorder to lifts on the twin side of a triangle; a model of how often such lifts succeed, given exclusion by nearby nodes, would turn the measured 20–29% into a calculated rate; (6) a formal version of the persistence postulate of Section 2, a justification of premises (P1) and (P2), and whether persistence

implies purification in full, which with the other five principles of Ref. [29] would select quantum theory; (7) a dynamics of repeated verification and repair under noise, which would turn the one-round preparation of Section 8.2 into persistence, and its threshold; and (8) whether exposed bonds can be handled as gauge qubits, and whether a code on this lattice with distance growing with size exists.

11. Conclusions

With nodes placed at exact close-packed sites, the SSM growth rules build the FCC lattice. From one tetrahedron, stitch generates only the FCC lattice spanned by its edges, so only lifts can create other stackings; the simulations measure the defects that lifts and growth add. At 10^6 nodes, about 98% of interior nodes of clusters are FCC (98.6% with the bare rule and 98.0% with the unit-bond condition, counting all nodes at depth 4 or more). The rules require every bond to have the same length, and this is motivated by what it gives: stability and error correction. Nodes placed off site give bonds of unequal length, and the network loses both the crystal and the code checks. New sheets grow from step edges by stitching on the triangles that join neighboring flat sheets, and this carries the stacking from sheet to sheet.

If operator weights are proportional to acceptance volumes, the lift rate follows from geometry as $P_{\text{lift}}/P_{\text{stitch}} = (2\sqrt{2}/\pi)\delta$, with the bond tolerance δ as a free parameter; every lift fraction below about 1% gives FCC. A least-squares unit-bond condition, the direct form of the Bell-pair rule, is needed to keep long-range order beyond about 100 layers (Appendix A). Each lift makes a tetrahedral unit, but the rules produce no global $K = 4$ stage, and if the SSM needs one, it must be assumed.

In a patch that grows one way, the void population is balanced: the two classes of tetrahedral void agree to within about 2×10^{-3} below the outermost layers. Selection of one class comes from trapping instead. Of 2400 remnant nodes that settled on rough simulated one-way fronts, none ended in a void pointing against growth; the seven sites with one node below and two above are ledge sites. The result is largely built into the settling rule. If remnants settle in this way, a patch grown one way would hold only one void class relative to its growth direction, which the SSM can read as matter only if matter is defined relative to that direction. The main open question is whether remnants settle onto the front as assumed, which a model of remnant motion could test. A patch started on one close-packed sheet grows as twinned FCC, with an interior FCC fraction of 84–92%; in 6 of 48 slabs a tilted grain takes over. The twin variants coarsen into ever wider single-variant domains, and the winner is random. On a three-layer ABC substrate the patch grows as a single variant, so the twins come from the single-sheet start.

Read as a network of Bell pairs, the grown lattice obeys an exact entanglement area law, $S = \ln 2$ times the number of bonds crossing the boundary; one parity check per node turns it into a code state with one qubit per bond, which lowers S by exactly $\ln 2$ for a connected region with a connected complement. The area-law coefficient lies 2.5–6% below that of an ideal FCC because of defects. The weight-12 FCC stabilizer code can be built on any grown graph, and its checks commute, as they must on any graph with complete octahedra. Its distance is at most 3 at every size, so it protects the lattice locally, cell by cell, rather than as one block. About 96% of interior nodes carry its full structure (89% when all defects are counted), 0.15% of interior bonds escape every X check (1.1% counting all defects), and in patches started on one sheet the code regains checks as twin domains coarsen.

Section 2 states the postulate behind the model, "what exists is what can persist", and makes it concrete with a verification map, one round of node checks and local repair, whose fixed points are exactly the checked states. It gives a reason, in the one-qubit-per-bond reading, to start from a triangle: the triangle is the smallest simple network whose bonds

check one another. Its checks fix a GHZ state, and each X check of the FCC code is a sum of four triangular loops. The growth rules follow its two conditions, checked bonds and equal lengths, and these are what make the simulated lattice error-correcting. The rest of that section is a proposal.

Funding: This research received no external funding.

Data Availability Statement: The original data presented in the study are openly available in Zenodo at https://doi.org/10.5281/zenodo.24096739 [16], in the archive `ssm-growth-sim.zip`, which is also available in GitHub at https://github.com/raghu91302/ssmtheory/blob/main/ssm-growth-sim.zip (accessed on 7 November 2026). The archive contains the simulator (`ssmsim`), tests, analysis scripts, raw results (aggregated for the robustness and lattice-check runs) and figure scripts. Two interactive views go with the paper. The growth viewer for both cases is at https://raghu91302.github.io/ssmtheory/lattice_growth.html (accessed on 7 November 2026) [14]. The formation of a new sheet at a step edge (Figure 5) is at https://raghu91302.github.io/ssmtheory/step_edge_sheet.html (accessed on 7 November 2026) [15].

