Homological percolation in characteristic two: the Ising gauge case

Tejasvi Singh Tomar. Version 7, 2026-10-27.

Abstract. Duncan and Schweinhart proved that the i-dimensional plaquette random-cluster model with cluster weight q on the 2i-dimensional torus has a sharp homological percolation transition at the self-dual point, for coefficients in a field of characteristic other than 2. Their proof needs the point group of the lattice to act irreducibly on the i-th homology of the torus. In characteristic 2 it does not, and the most important case is excluded: $q = 2$ with F_2 coefficients, which represents the $(i-1)$-form Ising lattice gauge theory. We prove the theorem for $q = 2^r$ with $r \geq 1$ and F_2 coefficients. The representation-theoretic step is replaced by a transfer-matrix inequality, the classical positivity of the partition functions of 't Hooft's electric-flux sectors; read in the random-cluster representation, this inequality compares the probabilities of homological events. As a corollary, for the ferromagnetic $(i-1)$-form Ising gauge theory on the 2i-dimensional torus (Wegner's model when $i = 2$), the free energy of every nontrivial 't Hooft twist tends to 0 below the self-dual coupling and to infinity above it as the torus grows. For Bernoulli plaquette percolation ($q = 1$) the characteristic-2 case remains open for $i \geq 2$; we describe obstructions to proving it, and on the 4-torus we prove a computer-assisted divisibility theorem for the classes of integral giant cycles just above $p = 1/2$.

hexagon:2611.00468v1

1. Introduction and statement of results

Plaquette percolation and the plaquette random-cluster model extend bond percolation and the random-cluster model to i-dimensional cells [HS16, DKS20, DS22], and on a torus the natural question is homological: when do the open i-cells carry cycles that are nontrivial in the homology of the torus? For $q = 2^r$ and coefficients in F_2, the random-cluster model represents a lattice gauge theory of $(i-1)$-forms with gauge group ${F_2}^r$, and the question is then one about 't Hooft's twisted boundary conditions ['tH79]: the ratio of the gauge theory's partition functions with and without a twist by a mod-2 cocycle is the probability that the twist vanishes on every cycle of open plaquettes. As the torus grows, a nontrivial twist therefore costs no free energy if, with probability tending to 1, these cycles are all trivial in homology, and infinite free energy if, with probability tending to 1, they fill every homology class. For $q = 2$ the gauge theory is the Ising lattice gauge theory, Wegner's Z_2 gauge theory when $i = 2$.

Duncan and Schweinhart [DS22] proved that on the 2i-dimensional torus the transition is sharp and sits at the self-dual point $p_{sd} = \sqrt{q}/(1+\sqrt{q})$: with probability tending to 1, below p_{sd} no cycle of open plaquettes is nontrivial in homology, and above p_{sd} every homology class contains one. Their theorem is stated for coefficient fields F with char $F \neq 2$, since the proof relies on the irreducibility of $H_i(T^{2i}; F)$ under the point group of the lattice (the signed permutations of the coordinates), and in characteristic 2 this representation is reducible (§4). The most important case not covered was therefore the Ising gauge theory itself, $q = 2$ with F_2 coefficients; for Bernoulli plaquette percolation ($q = 1$), the question whether characteristic 2 can be included was raised by Duncan, Kahle and Schweinhart [DKS20, §9]. Theorem A below settles the cases $q = 2^r$, $r \geq 1$, with F_2 coefficients, and Corollary B states the consequence for the Ising gauge theory.

1.1 Setting and main results

Throughout, $X = T^d_N$ is the cubical d-torus of side $N \geq 2$, with $d = 2i$ unless stated otherwise. For $x \in (Z/NZ)^d$ and $I \subset [d] = \{1, \dots, d\}$, the cell $c(x, I) = x + \Sigma_{a\in I}[0, 1]e_a$ has *type* I and dimension $|I|$, and $X^{(k)}$ is the set of k-cells. Unless stated otherwise, coefficient groups are F_2-vector spaces, so orientations play no role; Proposition 10 and §8 use other coefficients and fix an orientation of every cell.

A configuration $\omega \subset X^{(i)}$ is a set of open i-cells (plaquettes), and $P = P(\omega)$ is the full $(i-1)$-skeleton together with ω. For $q \geq 1$ and $p \in (0, 1)$, the plaquette random-cluster model with F_2 coefficients [HS16, DS22] is the measure

$$\mu(\omega) = \mu_{N,p,q}(\omega) \propto p^{|\omega|}\,(1-p)^{|X^{(i)}|-|\omega|}\,q^{b_{i-1}(P;\,F_2)}.$$

We write $H := H_i(X; F_2) \cong {F_2}^K$ with $K = C(d, i)$; its basis is $e_I = [T_I]$, where T_I is the coordinate i-torus in the directions I. Let $\varphi\colon P \to X$ be the inclusion. The giant cycles are the i-cycles of P whose classes in H are nonzero, and $V(P) := \mathrm{im}(\varphi_*\colon H_i(P; F_2) \to H)$ is the subspace they span. Consider the events $A := \{V(P) \neq 0\}$ and $S := \{V(P) = H\}$.

Theorem A. Let $i \geq 1$, $d = 2i$, $q = 2^r$ with $r \geq 1$, and $p_{sd} = \sqrt{q}/(1+\sqrt{q})$. Then, as $N \to \infty$,

$\mu_{N,p,q}(A) \to 0$ for $p < p_{sd}$, $\mu_{N,p,q}(S) \to 1$ for $p > p_{sd}$.

The analogue of Theorem A with coefficients in a field F of characteristic other than 2 is [DS22, Theorem 8]. For $i = 1$, Theorem A is the torus form of $p_c = p_{sd}$ for the planar random-cluster model [BDC12]. The new case is $F = F_2$, $i \geq 2$, and $q = 2^r$ with $r \geq 1$.

Corollary B. Consider the $(i-1)$-form Ising gauge theory on ${T^{2i}}_N$. Its spins are $\sigma \in \{\pm 1\}^{X^{(i-1)}}$, and its weight is $\exp(J\,\Sigma_c\,\sigma_c)$, where σ_c is the product of the spins on the boundary of the i-cell c. For a class $n \in H^i(X; F_2)$, let $Z_J(n)$ be the partition function with the couplings of the i-cells in a representative cocycle of n reversed. Then, with $J_{sd} = \frac{1}{2}\ln(1+\sqrt{2})$, for every $n \neq 0$:

$Z_J(n)/Z_J(0) \to 1$ for $0 \leq J < J_{sd}$, $Z_J(n)/Z_J(0) \to 0$ for $J > J_{sd}$.

As $N \to \infty$, then, the free energy $-\log(Z_J(n)/Z_J(0))$ of every nontrivial twist tends to 0 for $0 \leq J < J_{sd}$ and to ∞ for $J > J_{sd}$.

1.2 Idea of the proof

For $n \in H^i(X; F_2)$ let $R(n)$ be the probability that n vanishes on $V(P)$; for $q = 2$, $J \geq 0$ and $p = 1 - e^{-2J}$ it is the ratio $Z_J(n)/Z_J(0)$ of Corollary B (Lemma 3). The proof of [DS22] uses irreducibility to bound $\mu(S)$ from below by a constant times a power of $\mu(A)$ [DS22, Lemma 41]. For $q = 2^r$ we obtain such a bound from one inequality instead, Lemma 4: if n is written in the basis dual to (e_I), deleting its components on the coordinate tori that contain a fixed direction can only increase $R(n)$. Taking that direction as time, Lemma 4 says that adding a temporal twist to a spatial one cannot increase the twisted partition function. In 't Hooft's terms, the difference is, up to a factor 2, the partition function of the odd electric-flux sector, which is nonnegative; we prove this with a transfer matrix (§3), as in [BS83, TY85]. With the FKG inequality and the symmetry of the torus under coordinate permutations, Lemma 4 gives the spin-up bound of Lemma 5, $\mu(S) \geq (\mu(A)/K)^{2^K-1}$, without representation theory. The duality and influence inequalities of [DS22], which hold in every characteristic, then complete the proof (§5). For $q = 2^r$ and F_2 coefficients the same argument, with duality between degrees i and $d - i$, gives a sharp threshold for A and S in every dimension $d \geq 2$ and degree $1 \leq i \leq d - 1$, around the value $\lambda(N)$ of p at which $\mu(A) = 1/2$ (Theorem C).

1.3 Bernoulli plaquette percolation (q = 1)

For $q = 1$ there is no gauge theory and the inequality of Lemma 4 is not known, and for $i \geq 2$ it remains open whether the conclusion of Theorem A holds with F_2 coefficients. Section 8 proves a theorem about integral coefficients instead. For Bernoulli plaquette percolation P on T^4_N ($q = 1$, $i = 2$), let $\Lambda_P \subseteq H_2(T^4_N; Z) \cong Z^6$ be the lattice of classes of integral 2-cycles of P. Theorem W states that for every integer $M \geq 2$ there is $w_M > 0$ such that, for $p \in (1/2, 1/2 + w_M)$, with probability tending to 1, rank $\Lambda_P = 6$ and $\Lambda_P \subseteq MZ^6$; the rank statement comes from [DKS20], and the new part is the divisibility. The proof is an Aizenman–Grimmett strict-shift argument [AG91, BBR14] whose local step rests on finite computer checks and on an induction (§§8.6–8.7).

Sections 7 and 9 give bounds and obstructions for the F_2 question at $q = 1$. Among them is a family of product measures on 4-tori of side divisible by 8 that is invariant under coordinate permutations and exactly self-dual (§9.2). With high probability its giant-cycle space V equals one fixed subspace for all parameters in an interval around its self-dual point, and this subspace is Lagrangian for the intersection form. A proof at $q = 1$ therefore has to exploit a property that this family does not have, such as invariance under unit translations.

1.4 Organization

Sections 2–5 prove Theorem A, Corollary B and Theorem C; with the remarks of §6 they can be read on their own. Section 7 treats $q = 1$ and other real q, Section 8 proves Theorem W, and Section 9 studies how, with probability tending to 1 along a subsequence of N, V could contain a fixed nonzero subspace for all p in an interval below 1/2. A dagger † marks a statement that rests on a computation; the Appendix lists, for each, the program and the log in the supplementary material.

2. Twists and twist ratios

This section identifies twist ratios with probabilities of the events D_n (Lemma 3).

Coordinate cocycles. For an i-set I, let $\Sigma_I := \{c(x, I) : x_a = 0 \text{ for all } a \in I\}$ and $\eta_I := 1_{\Sigma_I} \in C^i(X; F_2)$. Then η_I is a cocycle: an (i+1)-cell of type $I \cup \{b\}$ has exactly two faces of type I, namely $c(x, I)$ and $c(x + e_b, I)$, and they are both in Σ_I or both outside it. Moreover $\langle \eta_I, T_J \rangle = \delta_{IJ}$, so $[\eta_I] = e_I^*$, the dual basis of $H^i := H^i(X; F_2)$. We write $n = \Sigma_I\, n_I\, e_I^*$ and $|n| = \#\{I : n_I = 1\}$.

Events D_n. For $n \in H^i$ let $D_n := \{\langle n, V(P) \rangle = 0\}$.

Lemma 1. (a) D_n is decreasing. (b) $D_m \cap D_n \subseteq D_{m+n}$. (c) $A = \bigcup_{n \neq 0} D_n^c$ and $S = \bigcap_{n \neq 0} D_n^c$.

Proof. (a) and (b) are immediate. For (c), the evaluation pairing $H^i \times H \to F_2$ is perfect, so $V = H$ if and only if $V^\perp = 0$, and $V \neq 0$ if and only if $V^\perp \neq H^i$. ■

Lemma 2 (FKG) [HS16; DS22, Thm 16]. For $q \geq 1$, μ satisfies the FKG lattice condition. Hence increasing events are positively correlated, and so are decreasing events. ■

The twisted gauge theory. Let $G = F_2^r$, so $|G| = q$, let $\beta \geq 0$, and let $k(x) := \exp(\beta \cdot 1\{x = 0\})$ for $x \in G$. For a cocycle $\tau \in Z^i(X; G)$ put

$$Z(\tau) := \Sigma_{f \in C^{i-1}(X; G)}\, \Pi_{\sigma \in X^{(i)}}\, k(\delta f(\sigma) + \tau(\sigma)).$$

Lemma 3. (a) $Z(\tau)$ depends only on the class $[\tau] \in H^i(X; G)$. (b) Let $p = 1 - e^{-\beta}$, let $\eta \in Z^i(X; F_2)$, and let $g \in G \setminus \{0\}$. Then $Z(\eta g)/Z(0) = \mu_{N,p,q}(D_{[\eta]})$.

Proof. (a) The substitution $f \mapsto f + h$ maps $Z(\tau)$ to $Z(\tau + \delta h)$.

(b) Since $k(x) = e^{\beta}[(1-p) + p\cdot 1\{x = 0\}]$, expanding the product gives

$Z(\tau) = e^{\beta|X^{(i)}|} \Sigma_{\omega} p^{|\omega|}(1-p)^{|X^{(i)}|-|\omega|} \cdot \#\{f : \delta f = \tau \text{ on } \omega\}$.

Fix ω and $P = P(\omega)$. Since $C^{i-1}(X; G) = C^{i-1}(P; G)$, the solution set of $\delta f = \tau$ on ω is either empty or a coset of $Z^{i-1}(P; G)$, and it is nonempty exactly when $[\tau|_P] = 0$ in $H^i(P; G)$. The coset has $|Z^{i-1}(P; G)| = |\delta C^{i-2}(X; G)| \cdot |H^{i-1}(P; G)|$ elements. The first factor does not depend on ω, and the second equals $|G|^{b_{i-1}(P;\, F_2)} = q^{b_{i-1}}$, coordinatewise in $G = {F_2}^r$. When $\tau = \eta g$, we have $[\tau|_P] = 0$ if and only if $[\eta|_P] = 0$ in $H^i(P; F_2)$: read off a coordinate j with $g_j = 1$ for one direction, and use $h = h'g$ for the other. Finally, by the universal coefficient theorem over a field, evaluation identifies $H^i(P; F_2)$ with $\mathrm{Hom}(H_i(P; F_2), F_2)$, and $\langle \eta, z\rangle = \langle [\eta], \varphi_*[z]\rangle$ for a cocycle η on X and a cycle z in P. Hence $[\eta|_P] = 0$ if and only if $D_{[\eta]}$ holds, and $Z(\eta g) = C\, \Sigma_{\omega}\, p^{|\omega|}(1-p)^{|X^{(i)}|-|\omega|}\, q^{b_{i-1}(P)}\, 1_{D_{[\eta]}}$ with C independent of ω and η. ■

From now on write $R(n) := \mu(D_n)$, which equals $Z(\eta g)/Z(0)$ for any representative η of n.

3. The transfer-matrix inequality

For a direction $\mu \in [d]$ split $n = s_\mu(n) + t_\mu(n)$, with $s_\mu(n) := \Sigma_{I \not\ni \mu} n_I e_I^*$ ("spatial") and $t_\mu(n) := \Sigma_{I \ni \mu} n_I e_I^*$ ("temporal").

Lemma 4. Let $q = 2^r$ with $r \geq 1$. For every $n \in H^i$ and every $\mu \in [d]$, $R(n) \leq R(s_\mu(n))$.

Proof. Let Y be the $(d-1)$-dimensional cubical torus $\{x_\mu = 0\}$, and identify each slice $\{x_\mu = t\}$ with Y. Spatial cells (of type $I \not\ni \mu$) correspond to $Z_N \times Y^{(|I|)}$, and temporal cells $c \times [t, t+1]_\mu$ to $Z_N \times Y^{(|I|-1)}$. A cochain $f \in C^{i-1}(X; G)$ becomes $(f_t, g_t)_{t \in Z_N}$, with $f_t \in C^{i-1}(Y; G)$ and $g_t \in C^{i-2}(Y; G)$; g_t is absent if $i = 1$. The boundary of $c \times [t, t+1]$ is $c \times \{t\} + c \times \{t+1\} + (\partial c) \times [t, t+1]$, so

$\delta f(t, c) = \delta_Y f_t(c)$ for $c \in Y^{(i)}$, $\delta f(c \times [t, t+1]) = f_t(c) + f_{t+1}(c) + \delta_Y g_t(c)$ for $c \in Y^{(i-1)}$.

To represent the twist, we fix $g \in G \setminus \{0\}$ and use $\tau := \Sigma_I n_I \eta_I g$. Its spatial part is the same in every slice: $\tau(t, c) = a(c)$ with $a := \Sigma_{I \not\ni \mu} n_I {\eta^Y}_I g \in C^i(Y; G)$. Its temporal part is supported on the cells $c \times [0, 1]$ and equals $\theta(c)$ there, with $\theta := \Sigma_{I \ni \mu} n_I {\eta^Y}_{I \setminus \{\mu\}} g \in C^{i-1}(Y; G)$. Here ${\eta^Y}_J := 1\{\text{cells of Y of type J with } x_J = 0\}$. By the argument for η_I applied inside Y, $\delta_Y \theta = 0$.

On $\mathcal{H} := \mathbb{C}^{C^{i-1}(Y;\, G)}$ we use three operators: the multiplication operator M by $S(\varphi) := \Pi_{c \in Y^{(i)}} k(\delta_Y \varphi(c) + a(c)) > 0$, the operator $\mathcal{K}$ with kernel $\mathcal{K}(\varphi, \varphi') := \Sigma_{\gamma \in C^{i-2}(Y;\, G)} \Pi_{c \in Y^{(i-1)}} k(\varphi(c) + \varphi'(c) + \delta_Y \gamma(c))$, and the shift $(U_\theta \psi)(\varphi) := \psi(\varphi + \theta)$. Summing over the temporal variables g_t gives

$Z(\tau) = \Sigma_{(f_t)} \Pi_t S(f_t) \Pi_{t \neq 0} \mathcal{K}(f_t, f_{t+1}) \cdot \mathcal{K}(f_0, f_1 + \theta) = \mathrm{Tr}(U_\theta\, (M\mathcal{K})^N)$.

These operators have two properties.

- (P1) $\mathcal{K} \geq 0$. Indeed, $\mathcal{K} = (\otimes_c \kappa)\, \Sigma_\gamma\, U_{\delta_Y \gamma}$, where κ is the operator on $\mathbb{C}^G$ with kernel $k(x + y)$, and κ is convolution by k because $x + y = x - y$ in G. Its eigenvalues are $\hat{k}(\chi) = \Sigma_x k(x)\chi(x) = |G|\cdot 1\{\chi = 1\} + (e^\beta - 1) \geq 0$, so $\otimes\kappa \geq 0$. The sum $\Sigma_\gamma U_{\delta\gamma}$ is a positive multiple of the orthogonal projection onto vectors invariant under the subgroup $\delta C^{i-2}(Y; G)$, and the two factors commute, because translations commute with convolutions. Hence $\mathcal{K} \geq 0$.
- (P2) U_θ commutes with $\mathcal{K}$ by translation invariance, and with M because $\delta_Y(\varphi + \theta) = \delta_Y \varphi$.

We now put $T := M^{1/2}\, \mathcal{K}\, M^{1/2} \geq 0$. By (P2), $\mathrm{Tr}(U_\theta\, (M\mathcal{K})^N) = \mathrm{Tr}(U_\theta\, T^N)$. The spatial twist $s_\mu(n)$ has the same a and $\theta = 0$, so $Z(\tau_{sp}) = \mathrm{Tr}(T^N)$. Since U_θ is unitary and $T^N \geq 0$,

$Z(\tau) = \mathrm{Tr}(U_\theta\, T^N) \leq \|U_\theta\| \cdot \|T^N\|_1 = \mathrm{Tr}(T^N) = Z(\tau_{sp})$.

Dividing by $Z(0)$ and using Lemma 3 gives $R(n) \le R(s_\mu(n))$. ■

Remark. $Z(\tau_{sp}) - Z(\tau) = 2\,\mathrm{Tr}(P_-\,T^N) \ge 0$, where $P_- = (1 - U_\theta)/2$: the difference is the partition function of the odd electric-flux sector. Lemma 4 is the classical positivity of 't Hooft's electric-flux partition functions ['tH79], proved with transfer matrices in [BS83, TY85] for lattice gauge theories. What is new here is reading it in the random-cluster representation (Lemma 3), where it becomes an inequality between homological events.

4. Spin-up

Lemma 4 and the FKG inequality give a lower bound on $\mu(S)$ in terms of $\mu(A)$.

Lemma 5. Let $q = 2^r$ with $r \ge 1$. For every N and p, with $K = C(d, i)$:

$\mu(S) \ge (\ \mu(A)\ /\ K\ )^{2^K - 1}$.

Proof.

- (i) Coordinate permutations are automorphisms of X that preserve μ and map Σ_I to $\Sigma_{\pi(I)}$. Hence $R(e_I^*) =: r$ is the same for all I.
- (ii) Let $n \ne 0$ and choose I with $n_I = 1$. Apply s_μ successively for the i directions $\mu \notin I$. Each intermediate twist still contains the I-component, and the last one is e_I^*, so Lemma 4 gives $R(n) \le r$.
- (iii) By Lemma 1(b), $D_n \supseteq \bigcap_{I:\ n_I=1} D_{e_I^*}$, and Lemma 2 for decreasing events gives $R(n) \ge r^{|n|} \ge r^K$.
- (iv) By Lemma 1(c), Lemma 2 for increasing events, and (ii), $\mu(S) = \mu(\bigcap_{n\ne 0} D_n^c) \ge \Pi_{n\ne 0}(1 - R(n)) \ge (1 - r)^{2^K - 1}$.
- (v) If $V \ne 0$, some nonzero $v \in V$ has a coordinate $v_I = 1$, so e_I^* does not vanish on V. Hence $A = \bigcup_I D^c_{e_I^*}$ and $\mu(A) \le K(1 - r)$.

Combining (iv) and (v) proves the lemma. We do not need step (iii) for the bound, but it shows that every $R(n)$ lies between r^K and r. ■

Lemma 5 replaces [DS22, Lemma 41] and [DKS20, Lemma 10 and Prop. 11], which give $\mu(S) \ge c\ \mu(A)^{c'}$ only when $H_i(T^d; F)$ is an irreducible representation of the point group. This fails for $F = F_2$: coordinate sign changes act trivially mod 2, so the point group acts on $\{e_I\}$ by permutations, and $\Sigma_I\, e_I$ spans an invariant line [DKS20, Prop. 11]. In characteristic 2 the symmetry argument therefore cannot exclude an "oblique" phase. In such a phase V(P) would be a Lagrangian subspace such as $\mathrm{span}\{e_{12} + e_{34},\ e_{13} + e_{24},\ e_{14} + e_{23}\}$, which is invariant under coordinate permutations and self-dual. Lemma 4 excludes it: there $R(e_{12}^* + e_{34}^*) \to 1$ while $R(e_{12}^*) \to 0$.

