

What the burp leaves behind: relic defects of fixed size in a graph model of emergent geometry, their kinds, their number, and how long they last

Emily Smith^{1,*}

¹*Independent researcher, Charlottesville, Virginia, USA*

(Dated: DRAFT, 13 October 2026)

In a companion paper a $4 \times L$ torus with one direction curled was shown to be metastable, and to convert after activation with a release of energy, in the family of graph energies obtained from combinatorial quantum gravity (CQG) by raising the coefficient λ of its local term above 1 (only $\lambda = 1$ is CQG; every result here is at $\lambda = 1.25$, next to CQG and not in it). Here we study the defects the conversion leaves. Read from the saved wiring there are three kinds, each of fixed size and each a local minimum of the energy: a column of the parent torus left curled (four vertices, $24\lambda - 16 = 14$ units), a double column (eight vertices, 8 units), and a knot of eight vertices of two kinds (20 to 30 units). Isolated, with a reservoir large enough to stay cold, a conversion started at one place leaves about one relic at every size from 64 to 288 vertices, at a position indistinguishable from uniform. The number rises with the number of nucleation events by less than one per event (1.10, 2.15, 2.50, 3.40 relics for 1, 2, 4, 8 planted seeds), and in proportion to length when a long torus nucleates spontaneously. In an external bath at fixed coupling a lone column anneals within 500 to 80,500 sweeps; under cooling it is arrested. When products of spontaneous nucleation at 1,024 vertices are cooled over 0 to 300,000 size-fair sweeps, the columns fall from about 2 to 0.1 per product while the other two kinds do not change, so that the excess energy falls only from 4.3% to 1.3% of the release. Two predictions of a gentler fall, the author's and ours, failed. Every verdict quoted was pre-registered before its run; the readings from wiring are exact.

I. INTRODUCTION

A change from one ordered arrangement to another releases energy and, in an isolated system, must keep the released energy in the graph or its reservoir. Conservation alone does not require a residual defect. In the companion paper [1] the starting arrangement is a $4 \times L$ torus, one direction curled to a single square, which is metastable once the local term of CQG [2–4] is weighted by a coefficient $\lambda > 1$ and which converts, after activation, to arrangements with two squares on every edge. That paper measured the escape and conversion clocks, the cost of the first exit and the release. This one measures what is left: its kinds, its number, where it sits, and how long it lasts.

The question is narrow and the answer is measurable. The reason for asking is the author's hypothesis that space emerged from an ordered precursor by a change of this kind, and that such a change leaves a remainder; she wants tested whether a remainder can last. Ref. [5] describes something in the same family for CQG itself: allotropes, long-lived metastable domains of the network in a different discrete arrangement. The relics studied here are not allotropes in that sense: they are small objects of fixed size, made by the conversion, at $\lambda > 1$.

The process has a second set of relatives, in the nonequilibrium dynamics of ordered phases in condensed-matter models, and we use that field's words where the concept is the same. There, order that forms in an isolated system does so by nucleation and growth and can

end in an arrested state whose defects no longer move [8]; in contact with a thermal bath the defects diffuse and annihilate in pairs [9]. Those works study other models and are context, not validation. Their defects are kinks, walls between two symmetry-related versions of one order. A relic is a small piece of the parent arrangement left in place, and we do not call it a kink (Sec. VIII).

Every kinetic statement below concerns one Markov chain, the local edge switch used to sample the model [3], with Metropolis acceptance, and the sealed (isolated, microcanonical) version of it built from Creutz demons [6]. Energies, the relics' costs and the classification of end states do not depend on the chain; waiting times, annealing times and counts per seed do.

II. MODEL, DYNAMICS, AND HOW A RELIC IS READ

The state is a four-regular bipartite graph on N vertices obeying the hard-core rule of CQG (no two vertices share more than two neighbors). With S the number of squares and $X = \sum_e (S_e - 2)_+$ the number of surplus squares on edges carrying more than two, the energy is

$$H = 16(N - S) + 4\lambda X, \quad (1)$$

which for $\lambda = 1$ is the CQG action [4] and for $\lambda \neq 1$ is not a curvature. Any arrangement with two squares on every edge has $H = 0$ and is a ground state for $\lambda \geq 1$; the curled torus sits $\varepsilon = 4(\lambda - 1)$ per vertex above it, 1 at $\lambda = 1.25$.