Conflicts of Interest: Raghu Kulkarni is an employee of SSMTheory Group, IDrive Inc. The views expressed in this article are those of the author and do not necessarily represent the views of IDrive Inc.

Abbreviations

The following abbreviations are used in this manuscript:

SSM	Selection-Stitch Model
3D	Three-dimensional
FCC	Face-centered cubic
hcp	Hexagonal close packing
CNA	Common neighbor analysis
CSS	Calderbank–Shor–Steane (code)
GHZ	Greenberger–Horne–Zeilinger (state)
GF(2)	Galois field with two elements

Appendix A. Instability of the Bare Growth Rule

Without the unit-bond condition, a new node's position comes straight from its triangle, and any error in the triangle carries over. Each node depends on several earlier nodes, so errors grow from layer to layer. In a 32×32 slab, the median deviation from the ideal lattice grows from 3×10^{-14} at layer 10 to 2×10^{-3} at layer 100. That is a factor of about 1.3 per layer (Figure A1). Position noise ε would therefore destroy order after about $\ln(\delta/\varepsilon)/\ln 1.3$ layers, or about 15 layers for $\varepsilon = 10^{-3}$ and $\delta = 0.05$; Section 9 tests this directly by placing nodes off site. Order breaks down at layer 109 on average (range 105–117 over 48 slabs; group means 107–111 for every width and orientation), and the crystal breaks into misoriented grains. That the height does not depend on width or orientation shows that it is an artifact. With the unit-bond condition, the deviation stays between 2×10^{-15} and 5×10^{-13} through 300 layers, and slabs stay ordered through 600 layers. Slab runs without the condition are kept in the archive but not used.

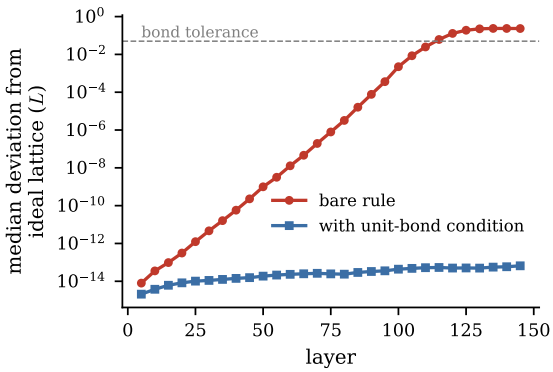


Figure A1. Median deviation of node positions from the ideal lattice against layer in a 32 × 32 slab, for the bare growth rule and with the unit-bond condition. Without the condition, rounding errors grow by about 30% per layer and reach the bond tolerance near layer 110.

Appendix B. Comparison with the Earlier Implementation

The earlier simulator [2] differs in three ways. It bonds nearby nodes in batches every 40 placements. It grows edges that belong to no triangle in a direction set by the global z axis. And it keeps a site active until both candidate positions fail. We changed one feature at a time ($N = 1500$, $p = 0.05$, 8 seeds; Table A1). Batched bonding explains about 9 of the 16 percentage points between the two simulators. The rest comes from site bookkeeping. The global-z rule has no measurable effect.

Table A1. Earlier implementation versus present simulator ($N = 1500$, $p = 0.05$, bare rule; 8 seeds unless stated; mean ± standard error).

Variant	Bulk FCC (CNA)
Earlier simulator as published	0.675 ± 0.019
Bonding at every placement	0.763 ± 0.029
No global-z growth of triangle-free edges	0.671 ± 0.022
Both changes	0.764 ± 0.024
Present simulator, legacy per-step rule (16 seeds)	0.839 ± 0.017
Present simulator, site-weighted rule (16 seeds)	0.883 ± 0.013