5. Proof of Theorem A

We use two facts from [DS22] that hold for every $q \ge 1$ and every field.

Fact D (duality) [DKS20, Lemma 9 (Duality Lemma); DS22, Thm 18]. Let $P\bullet$ be the dual complex in the dual torus $X\bullet \cong X$.

- (a) $\mathrm{rank}\ \varphi_* + \mathrm{rank}\ \psi_* = K$ for every configuration. Hence $A(P) \iff \neg S(P\bullet)$, and at least one of $A(P)$, $A(P\bullet)$ holds.

- (b) The balanced measure $\tilde{\mu}_{N,p,q} \propto q^{-\operatorname{rank}\varphi_*/2}\,\mu_{N,p,q}$ satisfies: $P\bullet$ has law $\tilde{\mu}_{N,p^*,q}$ when P has law $\tilde{\mu}_{N,p,q}$, where $p^* = (1-p)q/((1-p)q + p)$. Here $P\bullet$ is the (d−i)-dimensional model on $X\bullet$, which is i-dimensional because $d = 2i$. The translation $X\bullet \to X$ by $(½, \dots, ½)$ carries A and S to themselves.
- (c) $d\tilde{\mu}/d\mu \in [q^{-K/2}, q^{K/2}]$.

Fact T (sharp threshold) [GG06; DS22, Thm 39]. $\mu_{N,p,q}$ is monotonic, and the automorphism group of X acts transitively on $X^{(i)}$. Hence for every increasing automorphism-invariant event E,

$$d/dp\ \mu_p(E) \geq c(p) \cdot \min\{\mu_p(E), 1 - \mu_p(E)\} \cdot \log |X^{(i)}|,$$

with $c(p) > 0$ bounded below on compact subsets of (0, 1). Indeed, $\mu_p(\omega_\sigma = 1 \mid \text{rest}) \in [p/(p + q(1-p)), p]$, so the one-plaquette variance is bounded below by a constant depending only on p and q.

Proof of Theorem A.

1. At $p = p_{sd} = p_{sd}{}^*$, Fact D(a) and (b) give $1 \leq \tilde{\mu}(A(P)) + \tilde{\mu}(A(P\bullet)) = 2\tilde{\mu}(A)$. By Fact D(c), $\mu_{N,p_{sd}}(A) \geq q^{-K/2}/2$ for every N.
2. By Lemma 5, $\mu_{N,p_{sd}}(S) \geq \delta_0 := (q^{-K/2}/(2K))^{2^K-1} > 0$, uniformly in N.
3. S is increasing and automorphism-invariant, and $\mu_p(S)$ is nondecreasing in p. Integrating Fact T from p_{sd}, we see that $\mu_p(S)$ reaches 1/2 within $\Delta p \leq \log(1/(2\delta_0))/(c \log|X^{(i)}|)$, and that $1 - \mu_p(S)$ then decays like $\exp(-c \log|X^{(i)}| \cdot \Delta p)$. Hence $\mu_{N,p}(S) \to 1$ for every $p > p_{sd}$.
4. For $p < p_{sd}$ we have $p^* > p_{sd}$, so $\mu_{N,p^*}(S) \to 1$ by step 3, and also $\tilde{\mu}_{N,p^*}(S) \to 1$ by Fact D(c). By Fact D(b), $\tilde{\mu}_{N,p}(S(P\bullet)) \to 1$. By Fact D(a), $\tilde{\mu}_{N,p}(A) = 1 - \tilde{\mu}_{N,p}(S(P\bullet)) \to 0$. By Fact D(c), $\mu_{N,p}(A) \to 0$. ■

Theorem C (sharp threshold in every dimension). Let $d \geq 2$, $1 \leq i \leq d - 1$, $q = 2^r$ with $r \geq 1$, and F_2 coefficients. Define $\lambda(N)$ by $\mu_{N,\lambda(N),q}(A) = 1/2$. Then for every $\varepsilon > 0$, $\mu_{N,\lambda(N)-\varepsilon,q}(A) \to 0$ and $\mu_{N,\lambda(N)+\varepsilon,q}(S) \to 1$. This is the characteristic-2 analogue of [DS22, Prop. 42 and Thm 11].

Proof. Lemmas 3–5 do not use $d = 2i$. In Lemma 5(ii) the $d - i$ projections s_μ, $\mu \notin I$, leave only components $I' \subseteq I$ with $|I'| = i$, that is, $e_I{}^*$. Hence $\mu(S) \geq (\mu(A)/K)^{2^K-1}$ with $K = C(d, i)$. At $p = \lambda(N)$ this gives $\mu(S) \geq (2K)^{-(2^K-1)}$, and Fact T gives $\mu_{\lambda+\varepsilon}(S) \to 1$.

For the event A we pass to the dual complex $P\bullet$, which is the (d−i)-dimensional plaquette random-cluster model with the same q on $X\bullet$ (balanced form of [DS22, Thm 18]). The duality lemma $\operatorname{rank}\varphi_* + \operatorname{rank}\psi_* = K$ holds for every d and i, so $A(P) \iff \neg S(P\bullet)$. From $\mu_\lambda(A) = 1/2$ and Fact D(c), the dual model at λ^* has $\mu(S\bullet) \geq q^{-K}/2$, hence $\mu(A\bullet) \geq q^{-K}/2$. Lemma 5 for the dual model (dimension $d - i$, same q) gives $\mu(S\bullet) \geq \delta > 0$ at λ^*, and Fact T for the dual model then gives $\mu(S\bullet) \to 1$ at $(\lambda - \varepsilon)^* > \lambda^*$. By duality, $\mu_{\lambda-\varepsilon}(A) \to 0$. ■

Proof of Corollary B. Since $1\{\sigma = 1\} = (1+\sigma)/2$, the Ising weight $\exp(J\,\Sigma\,\sigma_c)$ is the gauge weight of §2 with $q = 2$ and $\beta = 2J$, so $p = 1 - e^{-2J}$ and p_{sd} corresponds to $J_{sd} = ½\ln(1+\sqrt{2})$. By Lemma 3, $Z_J(n)/Z_J(0) = \mu(D_n)$. For $0 \leq J < J_{sd}$ we have $\mu(A) \to 0$ and $A^c \subseteq D_n$; for $J > J_{sd}$ we have $\mu(S) \to 1$ and $S \subseteq D_n{}^c$. ■

6. Remarks

1. We indicate how the argument extends, with signs, to $q = \ell^r$ and $F = F_\ell$ for any prime ℓ, with gauge group $G = F_\ell{}^r$. The slicing formula becomes $\delta f(c \times [t, t+1]) = \pm(f_{t+1} - f_t)(c) + \delta_Y\, g_t(c)$,

and η_I needs a fixed orientation. The kernel k is even, so $\mathcal{K}$ is again a convolution with nonnegative Fourier coefficients. U_θ is unitary but no longer an involution, so the $\pm$ split of the remark after Lemma 4 becomes a Z_ℓ Fourier split. Finally $R(e_I^*) = R(-e_I^*)$, and $2^K - 1$ becomes $\ell^K - 1$. For odd ℓ this reproves [DS22, Thm 8] for these q without representation theory.

2. Theorem A concerns tori only; the infinite-volume area/perimeter law [DS26, Conjecture 1] remains open for d = 4.

7. Other values of q: reductions, partial results and obstructions

For real $q \geq 1$ with F_2 coefficients (in particular q = 1, Bernoulli plaquette percolation [DKS20]; whether the restriction on the characteristic can be removed there is asked in [DKS20, §9]) there is no gauge theory, and Lemma 4 is proved only for $q = 2^r$. The rest of the argument does not use the gauge theory. In this section we reduce the theorem for such q to a single inequality and study that inequality. For q = 1 we also prove partial results with coefficients F_2 and Z/4; §8 continues with integral coefficients.

7.1 Reduction to Lemma 4 and to a crossing estimate

Theorem D (reduction). Fix d, i, real $q \geq 1$, and F_2 coefficients. Suppose $R(n) \leq R(s_\mu(n))$ holds for $\mu_{N,p,q}$, for all N, p, n and μ. Then Lemma 5, Theorem C and, when d = 2i, Theorem A hold for this q.

Proof. Lemma 5 uses only Lemma 4, the FKG inequality and coordinate symmetry, all of which hold for every real $q \geq 1$ [HS16]. Section 5 uses Facts D and T, which hold for every real $q \geq 1$ [DS22]; for q = 1 they are exact self-duality at p = 1/2 and the Friedgut–Kalai theorem. ■

Proposition 6 (an unconditional window for Bernoulli). For Bernoulli plaquette percolation with F_2 coefficients on $T^{2i}{}_N$:

$P_p(A) \to 0$ for $p < \sqrt{2} - 1$, $P_p(S) \to 1$ for $p > 2 - \sqrt{2}$.

Proof. Given the other plaquettes, a plaquette is open with conditional probability p if it closes an i-cycle and with probability $p/(p + q(1-p))$ if it kills an (i−1)-cycle. Holley's inequality therefore gives $\text{Bernoulli}(p'/(p' + 2(1-p'))) \leq_{st} \mu_{p',2} \leq_{st} \text{Bernoulli}(p')$. Since A and S are increasing, for $p > p_{sd}(2) = 2 - \sqrt{2}$ we get $P_p(S) \geq \mu_{p,2}(S) \to 1$ by Theorem A. For $p < \sqrt{2} - 1$, put $p' = 2p/(1+p) < p_{sd}(2)$; then $P_p(A) \leq \mu_{p',2}(A) \to 0$. ■

In the proof of Theorem A, Lemma 4 is used only to show that every R(n) stays bounded away from 1 at $p = p_{sd}$. A lower bound on the probability of a coordinate giant cycle at the self-dual point gives the same conclusion, as we now show. For d = 2i we use the following duality. By Fact D(a), V(P) and $V(P^\bullet)$ have complementary dimensions, and they pair to zero: a primal i-cell meets a dual i-cell transversally only when the two cells are dual to each other, and $P^\bullet$ contains the duals of the closed i-cells only. Hence $V(P^\bullet) = V(P)^\perp$ under the intersection pairing, which after the translation of Fact D(b) is $e_I \cdot e_J = 1\{I \cap J = \emptyset\}$. Poincaré duality sends e_I^* to e_{I^c}, so for every $n \in H^i$ the event D_n means exactly that $PD(n) \in V(P^\bullet)$.

For d = 4 we read a class as a set of edges of the complete graph on {1, 2, 3, 4}, where the plane {a, b} is the edge ab. Let $u := \Sigma_I e_I$ be the full edge set, and let C be the span of the three 4-cycles, such as $e_{12} + e_{14} + e_{23} + e_{34}$. The three 4-cycles sum to 0, so dim C = 2.

Lemma 6′. Let d = 4, i = 2. Every nonzero subspace of H that is invariant under coordinate permutations contains u or contains C.

Proof. Let U be such a subspace and take $v \in U \setminus \{0, u\}$. Some transposition $\tau = (ab)$ moves v: the permutations act transitively on the six planes, so 0 and u are the only vectors fixed by all

transpositions. Let $\{c, d\} = [4] \setminus \{a, b\}$. Then τ fixes the edges ab and cd and swaps ac $\leftrightarrow$ bc and ad $\leftrightarrow$ bd. Hence $w := v + \tau v$ is a nonzero element of U, equal to one of ac + bc, ad + bd, or ac + bc + ad + bd. In the last case w is a 4-cycle. In the first case $w + (cd)w = ac + bc + ad + bd$ is a 4-cycle, and the second case is the same. Since coordinate permutations act transitively on the three 4-cycles, $C \subseteq U$. ■

Proposition 7 (reduction to a crossing estimate). Let $d = 4$, $i = 2$, $q \geq 1$, and F_2 coefficients. Suppose $r_0 := \inf_N \mu_{N,p_{sd},q}(e_{12} \in V) > 0$. Then Theorem A holds for this q.

Proof. Work at $p = p_{sd}$ and write $\tilde{\mu}$ for the balanced measure of Fact D.

- (i) If $a \cdot b = 1$, then $a \in V(P)$ and $b \in V(P\,\bullet)$ cannot both hold. $P\,\bullet$ has law $\tilde{\mu}$ when P does, so $\tilde{\mu}(a \in V) + \tilde{\mu}(b \in V) \leq 1$.
- (ii) The class e_{12} meets u in the single disjoint plane 34. It also meets the 4-cycle $e_{12} + e_{14} + e_{23} + e_{34} \in C$ in the plane 34. By (i) and Fact D(c), $\mu(u \notin V) \geq q^{-K} r_0$ and $\mu(C \not\subset V) \geq q^{-K} r_0$.
- (iii) Let $c \neq 0$ and $\varepsilon := \mu(c \notin V)$. By symmetry and a union bound, the span U_c of the 24 images of c satisfies $\mu(U_c \subseteq V) \geq 1 - 24\varepsilon$. By Lemma 6′, U_c contains u or C, so (ii) gives $\varepsilon \geq q^{-K} r_0/24 =: \delta$. Thus $\mu(c \in V) \leq 1 - \delta$ for every $c \neq 0$.
- (iv) For $n \neq 0$, $D_n{}^c = \{PD(n) \notin V(P\,\bullet)\}$. Two uses of Fact D(c) and the self-duality of $\tilde{\mu}$ give $\mu(D_n{}^c) \geq q^{-K}\,\mu(PD(n) \notin V) \geq q^{-K}\,\delta$.
- (v) By Lemma 1(c) and Lemma 2, $\mu(S) \geq \Pi_{n\neq 0}\,\mu(D_n{}^c) \geq (q^{-K}\,\delta)^{2^K-1}$, uniformly in N. Steps 3 and 4 of the proof of Theorem A now apply verbatim. ■

For $q = 1$, step (i) and the symmetry $e_{12} \leftrightarrow e_{34}$ give $\mu_{1/2}(e_{12} \in V) \leq 1/2$, while Monte Carlo at $p = 1/2$ (the samples of the $d = 4$ study in §7.5) gives $\mu(e_I \in V) = 0.424$, 0.436 and 0.452 for $N = 4$, 6 and 8, averaged over the six coordinate classes e_I.† The hypothesis of Proposition 7 is a box-crossing (Russo–Seymour–Welsh) bound for giant surfaces at the self-dual point; its $i = 1$ analogue is the classical crossing bound that underlies Kesten's proof of $p_c = 1/2$. A crossing probability that decays slower than every power of N would suffice. If $r_0(N) := \mu_{N,p_{sd},q}(e_{12} \in V)$ satisfies $\log(1/r_0(N)) = o(\log N)$, steps (i)–(v) give $\mu_{N,p_{sd}}(S) \geq (q^{-2K}\,r_0(N)/24)^{2^K-1} = N^{-o(1)}$, and step 3 of the proof of Theorem A still reaches $\mu_p(S) \to 1$ for every $p > p_{sd}$, because it needs only $\log(1/\mu_{p_{sd}}(S)) = o(\log |X^{(i)}|)$.

7.2 Partnerless classes

In this subsection $d = 4$ and $i = 2$. Call a nonzero class $z \in H$ *partnered* if $\langle z, gz\rangle = 1$ for some coordinate permutation g. At $p = 1/2$, partnered classes satisfy $\mu(z \in V) \leq 1/2$ by step (i) of the proof of Proposition 7, because $\mu(gz \in V) = \mu(z \in V)$.

Proposition 8 (the partnerless classes). A nonzero class is partnerless iff it lies in $C^\perp$. Read as an edge set of K_4, this means that its four vertex degrees have equal parity. $C^\perp$ is the 4-dimensional permutation-invariant subspace spanned by the triangles and the stars; its 15 nonzero elements are the matchings, triangles, stars, 4-cycles and u. For Bernoulli percolation, $P_p(V \not\subset C^\perp) \to 0$ for every $p < 1/2$ and $P_p(C \subseteq V) \to 1$ for every $p > 1/2$. Consequently, Theorem A at $q = 1$ holds iff $P_p(V \cap C^\perp \neq 0) \to 0$ for every $p < 1/2$.

Proof. For the 4-cycle c with diagonals 13 and 24, $\langle z, c\rangle = \deg_z(1) + \deg_z(3) \bmod 2$; the other two 4-cycles give $\deg_z(1) + \deg_z(2)$ and $\deg_z(1) + \deg_z(4)$. Thus $z \in C^\perp$ iff all degrees have equal parity, and a check of the ten permutation orbits of nonzero classes shows that this is exactly the set of partnerless classes.†

Each event $\{z \in V\}$ is increasing and invariant under translations, whose orbits on plaquettes have size N^4. By the Friedgut–Kalai theorem in the form of [BR06, Lemma 7], $\mu_{1/2}(z \in V) \leq 1/2$ forces $\mu_p(z \in V)$

$\to 0$ for every $p < 1/2$, and a union bound over the 48 partnered classes gives the first claim. For $p > 1/2$ the dual is Bernoulli$(1-p)$, so $V(P\,\bullet\,) \subseteq C^{\perp}$ whp, and $V = V(P\,\bullet\,)^{\perp} \supseteq C$. ■

The next proposition shows that the obstruction is not an artefact of the cubic lattice.

Proposition 9 (periodic structures). Let $\Lambda \subset R^4$ be a lattice and $G \subset GL(\Lambda)$ a finite group. Then $\Lambda^2(\Lambda/2\Lambda) \cong H_2(R^4/\Lambda;\, F_2)$ contains a nonzero z with $\langle z,\, gz\rangle = 0$ for all $g \in G$, where $\langle\cdot,\cdot\rangle$ is the mod-2 intersection form.

Proof. G preserves a positive definite integral form B; divide B by a power of 2 until $\bar{B} := B \bmod 2 \neq 0$. If $\bar{B}$ is alternating, it is a nonzero G-fixed element of $\Lambda^2(\Lambda/2\Lambda)^* \cong \Lambda^2(\Lambda/2\Lambda)$, and a fixed z has $\langle z,\, gz\rangle = \langle z,\, z\rangle = 0$. Otherwise $\ell(x) := \bar{B}(x,\, x)$ is a nonzero G-invariant linear form, and $\Lambda^2(\ker \ell)$ is G-invariant and isotropic, because $x \wedge y \wedge x' \wedge y' \in \Lambda^4(\ker \ell) = 0$. Any $z \in \Lambda^2(\ker \ell) \setminus \{0\}$ works. ■

For a periodic self-dual cell structure on R^4/Λ, apply Proposition 9 to the group generated by its point group and the linear part of its duality isometry. The bound of Proposition 7(i) then leaves some nonzero class unconstrained, so changing the lattice does not remove the obstruction. For the permutohedral lattice $A_4{}^*$, the second model of [DKS20], the point group S_5 fixes exactly one nonzero class $u_Q \in H_2(T^4;\, F_2)$, and u_Q is the only partnerless class.† [DKS20, Theorem 4] states that on permutohedral tori of even dimension the threshold $\lambda_{d/2}(N)$ tends to $1/2$ in every characteristic. Its proof uses the irreducibility of $H_{d/2}(T^d;\, F)$ under S_{d+1} [DKS20, Proposition 22], which is proved there only when char $F = 0$ or char F is an odd prime not dividing $d + 1$. For $d = 4$ it fails over F_2, since S_5 fixes u_Q. For F_2 and $d = 4$ that clause is therefore open, and the bound of Proposition 7(i) controls every class except u_Q.

Proposition 10 (the threshold with Z/4 coefficients). For Bernoulli percolation and every $p < 1/2$, $P_p(V_{Z/4} \not\subset 2(Z/4)^6) \to 0$, where $V_{Z/4}$ is the image of $H_2(P;\, Z/4)$ in $H_2(T^4;\, Z/4) \cong (Z/4)^6$. Equivalently, below $1/2$ no giant mod-2 cycle lifts to a mod-4 cycle.

Proof.

- (i) Enumerating all 4032 elements $x \in (Z/4)^6$ of order 4 shows that each has a signed permutation $g \in B_4$ with $\langle x,\, gx\rangle \not\equiv 0 \bmod 4$.†
- (ii) Disjoint Z/4-cycles have zero Z/4 intersection number, so $V_{Z/4}(P\,\bullet\,) \perp V_{Z/4}(P)$. At $p = 1/2$, $V_{Z/4}(P\,\bullet\,)$ has the law of $V_{Z/4}(P)$, because the half-translation acts trivially on homology. Signed permutations are induced by lattice automorphisms, so $\mu(gx \in V_{Z/4}) = \mu(x \in V_{Z/4})$. Hence $\mu_{1/2}(x \in V_{Z/4}) \leq 1/2$ for every x of order 4.
- (iii) Sharp thresholds as in Proposition 8, and a union bound over the 4032 elements. ■

In $Z/2^k$ every element of order 2 is partnerless, so the obstruction of Proposition 8 reappears at the 2-torsion level. Proposition 10 does not settle the problem; it shows only that a putative oblique phase below $1/2$ consists of mod-2 cycles whose Bockstein in $H_1(P;\, Z)$ is not divisible by 2. Proposition 20 in §8 extends the argument of Propositions 8 and 10 to all coefficient rings Z/m, and Theorem W shows that these constraints persist slightly above $1/2$: there no partnered class lifts to a Z/4-cycle and no nonzero mod-2 class lifts to a Z/8-cycle.

7.3 Box crossings and crumpled cycles

Proposition 11 (box crossings away from aspect ratio 1). Let $g_p(a,b,c,d)$ be the probability that the box B with 12-sides (a,b) and 34-sides (c,d) has an open 12-crossing: a 2-chain of open plaquettes in B, with boundary in the 12-faces $\partial_{12}B = \{x \in B : x_1 \text{ or } x_2 \text{ extremal}\}$, that is nontrivial in $H_2(B,\, \partial_{12}B;\, F_2) \cong F_2$.

- For Bernoulli percolation exactly one of {open 12-crossing, closed dual 34-crossing} occurs. Rotating by (13)(24) gives $g_p(a,b,c,d) + g_{1-p}(c+1, d+1, a-1, b-1) = 1$, and in particular $g_{1/2}(c+1, d+1, c, d) = 1/2$.
- Fix $\varepsilon \in (0, 1/10)$. For $p > 1/2$, $g_p((1-\varepsilon)L, (1-\varepsilon)L, (1+\varepsilon)L, (1+\varepsilon)L) \to 1$. For $p < 1/2$, $g_p((1+\varepsilon)L, (1+\varepsilon)L, (1-\varepsilon)L, (1-\varepsilon)L) \to 0$.

Proof. This is the argument of Bollobás and Riordan [BR06, Lemma 8], which they note extends to surfaces in higher dimensions.