Two letters need a warning for a reader from condensed matter, where both name interaction strengths. Here λ weights the local term. The coupling g of CQG plays

* emilysmithcreate@gmail.com

the role of a temperature: at fixed g a proposed switch is accepted with probability $\min[1, \exp(-\Delta H/g)]$, which puts the graph in contact with an external thermal bath.

Reading a relic. The square-count signature $d(v)$ of a vertex is the number of its six edge pairs that close no square: 2 where the graph is locally a square lattice, 1 on the curled torus, 0 in a 4-cube. It serves as a local order parameter and does not certify geometry. A relic is a localized defect: a connected piece of vertices not at $d = 2$ in a product that is otherwise at $d = 2$. Its excess energy is read exactly from Eq. (1) on the saved wiring.

What the product is. “Flat” below means $d = 2$, a local statement. The companion paper reconstructed 3,521 of its zero-energy end states: 3,446 are connected tori whose shortest noncontractible loop is six edges, the parent torus re-glued with its seam shifted by two sites; 73 are disconnected and two are Klein bottles [1]. The product is therefore not a wide sheet, and nothing here measures a width unless it says so.

Isolated (sealed) runs keep a reservoir of C stores beside the graph and exchange energy with nothing else. The stores play the part of the internal bath of Ref. [8], degrees of freedom inside the isolated system that take up and give back energy without dissipating it; here the reservoir is put in by hand and C is a declared parameter. A proposed switch picks one store at random, is refused if that store cannot pay for it, and otherwise pays or is paid; $H + \sum$ stores is conserved to the last unit in every run reported. The stores’ mean energy is not a calibrated temperature [1]. The isolated runs below use $C = 2N$, large enough that the release leaves the reservoir cold and the product at $d = 2$ [1]: in T10 and T11 one store starts with a 12-unit seed and the rest are empty; in T17, T19 and T25 every store starts empty and the change is started by planted seeds. A seed is a nucleus of the product arrangement, planted or spontaneous; a front is the moving boundary of the domain that grows from it; where two fronts meet, two domains coalesce. T37 and T51 run in an external bath.

III. THREE KINDS OF RELIC, READ FROM THE WIRING

Exact for each graph read. A leftover has been named from a count and been wrong before in this project, so each kind below was read from saved adjacency lists.

The column. A closed loop of four vertices at $d = 1$, each of its edges lying in three squares: one column of the parent torus that stayed curled while the columns around it converted. It has one extra square, and holds 4, 9 or 14 units at $\lambda = 1.25$ according to how many edges around it carry a third square ($S = N + 1$ and $X = 4, 5$ or 6). The 14-unit form, $24\lambda - 16$, is the common one: 1,070 of the 1,142 single columns in the end states of Sec. VII. Its cheapest single switch out costs 12.

The double column. Eight vertices at $d = 1$ with twelve edges among them, each vertex keeping three edges inside

the piece and one out: a 4×2 segment of the parent torus. It holds 8 units.

The knot of two kinds. Eight vertices: two joined pairs at $d = 1$, each vertex’s four edges carrying 2, 2, 3 and 3 squares, and two joined pairs at $d = 3$, each vertex’s edges carrying 2, 1, 1 and 2. It holds 20 units at its own vertices and 30 within one step. Counted by connected pieces it appears as four pieces of two vertices, two steps apart, which is how the registered analysis first counted it. A 20-unit “twist” seen in 9 of 80 isolated runs (Sec. IV) is of the same kind; whether it is the same arrangement was not settled.