References

1. Van Raamsdonk, M. Building up spacetime with quantum entanglement. *Gen. Relativ. Gravit.* **2010**, *42*, 2323–2329. [CrossRef]
2. Kulkarni, R. Matter as incomplete crystallization: Quark charges, color confinement, and the proton mass from a single extra node in the vacuum lattice. *Phys. Open* **2026**, *27*, 100423. [CrossRef]
3. Hales, T.C. A proof of the Kepler conjecture. *Ann. Math.* **2005**, *162*, 1065–1185. [CrossRef]
4. Bolhuis, P.G.; Frenkel, D.; Mau, S.-C.; Huse, D.A. Entropy difference between crystal phases. *Nature* **1997**, *388*, 235–236. [CrossRef] [PubMed]
5. Pusey, P.N.; van Megen, W.; Bartlett, P.; Ackerson, B.J.; Rarity, J.G.; Underwood, S.M. Structure of crystals of hard colloidal spheres. *Phys. Rev. Lett.* **1989**, *63*, 2753–2756. [CrossRef] [PubMed]
6. Satzinger, K.J.; Liu, Y.-J.; Smith, A.; Knapp, C.; Newman, M.; Jones, C.; Chen, Z.; Quintana, C.; Mi, X.; Dunsworth, A.; et al. Realizing topologically ordered states on a quantum processor. *Science* **2021**, *374*, 1237–1241. . [CrossRef] [PubMed]
7. Bluvstein, D.; Levine, H.; Semeghini, G.; Wang, T.T.; Ebadi, S.; Kalinowski, M.; Keesling, A.; Maskara, N.; Pichler, H.; Greiner, M.; et al. A quantum processor based on coherent transport of entangled atom arrays. *Nature* **2022**, *604*, 451–456. [CrossRef] [PubMed]
8. Iqbal, M.; Tantivasadakarn, N.; Verresen, R.; Campbell, S.L.; Dreiling, J.M.; Figgatt, C.; Gaebler, J.P.; Johansen, J.; Mills, M.; Moses, S.A.; et al. Non-Abelian topological order and anyons on a trapped-ion processor. *Nature* **2024**, *626*, 505–511. [CrossRef] [PubMed]
9. van Blaaderen, A.; Ruel, R.; Wiltzius, P. Template-directed colloidal crystallization. *Nature* **1997**, *385* 321–324. [CrossRef]
10. van der Drift, A. Evolutionary selection, a principle governing growth orientation in vapour-deposited layers. *Philips Res. Rep.* **1967**, *22* 267–288.