1. On T^4_{4L} let U be the union, over all translates and coordinate permutations, of the crossing event of the self-dual box B_{sd} with 12-sides L+1 and 34-sides L. U is increasing and invariant under a group acting transitively on plaquettes, and $\mu_{1/2}(U) \geq g_{1/2}(B_{sd}) = 1/2$. By [BR06, Lemma 7], $\mu_p(U) \to 1$ for every $p > 1/2$.
2. If B' has a smaller 12-square and a larger 34-square than B, suitably placed, a 12-crossing of B restricts to one of B': cut along the 12-hyperplanes of B'; the mod-2 number of 12-plaquettes over an interior point is unchanged.
3. Let $s = \lfloor \varepsilon L/2 \rfloor$. Rounding the corner of each box in U down to the grid sZ^4 gives $U \subseteq \bigcup_{i \leq k} A_i$, where each A_i is a translate or permutation of the crossing event of the box B° with 12-sides L+1−2s and 34-sides L+2s, and $k \leq 6(8/\varepsilon)^4$.
4. By Harris' inequality for the decreasing events A_i^c, $\mu_p(A(B^\circ)^c)^k \leq \mu_p(U^c) \to 0$. Containment gives the first claim.
5. The second claim is the first one for the dual, through the complementarity identity. ■

Proposition 11 needs no gluing of sheets, whereas the self-dual box itself, and the passage from boxes to tori, both need a lower bound at $p = 1/2$ for boxes that are slightly harder than self-dual, i.e. sheet gluing.

We explain in outline why Theorem A cannot be deduced from Proposition 11 by deterministic inclusions alone. Call a box δ-hard for I-crossings if each I-side exceeds a distinct I^c-side by δL. A flat 2-plane T that is a graph $x_{I^c} = A\, x_I$ contains an I-crossing of the box with I-sides s and I^c-sides t exactly when $|A|\, s \leq t$ entrywise ($|A|$ the entrywise absolute value). Call T *hardness-free* if no such box is δ-hard for any $\delta > 0$ and any I; by Collatz–Wielandt this is a condition on spectral radii of 2×2 matrices built from the Plücker coordinates of T. The plane $\{(x_1, x_2) = R_{\pi/4}(x_3, x_4)\}$ is hardness-free with margin (each spectral radius is at least $\sqrt{2}$),† so hardness-free planes form a set with interior. Every nonzero class of $H_2(T^4; F_2)$ is the class of a cellular approximation of one hardness-free rational plane, or of the sum of two, for instance $e_{34} \leftarrow \langle (2,2,1,0), (0,0,1,1) \rangle$ and $u \leftarrow \langle e_1+e_2, e_1+e_3 \rangle + \langle e_1+e_2+e_3, e_4 \rangle$.† These configurations contain no δ-hard crossing at scales $L \geq C/\delta$. Hence, for every $z \neq 0$, the event $\{z \in V\}$ is contained in no monotone combination of the events that Proposition 11 controls. Similarly, primal gluing of two boxes fails exactly when a dual crossing of their union is contained in neither box. That dual crossing can be a cellular hardness-free plane, so gluing must be proved probabilistically.

The obstruction just described is specific to the six coordinate orientations. Proposition 11 holds in tilted coordinates as well.

Lemma 12′. For every 2-plane $T \subset R^4$ there is an involution $g \in B_4$ with $gT \cap T = 0$.

Proof. Let $p = v_1 \wedge v_2$ be the Plücker vector of T. Then $\det[v_1, v_2, gv_1, gv_2] = p \wedge (\Lambda^2 g)p$, so $gT \cap T = 0$ iff this number is nonzero.

- For $g = \mathrm{diag}(1, 1, -1, -1)$, the Plücker relation $p_{12}p_{34} - p_{13}p_{24} + p_{14}p_{23} = 0$ gives $p \wedge (\Lambda^2 g)p = 4p_{12}p_{34}$. The other two double sign changes give $\pm 4p_{13}p_{24}$ and $\pm 4p_{14}p_{23}$. So we are done unless p_I

$p_{I^c} = 0$ for all three pairs; then the support of p, read as edges of K_4, is contained in a star or in a triangle.

- *Star.* After a coordinate permutation, $T = \mathrm{span}(e_1, w)$ with $w = (0, a, b, c) \neq 0$. For $j \in \{2, 3, 4\}$ let g swap $e_1 \leftrightarrow e_j$ and either fix one of the two remaining basis vectors and negate the other, or swap them. For $j = 2$ the determinant is $\pm 2bc$ or $\pm(b^2 - c^2)$, and both vanish only if $b = c = 0$. The cases $j = 3, 4$ are the same with (a, c) and (a, b). Since $w \neq 0$, some j works.
- *Triangle.* After a coordinate permutation, $T = \{x_4 = 0, n \cdot x = 0\}$ with $n = (n_1, n_2, n_3, 0) \neq 0$. For $i \in \{1, 2, 3\}$ let g swap $e_i \leftrightarrow e_4$ and act on the other two basis vectors as above. Then $gT \cap T = 0$ iff e_4, e_i, n, gn are linearly independent. With $\{j, k\} = \{1, 2, 3\} \setminus \{i\}$, the relevant determinant is $\pm 2 n_j n_k$ or $\pm(n_j^2 - n_k^2)$. Choose i with $(n_j, n_k) \neq 0$. ■

By compactness of the Grassmannian, the transversality constant $|\det[v_1, v_2, gv_1, gv_2]|$ (orthonormal v_1, v_2) of the best g is bounded below by a constant $c_0 > 0$, uniformly in T. Numerically $c_0 \approx 0.62$.†

Proposition 12 (tilted parallelepipeds). Let $T \subset R^4$ be a 2-plane and $g \in B_4$ an involution with $gT \cap T = 0$ (Lemma 12′). Choose a basis v_1, v_2 of T and write $x = s_1 v_1 + s_2 v_2 + t_1 g v_1 + t_2 g v_2$, so that g acts by $(s, t) \mapsto (t, s)$. For $a, b > 0$ let $\Pi(a, b)$ be the parallelepiped $\{s \in [0, a]^2, t \in [0, b]^2\}$, with T-boundary $\partial[0, a]^2 \times [0, b]^2$, and define a T-crossing of (a cubical approximation of) $\Pi(a, b)$ as for boxes. Fix $\varepsilon \in (0, 1/10)$. For $p > 1/2$, P_p(T-crossing of $\Pi((1-\varepsilon)L, (1+\varepsilon)L)) \to 1$, and for $p < 1/2$, P_p(T-crossing of $\Pi((1+\varepsilon)L, (1-\varepsilon)L)) \to 0$, uniformly in the position.

Proof. Let K be a union of lattice 4-cubes that is a 4-ball, with $\partial K = A \cup B$, where A and B are unions of boundary 3-faces, each a solid torus. Lefschetz duality for (K; A, B) gives, exactly as for boxes, that exactly one of {primal A-crossing, dual B-crossing} occurs. The dual configuration shifted by $h = (½, ½, ½, ½)$ has the law of the primal one at $p = 1/2$, and g maps T-crossings to gT-crossings and back. Therefore, for every such region, $P_{1/2}$(T-crossing of K) + $P_{1/2}$(T-crossing of K^♯) = 1, where K^♯ = g(K* + h) is the shifted dual region and lies within distance O(1) of $g(\Pi(a, b) + h) = \Pi(b, a) + gh$. Let U be the union, over all cubical approximations within distance C of translates of $\Pi(L, L)$, of the T-crossing events. For C large this family is closed under ♯, so U is an increasing, translation-invariant event with $\mu_{1/2}(U) \geq 1/2$. Steps 1–5 of the proof of Proposition 11 now apply: [BR06, Lemma 7], snapping to a grid of spacing $\varepsilon L/2$, containment (a T-crossing of $Q \times Q'$ restricts to one of $Q^\circ \times Q'^\circ$ when $Q^\circ \subseteq Q$ and $Q'^\circ \supseteq Q'$), and Harris' inequality. ■

A flat plane S crosses, in the S-direction, parallelepipeds $\Pi(L, \delta L)$ that are as thin as we like. In its own coordinates every flat piece is thus maximally hard, and by Proposition 12 flat pieces of macroscopic size are improbable below 1/2 in every direction. Concretely, fix K, δ and a finite net of directions T; a union bound over the finitely many parallelepipeds of size at least N/K with these directions and with corners on a grid of spacing N/K^2 shows that, for $p < 1/2$, w.h.p. no surface in P contains a δ-hard crossing of any of them.

Crumpled cycles, that is, cycles that at every macroscopic scale and position are not close to any flat plane, are the only remaining obstruction, in the following sense; we only sketch the proof. Fix p and a scale L, and suppose that for every grid box Q of side L there is an increasing event E_Q, depending only on the plaquettes within distance 2L + 2 of Q. Suppose also that E_Q occurs for every box Q visited by a connected cycle with nonzero class, and that $\sup_Q P_p(E_Q) \leq \eta_0$, an absolute constant. Then $P_p(V \neq 0) \leq 2(N/L)^4 e^{-6(N/L)^2} \to 0$. Indeed, a connected giant cycle visits a *-connected set of at least $(N/L)^2$ boxes, there are at most $(N/L)^4(80e)^m$ such sets of size m, and the events E_Q of a 6-separated subset of size $m/11^4$ are independent.

The coarse-grained entropy $(80e)^m$ does not depend on L, so near 1/2 the difficulty is coverage rather than entropy. With the events of Proposition 12, whose proof gives a polynomial rate, the same counting, which we do not write out, shows the following: for each $p < 1/2$ and $\theta > 0$, w.h.p. every giant

cycle carries a hard crossing in fewer than a θ-fraction of its visited boxes, at every dyadic scale between a fixed $L_0(p, \theta)$ and $N/16$. What is still missing is a local event that occurs in every visited box of a crumpled cycle and is rare below $1/2$. Its natural forms are subcritical arms of connected relative 2-cycles, or easy crossings (equivalently, by box duality, supercritical gluing). Numerically, in the ensemble in which exactly half of the plaquettes are open, the minimal representative of a coordinate giant class, computed by annealing, has area $\approx 1.06 \cdot N^{3.35}$ for $N = 4, \dots, 12$.† For $N \le 8$ the annealing was checked against exact integer programs on 40 test cases from both ensembles (Bernoulli and fixed density, densities 0.485 to 0.5, several classes): it found the exact minimum in 36 of them and was at most 7.3% above it in the other four.† These data indicate that critical giant cycles are far from flat. In runs at densities slightly below $1/2$, the giant cycles that occur are flatter and more often straight.†

For $p < 1/2$, Proposition 12 also excludes large graph pieces. Fix $\Lambda, \eta > 0$. We outline the argument. By compactness and quantitative differentiation (Arzelà–Ascoli plus Rademacher), a relative 2-cycle that lies within $O(1)$ of a Λ-Lipschitz graph over a square of side r is, on some sub-square of side $\delta(\Lambda, \eta)r$, within ηr of a plane. It then contains a hard crossing in that plane's coordinates. Thus, for $p < 1/2$, w.h.p. no surface of P contains a piece that is a bounded-slope graph over a square of side $\ge \eta N$. Below $1/2$, giant cycles, slab disks and gluing obstructions must fold back or be arbitrarily steep at every macroscopic scale.

The following construction shows a limitation of these events. Drill small holes in a flat sheet and cap each hole by a long thin finger, a tube closed at its far end. The result is homologous to the sheet. A finger that leaves a box through its transverse faces puts boundary on those faces, so the restricted chain is no longer a crossing. Placing fingers of length $\approx L$ at density $\approx L^{-2}$ at every dyadic scale L therefore destroys every cubic crossing at every scale above the shortest finger length ℓ_0. The area grows only by a factor $1 + O(1/\ell_0)$. The same construction avoids hard tilted crossings below a fixed fraction of the size of the sheet, and an adversarial search in a continuum model of the fingered sheet finds no hard tilted crossing at the largest scales either.† Fingers can thus make a sheet invisible to the events controlled by Propositions 11 and 12, at negligible cost in area.

We record one conditional route and sketch the implication. Call a pair of overlapping regions R_1, R_2 (two easy boxes, or two overlapping annular strips $T^1 \times [s, e] \times \square$) *uniquely crossed* if every 12-crossing of R_1 and every 12-crossing of R_2 have traces on the middle slice of the overlap that differ by a boundary in $P \cap (R_1 \cap R_2)$. If, for every $p > 1/2$, these pairs are uniquely crossed w.h.p., then Theorem A holds. Indeed, cyclically overlapping easy boxes give a crossing of an annular strip, and three cyclically overlapping strips give a 12-torus. In Monte Carlo runs at $p \ge 0.54$ the estimated probability of unique crossing increases towards 1 with size: for annular strips at $p = 0.56$ it is 0.22, 0.43, 0.66 and 0.86 for sides 4, 6, 8 and 10.† At $p = 1/2$ it is small and does not grow: 0.08, 0.03, 0.005 and 0.008 for the same sides.†

7.4 Reflection positivity

We now show that Lemma 4 at $q = 1$ does not follow from reflection positivity. Reflect the torus in direction μ through two planes cutting the temporal cells of two opposite gaps. Write a configuration as $X \cup H \cup \theta Y$, where X and Y are configurations of one half and H is the set of open plaquettes crossing the planes. Place the temporal twist e in one of the crossing layers.

For $q = 2$ one reflection suffices, and Lemma 4 then holds for every reflection-symmetric product of couplings ($\beta_{\theta\sigma} = \beta_\sigma$). We sketch the argument. In the gauge theory,

$$Z(\tau) = \Sigma_{a,b}\, F(a)\, M_\tau(a, b)\, F(b),$$

where a and b are the gauge fields on the two boundary layers, $F \ge 0$ is the partition function of one half with boundary value a, and M_τ is the weight of the crossing layers, after summing over the gauge

field on the crossing (i−1)-cells. Expanding the crossing weights in characters gives

$M_s - M_{s+e} = c\ \Sigma_{S:\ e(S)\,=\,1}\ (\Pi_{h\in S}\ t_h)\ \chi_S \otimes \chi_S$, with $t_h = p_h/(2 - p_h) \geq 0$ and $c > 0$.

Here S runs over the sets of crossing plaquettes that are cycles relative to the crossing (i−1)-cells, and $\chi_S(a) = (-1)^{a(\partial S)}$. This kernel is positive semidefinite, which gives $R(s+e) \leq R(s)$.

For $q = 1$ the analogous statement is false:

Proposition 13. Take the torus $Z_3 \times Z_2 \times Z_4$ ($d = 3$, $i = 2$), with $\mu = 3$ and the reflection $x_3 \mapsto 1 - x_3$. Its planes cut the temporal plaquettes of the gaps [0,1] and [2,3]. Let $s = [\eta_{12}]$ and $e = [\eta_{13}]$, and put η_{13} on the gap [0,1]. Let H consist of the following 14 of the 24 crossing plaquettes, given as (type; base point (x_1, x_2, x_3)):

- type {1,3}: (0,0,2), (0,1,0), (0,1,2), (1,0,0), (1,0,2), (1,1,2), (2,0,0), (2,0,2), (2,1,0);
- type {2,3}: (0,0,2), (0,1,0), (1,0,0), (2,0,2), (2,1,2).

Then $\Sigma_{X,Y}\ (1_{D_s} - 1_{D_{s+e}})(X \cup H \cup \theta Y) = -9\ 711\ 396$, the sum running over all $2^{24} \times 2^{24}$ pairs of half-configurations. Consequently, take the product measure with probability 1/2 on non-crossing plaquettes, $1 - \delta$ on H, and δ on the other crossing plaquettes. This measure is reflection-symmetric, and for all small $\delta > 0$ it satisfies $\mu(D_s) < \mu(D_{s+e})$.

The sum is computed exactly.† After scaling by 2^{48}, every intermediate quantity is an integer of absolute value at most 2^{48}. The computation reduces each half to its affine subspace of boundary values and works in $F_2{}^H$ modulo the crossing coboundaries. On 2×2×2 and 2×2×4, an independent brute-force computation reproduces its output to at least 10 significant digits, including a negative value on 2×2×4 for an inhomogeneous reflection-symmetric product measure.†

A proof of Lemma 4 at $q = 1$ must therefore use the homogeneity of the measure, which the proof for $q = 2^r$ never uses. In $d = 2$ the per-H inequality holds in every case we tested: all crossing configurations on the tori 2×4, 3×4, 4×4, 3×6 and 4×6, for nine homogeneous and 501 random inhomogeneous reflection-symmetric product measures.† There the obstruction is subtler. On the tori $Z_W \times Z_N$ with $W \in \{3, 4\}$ and $p = 1/2$, the twist gap is

$\Delta_N = R(s) - R(s+e) = \Sigma_j\ c_j\ \lambda_j{}^N$,

where all λ_j lie in (0, 1) and the c_j are integers in $\{4, 2, 1, -3\}$.† (These recurrences have orders 16 and 69 and are obtained by exact Berlekamp–Massey. For $W = 3$ at $p = 1/4$ and 3/4 the recurrences have order 17 and the same structure.) Some c_j are negative, so (Δ_N) is not a moment sequence. Hence $\Delta_N \neq \langle v, T^N v\rangle$ for any positive operator T, and Lemma 4 at $q = 1$ holds there because the leading positive mode dominates.

7.5 A cell-wise form of Lemma 4

Fix μ, a spatial twist $s \neq 0$ and a temporal twist $e \neq 0$. For a configuration ω let $W(\omega) := \mathrm{span}\{s, e\} \cap V(P(\omega))^{\perp}$, the twists in span{s, e} that vanish on every cycle of $P(\omega)$. Since $D_s \setminus D_{s+e} = \{W = \langle s\rangle\}$ and $D_{s+e} \setminus D_s = \{W = \langle s+e\rangle\}$,

$\mu_{N,p,q}(D_s) - \mu_{N,p,q}(D_{s+e}) = Z^{-1}\ \Sigma_{m,k}\ N_{m,k}\ p^m\ (1-p)^{|X^{(i)}|-m}\ q^k$, $N_{m,k} := \#\{\omega : |\omega| = m,\ b_{i-1}(P(\omega)) = k,\ W(\omega) = \langle s\rangle\} - \#\{\omega : |\omega| = m,\ b_{i-1}(P(\omega)) = k,\ W(\omega) = \langle s+e\rangle\}$.

Conjecture E (cell-wise twist inequality). $N_{m,k} \geq 0$ for all N, m, k, μ, s and e.

Conjecture E implies Lemma 4 for every $q > 0$ simultaneously. By Theorem D, it therefore implies the characteristic-2 sharp threshold for every $q \geq 1$, including the Bernoulli case. It has two equivalent forms.

Within a cell, $|Z^{i-1}(P(\omega))| = 2^{k+c}$ is constant, so $N_{m,k} \geq 0$ says that among pairs (ω, f) with ω in the cell, solutions of $\delta f = s$ on ω are at least as many as solutions of $\delta f = s + e$ on ω. Exchanging the sums turns this into $\Sigma_{c\in[s]} N^{(m,k)}(X \setminus \operatorname{supp} c) \geq \Sigma_{c\in[s+e]} N^{(m,k)}(X \setminus \operatorname{supp} c)$, where c runs over all cocycles in the class, equivalently over dual cycles, and $N^{(m,k)}(S)$ counts the configurations $\omega \subseteq S$ in the cell.

We enumerated all configurations on the following tori and found $N_{m,k} \geq 0$ in every cell, with no exceptions:† the cubical tori 2×2×2 and 2×2×3 ($d = 3$, $i = 2$; 2^{24} and 2^{36} configurations), for all nine admissible pairs (μ, n), and the tori 3×3, 3×4, 4×4, 3×5 and 4×5 ($d = 2$, $i = 1$; up to 2^{40} configurations), for $n = e_1^* + e_2^*$ and both transfer directions.

For Lemma 4 itself at $q = 1$ on longer tori, we use an exact transfer operator. Its states are the affine relations between the gauge fields of the first and the current layer, taken modulo coboundaries. On $Z_2 \times Z_2 \times Z_N$ with the long direction as μ, it gives $R(s) - R(s+e) > 0$ for every N up to the limit of double precision: $N \leq 25$ at $p = 1/2$, $N \leq 14$ at $p = 0.3$, and $N \leq 30$ at $p = 0.7$, for the temporal twist e_{13}^*, and $N \leq 25$ at $p = 1/2$ for the temporal twist $e_{13}^* + e_{23}^*$.† The gap decays geometrically, with ratio 0.2396 at $p = 1/2$.

In $d = 4$, Monte Carlo sampling of Bernoulli plaquette percolation on T^4_N ($N = 4, 6, 8$; up to 10^5 samples per point) is consistent with $R(n) \leq R(s_\mu(n))$ for all 63 twists: the smallest observed difference $R(s_\mu(n)) - R(n)$ is $-1{\cdot}10^{-5}$, well within the statistical error.† It also finds a transition at $p = 1/2$ that is already sharp at $N = 6$: $V = 0$ in every sample at $p = 0.45$ and $V = H$ in every sample at $p = 0.55$. The oblique Lagrangian $\operatorname{span}\{e_{12} + e_{34}, e_{13} + e_{24}, e_{14} + e_{23}\}$ occurs once in $2.2{\cdot}10^5$ samples at $p = 1/2$ ($N = 4, 6, 8$ combined).

7.6 Translation orbits of the layers

Proposition 13 shows that reflection symmetry alone does not give Lemma 4 at $q = 1$. The experiments below indicate that Lemma 4 depends on spatial homogeneity, and they point to the orbit averages on which it may rest.

Fix the transfer direction μ and let $G \cong (Z/NZ)^{d-1}$ be the group of translations orthogonal to μ. A *class* is the set of i-cells of a given type I in a given time layer: spatial cells of type I in the layer $x_\mu = t$, or temporal cells of type I in the gap $[t, t+1]$. Each class is a single free G-orbit. For $g = (g_c) \in G^C$, where C is the set of classes, $g{\cdot}\omega$ translates the part of ω in each class c by g_c. Let R be the group generated by G^C and, in each class separately, the reflections of the spatial coordinates. These maps are not automorphisms of X, but they are bijections of the configuration space that preserve every class-wise homogeneous product measure.

Our tests of Lemma 4 under random product measures are summarized in the following table.† In all tests with measures homogeneous within each class Lemma 4 at $q = 1$ held, and the Ising gauge theory ($q = 2$) behaved the same way.

q	Product measure	Torus	Trials	Result
1	homogeneous within each class, with arbitrary dependence on t and on the cell type; all three temporal twists	$Z_2 \times Z_2 \times Z_N$, $N = 2, 3$	210	no violation
1	parameters varying from cell to cell	$Z_2 \times Z_2 \times Z_2$	240	2 violations, up to $3{\cdot}10^{-2}$
2	time-dependent, spatially homogeneous couplings	—	800	no violation (negative values only of order 10^{-15}, i.e. round-off)
2	couplings that also vary from plaquette to plaquette	—	—	violations of relative size up to $1.3{\cdot}10^{-4}$

Proposition 14. There are configurations ω, in d = 2, 3 and 4, such that

$\Sigma_{g \in G^C} (1_{D_s} - 1_{D_{s+e}})(g \cdot \omega) < 0$.