Pre-registered as T55, on the 600 end states of Sec. VII, before any of them was read this way. Three readings of the pieces that cooling does not remove were registered: SCAR (a single switch removes the piece and lowers the energy), SEAM (the piece carries an edge joining places at least seven steps apart by any other route, the stitching between domains that grew from different seeds) and KNOT (neither). The verdict was KNOT for both piece sizes scored: all 652 pieces of two and all 103 pieces of eight at $L = 256$ are local minima, with a cheapest exit of 11 in nearly all, and 4 of 2,121 pieces of any size carry such an edge. The author’s prediction, SEAM, failed; ours held in part (not SCAR; the pieces of eight are not SEAM as we had guessed).

Rarer resting states, 8 to 45 units, were read from twelve of the companion paper’s decays: combinations of small defects, among them a two-vertex twist (20 units), a 3-cube, an open line of three vertices, single edges and single vertices; eleven matched their recorded release exactly, and the twelfth changed state inside its last measuring window. Whether the list of kinds is complete is not known.

The share of the release kept by one 14-unit column is exact, $(24\lambda - 16)/[4(\lambda - 1)N]$: 22% at $N = 64$, falling as $1/N$.

IV. ONE SEED, ABOUT ONE RELIC, AT EVERY SIZE

Pre-registered as T10. Twenty isolated conversions at each of $N = 64, 96, 192, 288$ ($C = 2N$, a 12-unit seed, 30,000 sweeps). The prediction on record was that the mean number of relics rises with N . The verdict was ONE RING, HOWEVER LARGE: at every size the product holds about one column, the mean count 0.85 to 1.05, with a slope of 0.0003 ± 0.0003 per vertex. (The registered label keeps the word “ring”; the four vertices lie along the torus, spanning one to four of its original columns.) A second check, that every relic holds 14 units, failed in 9 of the 80 runs, which held the 20-unit twist instead.

Pre-registered as T11, at $N = 64$ and 96. Two readings were possible: the relic marks the seed, or it marks the place opposite the seed, where the two fronts from one nucleus would meet. In the 36 and 28 runs holding one

relic its position is indistinguishable from uniform along the torus: NEITHER. Where the fronts actually met was not measured.

V. HOW MANY

Planted seeds; pre-registered as T17. A 96×4 torus ($N = 384$), isolated and cold, with k seeds planted as the cheapest move out of the torus at evenly spaced columns; twenty runs per k , all converting. The author's hypothesis corresponds to one relic per seed, though she did not pick among the registered options; ours leaned toward one per torus whatever k . Neither held: 1.10, 2.15, 2.50 and 3.40 relics per torus for $k = 1, 2, 4, 8$, a slope of 0.29 ± 0.04 per extra seed, verdict BETWEEN. The slope is only about one standard error above the boundary with one per torus, and the rises (1.05, then 0.35, then 0.90) are poorly described by a line. The counts are more regular than independent events would give (sub-Poissonian: variance over mean 0.09 to 0.50). Where domains grown from neighboring seeds coalesce, relics may merge or anneal; that was not measured.

Spontaneous nucleation; pre-registered as T37. The chance of a seed per sweep does not grow with the torus while a front's speed per sweep falls with its size, so a long enough torus nucleates in several places before any one domain has grown across it. Tori $4 \times L$ were started in an external bath at $g = 1.25$ ($L = 1024$) and at $g = 1.5$ and 1.75 ($L = 64$ to 1024), 12 to 40 replicas a cell. The registered predictions were ours, with the same verdict inferred for the author: that the seed count grows with length, that the converted fraction follows the one-dimensional law of nucleation and growth [7] (Avrami exponent 2), and that the energy left grows at least threefold from $L = 64$ to the largest L . All three held: the seed count grows as $L^{0.92}$ and $L^{0.78}$ at $g = 1.5$ and 1.75 , and the Avrami exponent is 2.02, 2.14 and 2.03 at the three couplings. A fourth, the seed count computed from the measured nucleation rate and front speed with nothing fitted, holds at $g = 1.25$ (0.71 of the observed) and fails at the warmer couplings (0.41 and 0.16 at $L = 1024$).