11. Kulkarni, R. A 67%-rate CSS code on the FCC lattice: $[[192, 130, 3]]$ from weight-12 stabilizers. *arXiv* **2026**, arXiv:2603.20294. [CrossRef]
12. Verstraete, F.; Wolf, M.M.; Pérez-García, D.; Cirac, J.I. Criticality, the area law, and the computational power of projected entangled pair states. *Phys. Rev. Lett.* **2006**, *96*, 220601. [CrossRef] [PubMed]
13. Gu, Z.-C.; Levin, M.; Swingle, B.; Wen, X.-G. Tensor-product representations for string-net condensed states. *Phys. Rev. B* **2009**, *79*, 085118. [CrossRef]
14. SSMTheory. Lattice Growth Viewer (Interactive 3D Visualization of Cluster and Patch Growth). 2026. Available online: https://raghu91302.github.io/ssmtheory/lattice_growth.html (accessed on 7 November 2026).
15. SSMTheory. Step-Edge Sheet Formation (Interactive 3D Visualization). 2026. Available online: https://raghu91302.github.io/ssmtheory/step_edge_sheet.html (accessed on 7 November 2026).
16. Kulkarni, R. ssm-growth-sim: Simulator, Analysis Scripts and Data for This Paper. *Zenodo* **2026**. [CrossRef]
17. Wheeler, J.A. Information, physics, quantum: The search for links. In *Proceedings of the 3rd International Symposium on Foundations of Quantum Mechanics in the Light of New Technology, Tokyo, Japan, 1989*; Kobayashi, S., Ezawa, H., Murayama, Y., Nomura, S., Eds.; Physical Society of Japan: Tokyo, Japan, 1990; pp. 354–368.
18. Tryon, E.P. Is the Universe a vacuum fluctuation? *Nature* **1973**, *246*, 396–397. [CrossRef]
19. Zurek, W.H. Quantum Darwinism. *Nat. Phys.* **2009**, *5*, 181–188. [CrossRef]
20. Deutsch, D.; Marletto, C. Constructor theory of information. *Proc. R. Soc. A* **2015**, *471*, 20140540. [CrossRef] [PubMed]
21. Swingle, B. Entanglement renormalization and holography. *Phys. Rev. D* **2012**, *86*, 065007. [CrossRef]
22. Cao, C.; Carroll, S.M.; Michalakis, S. Space from Hilbert space: Recovering geometry from bulk entanglement. *Phys. Rev. D* **2017**, *95*, 024031. [CrossRef]
23. Barrett, J. Information processing in generalized probabilistic theories. *Phys. Rev. A* **2007**, *75*, 032304. [CrossRef]
24. Popescu, S.; Rohrlich, D. Quantum nonlocality as an axiom. *Found. Phys.* **1994**, *24*, 379–385. [CrossRef]
25. Gross, D.; Müller, M.; Colbeck, R.; Dahlsten, O.C.O. All reversible dynamics in maximally nonlocal theories are trivial. *Phys. Rev. Lett.* **2010**, *104*, 080402. [CrossRef] [PubMed]
26. Coffman, V.; Kundu, J.; Wootters, W.K. Distributed entanglement. *Phys. Rev. A* **2000**, *61*, 052306. [CrossRef]
27. Terhal, B.M. Is entanglement monogamous? *IBM J. Res. Dev.* **2004**, *48*, 71–78. [CrossRef]
28. Hardy, L. Quantum Theory from Five Reasonable Axioms. *arXiv* **2001**, arXiv:quant-ph/0101012.. [CrossRef]
29. Chiribella, G.; D'Ariano, G.M.; Perinotti, P. Informational derivation of quantum theory. *Phys. Rev. A* **2011**, *84*, 012311. [CrossRef]
30. Masanes, L.; Müller, M.P. A derivation of quantum theory from physical requirements. *New J. Phys.* **2011**, *13*, 063001. [CrossRef]
31. Plávala, M. General probabilistic theories: An introduction. *Phys. Rep.* **2023**, *1033*, 1–64. [CrossRef]
32. Renou, M.-O.; Trillo, D.; Weilenmann, M.; Le, T.P.; Tavakoli, A.; Gisin, N.; Acín, A.; Navascués, M. Quantum theory based on real numbers can be experimentally falsified. *Nature* **2021**, *600*, 625–629. [CrossRef] [PubMed]
33. Chen, M.-C.; Wang, C.; Liu, F.-M.; Wang, J.-W.; Ying, C.; Shang, Z.-X.; Wu, Y.; Gong, M.; Deng, H.; Liang, F.; et al. Ruling out real-valued standard formalism of quantum theory. *Phys. Rev. Lett.* **2022**, *128*, 040403. [CrossRef] [PubMed]
34. Sinha, U.; Couteau, C.; Jennewein, T.; Laflamme, R.; Weihs, G. Ruling out multi-order interference in quantum mechanics. *Science* **2010**, *329*, 418–421. [CrossRef] [PubMed]
35. Page, D.N.; Wootters, W.K. Evolution without evolution: Dynamics described by stationary observables. *Phys. Rev. D* **1983**, *27*, 2885–2892. [CrossRef]
36. Moreva, E.; Brida, G.; Gramegna, M.; Giovannetti, V.; Maccone, L.; Genovese, M. Time from quantum entanglement: An experimental illustration. *Phys. Rev. A* **2014**, *89*, 052122. [CrossRef]
37. Regge, T. General relativity without coordinates. *Nuovo Cim.* **1961**, *19*, 558–571. [CrossRef]
38. Lam, S.K.; Pitrou, A.; Seibert, S. Numba: A LLVM-based Python JIT compiler. In *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC (LLVM-HPC 2015)*; ACM: New York, NY, USA, 2015; article 7. [CrossRef]
39. Honeycutt, J.D.; Andersen, H.C. Molecular dynamics study of melting and freezing of small Lennard-Jones clusters. *J. Phys. Chem.* **1987**, *91*, 4950–4963. [CrossRef]
40. Faken, D.; Jónsson, H. Systematic analysis of local atomic structure combined with 3D computer graphics. *Comput. Mater. Sci.* **1994**, *2*, 279–286. [CrossRef]
41. Calderbank, A.R.; Shor, P.W. Good quantum error-correcting codes exist. *Phys. Rev. A* **1996**, *54*, 1098–1105. [CrossRef] [PubMed]
42. Steane, A.M. Error correcting codes in quantum theory. *Phys. Rev. Lett.* **1996**, *77*, 793–797. [CrossRef] [PubMed]
43. Gottesman, D. Stabilizer Codes and Quantum Error Correction. Ph.D Thesis, California Institute of Technology, Pasadena, CA, USA, 1997. [CrossRef]
44. Eisert, J.; Cramer, M.; Plenio, M.B. Colloquium: Area laws for the entanglement entropy. *Rev. Mod. Phys.* **2010**, *82*, 277–306. [CrossRef]
45. Kitaev, A.; Preskill, J. Topological entanglement entropy. *Phys. Rev. Lett.* **2006**, *96*, 110404. [CrossRef] [PubMed]

46. Levin, M.; Wen, X.-G. Detecting topological order in a ground state wave function. *Phys. Rev. Lett.* **2006**, *96*, 110405. [CrossRef] [PubMed]
47. Kitaev, A.Y. Fault-tolerant quantum computation by anyons. *Ann. Phys.* **2003**, *303*, 2–30. [CrossRef]
48. Bennett, C.H. Serially deposited amorphous aggregates of hard spheres. *J. Appl. Phys.* **1972**, *43*, 2727–2734. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.