- *d = 3, i = 1 (bonds).* Take the torus $Z_3 \times Z_3 \times Z_2$, with time $t \in \{0, 1\}$, $s = e_1^*$ and $e = e_t^*$. In time slice 0, let the bonds of direction 1 be all sites off the diagonal {(0,0), (1,1), (2,2)} and the bonds of direction 2 the diagonal, with temporal bonds at {(0,1), (1,0)}. In slice 1, let the direction-1 bonds be at {(0,0), (0,1), (0,2), (2,0), (2,2)} and the direction-2 bonds at {(0,1), (0,2), (1,1), (1,2), (2,0), (2,1)}, with temporal bonds at {(0,1), (1,0), (1,1), (2,0)}. With one class pinned, the orbit sum over the remaining 9^5 translations is −243 (486 positive and 729 negative terms).†
- *d = 4, i = 2.* Thicken this example by a circle S^1_u of length 2: replace every open bond b by the two plaquettes b × f, f an edge of S^1_u. A 2-chain $\Sigma\, c_{b,f}\, b \times f$ is a cycle iff $c_{b,f} = c_b$ does not depend on f and $\{b : c_b = 1\}$ is a 1-cycle. Hence the 2-cycles are exactly $\gamma \times S^1_u$ for 1-cycles γ, and $D_{n \wedge e_u^*}$ on the thickened torus equals D_n on the original one. Translations along S^1_u act trivially on the nonempty classes, so on $Z_3 \times Z_3 \times Z_2 \times Z_2$ the sum over the translations of the nonempty classes, one of them pinned, is $-243 \cdot 2^5 = -7776$.†
- *d = 2, i = 1.* On $Z_6 \times Z_{39}$ ($\mu = 2$), take 13 copies of the three-row block with rows {1,3,5}, {0,1,3,4}, {1,3,5} and gaps {0,3,5}, {2,3,5}, {0}. The uniform measure on its orbit has $\mu(D_s) - \mu(D_{s+e}) = -6.98 \cdot 10^{-9}$. Two programs compute this value exactly with big integers, and a third in floating point.†

The mechanism behind these examples is the following. Between two gaps with one open cell each, a block of layers contributes a displacement law ν on Z_{2W}, where the extra factor 2 records the winding parity. The orbit sum of m stacked copies of a block is the sum of the antiperiodic Fourier coefficients of ν^{*m}. Reflection sends $\nu(x)$ to $\nu(-x)$. If the class patterns are chiral, ν is in general not symmetric, its leading antiperiodic coefficients are $r\, e^{\pm i\alpha}$ with $\alpha \neq 0$, and m copies contribute $(2/W)\, r^m \cos(m\alpha) + O(r'^m)$ with $r' < r$. For the d = 2 block above, $\alpha = 0.1286$ and the sum first becomes negative at m = 13.† Every subset of Z_W with $W \leq 5$ is a translate of its reflection, so in d = 2 chirality first occurs at W = 6.† The same mechanism gives negative q-weighted orbit sums at q = 2 (on $Z_7 \times Z_{232}$) and q = 4 (on $Z_7 \times Z_{378}$).† Since at q = 2 Lemma 4 holds for every reflection-symmetric product measure (§7.4), the obstruction comes from chirality rather than from the value of q.

The uniform measure on such an orbit is independent across classes and translation-invariant within each class, and it violates Lemma 4. Thus, at q = 1, Lemma 4 can fail for class-wise translation-invariant measures. Homogeneous product measures are also invariant under the reflections of each class, and the conjecture below keeps that symmetry. Conjecture E is not affected, because its sums run over all configurations with given $|\omega|$ and b_{i-1}.

Conjecture F (reflection-orbit twist inequality). Let $q \geq 1$. For every configuration ω,

$\Sigma_{r \in R}\, q^{b_{i-1}(P(r \cdot \omega))}\, (\, 1_{D_s} - 1_{D_{s+e}}\,)(r \cdot \omega) \geq 0$.

Proposition 15. Conjecture F implies Lemma 4 for the random-cluster weight $q^{b_{i-1}}$ applied to any product measure that is homogeneous within each class. In particular it implies Lemma 4 for $\mu_{N,p,q}$, hence Theorem A at that q by Theorem D; at q = 1 this is the Bernoulli case.

Proof. A law on the subsets of a class that is invariant under the translations and reflections of the class is a mixture of uniform laws on single orbits of these maps, and a homogeneous product law is such a law. By independence across classes, the base measure is then a mixture of uniform measures on R-orbits. The weight $q^{b_{i-1}}$ is applied configuration by configuration, so $\mu(D_s) - \mu(D_{s+e})$ is a nonnegative combination of the sums in Conjecture F. ■

Every R-orbit is a union of translation orbits, so a nonnegative sum over every translation orbit implies Conjecture F on that torus. By exhaustive enumeration we found no negative translation-orbit sum on the following tori:†

Torus (d)	**Twists**	**q tested (no negative orbit)**	**Nonzero orbit sums at q = 1**
3×4 (d = 2)	—	0.5, 1, 1.5, 2, 4	6 561
4×4 (d = 2)	—	0.1, 0.5, 1, 1.5, 2, 4, 8	390 625
2×2×2 (d = 3)	e_{13}^*	0.5, 1, 1.5, 2, 4	20 696
2×2×3 (d = 3)	e_{13}^*	0.1, 0.5, 1, 1.5, 2, 4, 8	2 868 660
2×2×3 (d = 3)	e_{23}^*	1	2 868 660
2×3×2 (d = 3)	e_{13}^*	0.1, 0.5, 1, 1.5, 2, 4, 8	2 762 976
2×3×2 (d = 3)	e_{23}^*	0.5, 1, 1.5, 2, 4, 8	988 624

Conjecture F cannot be refined to single values of b_{i-1}, and it fails for every $q < 1$, even after reflection averaging. The smallest example has d = 2, i = 1, on $Z_6 \times Z_2$. Let every row and every gap be the alternating pattern {0, 2, 4}; a translation by 1 or a reflection maps it to {1, 3, 5}, so the R-orbit has 16 configurations. In 12 of them the open edges form three vertical cycles, with $W = \langle s \rangle$ and $b_0 = 3$. In the other 4 (rows different and gaps different) they form one Hamiltonian cycle of class (1, 1) mod 2, with $W = \langle s+e \rangle$ and $b_0 = 1$. The q-weighted orbit sum is therefore $12q^3 - 4q$, which is negative for $q < 1/\sqrt{3}$.† On $Z_6 \times Z_N$ the same pattern gives a threshold value of q that tends to 1 as $N \to \infty$.† Thickening by circles (as in Proposition 14) transfers these orbits to d = 3 and d = 4.†

In d = 3, on 2×3×2 with $e = e_{23}^*$, exhaustive enumeration of the 714 136 R-orbits with nonzero sum finds none with a negative total, but 20 with a negative sum at a fixed value of b_{i-1}; these 20 orbits are achiral.† With $e = e_{13}^*$ on the same torus, and on 2×2×2 and 2×2×3, every rank part of every R-orbit is nonnegative.† There the rank failures of translation orbits appear to be effects of chirality: on 2×3×2 the translation orbit of the configuration 0x3f719edb6 (in the cell order of the program `orbtr.c` of the supplement) has rank sums +1584, +3264 and −192, and its R-orbit has rank sums +6336, +12672 and 0.†

The following argument, given in outline, explains why q = 1 is the boundary case. Let k be odd, and lift ω to the k-fold cover of the torus in one spatial direction. The covering map and the transfer map act on $H_i(\cdot; Z)$ by matrices that are the identity mod 2. Cycles project to cycles and preimages of cycles are cycles, so the lift has the same V, hence the same values of $1_{D_s} - 1_{D_{s+e}}$. The R-orbit of the lift consists of the lifts of the R-orbit of ω, with equal multiplicities. Consequently the orbit sum at q = 1 is invariant under odd covers, while b_{i-1} is not: in d = 2, a cluster whose cycles generate the horizontal winding a lifts to gcd(k, a) clusters. The orbits on $Z_6 \times Z_N$ above are 3-fold covers of configurations on $Z_2 \times Z_N$ with one open edge in every row and every gap, where Proposition 16 applies.

The reflection averages of the counterexamples of Proposition 14 are positive. For the d = 3 example, the sum over the translations and the reflections of the first spatial coordinate of each class (one class pinned) is +69 984, and with the reflections of both coordinates it is +2 239 488.† In d = 4 the orbits are too large to enumerate directly (8^{18} translations on 2×2×2×3). On the spatial torus $Z_2{}^3$ every reflection acts as a translation, so there R-orbits are translation orbits, and we test random orbits and sum them exactly where we can.

To sum orbits exactly in d = 4, fix the translations of the temporal classes. A 2-cycle of P is then $z = \Sigma_t\, c_t + \Sigma_t\, \gamma_t \times [t, t+1]$, where γ_t is a 1-cycle of the slice formed by edges b for which $b \times [t, t+1]$ is open, and c_t is a chain of open spatial plaquettes of the layer t with $\partial c_t = \gamma_{t-1} + \gamma_t$. Moreover $\langle s, z \rangle = \Sigma_t \langle s, c_t \rangle$, while $\langle e, z \rangle$ depends only on γ_0. The layer t therefore enters only through the relation it induces between the cycle spaces of the two adjacent gaps, together with the values of s on the chains c_t. Grouping the translations of each layer by this relation reduces the orbit sum to a sum over distinct

relations. Our exact computation does this, and it also skips the translations of the temporal classes for which $\langle e, \cdot \rangle$ vanishes on every cycle that any choice of layers allows. It agrees with a direct cochain computation on every configuration tested, and with direct enumeration on a whole orbit of 2×2×2×2.†

The orbits we tested in d = 4 are listed in the following table.† No orbit sum was found to be negative.

Torus	**Twist pairs**	**Random orbits per pair**	**Orbit sums**
2×2×2×3	(e_{12}^*, e_{34}^*), (e_{12}^*, e_{14}^*), and the spatial and temporal halves of $\Sigma_I e_I^*$	1500	3705 of the 4500 computed exactly: 3605 are 0, 100 positive, none negative; of the other 795, 792 positive by Monte Carlo (by more than 5σ) and 3 zero in every one of $2{\cdot}10^8$ samples
2×2×2×2	the three above and $(e_{12}^* + e_{13}^*, e_{24}^*)$	3000	Monte Carlo with $3{\cdot}10^5$ samples per orbit: more than 3σ above zero or zero in every sample, except for 3 orbits, which are 0 when summed exactly
3×2×2×2	the first three above	1500	Monte Carlo with $4{\cdot}10^5$ samples per orbit: more than 3σ above zero or zero in every sample, except for 12 orbits, which are positive when summed exactly

One of the exact zeros on 2×2×2×3 is the orbit whose Monte Carlo mean is $-1.2{\cdot}10^{-7} \pm 0.9{\cdot}10^{-7}$ after 10^{10} samples.†

The coarser groupings that we tested fail. Translating all classes of a time step together gives negative orbit sums on 2×2×2 and 3×4. In d = 4, translating all spatial classes of a layer together and all temporal classes of a gap together gives negative orbit sums on 2×2×2×2, for 19 to 40 of 20 000 random patterns per twist pair; the exact per-class sums of eight of these patterns are positive.†

On the 3×4 torus, which has eight classes, averaging over the translations of any six classes (with the other two fixed) gives a nonnegative sum for every configuration. For every choice of five classes, some configuration gives a negative sum.†

Proposition 16 (one temporal cell per gap, d = 2). Let d = 2, i = 1, and work on $Z_W \times Z_N$ with $\mu = 2$. Fix arbitrary patterns H_t of open edges in the rows and exactly one open vertical edge in each gap. Translate each row and each gap by independent uniform elements of Z_W. Then $E[1_{D_s} - 1_{D_{s+e}}] \geq 0$.

Proof.

1. If some row is full, it is an s-odd, e-even cycle, so both events fail. Assume no row is full.
2. Every cycle crosses each gap exactly once, so there is at most one cycle. It exists exactly when, for every t, the columns c_{t-1} and c_t of the vertical edges below and above row t lie in a common arc of row t.
3. Given the columns, the rows are independent. With $r_t = (c_t - c_{t-1}) \bmod W$, row t joins them to the right with probability $f_t(r_t)$ and to the left with probability $f_t(W - r_t)$. Here $f_t(\ell)$ is the probability that ℓ consecutive edges are open after the random translation, with $f_t(0) = 1$ and $f_t(W) = 0$. If $\Sigma_t\, r_t = mW$, the winding of the cycle in the first direction is $m - \#\{\text{left}\}$. So $E[1_{D_s} - 1_{D_{s+e}} \mid c] = (-1)^m \prod_t u_t(r_t)$, where $u_t(r) = f_t(r) - f_t(W - r)$.
4. The columns are independent and uniform, and $(-1)^m\, 1\{W \mid \Sigma\, r_t\} = W^{-1} \sum_{k=0}^{W-1} \exp(i\pi(2k+1)\Sigma_t\, r_t/W)$. Hence $E[1_{D_s} - 1_{D_{s+e}}] = W^{-N} \sum_{k=0}^{W-1} \prod_t \hat{U}_t(k)$, where $\hat{U}_t(k) = \sum_{|r|<W} f_t(|r|)\, e^{i\pi(2k+1)r/W}$.
5. Each factor is nonnegative. For $r \geq 1$, $f_t(r) = W^{-1} \sum_A (|A| + 1 - r)_+$, the sum running over the arcs A of row t. Write $F_n(\theta) = |\sum_{j=0}^{n} e^{ij\theta}|^2 \geq 0$ for the Fejér kernel. Then

$\hat{U}_t(k) = 1 - W^{-1}\,\Sigma_A\,(|A|+1) + W^{-1}\,\Sigma_A\,F_{|A|}(\pi(2k+1)/W) \geq 0,$
because each arc is followed by a closed edge, so $\Sigma_A\,(|A|+1) \leq W$. ■

The proof shows how the assumption of one temporal cell per gap is used. After averaging over the translations of each class, every row contributes a translation-invariant kernel whose Fourier transform is nonnegative at every frequency, in particular at the antiperiodic modes selected by the temporal twist. The same proof applies, with changes we do not write out, in $d = 3$, $i = 2$ when the open temporal cells of every gap form one full line of type $\{1, 3\}$; the role of the arcs is then played by maximal runs of fully open rows of spatial plaquettes. It also gives the inequality rank by rank in a larger class.

Proposition 17 (thin or full gaps, d = 2). Let $d = 2$, $i = 1$, $\mu = 2$, and suppose every gap of ω is thin (one open edge) or full (all W edges open), and at least one gap is thin. Then for every k,

$\Sigma_{r \in R\,:\,b_0(P(r\omega)) = k}\,(1_{D_s} - 1_{D_{s+e}})(r\omega) \geq 0,$

and the same holds with R replaced by the translations G^C.

Proof.

1. A full row is a horizontal cycle, and then both events fail. Cut the time circle at the thin gaps into m blocks of consecutive rows. Inside a block the full gaps join (x, t) to (x, t+1) for every x, so the components of the block are $\{\alpha\} \times$ (rows of the block), where α runs over the arcs, isolated vertices included, of the union U of the row patterns of the block. If $U = Z_W$ there is a horizontal cycle and both events fail again. Assume neither happens.
2. Every cycle with odd vertical winding uses the edge of each thin gap exactly once. So there is a cycle through all thin edges, its class (h, 1) mod 2 spans V, and the events give $+1$ if $h = 0$ and -1 if $h = 1$. Such a cycle exists exactly when, in every block, the columns of the thin edges below and above lie in one arc of U. Its horizontal winding is D/W, where D is the sum over the blocks of the displacement inside the arc. On these configurations, $b_0 = \Sigma_j\,c_j - m + 1$, where c_j is the number of arcs of the union of block j.
3. For each block j and each c, let $\nu_j^{(c)}(d)$, $d \in Z_{2W}$, count the images of the rows of block j with c arcs in which the columns 0 and d mod W lie in one arc with lifted displacement $\equiv d$ mod 2W. Summing over the translations of the thin gaps gives
   $\Sigma_{r\,:\,b_0 = k}\,(1_{D_s} - 1_{D_{s+e}})(r\omega) = W\,\Sigma_{c_1+\ldots+c_m = k+m-1}\,[(\nu_1^{(c_1)} \circledast \ldots \circledast \nu_m^{(c_m)})(0) - (\nu_1^{(c_1)} \circledast \ldots \circledast \nu_m^{(c_m)})(W)],$
   with $\circledast$ the cyclic convolution on Z_{2W}.
4. Group the images of the rows of block j into orbits of the diagonal translation, which moves all rows of the block together. On such an orbit the arcs of U are translated, so c is constant, and summing over the orbit gives $\Sigma_\alpha\,(|\alpha| - |d|)_+$ for $|d| < W$, where $|\alpha|$ is the number of vertices of the arc α. Every $\nu_j^{(c)}$ is thus a nonnegative combination of tents supported in $|d| \leq W - 1$, hence symmetric and nonincreasing in $|d|$ on Z_{2W}.
5. Cyclic convolution preserves this property, so each bracket in step 3 is ≥ 0. ■

A gap is thick if it has at least two open edges. Full gaps are thick, but Proposition 17 does not extend to the other thick gaps: the rank-refined inequality fails there (the example on $Z_6 \times Z_2$ above, which has no thin gap), and the single-block argument fails too, because the averaged displacement law of a block between two thin gaps need not be nonincreasing in $|d|$. For the two-row block on Z_8 with rows $\{0, 1, 4, 5\}$ and internal gap $\{0, 2, 4\}$, the law averaged over translations and reflections takes the values 64 and 96 at $|d| = 5$ and 6.† Its antiperiodic Fourier coefficients are positive, but other blocks have negative ones.†

For an arbitrary configuration with a thin gap, the proof of Proposition 17 still reduces the orbit sum to a convolution of block laws. Cut the time circle at the thin gaps into blocks of consecutive rows. For the j-th block, let $\bar\nu_j(d)$, $d \in Z_{2W}$, be the number of elements of the product of the dihedral groups of its rows and internal gaps for which:

- the block has no cycle of odd horizontal winding; and
- the column-0 vertex of its bottom row is joined, inside the block, to its top row with lifted horizontal displacement $\equiv d \bmod 2W$.

Reflecting all classes of the block at once shows that $\bar\nu_j(-d) = \bar\nu_j(d)$. Put $h_j(d) := \bar\nu_j(d) - \bar\nu_j(d + W)$. This function is symmetric and antiperiodic ($h_j(d + W) = -h_j(d)$).

Proposition 17′. Let $d = 2$, $i = 1$, $\mu = 2$, and let ω have $m \geq 1$ thin gaps.

- (i) $\Sigma_{r\in R}\,(1_{D_s} - 1_{D_{s+e}})(r\omega) = 2W \cdot (h_1 \circledast \cdots \circledast h_m)(0)$, where $\circledast$ is the cyclic convolution on Z_{2W}.
- (ii) For $W \geq 3$ and a symmetric antiperiodic h on Z_{2W}, put $F_j(h) := \Sigma_{d\in Z_{2W}}\, h(d)\cos(\pi(2j + 1)d/W)$ for $j = 0, \ldots, W - 1$. Let K_W be the set of such h with $F_0 > 0$, $|F_j| \leq F_0$ and $F_j \geq -2F_0/(W - 2)$ for every j. Then K_W is closed under $\circledast$, and $h(0) \geq 0$ for every $h \in K_W$.
- (iii) For $W \geq 3$ and every block with n rows, $h_j = 0$ or $h_j \in K_W$, whenever $n = 1$ and $W \leq 16$, $n = 2$ and $W \leq 12$, $n = 3$ and $W \leq 7$, or $n = 4$ and $W \leq 5$.†

Consequently, Conjecture F at $q = 1$ holds on $Z_W \times Z_N$ for every configuration with a thin gap whose blocks all have at most n_W rows. Here $n_W = 4$ for $3 \leq W \leq 5$, $n_W = 3$ for $W = 6, 7$, $n_W = 2$ for $8 \leq W \leq 12$, and $n_W = 1$ for $13 \leq W \leq 16$.

Proof.

1. As in step 1 of the proof of Proposition 17, a cycle of odd horizontal and even vertical winding makes both events fail. Such a cycle exists if and only if some block contains a cycle of odd horizontal winding.
2. Assume no block contains one. As in step 2, every cycle with odd vertical winding uses the edge of each thin gap exactly once, and cycles with even vertical winding lie inside the blocks. Every cycle with odd vertical winding therefore has horizontal winding $\equiv \Sigma_j D_j/W \bmod 2$, where D_j is the lifted displacement across block j, and the events give $+1$ or -1 according to this parity.
3. Fix the positions of the thin edges below and above block j. The images of the block in which these columns are joined with displacement d are counted by $\nu_j(d)\cdot 1\{d \equiv c' - c \bmod W\}$, where c and c' are the two columns and ν_j is the law over translations only.
4. Summing over the translations of the thin gaps, as in step 3, gives $\Sigma_{r\in G^c}\,(1_{D_s} - 1_{D_{s+e}})(r\omega) = W \cdot (\nu_1 \circledast \cdots \circledast \nu_m)^-(0)$, where $f^-(d) := f(d) - f(d + W)$.
5. Adding the reflections of the classes replaces each ν_j by $\bar\nu_j$, by multilinearity. A reflection of a thin gap is a translation of its single edge, which gives a factor 2 for each of the m thin gaps.
6. Finally, $(f \circledast g)^- = \tfrac12\, f^- \circledast g^-$. To see this, write $f = f_+ + \tfrac12\, f^-$, where $f_+(d) := \tfrac12(f(d) + f(d + W))$ is W-periodic; then $f_+ \circledast g^- = 0$, because g^- is antiperiodic. Applying this identity repeatedly gives $2^m\, W\, (\bar\nu_1 \circledast \cdots \circledast \bar\nu_m)^-(0) = 2W\, (h_1 \circledast \cdots \circledast h_m)(0)$. This proves (i).

For (ii), convolution multiplies the coefficients F_j, which are real.

- If $F_j(h) < 0 < F_j(g)$, then $F_j(h)F_j(g) \geq -2F_0(h)F_j(g)/(W - 2) \geq -2F_0(h)F_0(g)/(W - 2)$. The other sign cases are immediate, so K_W is closed under $\circledast$.