Read afterwards from the saved graphs, the cells divide in two. Where the product stays at $d = 2$ (about 95% of vertices), columns are left in proportion to length: one per 57 to 72 columns of torus at $g = 1.5$ for $L \leq 256$, and at $g = 1.25$ about one for every three seeds, close to the 0.29 per planted seed above. At $L = 1024$ and $g = 1.5$ or 1.75 , 55% and 74% of vertices ended at $d > 2$, the squares per vertex fell to 0.90 and 0.83, and the graph had absorbed 1.75 and 2.49 per vertex from the bath: thermal melting, with energy supplied from outside. At this λ an isolated torus's own release is too small to do that [1]. The coupling at which the product gives way to the random phase falls as the system grows (the published curves for $N = 100$ to 200 in the parent model do not collapse either; Ref. [3], Fig. 8a, our reading), so the test

of Sec. VII nucleates at the coldest of these couplings.

VI. LIFETIME IN AN EXTERNAL BATH

Pre-registered as T19. Products with exactly one column, made as in T17 at $k = 1$ from a 24×4 torus ($N = 96$), were held at fixed coupling $g = 1.0, 1.25$ and 1.5 for 100,000 sweeps, the relic's position read every 500. The options on record were that it moves, that it anneals at the warmer couplings and stays at the colder, or that it stays. It anneals at every coupling: at $g = 1.0$ all nine were gone between 4,000 and 80,500 sweeps; at 1.25 all ten within 500 to 13,000; at 1.5 all ten within 11,000, two after moving first. Our prediction, that it would stay at $g = 1.0$, was wrong. In T10's cold isolated runs the relic lasted the whole 30,000 sweeps.

In the terms of defect kinetics: each product held one relic, so each annealed singly, with no partner to meet, and 27 of the 29 did so without first being seen to move. A column has low mobility and is not removed by pair annihilation, which is how the kinks of Refs. [8, 9] are removed.

VII. COOLING

Reality cools as it expands, so what matters for a long-lived relic is a race between annealing and cooling.

One relic; pre-registered as T25. A product with one column ($N = 96$) was cooled from $g = 1.25$, where every column anneals within 13,000 sweeps, to $g = 0.25$, where no move is accepted, by the same factor each block of 500 sweeps over $t_{\text{cool}} = 300$ to 100,000 sweeps, then held cold for 20,000 sweeps. The registered verdict was FREEZES IN if some cooling time keeps a majority of relics and the fastest keeps the most. The prediction recorded for the author was FREEZES IN; it was inferred from her stated positions, and she confirmed it only after the reading. Ours was FREEZES IN with a threshold between 1,000 and 10,000 sweeps. The relics surviving were 20/20, 17/17, 19/19, 15/19, 14/18 and 7/19 at $t_{\text{cool}} = 300, 1,000, 3,000, 10,000, 30,000$ and 100,000: FREEZES IN, with the registered "freeze-out time", the longest cooling that keeps a majority, at 30,000 sweeps, three to thirty times longer than we predicted. (That label is a threshold cooling duration. It is not the freeze-out time of the Kibble-Zurek literature, the instant at which a system can no longer keep up with a cooling of finite rate [11].) A relic that anneals at every fixed coupling tried is therefore kept by cooling. This is dynamical arrest in the sense of Ref. [8], where defects become immobile once the energy available to them falls below the local barrier; there the arrest follows a sudden change in an isolated system, and here it follows cooling by an external bath (our reading).

Many relics; pre-registered as T51. Tori $4 \times L$ were started at $g = 1.25$ and left to nucleate spontaneously

until at least 90% of vertices were at $d = 2$ and no stretch of unconverted torus remained. Each product was then cooled to $g = 0.25$ over t_{cool} and held, every cooling time starting from the same product. Times are in size-fair sweeps, $N/96$ sweeps of the chain, so that each local pair of edges is offered as often per unit time as at T25's 96 vertices. At $L = 256$ ($N = 1,024$) there were 80 products for $t_{\text{cool}} = 0$ to 100,000 and 40 for 300,000; at $L = 128$ and 512, 40 each at 10,000. All 160 tori converted, and at the end of the nucleation stage they held 1.27 ± 0.21 , 2.14 ± 0.15 and 4.90 ± 0.33 columns at $L = 128$, 256 and 512, one per 100 to 120 columns of torus, with 1.75, 2.67 and 3.90 pieces of other kinds.