- By Fourier inversion, $2W \cdot h(0) = \Sigma_{j=0}{}^{W-1} F_j$, since only the antiperiodic frequencies contribute. The terms $j = 0$ and $j = W - 1$ both equal F_0, and each of the other $W - 2$ terms is at least $-2F_0/(W - 2)$. Hence $h(0) \geq 0$.

Part (iii) is a finite computation over all block types up to the symmetries of each class, in floating point.† Apart from the blocks with $h(d) = 0$ for $0 < d < W$, for which $|F_j| = F_0$ for every j, the largest ratio $|F_j|/F_0$ for $j \neq 0, W - 1$ is 0.983. It is attained by a row with one edge and tends to 1 as W grows. The smallest ratio F_j/F_0 is -0.022, against a threshold $-2/(W - 2) \leq -1/7$.

The consequence follows from (i)–(iii). If some $h_j = 0$, the orbit sum is 0. ■

Neither of the two natural cones contains the antiperiodic part of every block law, which is why K_W is needed.

- Nonnegative antiperiodic Fourier coefficients: this cone fails for the blocks with some $F_j < 0$, such as the two-row block on Z_{10} with ratio -0.022.
- Antiperiodic parts that are nonincreasing on $\{0, \ldots, W\}$: this cone fails for 17 two-row block types on Z_{11}, all at the middle step $d = 5$.†

Proposition 17′ covers only blocks of at most n_W rows, and the cone K_W shrinks as W grows, so it does not settle the thin-gap case.

For thick gaps one can ask for positivity frequency by frequency. Call $[N - 1, N]$ the closing gap. Let E be an event that depends only on which gaps are thick and on the pattern of the closing gap up to translations and reflections. For $a \in Z_{2W}$, let X_a be the torus obtained from $Z_W \times [0, N]$ by gluing (x, N) to $(x + a, 0)$, let $s_a \in H^1(X_a; F_2)$ be the class of the double cover obtained in the same way from $Z_{2W} \times [0, N]$, and let $\Psi^E(a)$ be the probability of $E \cap D_{s_a}$ for Bernoulli bond percolation P_p on X_a. Then $X_W = X_0 = Z_W \times Z_N$, $s_0 = s$ and $s_W = s + e$, so D_{s_a} is D_s for $a = 0$ and D_{s+e} for $a = W$. For odd $k \in Z_{2W}$ put $\Delta^E{}_k := W^{-1} \Sigma_{a \in Z_{2W}} \Psi^E(a)\, e^{-i\pi ka/W}$. Reflecting every class maps X_a to X_{-a} and preserves E, so $\Psi^E(-a) = \Psi^E(a)$ and $\Delta^E{}_k$ is real. Then

$$P_p(D_s \cap E) - P_p(D_{s+e} \cap E) = \Sigma_{k \text{ odd}}\, \Delta^E{}_k,$$

where k runs over the W odd residues mod 2W. Moreover $\Delta^E{}_{2j+1} = F_j(\Psi^E{}_-)/(2W)$, where $\Psi^E{}_-(a) := \Psi^E(a) - \Psi^E(a + W)$ and F_j is as in Proposition 17′(ii). Convolution multiplies the F_j, so positivity of every $\Delta^E{}_k$ is preserved under the convolutions of the block decomposition. When the closing gap is thin, every $\Delta^E{}_k$ was positive in all the cases we computed,† although single block laws can have negative coefficients F_j (see above). When the gaps are thick, it can fail:

- On $Z_9 \times Z_2$ at $p = 1/2$, $\Delta^E{}_7 < 0$ for each of three events: E requiring all gaps to be thick, E requiring the closing gap to be thick (the other gap is free), and E the whole space (the total over all closing patterns). For the total, $\Delta^E{}_7 \approx -1.17 \cdot 10^{-3}$.†
- For these events the sum $\Sigma_k\, \Delta^E{}_k = P_p(D_s \cap E) - P_p(D_{s+e} \cap E)$ stays positive in every case we computed.†
- Among the tori with $W \geq 3$ and $N \geq 2$, $Z_9 \times Z_2$ is the smallest such example by area: for these three events and for E requiring the closing gap to be thin, every other such torus with $WN \leq 18$ has all $\Delta^E{}_k > 0$ for every $p \in (0, 1)$, and every value we computed with $W \leq 8$ and $N \geq 2$ is positive (for $W \leq 7$ in floating point, at $p = 0.1, 0.3, 0.5, 0.7, 0.9$ and $N \leq 30$, or $N \leq 25$ in some runs at $W = 7$).†

The failures are not confined to very wide tori. At $p = 1/2$, $\Delta^E{}_7 < 0$ for each of the three events also on $Z_{11} \times Z_3$, $Z_{12} \times Z_3$, $Z_{13} \times Z_4$ and $Z_{14} \times Z_5$, where in the last case $N > W/3$.† On these tori every other

Δ^{E}_{k} is positive, except Δ^{E}_{5} on $Z_{12} \times Z_3$ for E requiring all gaps to be thick.† At $p = 1/2$ and for $W \geq 3$, the smallest failing width is W = 6, 9, 11, 13 and 14 for N = 1, ..., 5: for $N \leq 4$ every narrower torus is positive, and for N = 5 so is every torus with $W \leq 12$, while the positivity of $Z_{13} \times Z_5$ rests on a Monte Carlo estimate.† On the square tori $Z_N \times Z_N$ with $3 \leq N \leq 11$, every Δ^{E}_{k} is positive for the three events at $p = 1/2$.† Apart from the Monte Carlo estimate and the whole space on $Z_{11} \times Z_3$ and $Z_{12} \times Z_3$, which we computed exactly, these values were computed in floating point with a first-order bound on the rounding error. Since P_p restricted to such an E is a mixture of uniform measures on R-orbits (proof of Proposition 15), on each failing torus the failure for E requiring all gaps to be thick gives an R-orbit of configurations with all gaps thick that has a negative coefficient. So, on the failing tori, positivity frequency by frequency does not yield the thick-gap case of Conjecture F.

In the scaling limit there are failures on wide, square and long tori. Assume that at $p = 1/2$, as $W \to \infty$ with $N/W \to \rho$, the probabilities of the homology classes of the open clusters on X_a converge, uniformly in a, to their values in the Coulomb-gas description of critical percolation on the torus [DSZ87, Pin94] with modular parameter $\tau = (-a + iN)/W$. This convergence is not proved, and no rate is known. In this limit each Δ^{E}_{k} converges, and the three events have the same limits Δ_k, since they differ only on configurations with a gap of at most one open edge, which have probability at most $N(W + 1)2^{-W}$. Then some limiting sector is negative at each of the following ratios:

- Δ_7 on a grid of step 0.01 in $0.03 \leq \rho < 0.9768$, and Δ_{11} and Δ_{13} at $\rho = 0.02$;†
- Δ_5 on a grid of step 0.02 in $2.8082 \leq \rho \leq 12$;†
- a negative tail at $\rho = 0.98$, 1, 1.2, 1.5, 2, 2.5, 3 and 4: every computed Δ_k from some $k_0(\rho)$ on is negative, with $k_0 = 1031$, 987, 665, 405, 209, 123, 77 and 35.† On the square $\rho = 1$ every computed Δ_k with $k \geq 987$ is negative (all odd k up to 1201, and k = 1501, 1503, 2001, 3001 and 4001), beginning with $\Delta_{987} \approx -5.2{\cdot}10^{-1333}$.†

We have no proof that every ratio fails. The lattice and the limit agree in sign only in part: on $Z_9 \times Z_3$ and $Z_{12} \times Z_4$ every lattice sector is positive, although the limiting Δ_5 and Δ_7 at $\rho = 1/3$ are negative, and no lattice failure with $N \geq W$ is seen at computable sizes.† On square tori, then, positivity frequency by frequency fails in the limit. But the negative sectors we found on the square have $k \geq 987$, and the sectors satisfy $k \leq W$, so this failure needs $W \geq 987$, and possibly much larger widths.

The odd-cover argument sketched above gives Conjecture F at q = 1 for the odd covers of the configurations of Propositions 16 and 17. These covers have thick gaps.

8. Divisibility of the classes of giant cycles near p = 1/2

In this section P is Bernoulli(p) plaquette percolation on $X = T^4_N$ (q = 1, i = 2), and $P\bullet$ is its dual complex, which is Bernoulli(1 − p) plaquette percolation on the dual torus $X\bullet$. Propositions 8 and 10 constrain the giant cycles below 1/2 with coefficients F_2 and Z/4. Here we study integral giant cycles. Above 1/2 they exist in every direction [DKS20, Theorem 1 with F = Q]. We show that just above 1/2 the classes of all integral giant cycles are divisible by any fixed integer M. As in §1, we use the strict-shift argument of [AG91] in the finite-volume form of [BBR14]. Its input is the local conversion lemma of §8.3, which §§8.4–8.7 prove for all configurations.

Conventions.

- All cells carry fixed orientations. For an abelian group Γ put $V_\Gamma(P) := \mathrm{im}(H_2(P; \Gamma) \to H_2(X; \Gamma))$, so $V = V_{F_2}$ and $\Lambda_P := V_Z(P) \subseteq H_2(X; Z) = Z^6$. For $m = 2^k$ put $L_m(P) := \mathrm{red}_2\, V_{Z/m}(P) \subseteq {F_2}^6$.

Here red_r is reduction mod r of coefficients and of classes, and v_ℓ is the ℓ-adic valuation. Reduction of coefficients maps cycles to cycles, so $\mathrm{red}_2\Lambda_P \subseteq L_8(P) \subseteq L_4(P) \subseteq V(P)$.

- Through the translation by (½, ½, ½, ½) we identify $H_2(X\bullet; \Gamma)$ with $H_2(X; \Gamma) = \Lambda^2\Gamma^4 = \Gamma^6$. The translation acts trivially on homology. The cellular intersection number of a 2-cycle of X with a 2-cycle of $X\bullet$ is then the intersection form $\langle z, w\rangle$, defined by $z \wedge w = \langle z, w\rangle\, e_1\wedge e_2\wedge e_3\wedge e_4$. It is unimodular, and $\langle e_I, e_J\rangle = \pm 1\{I \cap J = \emptyset\}$.
- The signed permutations $g \in B_4$ are automorphisms of the lattice Z^4. They act on X, on $\Gamma^6 = \Lambda^2\Gamma^4$ by $\Lambda^2 g$, and they preserve the law of P. An element $x \in (Z/m)^6$ is *partnered* if $\langle x, gx\rangle \neq 0$ in Z/m for some $g \in B_4$. For $m = 2$ this is the notion of §7, because signs act trivially mod 2.
- $B_t(y)$ is the set of plaquettes whose centres lie within $l\infty$ distance t of the centre of the plaquette y, and $|B_t|$ is its cardinality.
- In this section the letter d denotes an ℓ-adic exponent, not the dimension.

8.1 Statement

Theorem W (divisibility near p = 1/2). For every integer $M \geq 2$ there is $w_M > 0$ such that, as $N \to \infty$:

- (i) $P_p(\Lambda_P \not\subset MZ^6) \to 0$ for all $p < 1/2 + w_M$. More precisely, for each prime power ℓ^d dividing M, $P_p(V_{Z/\ell^n}(P) \not\subset \ell^d(Z/\ell^n)^6) \to 0$, where $n := 3d$ if $\ell = 2$ and $n := 3d - 1$ if ℓ is odd.
- (ii) For $p \in (1/2, 1/2 + w_M)$, with probability tending to 1, $\Lambda_P \subseteq MZ^6$ and rank $\Lambda_P = 6$. In particular M^6 divides the index $[Z^6 : \Lambda_P]$.
- (iii) If M is even, $P_p(L_4(P) \not\subset C^\perp) \to 0$ for all $p < 1/2 + w_M$.
- (iv) For $p > 1/2 - w_M$ and each prime power ℓ^d dividing M, with probability tending to 1, $V_{Z/\ell^n}(P) \supseteq \ell^{n-d}(Z/\ell^n)^6$.

Since $1/2 < 1/2 + w_M$, statements (i) and (iii) hold at $p = 1/2$. For $M = 2$, part (i) says that $L_8(P) = 0$ with probability tending to 1 for $p < 1/2 + w_2$: no giant mod-2 cycle lifts to a Z/8-cycle. Part (ii) says that just above 1/2 the integral giant lattice has full rank, while all integral giant cycles have even classes. Part (iii) says that no partnered class lifts to a Z/4-cycle, and part (iv) says that for $p > 1/2 - w_2$, in particular just below 1/2, each class $4\tilde{s}$ ($s \in {F_2}^6$, $\tilde{s}$ a lift) is the class of a giant Z/8-cycle. For each M the index of Λ_P is a multiple of M^6 for $p \in (1/2, 1/2 + w_M)$, with probability tending to 1, so it tends to infinity as $p \downarrow 1/2$ in this sense. However, Theorem W does not settle the case $q = 1$ of Theorem A, which concerns V(P) below 1/2. The constant w_M is explicit (Remark 1 in §8.8) and very small.

Corollary W1 (an F_ℓ-only window). Let ℓ be an odd prime. For $p \in (1/2 - w_\ell, 1/2 + w_\ell)$, with probability tending to 1, $V_{Z/\ell^2}(P) = \ell(Z/\ell^2)^6$. For $p \in (1/2, 1/2 + w_\ell)$, moreover $V_{F_\ell}(P) = {F_\ell}^6$ with probability tending to 1.

In particular, just above 1/2 there are giant F_ℓ-cycles in all classes, and none of them lifts to a Z/ℓ^2-cycle.

Corollary W2 (torsion just below 1/2). For $p \in (1/2 - w_M, 1/2)$, with probability tending to 1, the torsion subgroup T of $H_1(P; Z)$ has order at least M^6 and exponent divisible by M, and the linking map of Proposition 21 maps T onto a subgroup of Q^6/Z^6 that contains $(M^{-1}Z/Z)^6$.

8.2 Algebraic inputs

We collect the facts about coefficients that the proof uses: the partner rule at every level, exact duality, the subcritical bound, and the linking map of the torsion.

Lemma 18 (partner rule).

- (a) Let ℓ be an odd prime. Every nonzero $z \in (\mathbb{Z}/\ell)^6$ is partnered. Hence each $x \in (\mathbb{Z}/\ell^a)^6$ of order ℓ^b with $2b > a$ is partnered.
- (b) Each $x \in (\mathbb{Z}/2^a)^6$ of order 2^b with $2b \geq a + 2$ is partnered.

Proof. (a) Write $z = \Sigma\, z_{ij}\, e_i \wedge e_j$ and put $\alpha_1 := z_{12}z_{34}$, $\alpha_2 := z_{13}z_{24}$, $\alpha_3 := z_{14}z_{23}$. Let g_{jk} be the diagonal sign change of coordinates j and k. A direct computation gives

$\langle z, z\rangle = 2(\alpha_1 - \alpha_2 + \alpha_3)$, $\langle z, g_{34}z\rangle = 2(\alpha_1 + \alpha_2 - \alpha_3)$, $\langle z, g_{24}z\rangle = -2(\alpha_1 + \alpha_2 + \alpha_3)$, $\langle z, g_{23}z\rangle = 2(-\alpha_1 + \alpha_2 + \alpha_3)$.

For odd ℓ these four linear forms span all linear forms in $(\alpha_1, \alpha_2, \alpha_3)$ mod ℓ (the 3×3 minors of the coefficient matrix of $\alpha_1 - \alpha_2 + \alpha_3$, $\alpha_1 + \alpha_2 - \alpha_3$, $\alpha_1 + \alpha_2 + \alpha_3$, $-\alpha_1 + \alpha_2 + \alpha_3$ are ± 4). So z is partnered unless $\alpha_1 \equiv \alpha_2 \equiv \alpha_3 \equiv 0$. In that case the support of z, read as a set of edges of K_4, contains no two disjoint edges, so it lies in a star or in a triangle. Then $z = v_1 \wedge v_2$ is decomposable over $\mathbb{F}_\ell$. With $T := \mathrm{span}(v_1, v_2)$, the number $\langle z, gz\rangle = \det[v_1, v_2, gv_1, gv_2]$ is nonzero if and only if $gT \cap T = 0$. After a coordinate permutation, $T = \mathrm{span}(e_1, w)$ with $w = (0, w_2, w_3, w_4) \neq 0$ (star), or $T = \{x_4 = 0,\ n\cdot x = 0\}$ with $n = (n_1, n_2, n_3, 0) \neq 0$ (triangle). The proof of Lemma 12′ uses only linear algebra and polynomial identities, so it holds over $\mathbb{F}_\ell$. For a suitable pair $(x, y) \neq (0, 0)$ of coordinates of w, resp. of n, it gives $g, g' \in B_4$ with $gT \cap T = 0$ if and only if $2xy \neq 0$, and $g'T \cap T = 0$ if and only if $x^2 - y^2 \neq 0$. For odd ℓ these two numbers do not both vanish.

For the second claim of (a), write $x = \ell^{a-b}y$ with $y \not\equiv 0$ mod ℓ. Then $\langle x, gx\rangle = \ell^{2(a-b)}\langle y, gy\rangle$. Choosing g with $\langle y, gy\rangle \not\equiv 0$ mod ℓ gives $\langle x, gx\rangle \not\equiv 0$ mod $\ell^{2(a-b)+1}$, and $2(a - b) + 1 \leq a$. For (b), write $x = 2^{a-b}y$ with $y \not\equiv 0$ mod 2. The reduction of y mod 4 has order 4, so Proposition 10(i) gives g with $\langle y, gy\rangle \not\equiv 0$ mod 4. Hence $\langle x, gx\rangle \not\equiv 0$ mod $2^{2(a-b)+2}$, and $2(a - b) + 2 \leq a$. ■

The identities in the proof, and the conclusion of (a) for $\ell = 3, 5, 7$, are also checked by computer.†

Lemma 19 (exact duality). Let $m \geq 2$ and $A := \mathbb{Z}/m$. For any configuration, $V_A(P\bullet) = V_A(P)^\perp$ with respect to $\langle\cdot,\cdot\rangle$ on A^6.

Proof. The cells of X that are not in P are the closed plaquettes and all 3- and 4-cells. Their duals are exactly the cells of $P\bullet$. Sending a cell σ to the cochain dual to σ^* identifies $C_k(X; A)/C_k(P; A)$ with $C^{4-k}(P\bullet; A)$, and the boundary map with the coboundary map up to signs. Hence $H_2(X, P; A) \cong H^2(P\bullet; A)$. The ring A is self-injective, so $\mathrm{Hom}_A(\cdot, A)$ is exact and $H^2(P\bullet; A) \cong \mathrm{Hom}_A(H_2(P\bullet; A), A)$. The composite $H_2(X; A) \to H_2(X, P; A) \to \mathrm{Hom}_A(H_2(P\bullet; A), A)$ sends x to $y \mapsto \langle x, y\rangle$. Exactness of $H_2(P; A) \to H_2(X; A) \to H_2(X, P; A)$ gives $V_A(P) = V_A(P\bullet)^\perp$. The dual complex of $P\bullet$ is P, up to the translation, so the same argument gives $V_A(P\bullet) = V_A(P)^\perp$. ■

With $A = \mathbb{Q}$ in place of $\mathbb{Z}/m$, the same argument gives $\dim V_{\mathbb{Q}}(P) + \dim V_{\mathbb{Q}}(P\bullet) = 6$.

Proposition 20 (the subcritical bound at all levels). Let $m \geq 2$ and let $x \in (\mathbb{Z}/m)^6$ be partnered. Then $P_p(x \in V_{\mathbb{Z}/m}(P)) \to 0$ for all $p < 1/2$.

Proof. Choose $g \in B_4$ with $\langle x, gx\rangle \neq 0$. Disjoint cycles have intersection number 0, so $x \in V_{\mathbb{Z}/m}(P)$ and $gx \in V_{\mathbb{Z}/m}(P\bullet)$ cannot both hold. At $p = 1/2$, $V_{\mathbb{Z}/m}(P\bullet)$ has the law of $V_{\mathbb{Z}/m}(P)$, and g preserves this law. Hence $P_{1/2}(x \in V_{\mathbb{Z}/m}(P)) \leq 1/2$. The event $\{x \in V_{\mathbb{Z}/m}(P)\}$ is increasing and invariant under translations, whose orbits on plaquettes have size N^4. By [BR06, Lemma 7], $P_p(x \in V_{\mathbb{Z}/m}(P)) \to 0$ for all $p < 1/2$. ■

Propositions 8 and 10 are the cases $m = 2$ and $m = 4$ of Proposition 20, together with the classification of the partnered elements.

Proposition 21 (torsion linking). Let P be a configuration with rank $\Lambda_P = 6$, and let $T\bullet$ be the torsion subgroup of $H_1(P\bullet; Z)$. For $[\gamma] \in T\bullet$ choose $m \geq 1$ and an integral 2-chain W on $P\bullet$ with $\partial W = m\gamma$, and put $\Psi([\gamma]) := m^{-1}[W \bmod m] \in Q^6/Z^6$, where $[W \bmod m] \in (Z/m)^6$ is the class of the Z/m-cycle W mod m and m^{-1} maps $(Z/m)^6$ onto $(m^{-1}Z/Z)^6$. Then Ψ is a well-defined homomorphism $T\bullet \to Q^6/Z^6$ with image ${\Lambda_P}^{\#}/Z^6$, where ${\Lambda_P}^{\#} := \{y \in Q^6 : \langle y, \Lambda_P\rangle \subseteq Z\}$. In particular $|T\bullet| \geq [Z^6 : \Lambda_P]$. If $\Lambda_P \subseteq MZ^6$, the image contains $(M^{-1}Z/Z)^6$ and M divides the exponent of $T\bullet$.

Proof.

1. By the remark after Lemma 19, $V_Q(P\bullet) = 0$. Since $H_2(X\bullet; Z)$ is torsion-free, all integral 2-cycles of $P\bullet$ are null-homologous in $X\bullet$. Given (γ, m), the chain W is unique up to such cycles, which do not change $[W \bmod m]$. Replacing γ by $\gamma + \partial c$ replaces W by $W + mc$, and replacing m by km replaces W by kW. Ψ is additive.
2. As in the proof of Lemma 19, $H_2(X, P; Z) \cong H^2(P\bullet; Z)$. The class $x \in H_2(X; Z)$ maps to the class of a cochain f_x with $f_x(W) = \langle x, W\rangle$ for integral 2-chains W on $P\bullet$. The kernel of $H_2(X; Z) \to H_2(X, P; Z)$ is Λ_P.
3. The cochain f_x vanishes on the cycles of $P\bullet$, by step 1. Its class therefore lies in the subgroup $\mathrm{Ext}(H_1(P\bullet; Z), Z)$ of $H^2(P\bullet; Z)$, which is identified with $\mathrm{Hom}(T\bullet, Q/Z)$ through $\chi_f([\gamma]) := f(W)/m \bmod 1$, for $m\gamma = \partial W$.
4. Hence $x \mapsto \chi_{f_x}$ is injective on Z^6/Λ_P. Since $\langle x, W\rangle \bmod m$ is the Z/m intersection number of x with the cycle W mod m, $\chi_{f_x}([\gamma]) = \langle x, \Psi([\gamma])\rangle \bmod 1$. For $x \in \Lambda_P$ the cochain f_x is a coboundary, so $\langle x, \Psi(T\bullet)\rangle \subseteq Z$ and $\Psi(T\bullet) \subseteq {\Lambda_P}^{\#}/Z^6$.
5. The group Z^6/Λ_P is finite, and $\langle\cdot,\cdot\rangle$ induces a perfect pairing between Z^6/Λ_P and ${\Lambda_P}^{\#}/Z^6$, because the form is unimodular. By step 4 the annihilator of $\Psi(T\bullet)$ in Z^6/Λ_P is trivial, so $\Psi(T\bullet) = {\Lambda_P}^{\#}/Z^6$.

Finally $|{\Lambda_P}^{\#}/Z^6| = [Z^6 : \Lambda_P]$. If $\Lambda_P \subseteq MZ^6$ then ${\Lambda_P}^{\#} \supseteq M^{-1}Z^6$, and any preimage of an element of order M has order divisible by M. ■

8.3 The strict-shift argument

We first show that a local conversion lemma rules out certain giant cycles slightly above 1/2.

Setting. Fix a prime ℓ, integers $a, d \geq 1$, $n := a + d$, the ring $A := Z/\ell^n$, and a submodule $\Theta \subseteq A^6$ such that

- (Θ1) $\ell^d A^6 \subseteq \Theta$;
- (Θ2) for each $v \in A^6 \setminus \Theta$, the reduction $\mathrm{red}_{\ell^a}(v) \in (Z/\ell^a)^6$ is partnered.

An A-cycle is Θ-*odd* if its class does not lie in Θ. Let $E_\Theta(\omega)$ be the event that the plaquette set ω carries a Θ-odd A-cycle, that is, $V_A(\omega) \not\subset \Theta$. It is increasing. Theorem W uses two instances.

- (I1) Given a prime power ℓ^d, take $a := 2d$ if $\ell = 2$ and $a := 2d - 1$ if ℓ is odd, so $n = 3d$ or $3d - 1$, and $\Theta := \ell^d A^6$. Then $E_\Theta = \{V_{Z/\ell^n}(P) \not\subset \ell^d(Z/\ell^n)^6\}$. Condition ($\Theta$1) is trivial. For ($\Theta$2), let $v \notin \ell^d A^6$, so $v_\ell(v) = j < d$. Then $\mathrm{red}_{\ell^a}(v)$ has order ℓ^b with $b = a - j \geq a - d + 1$. For $\ell = 2$ this gives $2b \geq 2d + 2 = a + 2$, and for odd ℓ it gives $2b \geq 2d = a + 1$. In both cases Lemma 18 applies.

- (I2) Take $\ell = 2$, $a = d = 1$ and $\Theta := \mathrm{red}_2^{-1}(C^\perp) \subseteq (\mathbb{Z}/4)^6$. Then $E_\Theta = \{L_4(P) \not\subset C^\perp\}$. Condition ($\Theta$1) holds because $\mathrm{red}_2(2A^6) = 0$, and (Θ2) is Proposition 8.

Enhancement. Fix $\rho \geq 1$. Give each plaquette y an independent bit ξ_y of mean $s \in [0, 1]$, independent of P. For a plaquette set ω let

$\omega^e := \omega \cup \{y : \xi_y = 1 \text{ and } \partial y \in \partial C_2(\omega \cap B_\rho(y); \mathbb{Z}/\ell^a)\}$.

The condition says that ∂y bounds a $\mathbb{Z}/\ell^a$-chain on the plaquettes of ω within distance ρ of y. The configuration ω^e is increasing in ω and in ξ.

Lemma 22. Let $N > 2\rho + 2$. Then $V_{\mathbb{Z}/\ell^a}(\omega^e) = V_{\mathbb{Z}/\ell^a}(\omega)$. Hence, if $E_\Theta(\omega^e)$ holds, then $V_{\mathbb{Z}/\ell^a}(\omega)$ contains a partnered element.

Proof. For each added y choose a $\mathbb{Z}/\ell^a$-chain g_y on $\omega \cap B_\rho(y)$ with $\partial g_y = \partial y$. Then $y - g_y$ is a $\mathbb{Z}/\ell^a$-cycle in a box of side less than N, hence null-homologous. For a $\mathbb{Z}/\ell^a$-cycle z of ω^e, the cycle $z - \Sigma_y\, z(y)(y - g_y)$ lies on ω and has the class of z. For the second claim, let $v \in V_A(\omega^e) \setminus \Theta$. Then $\mathrm{red}_{\ell^a}(v) \in V_{\mathbb{Z}/\ell^a}(\omega^e) = V_{\mathbb{Z}/\ell^a}(\omega)$, and it is partnered by (Θ2). ■

A plaquette τ is *E_Θ-pivotal* for ω if $E_\Theta(\omega \cup \tau)$ holds and $E_\Theta(\omega \setminus \tau)$ fails.

Local conversion lemma $\mathrm{LCL}_\Theta(\rho, \kappa)$. There is N_0 such that the following holds for all $N \geq N_0$. Let ω be any set of plaquettes of T^4_N and let τ be E_Θ-pivotal for ω. Then there are a plaquette set ω' with $\omega' \,\Delta\, \omega \subseteq B_\kappa(\tau)$ and a plaquette $x \in B_\kappa(\tau) \setminus \omega'$ such that:

- $E_\Theta(\omega' \cup x)$ holds;
- $E_\Theta(\omega')$ fails;
- $\partial x \in \partial C_2(\omega' \cap B_\rho(x); \mathbb{Z}/\ell^a)$.

Proposition 23 (strict shift). Assume (Θ2) and $\mathrm{LCL}_\Theta(\rho, \kappa)$. Put $\kappa' := \kappa + 2\rho + 1$ and $\delta := 1/(|B_{\kappa'}| \cdot 16^{|B_{\kappa'}|})$. Then $P_p(E_\Theta(P)) \to 0$ for all $p < 1/2 + \delta/2$.

Proof. We follow [AG91] in the finite-volume form of [BBR14]. Let $N \geq \max(N_0, 2\kappa' + 2)$.

1. *Russo.* Put $\theta_N(p, s) := P_{p,s}(E_\Theta(P^e))$. A plaquette is *p-pivotal* in (ω, ξ) if changing its state changes $E_\Theta(\omega^e)$, and a bit is *s-pivotal* if changing it changes $E_\Theta(\omega^e)$. Let Piv_p and Piv_s count them. The event is increasing in finitely many independent bits, so $\partial_p \theta_N = \mathrm{E}\,\mathrm{Piv}_p$ and $\partial_s \theta_N = \mathrm{E}\,\mathrm{Piv}_s$.
2. *Surgery.* Let τ be p-pivotal in (ω, ξ). We keep τ open and switch off the bits in $B_{\kappa'}(\tau)$ one at a time, in a fixed order.
   - Suppose E_Θ first fails when the bit of y is switched off. We output y, together with the configuration just before that step. The bit of y is s-pivotal there.
   - Otherwise let $\tilde\xi$ be ξ with all bits in $B_{\kappa'}(\tau)$ switched off. Then τ is still p-pivotal in $(\omega, \tilde\xi)$, because $E_\Theta((\omega \setminus \tau)^e)$ is increasing in the bits. Let Y_{far} be the set of plaquettes outside ω that are added by bits at distance more than κ' from τ. Their balls B_ρ avoid $B_{\kappa+\rho}(\tau)$, so Y_{far} does not change when ω changes inside $B_\kappa(\tau)$. Now τ is E_Θ-pivotal for $\omega \cup Y_{\mathrm{far}}$, and LCL_Θ gives $\hat\omega'$ and x. Put $\omega' := \hat\omega' \setminus Y_{\mathrm{far}}$, which agrees with ω outside $B_\kappa(\tau)$. In $(\omega', \tilde\xi)$ the enhanced configuration is $\hat\omega' \cup x$ if the bit of x is on and $\hat\omega'$ if it is off, because $\hat\omega' \cap B_\rho(x) = \omega' \cap B_\rho(x)$. So the bit of x is s-pivotal. We output $(\omega', \tilde\xi)$ and x. Pivotality of a bit does not depend on its own value, and $\tilde\xi_x = 0$.
3. *Distortion.* An output differs from (ω, ξ) in at most $|B_{\kappa'}|$ plaquette states and $|B_{\kappa'}|$ bits, all in $B_{\kappa'}(\tau)$, and bits change only from 1 to 0. For $s \leq 1/2$ its probability is therefore at least $\min(p, 1 - p)^{|B_{\kappa'}|}$ times that of (ω, ξ). The output site lies in $B_{\kappa'}(\tau)$, so an output arises from at most $|B_{\kappa'}|$

choices of τ and $4^{|\mathrm{B}_{\kappa'}|}$ original values in the ball. Hence E $\mathrm{Piv}_{\mathrm{s}} \geq \delta$ E $\mathrm{Piv}_{\mathrm{p}}$ for $\mathrm{p} \in [1/4, 3/4]$ and $\mathrm{s} \in [0, 1/2]$.

4. *Integration.* By steps 1 and 3, $\mathrm{t} \mapsto \theta_{\mathrm{N}}(\mathrm{p}_0 - \delta\mathrm{t}, \mathrm{t})$ is nondecreasing on $[0, 1/2]$ for $\mathrm{p}_0 \in [1/2, 1/2 + \delta/2)$. Hence $\mathrm{P}_{\mathrm{p}_0}(\mathrm{E}_\Theta(\mathrm{P})) = \theta_{\mathrm{N}}(\mathrm{p}_0, 0) \leq \theta_{\mathrm{N}}(\mathrm{p}_0 - \delta/2, 1/2)$. By Lemma 22, the right side is at most the probability that $\mathrm{V}_{\mathrm{Z}/\ell^{\mathrm{a}}}(\mathrm{P})$ contains a partnered element, for P with parameter $\mathrm{p}_0 - \delta/2 < 1/2$. This tends to 0 by Proposition 20 and a union bound over the finitely many partnered elements. For $\mathrm{p}_0 < 1/2$ the same bound applies directly, with $\mathrm{s} = 0$. ■

8.4 Descending funnels

We now reduce LCL_Θ to the existence of a design inside a fixed region. Section 8.5 will reduce the existence of designs to three finite conditions on the region, and §8.7 will verify them. Fix $(\ell, \mathrm{a}, \mathrm{d}, \Theta)$ as in §8.3, satisfying $(\Theta 1)$.

Regions. A *region* is a finite union $\mathfrak{D}$ of closed unit 4-cubes of the lattice that is star-shaped about one of its vertices o, hence contractible. It is placed with o a vertex of the pivotal plaquette τ, and N is large enough that $\mathfrak{D}$ embeds in X and that all A-cycles supported in $\mathfrak{D}$ are null-homologous in X. Let $\mathcal{P}(\mathfrak{D})$ be the set of plaquettes that are faces of cubes of $\mathfrak{D}$. We assume $\tau \in \mathcal{P}(\mathfrak{D})$.

- An edge of $\mathfrak{D}$ is *safe* if all six plaquettes containing it lie in $\mathcal{P}(\mathfrak{D})$. Let NS be the set of edges of $\mathfrak{D}$ that are not safe. A plaquette outside $\mathcal{P}(\mathfrak{D})$ can contain an edge of $\mathfrak{D}$ only if that edge lies in NS.
- NS carries a fixed total order.
- Let $\mathrm{O} := \omega \setminus \mathcal{P}(\mathfrak{D})$ be the *outside*. Since $\mathrm{O} \subseteq \omega \setminus \tau$ and τ is pivotal, O carries no Θ-odd cycle.

Relative cycles. A *relative cycle* on O is an A-chain U on O whose boundary lies in $\mathfrak{D}$. Its *trace* $\gamma := -\partial\mathrm{U}$ is a 1-cycle on NS. Since $\mathfrak{D}$ is contractible, $\gamma = \partial\mathrm{G}$ for a chain G in $\mathcal{P}(\mathfrak{D})$, and two such G differ by a cycle in $\mathfrak{D}$, which is null-homologous. We call U Θ-*odd* if the cycle U + G is Θ-odd for one, equivalently every, such G. The restriction to O of a Θ-odd cycle of $\omega \cup \tau$ is a Θ-odd relative cycle.

Lemma 24 (the min-max edge). Let e^* be the least edge $\mathrm{e} \in \mathrm{NS}$ for which some Θ-odd relative cycle on O has trace supported on $\{\mathrm{f} \in \mathrm{NS} : \mathrm{f} \leq \mathrm{e}\}$. Fix such a U and put $\gamma := -\partial\mathrm{U}$. Then:

- (a) no Θ-odd relative cycle on O has trace supported on $\{\mathrm{f} \in \mathrm{NS} : \mathrm{f} < \mathrm{e}^*\}$;
- (b) $\mathrm{e}^* \in \operatorname{supp} \gamma$;
- (c) each endpoint of e^* lies on an edge of $\operatorname{supp} \gamma$ other than e^*, and that edge is $< \mathrm{e}^*$.

Proof. (a) A nonzero such trace contradicts the minimality of e^*, and a zero trace gives a Θ-odd cycle on O. (b) follows from (a). For (c), $\partial\gamma = 0$ and $\gamma(\mathrm{e}^*) \neq 0$, so each endpoint of e^* lies on another edge of $\operatorname{supp} \gamma$, and all edges of $\operatorname{supp} \gamma$ are $\leq \mathrm{e}^*$. ■

Port networks. Let $\mathrm{t} \in \{0, \ldots, \mathrm{n} - 1\}$. A *port network of type t* at a plaquette π is a set $\mathcal{N}$ of plaquettes with $\pi \notin \mathcal{N}$, together with a plaquette $\mathrm{x} \notin \mathcal{N} \cup \{\pi\}$, such that:

- (N1) each A-chain on $\mathcal{N}$ whose boundary lies on the four edges of $\partial\pi$ has boundary $\mathrm{b}\partial\pi$ with $\mathrm{b} \in \ell^{\mathrm{t}+\mathrm{d}}\mathrm{A}$;
- (N2) for every $\mathrm{b} \in \ell^{\mathrm{t}}$ A there is an A-chain on $\mathcal{N} \cup \{\mathrm{x}\}$ with boundary $\mathrm{b}\partial\pi$;
- (N3) $\partial\mathrm{x} \in \partial\mathrm{C}_2(\mathcal{N} \cap \mathrm{B}_\rho(\mathrm{x}); \mathrm{Z}/\ell^{\mathrm{a}})$.

Here $\ell^{\mathrm{t}+\mathrm{d}}\mathrm{A} = 0$ when $\mathrm{t} + \mathrm{d} \geq \mathrm{n}$.

Funnel designs. Let $\mathrm{t} := \mathrm{v}_\ell(\gamma(\mathrm{e}^*))$. A *funnel design* for (γ, e^*) is a 2-chain F in $\mathcal{P}(\mathfrak{D})$ together with a port network $(\mathcal{N}, \mathrm{x})$ of type t at a plaquette π of supp F, such that:

- (F1) $\partial \mathrm{F} = \gamma$;
- (F2) the NS edges of the plaquettes of supp F are all $\leq \mathrm{e}^*$;
- (F3) e^* lies in exactly one plaquette Q_0 of supp F;
- (F4) there are pairwise distinct plaquettes $\mathrm{Q}_0, \mathrm{Q}_1, \ldots, \mathrm{Q}_\mathrm{J} = \pi$ of supp F such that $\mathrm{Q}_{\mathrm{k}-1}$ and Q_k share a safe edge σ_k, and σ_k lies in no other plaquette of supp F;
- (F5) $\mathcal{N} \cup \{\mathrm{x}\} \subseteq \mathcal{P}(\mathfrak{D})$, and no edge of $\mathcal{N} \cup \{\mathrm{x}\}$ other than the four edges of $\partial\pi$ lies in a plaquette of $\mathrm{O} \cup (\mathrm{supp\ F} \setminus \pi)$.

Lemma 25. Assume (Θ1) and $\mathcal{P}(\mathfrak{D}) \subseteq \mathrm{B}_\kappa(\tau)$. If a funnel design exists, then $\omega' := \mathrm{O} \cup (\mathrm{supp\ F} \setminus \pi) \cup \mathcal{N}$ and x satisfy the three conditions of $\mathrm{LCL}_\Theta(\rho, \kappa)$ at τ.

Proof. Each plaquette of $\mathcal{N}$ has an edge outside $\partial\pi$, so by (F5) $\mathcal{N}$ is disjoint from $\mathrm{O} \cup (\mathrm{supp\ F} \setminus \pi)$. For the same reason $\mathrm{x} \notin \omega'$. Moreover $\omega' \,\Delta\, \omega \subseteq \mathcal{P}(\mathfrak{D}) \subseteq \mathrm{B}_\kappa(\tau)$.

- *$E_\Theta(\omega' \cup x)$.* $\Xi := \mathrm{U} + \mathrm{F}$ is a Θ-odd cycle on $\mathrm{O} \cup \mathrm{supp\ F}$. By (F3), $\mathrm{F}(\mathrm{Q}_0) = \pm\gamma(\mathrm{e}^*)$. At σ_k only $\mathrm{Q}_{\mathrm{k}-1}$ and Q_k carry F, by (F4), and $\partial\mathrm{F}(\sigma_\mathrm{k}) = \gamma(\sigma_\mathrm{k}) = 0$ because σ_k is safe. Hence $\mathrm{F}(\pi) = \pm\gamma(\mathrm{e}^*)$, which has valuation t. By (N2) there is a chain n^* on $\mathcal{N} \cup \{\mathrm{x}\}$ with $\partial\mathrm{n}^* = \Xi(\pi)\partial\pi$. Then $\Xi - \Xi(\pi)\pi + \mathrm{n}^*$ is a cycle on $\omega' \cup \mathrm{x}$. It differs from Ξ by a cycle in $\mathcal{P}(\mathfrak{D})$, which is null-homologous, so it is Θ-odd.
- *The enhancement condition* is (N3), since $\mathcal{N} \subseteq \omega'$.
- *$\neg E_\Theta(\omega')$.* Suppose Ω is a Θ-odd cycle on ω'.

  1. Write $\Omega = \Omega_1 + \mathrm{n}$ with n on $\mathcal{N}$. By (F5), $\partial\mathrm{n}$ lies on the edges of $\partial\pi$, and a 1-cycle on these four edges is a multiple of $\partial\pi$. So $\partial\mathrm{n} = \mathrm{b}\partial\pi$ with $\mathrm{b} \in \ell^{\mathrm{t}+\mathrm{d}}\mathrm{A}$ by (N1). Then $\Omega' := \Omega_1 + \mathrm{b}\pi$ is a cycle on $\mathrm{O} \cup \mathrm{supp\ F}$ with the class of Ω, because $\mathrm{n} - \mathrm{b}\pi$ is a cycle in $\mathcal{P}(\mathfrak{D})$. Hence Ω' is Θ-odd, and $\Omega'(\pi) = \mathrm{b}$.
  2. Write $\mathrm{F}(\pi) = \ell^\mathrm{t}\,\nu$ with ν a unit, and $\mathrm{b} = \ell^{\mathrm{t}+\mathrm{d}}\mu$. Put $\zeta := \ell^\mathrm{d}\,\mu\,\nu^{-1} \in \ell^\mathrm{d}\,\mathrm{A}$ and $\Omega'' := \Omega' - \zeta\Xi$. Then $\Omega''(\pi) = 0$. Its class differs from that of Ω' by an element of $\ell^\mathrm{d}\,\mathrm{A}^6 \subseteq \Theta$, by ($\Theta$1), so Ω'' is Θ-odd.
  3. Write $\Omega'' = \mathrm{U}'' + \mathrm{Y}$ with U'' on O and Y on supp F. No plaquette of O contains a safe edge, so $\partial\mathrm{Y}$ vanishes on safe edges. By (F4), $\mathrm{Y}(\mathrm{Q}_{\mathrm{k}-1}) = \pm\mathrm{Y}(\mathrm{Q}_\mathrm{k})$ for all k. Hence $\mathrm{Y}(\mathrm{Q}_0) = \pm\mathrm{Y}(\pi) = 0$.
  4. The trace of U'' is $\partial\mathrm{Y}$. It lies on NS edges of plaquettes of supp F, which are $\leq \mathrm{e}^*$ by (F2), and it vanishes at e^* by (F3), since $\mathrm{Y}(\mathrm{Q}_0) = 0$. Also U'' is Θ-odd, since $\mathrm{U}'' + \mathrm{Y} = \Omega''$. This contradicts Lemma 24(a). ■

8.5 Existence of designs

Holes and sublevel complexes. Fix an axis box $\mathfrak{B} \subseteq \mathfrak{D}$ with $\mathfrak{B} \pm 1 \subseteq \mathfrak{D}$ (the box enlarged by 1 in each coordinate), and a plaquette $\pi_\mathfrak{B}$ on a face of $\mathfrak{B}$. A port network $(\mathcal{N}, \mathrm{x})$ at $\pi_\mathfrak{B}$ *fits* in $\mathfrak{B}$ if all vertices of $\mathcal{N} \cup \{\mathrm{x}\}$ other than the four vertices of $\pi_\mathfrak{B}$ lie in the interior of $\mathfrak{B}$. Call a vertex, an edge or a plaquette *available* if none of its vertices lies in the interior of $\mathfrak{B}$. All NS edges are available, since an edge with a vertex in the interior of $\mathfrak{B}$ has all its plaquettes in $\mathcal{P}(\mathfrak{B} \pm 1)$ and is therefore safe. For $\mathrm{h} \in \mathrm{NS} \cup \{\infty\}$, let $\mathcal{K}_{<\mathrm{h}}$ be the 2-complex whose vertices are the available vertices, whose edges are the

available safe edges and the NS edges below h, and whose 2-cells are the available plaquettes of $\mathcal{P}(\mathfrak{D})$ all of whose edges are edges of $\mathcal{K}_{<\mathrm{h}}$. Removing an edge from a 2-complex means removing it together with the 2-cells that contain it.

Lemma 26 (expansions). Suppose the available safe edges, followed by NS in its order, can be listed so that each edge either joins two components of the graph formed by the available vertices and the earlier edges, or lies in an available plaquette of $\mathcal{P}(\mathfrak{D})$ whose other three edges come earlier. Then $\mathrm{H}_1(\mathcal{K}_{<\mathrm{h}}; \mathrm{Z}) = 0$ for all h.