The registered question was the shape of the fall of the column count C with cooling time, scored on $R(t) = C(10t)/C(t)$ at $t = 10,000$ and 30,000: GENTLE if both lie between 0.25 and 0.85, CLIFF if either is below 0.25. The author predicted GENTLE and so did we. The result is $R(10,000) = 0.347 \pm 0.040$ and $R(30,000) = 0.079 \pm 0.035$: CLIFF. Both predictions failed.

Three further results are in Table I and beside it. *The energy that stays is not in the columns.* The columns fall seventeenfold while the excess energy falls threefold, because the knots and double columns do not change with cooling time at all over the range run. Fitted from 10,000 upward, as registered before the reading, the columns fall as $t_{\text{cool}}^{-0.82 \pm 0.13}$ and the excess energy as $t_{\text{cool}}^{-0.34 \pm 0.06}$. These are apparent exponents over a decade and a half [12], and a constant population beside a vanishing one implies a floor that a power law does not have; a form with a floor has not been tested on registered data. *The size-fair clock missed one of its two points.* Its registered prediction was that the survival at 1,024 vertices would repeat T25's at 96: against T25's $C(30,000)/C(10,000) = 0.99$ and $C(100,000)/C(10,000) = 0.47$ we find 0.78 ± 0.05 and 0.35 ± 0.04 , the first outside the registered window of 0.20. On that clock the larger product loses its columns sooner. *The count per length does not depend on size* at the one cooling time tested: 0.0107, 0.0083 and 0.0093 columns per column of torus at $L = 128$, 256 and 512, each pair within two standard errors. At $L = 512$, 7 of 40 products ended as two separate pieces.

VIII. DISCUSSION

What is on the record for this model. The conversion leaves three kinds of small defect of fixed size, each a local minimum. Started at one place it leaves about one, whatever the size. Their number is set by how the conversion started: it grows with the number of nucleation events by less than one per event, and in proportion to the length of a torus that nucleates spontaneously. A lone column anneals at each fixed coupling tried and is arrested by cooling. Among many, slow cooling removes the columns and leaves the knots and double columns, which hold most of the energy that stays.

For the author's hypothesis, which asks whether a re-

mainder of such a change can last: in this model the remainder that lasts is not the common relic. It is the rarer kinds, whose number over the range run depends on how the change started and not on how slowly its product cooled. Whether they anneal on longer clocks is not known: the range ends at 300,000 size-fair sweeps. A first exit of nearly the same cost separates the kind that anneals (12 units) from the kinds that do not (11), so the first exit is not what decides; the barrier that would be the highest point on the lowest route to the product, which has not been computed for any of the three.

In the vocabulary of nonequilibrium defect kinetics [8–10] (our reading; those works study other models): the conversion is the decay of a metastable ordered state by nucleation and growth; a relic is a localized defect trapped in a local minimum; its number is set by the nucleation events, with a yield below one per event; in an external bath a column anneals singly and has low mobility; under cooling it is dynamically arrested. Two differences from the kinks studied there matter to anyone carrying results across. A kink separates two symmetry-related versions of one order and is removed only by meeting another, while a lone column anneals by itself. And the dynamics there is a physical equation of motion, while ours is a sampling chain whose rates are its own.

Is a relic a wall between two regions? (exploratory). One point cuts the other way. The product is not unique: every arrangement with two squares on every edge has zero energy, whatever its short period. A reading of the companion paper's 2,014 saved end states that keep a defect (the project record, entry O116) finds that about half hold a stretch wider than six edges around, and that in the 440 where a stretch six around and a wider one are both present the two never meet directly: a defect always lies between them. A relic of that sort is the joint between two regions of the product, the nearer relative of a domain wall, and T55's SEAM did not test for it, since a joint need carry no long-range edge. Which of the three kinds are joints, and whether the relics followed in Secs. VI and VII were, has not been read. Regions of differing short period joined at defects are also the majority outcome of the companion paper's isolated runs when the reservoir is small or the curling cost larger [1].