Proof. Start from the available vertices and no edges. A joining edge does not change H_1. An edge that lies in a plaquette whose other three edges come earlier is attached together with that plaquette, which has it as a free face, and this does not change the homotopy type. The remaining 2-cells only kill H_1. The 1-skeleton of $\mathcal{K}_{<\mathrm{h}}$ is an initial segment of the list. ■

Lemma 27 (integral link criterion). Let $\mathcal{K}'$ be a 2-complex with $\mathrm{H}_1(\mathcal{K}'; \mathrm{Z}) = 0$, let σ be an edge of $\mathcal{K}'$, and let $\pi_1, \ldots, \pi_\mathrm{r}$ be the 2-cells of $\mathcal{K}'$ that contain σ, with incidence numbers $\varepsilon_\mathrm{i} \in \{\pm 1\}$. If each 1-cycle $\partial(\pi_1 - \varepsilon_1\varepsilon_\mathrm{j}\pi_\mathrm{j})$, $\mathrm{j} = 2, \ldots, \mathrm{r}$, bounds an integral 2-chain in $\mathcal{K}' \setminus \sigma$, then $\mathrm{H}_1(\mathcal{K}' \setminus \sigma; \mathrm{Z}) = 0$. The hypothesis holds if the π_j are connected to π_1 through lattice 3-cubes whose two faces at σ are 2-cells of $\mathcal{K}'$ and whose other four faces are 2-cells of $\mathcal{K}' \setminus \sigma$.

Proof. Let z be a 1-cycle of $\mathcal{K}' \setminus \sigma$, and write $\mathrm{z} = \partial(\mathrm{c}' + \Sigma\, \mathrm{a}_\mathrm{i}\pi_\mathrm{i})$ in $\mathcal{K}'$, with c' supported away from σ. The coefficient of σ gives $\Sigma\, \mathrm{a}_\mathrm{i}\varepsilon_\mathrm{i} = 0$, so $\Sigma\, \mathrm{a}_\mathrm{i}\pi_\mathrm{i} = \Sigma_{\mathrm{j}\geq 2}\, \mathrm{a}_\mathrm{j}(\pi_\mathrm{j} - \varepsilon_\mathrm{j}\varepsilon_1\pi_1) = -\Sigma_{\mathrm{j}\geq 2}\, \mathrm{a}_\mathrm{j}\varepsilon_\mathrm{j}\varepsilon_1(\pi_1 - \varepsilon_1\varepsilon_\mathrm{j}\pi_\mathrm{j})$. Each $\partial(\pi_1 - \varepsilon_1\varepsilon_\mathrm{j}\pi_\mathrm{j})$ avoids σ and bounds in $\mathcal{K}' \setminus \sigma$, so z bounds there. For the second claim, a 3-cube with faces π_i, π_j at σ has a boundary 2-cycle whose coefficients at π_i and π_j are proportional to 1 and $-\varepsilon_\mathrm{i}\varepsilon_\mathrm{j}$, so $\partial(\pi_\mathrm{i} - \varepsilon_\mathrm{i}\varepsilon_\mathrm{j}\pi_\mathrm{j})$ is $\pm$ the boundary of its other four faces. These relations compose along chains of cubes. ■

Proposition 28 (existence of designs). Let $\mathfrak{D}$, $\mathfrak{B}$, $\pi_\mathfrak{B}$ be as above, with $\mathfrak{B} \pm 1 \subseteq \mathfrak{D}$, and suppose:

- (C1) the hypothesis of Lemma 26 holds;
- (C2) each NS edge e that is not a joining edge of the list in (C1), and whose endpoints both lie on NS edges below e, lies in an available plaquette Q_0 of $\mathcal{P}(\mathfrak{D})$, the *collar plaquette*, whose other three edges are edges of $\mathcal{K}_{<\mathrm{e}}$, at least one of them safe;
- (C3) for each such e, with Q_0 as in (C2), there are pairwise distinct plaquettes $\mathrm{Q}_0, \mathrm{Q}_1, \ldots, \mathrm{Q}_\mathrm{J} = \pi_\mathfrak{B}$ (the *strip*) with $\mathrm{Q}_1, \ldots, \mathrm{Q}_\mathrm{J}$ 2-cells of $\mathcal{K}_{<\mathrm{e}}$, and pairwise distinct safe edges $\sigma_\mathrm{k} \subseteq \mathrm{Q}_{\mathrm{k}-1} \cap \mathrm{Q}_\mathrm{k}$ (the *spine edges*) lying in no other Q_j, such that removing $\sigma_\mathrm{J}, \sigma_{\mathrm{J}-1}, \ldots, \sigma_1$ one at a time from $\mathcal{K}_{<\mathrm{e}}$ keeps $\mathrm{H}_1(\cdot; \mathrm{Z}) = 0$.

Then for each port network $(\mathcal{N}, \mathrm{x})$ of type t that fits in $\mathfrak{B}$ at $\pi_\mathfrak{B}$, and each (γ, e^*) of Lemma 24 with $\mathrm{v}_\ell(\gamma(\mathrm{e}^*)) = \mathrm{t}$, a funnel design exists.

Proof.

- *The edge e* is not a joining edge.* All edges of supp γ are NS edges $\leq \mathrm{e}^*$. If e^* joined two components of the earlier graph, then $\gamma - \gamma(\mathrm{e}^*)\mathrm{e}^*$ would be a chain on that graph with boundary $-\gamma(\mathrm{e}^*)\partial\mathrm{e}^* \neq 0$, which is impossible.
- By Lemma 24(c), both endpoints of e^* lie on NS edges below e^*. So (C2) gives the collar plaquette Q_0, and (C3) gives the strip. Put $\mathrm{F} := \Sigma_{\mathrm{k}=0}{}^{\mathrm{J}}\, \alpha_\mathrm{k}\, \mathrm{Q}_\mathrm{k} + \mathrm{F}'$, where $\alpha_0 = \pm\gamma(\mathrm{e}^*)$ is chosen so that $\partial(\alpha_0\mathrm{Q}_0)(\mathrm{e}^*) = \gamma(\mathrm{e}^*)$, and $\alpha_\mathrm{k} = \pm\alpha_{\mathrm{k}-1}$ is chosen so that $\alpha_{\mathrm{k}-1}\mathrm{Q}_{\mathrm{k}-1} + \alpha_\mathrm{k}\, \mathrm{Q}_\mathrm{k}$ has no boundary at σ_k.
- Let $\mathcal{K}''$ be $\mathcal{K}_{<\mathrm{e}^*}$ with $\sigma_1, \ldots, \sigma_\mathrm{J}$ removed. Then $\gamma' := \gamma - \partial(\Sigma\, \alpha_\mathrm{k}\, \mathrm{Q}_\mathrm{k})$ is a 1-cycle of $\mathcal{K}''$: it vanishes at e^* and at each σ_k, the other edges of Q_0 are edges of $\mathcal{K}_{<\mathrm{e}^*}$ by (C2), and the Q_k with $\mathrm{k} \geq 1$ are 2-cells of $\mathcal{K}_{<\mathrm{e}^*}$.

- By Lemma 26 and (C3), $H_1(\mathcal{K}''; Z) = 0$. Its H_0 is free, so $H_1(\mathcal{K}''; A) = 0$ by the universal coefficient theorem. Hence $\gamma' = \partial F'$ for an A-chain F' on $\mathcal{K}''$.
- (F1) holds by construction.
- (F2): Q_0 satisfies it by (C2), and all other plaquettes used are 2-cells of $\mathcal{K}_{<e^*}$.
- (F3): no 2-cell of $\mathcal{K}_{<e^*}$ contains e^*, so Q_0 is the only plaquette of supp F at e^*, and $F(Q_0) = \alpha_0 \neq 0$.
- (F4): $\mathcal{K}''$ has no 2-cell containing a spine edge σ_k, and Q_k contains σ_k. Therefore $F(Q_k) = \alpha_k \neq 0$, and σ_k lies in no plaquette of supp F other than Q_{k-1} and Q_k.
- (F5): $\mathcal{N} \cup \{x\} \subseteq \mathcal{P}(\mathfrak{B}) \subseteq \mathcal{P}(\mathfrak{D})$. Each edge of $\mathcal{N} \cup \{x\}$ other than the edges of $\partial\pi_{\mathfrak{B}}$ has a vertex v in the interior of $\mathfrak{B}$, because the only lattice edges spanned by the vertices of $\pi_{\mathfrak{B}}$ are those of $\partial\pi_{\mathfrak{B}}$. All sixteen unit cubes at v lie in $\mathfrak{B} \subseteq \mathfrak{D}$, so all plaquettes containing the edge lie in $\mathcal{P}(\mathfrak{D})$, and none lies in O. The plaquettes of supp F are available, so none of them contains the edge either. ■

By Lemma 25, (C1)–(C3) for a region with $\mathcal{P}(\mathfrak{D}) \subseteq B_\kappa(\tau)$ and a family of fitting port networks of all types $t \in \{0, \ldots, n-1\}$ with $\mathcal{N} \subseteq B_\rho(x)$ prove $LCL_\Theta(\rho, \kappa)$, with $N_0 := 2\kappa' + 2$. The conditions (C1)–(C3) concern only $\mathfrak{D}$, $\mathfrak{B}$ and $\pi_{\mathfrak{B}}$: one certificate serves all networks that fit, and all (ℓ, a, d, Θ).

8.6 Comb networks

All port networks we use share one cross-section and differ only in length. We describe them in a local frame (u_1, u_2, u_3, u_4) with long axis u_1. Let sq(r) be the unit square spanning u_1, u_2 with lower corner $(r, 0, 0, 0)$.

- A *sheet box* over [lo, hi] with holes at $z_1 < \ldots < z_s$ in [lo, hi] is the boundary of the 3-dimensional box $[lo, hi + 1] \times [-2, 3] \times [0, 1] \times \{0\}$, minus the squares $sq(z_i)$, which lie in its face $u_3 = 0$. It is a 2-sphere with s holes, oriented so that its boundary is $-\Sigma_i\, \partial sq(z_i)$.
- An *around-tube* from $sq(r_1)$ to $sq(r_2)$, $r_2 \geq r_1 + 3$, is the lateral surface of a column of unit 3-cubes. The column goes from $sq(r_1)$ down to $u_4 = -2$, out to $u_3 = -2$, along u_1 by 2, up to $u_4 = 2$, along u_1 at height $1 \leq u_4 \leq 2$ to r_2, back to $u_3 = 0$, and down into $sq(r_2)$. A *mirrored* tube is its image under $u_4 \mapsto -u_4$. It leaves upwards and arrives from below.
- The *tail* at $sq(r_1)$ is the lateral surface of the column of three unit cubes below $sq(r_1)$. Its bottom face, at $u_4 = -3$, is the port π.

Fix $L_c \geq 0$ and $L_g \geq 2$. The network $\mathcal{N}(L_c, L_g)$ consists of the following pieces.

- *Converter* (only if $L_c \geq 1$): a sheet box over $[-2, 4L_c + 2]$ with holes $h_j := sq(4j)$ for $0 \leq j < L_c$ and $x := sq(4L_c)$, mirrored around-tubes from h_j to h_{j+1} for $0 \leq j \leq L_c - 2$, and the tail at h_0.
- If $L_c = 0$ there is no converter: $x := sq(0)$, and the tail hangs below x.
- *Gadget*: with $g_0 := 4L_c + 8$, a sheet box over $[g_0 - 2, g_0 + 4L_g + 2]$ with holes $y_0 := sq(g_0)$ and $y_k := sq(g_0 + 4k)$ for $1 \leq k \leq L_g$, and around-tubes from y_k to y_{k+1} for $1 \leq k \leq L_g - 1$.
- The *X-tube*: an around-tube from x to y_0, mirrored if $L_c = 0$.

The enhanced plaquette is x, and $x, \pi \notin \mathcal{N}$. Each piece carries its fundamental chain with the orientation induced by the cubes, written $[\Pi]$ for the piece Π. The *hole loops* are the boundaries of the squares h_j, x, y_k and of π.

Lemma 29 (structure). For all $L_c \geq 0$ and $L_g \geq 2$:

- (S0) the fundamental chain of a tube with entry square $\square$ and exit square $\square'$ (the port π for the tail) has boundary $\varepsilon(\partial\square' - \partial\square)$ for a sign ε, and each sheet box has boundary $-\Sigma_i\, \partial \mathrm{sq}(z_i)$;
- (U) the fundamental chains have coefficients ± 1;
- (S1) the pieces are pairwise plaquette-disjoint;
- (S2) the edges lying in exactly one plaquette of $\mathcal{N}$ are the four edges of $\partial\pi$, and each edge that lies in plaquettes of two different pieces lies on a hole loop;
- (S3) each piece is connected through its *interior* edges, the edges that lie in exactly two plaquettes of $\mathcal{N}$, both in that piece;
- (S4) the hole loops are pairwise edge-disjoint;
- (S5) all vertices of $\mathcal{N} \cup \{x\}$ other than the four vertices of π lie in the open box $(-3,\, 4L_c + 4L_g + 15) \times (-3,\, 4) \times (-3,\, 2) \times (-3,\, 3)$.

Proof. These are finite checks for $L_c \le 7$ and $2 \le L_g \le 7$.† Each condition concerns a single piece, or the pieces that meet a single edge or a single hole loop. Within a comb the tubes are translates of each other by multiples of 4 along u_1, the X-tube and the tail have fixed shapes, and the sheet boxes are periodic in u_1 away from their ends. Tubes T_j and $T_{j'}$ with $|j - j'| \ge 2$ have disjoint u_1-ranges separated by a gap, and a tube meets its sheet box only near its two holes. Hence each configuration of pieces near an edge that occurs for some (L_c, L_g) already occurs, up to translation, for $L_c, L_g \le 4$. ■

Lemma 30 (comb modules). Let $A = \mathbb{Z}/\ell^n$ with $n = a + d$, let $t \in \{0, \dots, n-1\}$, and put $L_c := \ell^t$ if $t \ge 1$, $L_c := 0$ if $t = 0$, and $L_g := \ell^a$. Then $(\mathcal{N}(L_c, L_g), x)$ is a port network of type t at π in the sense of §8.4, for all $\rho \ge 4L_c + 4L_g + 18$. More precisely:

- the A-chains on $\mathcal{N}$ whose boundary lies on $\partial\pi$ have boundaries exactly $\ell^{t+d}A\cdot\partial\pi$;
- the A-chains on $\mathcal{N} \cup \{x\}$ whose boundary lies on $\partial\pi$ have boundaries exactly $\ell^t A\cdot\partial\pi$;
- ∂x bounds a $\mathbb{Z}/\ell^a$-chain on the gadget and the X-tube.

Proof.

1. *Chains are piecewise constant.* Let c be an A-chain on $\mathcal{N} \cup \{x\}$ with ∂c supported on $\partial\pi$. Let e be an interior edge of a piece Π, lying in the plaquettes ϖ, ϖ' of Π. Then $e \notin \partial\pi$ by (S2), and $e \notin \partial x$, because by (S0) each edge of ∂x lies in plaquettes of two different pieces (the X-tube, and the converter sheet box or the tail). Thus c and $[\Pi]$ both have zero boundary at e. By (U) this gives $c(\varpi)/[\Pi](\varpi) = c(\varpi')/[\Pi](\varpi')$. By (S3), $c = \beta_\Pi[\Pi]$ on Π for some $\beta_\Pi \in A$.
2. *Loop equations.* So $c = \Sigma_\Pi\, \beta_\Pi[\Pi] + \beta_x\, x$. By (S0), ∂c is a combination of hole loops, and by (S4) $\partial c = b\partial\pi$ holds if and only if the coefficients of all hole loops vanish and that of $\partial\pi$ equals b. Replace the coefficient of each tube by its product with the sign ε of (S0). Write β_c and β_g for the coefficients of the converter and gadget sheet boxes, $\beta^c{}_j$ for the converter tube from h_j to h_{j+1}, $\beta^g{}_k$ for the gadget tube from y_k to y_{k+1}, β_π for the tail and β_X for the X-tube. The equations read:
   - gadget: $-\beta_g + \beta_X = 0$ at y_0; $-\beta_g - \beta^g{}_1 = 0$ at y_1; $-\beta_g + \beta^g{}_{k-1} - \beta^g{}_k = 0$ at y_k for $2 \le k \le L_g - 1$; and $-\beta_g + \beta^g{}_{L_g-1} = 0$ at y_{L_g};
   - converter ($L_c \ge 1$): $-\beta_c - \beta^c{}_0 - \beta_\pi = 0$ at h_0; $-\beta_c + \beta^c{}_{j-1} - \beta^c{}_j = 0$ at h_j for $1 \le j \le L_c - 2$; $-\beta_c + \beta^c{}_{L_c-2} = 0$ at h_{L_c-1} (for $L_c = 1$ the equation at h_0 is $-\beta_c - \beta_\pi = 0$); and $-\beta_c - \beta_X + \beta_x = 0$ at x;
   - without converter: $-\beta_\pi - \beta_X + \beta_x = 0$ at x;
   - at $\partial\pi$: $\beta_\pi = b$.

3. The gadget equations give $\beta^{g}{}_{k} = -k\beta_{g}$ and $L_{g}\beta_{g} = 0$, so $\beta_{X} = \beta_{g}$ satisfies $\ell^{a}\,\beta_{X} = 0$, that is, $\beta_{X} \in \ell^{d}A$. Each $\beta_{X} \in \ell^{d}A$ occurs.
4. The converter equations give $\beta^{c}{}_{j} = (L_{c} - 1 - j)\beta_{c}$ and $\beta_{\pi} = -L_{c}\beta_{c}$, with $\beta_{c} = \beta_{x} - \beta_{X}$. Hence $b = -L_{c}(\beta_{x} - \beta_{X})$. Without converter, $b = \beta_{x} - \beta_{X}$.
5. For $t \geq 1$, $\beta_{x} = 0$ gives $b = L_{c}\beta_{X} \in L_{c}\ell^{d}A = \ell^{t+d}A$, and arbitrary β_{x} gives $b \in L_{cA} = \ell^{t}A$. For $t = 0$, $b = \beta_{x} - \beta_{X}$ lies in $\ell^{d}A$ when $\beta_{x} = 0$ and ranges over A otherwise. All such b occur.
6. Over Z/ℓ^{a}, $L_{g}\beta_{g} = \ell^{a}\beta_{g} = 0$ for all β_{g}. So $\beta_{g} = \beta_{X} = 1$, $\beta^{g}{}_{k} = -k$ and all other coefficients 0 solve all equations except the one at x, and this chain on the gadget and the X-tube has boundary $-\partial x$.

Finally, by (S5) the network lies in a box of $l\infty$ diameter $4L_{c} + 4L_{g} + 18$, so $\mathcal{N} \subseteq B_{\rho}(x)$. ■

The modules of Lemma 30 are also computed directly, by exact elimination, for several (ℓ, a, d) and all types.† For (ℓ, a, d) = (2, 1, 1), (2, 2, 1), (3, 1, 1), (5, 1, 1) and every type t, the first two module statements, and the third with the chain allowed anywhere on the network, are also proved in Lean 4, with Mathlib [Mat20], for the networks $\mathcal{N}(L_{c}, L_{g})$ of the supplement. The proof checks explicit chain and cochain certificates with compiled code (`native_decide`). In `lean_kernel_nets/` the same checks run in Lean's kernel alone (`decide +kernel`) on an explicit copy of the networks and certificates, which is matched to the original data by a generator and a compiled comparison.†

8.7 Capsule regions

We now define the regions of the proof and verify (C1)–(C3) for them. Fix $o \in Z^{4}$ and $L \geq 1$, and put $s_{L}(y) := \max(0, |y_{3} - o_{3}| - L)$. For a cell or an edge e, $x_{3}(e)$ is the third coordinate of its lower corner; the *axial distance* of two sets is the distance of their projections to the x_{3}-axis; and the *window of radius r* of a cell is the set of cells within $l\infty$ distance r.

- The *capsule* $\mathfrak{D}_{L}$ is the union of the closed unit 4-cubes all of whose corners y satisfy $|y_{1} - o_{1}| + |y_{2} - o_{2}| + s_{L}(y) + |y_{4} - o_{4}| \leq 19$ and $\max(|y_{1} - o_{1}|, |y_{2} - o_{2}|, s_{L}(y), |y_{4} - o_{4}|) \leq 10$. The set of points satisfying these inequalities is convex and invariant under the reflections fixing o, so with each point it contains the box spanned by that point's reflections. Hence all lattice cubes between o and a cube of $\mathfrak{D}_{L}$ lie in $\mathfrak{D}_{L}$, and $\mathfrak{D}_{L}$ is star-shaped about o. It consists of a prism of length 2L in the direction e_{3} and two end caps, each a translate of half of the region obtained for $L = 0$.
- The *hole* is $\mathfrak{B}_{L} := [o_{1} - 3, o_{1} + 4] \times [o_{2} - 3, o_{2} + 2] \times [o_{3} - L, o_{3} + L] \times [o_{4} - 3, o_{4} + 3]$, and the *port* is $\pi_{L} := c((o_{1}, o_{2}, o_{3} - L + 3, o_{4} - 3), \{1, 3\})$, on the face $x_{4} = o_{4} - 3$ of $\mathfrak{B}_{L}$.
- The corners of $\mathfrak{B}_{L} \pm 1$ satisfy the defining inequalities (the left sides are at most 14 and 5), so $\mathfrak{B}_{L} \pm 1 \subseteq \mathfrak{D}_{L}$ by convexity. The vertex o lies in the interior of $\mathfrak{B}_{L}$, so the pivotal plaquette τ lies in $\mathcal{P}(\mathfrak{D}_{L})$.
- *Placement.* The map $(u_{1}, u_{2}, u_{3}, u_{4}) \mapsto (o_{1} + u_{2}, o_{2} + u_{3}, o_{3} - L + 3 + u_{1}, o_{4} + u_{4})$ sends the port of §8.6 to π_{L}. By (S5), $\mathcal{N}(L_{c}, L_{g})$ then fits in $\mathfrak{B}_{L}$ at π_{L} whenever $4L_{c} + 4L_{g} + 15 \leq 2L - 3$.
- *Radii.* $\mathfrak{D}_{L}$ lies in the $l\infty$ ball of radius $L + 10$ about o, so $\mathcal{P}(\mathfrak{D}_{L}) \subseteq B_{L+11}(\tau)$. A fitting network lies in $\mathfrak{B}_{L}$, whose $l\infty$ diameter is 2L, so $\mathcal{N} \subseteq B_{2L+1}(x)$.