Not claimed: anything about the real universe; a relation between the model's sweep and a cosmological clock; that the knots and double columns never anneal; that the list of kinds is complete; a law for the fall with cooling time; or that a relic in three dimensions, or with the model's vertices treated as interchangeable, behaves the same way. Under a rule that also charges for loops of six, tested since, the product widens and holds flaws throughout; there the defects are the main obstacle to a clean product and not a rare remainder, and that rule is the subject of another paper.

t_{cool} (size-fair sweeps)	0	1,000	3,000	10,000	30,000	100,000	300,000
columns, relative to the end of nucleation	1.04	1.09	1.08	0.99	0.78	0.35	0.06
excess energy, share of the release	0.043	0.044	0.042	0.040	0.034	0.021	0.013
pieces of two vertices, per 80 products	91	101	101	102	101	103	106
pieces of eight vertices, per 80 products	16	15	16	16	16	15	18

TABLE I. Cooling 1,024-vertex products of spontaneous nucleation ($T51$, $L = 256$). The column count can exceed its value at the end of nucleation because columns can still form after that count. Standard errors are 0.03 to 0.06 on the first row and 0.003 on the second. The last column is from 40 products, scaled to 80. The pieces of two are the quarters of the knot of Sec. III; the pieces of eight are double columns.

ACKNOWLEDGMENTS

Code, analysis and text were drafted with Claude (Anthropic), with a scope review assisted by Codex (OpenAI), in conversation with the author, who directed and checked the work; no physicist has reviewed it. The repository holds every configuration, seed, result and pre-registration: <https://github.com/EmilySmithCreate/RecreationalPhysics>.

nAI), in conversation with the author, who directed and checked the work; no physicist has reviewed it. The repository holds every configuration, seed, result and pre-registration: <https://github.com/EmilySmithCreate/RecreationalPhysics>.

-
- [1] E. Smith, The curled torus burps: activated escape and conversion in a graph model of emergent geometry, technical manuscript (2026), [docs/papers/curled_torus/](#) in the repository above.
 - [2] C. A. Trugenberger, Combinatorial quantum gravity: geometry from random bits, JHEP 09, 045 (2017).
 - [3] C. Kelly, C. A. Trugenberger and F. Biancalana, Self-assembly of geometric space from random graphs, Class. Quantum Grav. 36, 125012 (2019).
 - [4] C. A. Trugenberger, Networks as the fundamental constituents of the universe, J. Phys. Complex. 6, 042001 (2025); arXiv:2512.17676.
 - [5] C. A. Trugenberger, Dark matter and dark energy in combinatorial quantum gravity, arXiv:2409.09385 (2024).
 - [6] M. Creutz, Microcanonical Monte Carlo simulation, Phys. Rev. Lett. 50, 1411 (1983). (Not read by us.)
 - [7] A. N. Kolmogorov, Izv. Akad. Nauk SSSR, Ser. Mat. 1, 355 (1937); W. A. Johnson and R. F. Mehl, Trans. AIME 135, 416 (1939); M. Avrami, J. Chem. Phys. 7, 1103 (1939). (Not read by us; details to verify.)
 - [8] H. Jang and G.-W. Chern, Suppressed coarsening after an interaction quench in the Holstein chain, arXiv:2602.05815 (2026).
 - [9] H. Jang, Y. Yang and G.-W. Chern, Anomalous coarsening and nonlinear diffusion of kinks in a one-dimensional quasi-classical Holstein model, arXiv:2512.07744 (2025). (Read in part.)
 - [10] Y. Fan, S. Zhang and G.-W. Chern, Machine learning nonequilibrium phase transitions in charge-density wave insulators, arXiv:2601.07583 (2026). (Read in part.)
 - [11] A. del Campo and W. H. Zurek, Universality of phase transition dynamics: topological defects from symmetry breaking, Int. J. Mod. Phys. A 29, 1430018 (2014); arXiv:1310.1600. (Read in part; journal details to verify.)
 - [12] A. Bora, S. S. Bakshi, H. Jang and G.-W. Chern, Cooperative domain-wall dynamics in a two-dimensional quasi-classical Holstein model, arXiv:2609.28436 (2026). (Read in part.)