Orders and strips. The orders and the strips are given by explicit rules, implemented in the supplement. For the list of (C1), available vertices y are ordered by their $l\infty$ distance to $\mathfrak{B}_{L}$, then by the number of coordinates attaining that distance, then by $\langle G, y - o\rangle$, where $G := (0.618034, 0.414214, 0.302776, 0.236068)$. Each vertex then contributes its safe edges to earlier vertices, in an explicit order. NS is ordered lexicographically by the key of the endpoint with the larger key, then by the key of the

midpoint, where $\mathrm{key}(y) := (\|\varsigma(y-o)\|_1, \|\varsigma(y-o)\|_2^2, \langle\nu_0, y-o\rangle)$, $\nu_0 := (1, 1/\sqrt{2}, 1/\sqrt{3}, 1/\sqrt{5})$, and ς replaces the third coordinate z_3 by $\mathrm{sign}(z_3)\cdot\max(0, |z_3| - L)$.

Since the entries of ν_0 are linearly independent over $\mathbb{Q}$, $\langle\nu_0, y-o\rangle$ separates the points of $(\tfrac{1}{2}\mathbb{Z})^4$. Hence distinct NS edges have distinct keys, and the NS order is the exact lexicographic order of the keys. The floating-point sort of the computation produces exactly this order: for all 284,752 NS edges of $\mathfrak{D}_{24}$ it agrees with a 50-digit evaluation of the keys, in which no two keys are equal; locally the keys are also well separated, since the minimum of $|\langle\nu_0, \Delta y\rangle|$ over nonzero $\Delta y \in (\tfrac{1}{2}\mathbb{Z})^4$ of $l\infty$ norm at most 17 is $1.86\cdot10^{-6}$, far above double-precision rounding.† The vertex key may tie for distant vertices, but the checks of the local form of (C1) in the proof below compare keys only of vertices at $l\infty$ distance at most 2, whose difference Δy satisfies $|\langle G, \Delta y\rangle| \geq 0.012$;† so the floating-point order and every total order refining the vertex key give the same checks.

For the strips of (C3), let the *zone* be the set of plaquettes with $o_3 - L + 8 \leq x_3 \leq o_3 + L - 8$, and let the *lane plaquettes* be $\lambda_m := c((o_1, o_2, m, o_4 - 4), \{1, 3\})$. The strip begins with the collar plaquette Q_0 of (C2) and continues by a depth-first search through 2-cells of $\mathcal{K}_{<e}$, until it reaches a safe edge σ such that all faces of the lattice 3-cubes containing σ are available plaquettes of $\mathcal{P}(\mathfrak{D}_L)$ with only safe edges. From there it follows the parent pointers of a breadth-first search tree rooted at π_L. The tree is built on the available plaquettes of $\mathcal{P}(\mathfrak{D}_L)$ all of whose edges are safe, from which the plaquettes of the zone in planes containing the direction 3, other than the lane plaquettes, are excluded. The *start part* of the strip is its initial segment before its first lane plaquette in the zone, the *lane run* is the following run of lane plaquettes, and the *end path* is the rest.

Proposition 31 (capsule certificate). For all $L \geq 24$, the triple $(\mathfrak{D}_L, \mathfrak{B}_L, \pi_L)$ satisfies (C1)–(C3) of Proposition 28.

Proof. We prove by induction on $L \geq 24$ that (C1) holds in the following local form, and that each edge e of (C2) has a certificate of the following shape (H_L).

- *(C1), local form.* Each edge of the list is either an expansion (it lies in an available plaquette of $\mathcal{P}(\mathfrak{D}_L)$ whose other three edges come earlier) or has an endpoint that lies on no earlier edge of the list. Such an edge joins two components, so this implies (C1).
- *(H_L), route structure and window trichotomy.* The strip of e consists of its start part, followed, unless the whole strip lies in the start window, by a contiguous lane run $\lambda_{m_0}, \lambda_{m_0-1}, \ldots, \lambda_{o_3-L+8}$ for some m_0 and by the end path from λ_{o_3-L+8} to π_L, which is the same for all such e and has a fixed position relative to the lower end of $\mathfrak{D}_L$. The *start window* W(e) is the axial range of the start part, or of the whole strip if the strip uses no lane plaquette of the zone. It contains the slice of e, because the collar plaquette does, and it is of one of three types:
  - *zone type*: W(e) has axial extent ≤ 2;
  - *lower type*: $W(e) \subseteq \{x_3 \leq o_3 - L + 10\}$;
  - *upper type*: $W(e) \subseteq \{x_3 \geq o_3 + L - 8\}$.
- *(F) Far windows.* No NS edge lies within $l\infty$ distance 3 of the lane plaquettes in the zone or of the end path.

Base case L = 24.† (C1) holds with no failure. Apart from the first edge, each joining edge of the list has exactly one endpoint that lies on no earlier edge, so the local form holds. (C2) and (C3), with the link criterion of Lemma 27, hold at all 168,712 edges e as in (C2). The route structure holds, and all start windows are typed: 107,378, 45,544 and 39,876 of them satisfy the zone, lower and upper conditions (the types overlap). Finally (F) holds.

Formal verification. The base case is also checked in Lean 4 [dMU21] by the kernel alone.† The formalization defines $\mathfrak{D}_{24}$, $\mathfrak{B}_{24}$, π_{24}, the safe and NS edges, availability, the complexes $\mathcal{K}_{<\mathrm{h}}$, the lane, the end path and (H_{24}) as above. Its NS order is the lexicographic order of an integer form of the key, with doubled coordinates and ν_0 replaced by $\mathrm{round}(2^{50}\nu_0)$. On the NS edges of $\mathfrak{D}_{24}$ this is the order of the keys above.† Certificates produced by an untrusted program are checked by the Lean kernel: (C1) in the local form; (C2) and (C3) at all 168,712 edges, with an explicit integral 2-chain for every relation of Lemma 27 that is not a 3-cube relation; (H_{24}) for these certificates; and (F). From these checks Lean proves, for every commutative ring A:

- $\mathrm{H}_1(\mathcal{K}_{<\mathrm{h}}; \mathrm{A}) = 0$ for every h;
- each edge e as in (C2) has a certificate for (C2) and (C3) of the shape (H_{24}), and $\mathrm{H}_1(\mathcal{K}''; \mathrm{A}) = 0$, where $\mathcal{K}''$ is $\mathcal{K}_{<\mathrm{e}}$ with the spine edges removed;
- for every NS edge e* and every A-cycle γ on NS edges $\le$ e* with $\gamma(\mathrm{e}^*) \neq 0$, there is a 2-chain F on available plaquettes of $\mathcal{P}(\mathfrak{D}_{24})$ with (F1)–(F4) and $\mathrm{Q_J} = \pi_{24}$.

These proofs use only the axioms propext, Classical.choice and Quot.sound of Lean. No step uses `native_decide` or other code generated by Lean's compiler; the trusted base is Lean's kernel, including its built-in arithmetic on natural numbers. The induction step, Lemmas 24, 25 and 29, the placement of the networks and §8.3 are not formalized.

Locality. On $|\mathrm{x}_3 - \mathrm{o}_3| \le \mathrm{L} - 3$ the region, the hole, the safe and NS classes and availability are invariant under translation by e_3. A comparison of keys between two cells there is unchanged when both are translated by the same multiple of e_3, because the keys are functions that do not depend on x_3 there, plus linear forms. The two end caps of $\mathfrak{D}_{\mathrm{L}+1}$, with their keys up to an additive constant in the linear forms, are the translates by $\pm\mathrm{e}_3$ of those of $\mathfrak{D}_{\mathrm{L}}$, because ς measures the third coordinate from the end. For a slice $\hat{\mathrm{t}}$ with $|\hat{\mathrm{t}} - \mathrm{o}_3| \le \mathrm{L} - 11$, let $\iota_{\hat{\mathrm{t}}}$ map the cells of $\mathfrak{D}_{\mathrm{L}}$ below $\hat{\mathrm{t}}$ by $-\mathrm{e}_3$ and those above $\hat{\mathrm{t}}$ by $+\mathrm{e}_3$. This inserts a slab of width 2 at $\hat{\mathrm{t}}$ and identifies $\mathfrak{D}_{\mathrm{L}}$, away from the slab, with $\mathfrak{D}_{\mathrm{L}+1}$, away from the inserted slab. Consequently $\iota_{\hat{\mathrm{t}}}$ preserves the safe and NS classes, availability, and every comparison between two cells on the same side of the slab. In particular each window of radius 3 of $\mathfrak{D}_{\mathrm{L}+1}$ is a translate of a window of $\mathfrak{D}_{\mathrm{L}}$, with the same comparisons among its cells. Comparisons between cells on opposite sides of the slab may change.

(C1) for L + 1. The local form at an edge is decided by the window of radius 2 of the edge, so it transfers from $\mathfrak{D}_{\mathrm{L}}$ to $\mathfrak{D}_{\mathrm{L}+1}$. The same holds for (F).

(C2), (C3) for L + 1: choice of $\hat{t}$. Let e_0 be an edge of $\mathfrak{D}_{\mathrm{L}+1}$ as in (C2), at slice m. We need an edge $\mathrm{e}_0{}'$ of $\mathfrak{D}_{\mathrm{L}}$ as in (C2) and a slice $\hat{\mathrm{t}}$ with $|\hat{\mathrm{t}} - \mathrm{o}_3| \le \mathrm{L} - 11$ such that either $\mathrm{e}_0 = \mathrm{e}_0{}' + \mathrm{e}_3$ and $\mathrm{W}(\mathrm{e}_0{}')$ lies above $\hat{\mathrm{t}} + 3$ (case A), or $\mathrm{e}_0 = \mathrm{e}_0{}' - \mathrm{e}_3$ and $\mathrm{W}(\mathrm{e}_0{}')$ lies below $\hat{\mathrm{t}} - 3$ (case B). We call such a slice $\hat{\mathrm{t}}$ admissible. Assume $\mathrm{L} \ge 18$.

- If $\mathrm{m} \ge \mathrm{o}_3 - \mathrm{L} + 18$, put $\hat{\mathrm{t}} := \mathrm{o}_3 - \mathrm{L} + 11$ and $\mathrm{e}_0{}' := \mathrm{e}_0 - \mathrm{e}_3$. The window of radius 3 of e_0 lies above a slab inserted at $\hat{\mathrm{t}}$ and is the translate by $+\mathrm{e}_3$ of that of $\mathrm{e}_0{}'$, so $\mathrm{e}_0{}'$ is an edge of $\mathfrak{D}_{\mathrm{L}}$ as in (C2), by locality. $\mathrm{W}(\mathrm{e}_0{}')$ contains the slice $\mathrm{m} - 1 \ge \mathrm{o}_3 - \mathrm{L} + 17 > \mathrm{o}_3 - \mathrm{L} + 10$, so it is of zone type (then $\mathrm{W}(\mathrm{e}_0{}') \subseteq [\mathrm{m} - 3, \mathrm{m} + 1]$) or of upper type. In both cases it lies above $\hat{\mathrm{t}} + 3$: case A.
- Otherwise $\mathrm{m} \le \mathrm{o}_3 + \mathrm{L} - 18$, since $\mathrm{L} \ge 18$. Put $\hat{\mathrm{t}} := \mathrm{o}_3 + \mathrm{L} - 11$ and $\mathrm{e}_0{}' := \mathrm{e}_0 + \mathrm{e}_3$. Symmetrically, $\mathrm{e}_0{}'$ is an edge of $\mathfrak{D}_{\mathrm{L}}$ as in (C2), and $\mathrm{W}(\mathrm{e}_0{}')$ contains the slice $\mathrm{m} + 1 \le \mathrm{o}_3 + \mathrm{L} - 17 < \mathrm{o}_3 + \mathrm{L} - 8$, so it is of zone type ($\mathrm{W}(\mathrm{e}_0{}') \subseteq [\mathrm{m} - 1, \mathrm{m} + 3]$) or of lower type. In both cases it lies below $\hat{\mathrm{t}} - 3$: case B.

The other translate need not correspond: near the lower end $\mathrm{e}_0 - \mathrm{e}_3$, and near the upper end $\mathrm{e}_0 + \mathrm{e}_3$, is in general not an edge of $\mathfrak{D}_{\mathrm{L}}$ as in (C2) (17,316 and 17,356 edges of $\mathfrak{D}_{25}$).† In case A the strip of $\mathrm{e}_0{}'$ uses

the lane, because its window does not contain the slice of π_{L}.

Transport. In case A, translate the start part by $+\mathrm{e}_3$, extend the lane run by two lane plaquettes, and translate the end path by $-\mathrm{e}_3$. In case B, translate the whole certificate of e_0' by $-\mathrm{e}_3$. It lies below $\hat{\mathrm{t}} - 3$: its window does, and its lane run descends from the window to the lower end, followed by the end path.

We check that the result is a certificate for e_0 in $\mathfrak{D}_{\mathrm{L}+1}$. Each condition of (C2) and (C3) is local. The conditions on Q_0 and the statement that Q_{k} is a 2-cell of $\mathcal{K}_{<\mathrm{e}_0}$ are read in windows of radius 1. The link criterion at σ_{k} is read in the box of radius 2 about σ_{k}, in the complex from which the later spine edges have been removed. Near the start part, these windows and their comparisons with e_0 are images under $\iota_{\hat{\mathrm{t}}}$, together with e_0, of those of e_0'. By (F), the windows along the lane run and the end path contain no NS edge, so their content does not depend on e_0. Lane windows are translates of each other, and the windows at the two ends of the lane run are translates of those for e_0'. The start part, the lane run and the end path occupy disjoint stretches of the strip, so its plaquettes and spine edges remain pairwise distinct, and each σ_{k} lies in no strip plaquette other than $\mathrm{Q}_{\mathrm{k}-1}$ and Q_{k}.

Finally the transported certificates have the route structure and the window types of $(\mathrm{H}_{\mathrm{L}+1})$: lower windows move by $-\mathrm{e}_3$ with the lower end, upper windows by $+\mathrm{e}_3$ with the upper end, zone windows keep their extent, and the lane runs end at $\mathrm{o}_3 - (\mathrm{L}+1) + 8$. Thus the induction hypothesis holds at $\mathrm{L} = 24$, and by the induction step, which needs $\mathrm{L} \geq 18$, Proposition 31 holds for all $\mathrm{L} \geq 24$. ■

The induction step transports certificates and makes no use of the search that produced them; we have also tested it directly. Every edge of $\mathfrak{D}_{25}$ and of $\mathfrak{D}_{26}$ as in (C2) is reached by the rule above. Transporting a complete set of certificates for $\mathfrak{D}_{24}$ to $\mathfrak{D}_{25}$ and then to $\mathfrak{D}_{26}$ (with case A whenever it holds for some admissible $\hat{\mathrm{t}}$) gives certificates that pass an independent checker, and transported certificates are checked in $\mathfrak{D}_{25}$, $\mathfrak{D}_{26}$ and $\mathfrak{D}_{30}$ by a second independent implementation.† For $\mathrm{L} = 16$ the coverage fails even with any admissible $\hat{\mathrm{t}}$: 56 edges of $\mathfrak{D}_{17}$ are reached by no transport.†

8.8 Proofs of Theorem W and its corollaries

Proof of Theorem W. Let ℓ^{d} be a prime power dividing M, and take $(\mathrm{a}, \mathrm{n}, \Theta)$ as in (I1). Put $\mathrm{L} := \max(24, 2(\ell^{\mathrm{n}-1} + \ell^{\mathrm{a}}) + 9)$, $\kappa := \mathrm{L} + 11$ and $\rho := 2\mathrm{L} + 1$. For $\mathrm{t} \in \{0, \ldots, \mathrm{n}-1\}$ the network of Lemma 30 has $\mathrm{L}_{\mathrm{c}} \leq \ell^{\mathrm{n}-1}$, so it fits in $\mathfrak{B}_{\mathrm{L}}$ at π_{L} (§8.7), lies in $\mathrm{B}_{\rho}(\mathrm{x})$, and is a port network of type t. By Proposition 31, (C1)–(C3) hold for $(\mathfrak{D}_{\mathrm{L}}, \mathfrak{B}_{\mathrm{L}}, \pi_{\mathrm{L}})$, and $\mathcal{P}(\mathfrak{D}_{\mathrm{L}}) \subseteq \mathrm{B}_{\kappa}(\tau)$. By Proposition 28 and Lemma 25, $\mathrm{LCL}_{\Theta}(\rho, \kappa)$ holds, and Proposition 23 gives $\mathrm{P}_{\mathrm{p}}(\mathrm{E}_{\Theta}(\mathrm{P})) \to 0$ for $\mathrm{p} < 1/2 + \delta_{\ell,\mathrm{d}}/2$, where $\delta_{\ell,\mathrm{d}}$ is the constant of Proposition 23 for these radii. Put $\mathrm{w}_{\mathrm{M}} := \tfrac{1}{2} \min \delta_{\ell,\mathrm{d}}$, the minimum over all prime powers ℓ^{d} dividing M.

- (i) The precise form is the statement just proved. If $\Lambda_{\mathrm{P}} \not\subset \mathrm{M}\mathbb{Z}^6$, then $\Lambda_{\mathrm{P}} \not\subset \ell^{\mathrm{d}}\mathbb{Z}^6$ for some prime power ℓ^{d} dividing M. The reduction mod ℓ^{n} of an integral giant cycle whose class is not in $\ell^{\mathrm{d}}\mathbb{Z}^6$ is a giant $\mathbb{Z}/\ell^{\mathrm{n}}$-cycle whose class is not in $\ell^{\mathrm{d}}(\mathbb{Z}/\ell^{\mathrm{n}})^6$, since $\mathrm{d} \leq \mathrm{n}$.
- (ii) $\operatorname{rank} \Lambda_{\mathrm{P}} = \dim \mathrm{V}_{\mathrm{Q}}(\mathrm{P})$, because $\mathrm{H}_2(\mathrm{P}; \mathrm{Q}) = \mathrm{H}_2(\mathrm{P}; \mathbb{Z}) \otimes \mathrm{Q}$. For $\mathrm{p} > 1/2$ this dimension is 6 with probability tending to 1, by [DKS20, Theorem 1] with $\mathrm{F} = \mathrm{Q}$. A sublattice of $\mathrm{M}\mathbb{Z}^6$ of rank 6 has index divisible by M^6.
- (iii) Instance (I2) with the networks of Lemma 30 for $(\ell, \mathrm{a}, \mathrm{d}) = (2, 1, 1)$ needs only $\mathrm{L} = 24$, and the constant δ of Proposition 23 decreases as L grows. The argument above therefore gives (iii) for $\mathrm{p} < 1/2 + \mathrm{w}_{\mathrm{M}}$ when M is even.
- (iv) $\mathrm{P}\bullet$ is Bernoulli$(1-\mathrm{p})$ plaquette percolation on the cubical torus $\mathrm{X}\bullet$, and $1 - \mathrm{p} < 1/2 + \mathrm{w}_{\mathrm{M}}$. By (i) for $\mathrm{P}\bullet$, $\mathrm{V}_{\mathbb{Z}/\ell^{\mathrm{n}}}(\mathrm{P}\bullet) \subseteq \ell^{\mathrm{d}}(\mathbb{Z}/\ell^{\mathrm{n}})^6$ with probability tending to 1. By Lemma 19 and unimodularity, $\mathrm{V}_{\mathbb{Z}/\ell^{\mathrm{n}}}(\mathrm{P}) = \mathrm{V}_{\mathbb{Z}/\ell^{\mathrm{n}}}(\mathrm{P}\bullet)^{\perp} \supseteq (\ell^{\mathrm{d}}(\mathbb{Z}/\ell^{\mathrm{n}})^6)^{\perp} = \ell^{\mathrm{n}-\mathrm{d}}(\mathbb{Z}/\ell^{\mathrm{n}})^6$. ■

Proof of Corollary W1. Take $M = \ell$, so $n = 2$ and $d = 1$. The inclusion $V_{\mathbb{Z}/\ell^2}(P) \subseteq \ell(\mathbb{Z}/\ell^2)^6$ for $p < 1/2 + w_\ell$ is Theorem W(i). The reverse inclusion for $p > 1/2 - w_\ell$ is Theorem W(iv). For $p > 1/2$, since all nonzero classes mod ℓ are partnered, $V_{\mathbb{F}_\ell}(P\bullet) = 0$ with probability tending to 1 by Proposition 20, and $V_{\mathbb{F}_\ell}(P) = V_{\mathbb{F}_\ell}(P\bullet)^\perp = {\mathbb{F}_\ell}^6$ by Lemma 19. ■

Proof of Corollary W2. $P\bullet$ is Bernoulli$(1-p)$ plaquette percolation with $1 - p \in (1/2, 1/2 + w_M)$. By Theorem W(ii) for $P\bullet$, with probability tending to 1, $\Lambda_{P\bullet} \subseteq M\mathbb{Z}^6$ and rank $\Lambda_{P\bullet} = 6$. The dual complex of $P\bullet$ is P, up to the translation. Proposition 21, applied to $P\bullet$ in place of P, gives the claim for $T = \mathrm{Tors}\, H_1(P; \mathbb{Z})$. ■

Remarks.

1. *Constants.* For $M = 2$ the proof uses $L = 25$, $\kappa = 36$, $\rho = 51$ and $\kappa' = 139$, and $|B_{139}| =$ 36,095,493,313, so $\log_2(1/w_2) \approx 1.44{\cdot}10^{11}$;† in this sense Theorem W is qualitative. Theorem W(iii) and the prime power 3 use $L = 24$, so for them the region certificate is exactly the one verified in Lean.
2. *A second certificate for M = 2.* Parts (i) and (iii) for $M = 2$ also follow from a different region and different networks. The region is the union of the unit 4-cubes all of whose corners satisfy $|y - o|_1 \le 25$ and $|y - o|_\infty \le 13$. The networks are built from a tube, lattice Möbius bands and converters of degree 2 and 4. Since only $\mathbb{Z}/8$ and $\mathbb{Z}/4$ coefficients occur, the $\mathbb{F}_2$ form of the link criterion suffices there: $H_1(\mathcal{K}; \mathbb{F}_2) = 0$ implies that $H_1(\mathcal{K}; \mathbb{Z})$ is finite of odd order, so $H_1(\mathcal{K}; \mathbb{Z}/2^k) = 0$. (C1)–(C3) hold at all 77,712 relevant edges for each of the five networks ($\mathbb{Z}/8$-types 0, 1, 2 and $\mathbb{Z}/4$-types 0, 1).† (C1)–(C3) for this region are also checked in Lean 4 by the kernel alone, for the hole of each of the five networks and with the integral link criterion, on certificates produced by an untrusted program. From these checks Lean proves Lemma 26 and the 2-chain of Proposition 28 for this region, over every commutative ring and using only the standard axioms.† As for $\mathfrak{D}_{24}$, the formalization orders the NS edges by an integer form of the key; on the NS edges of this region this order agrees with the floating-point order of the certificate above, and that agreement is checked outside Lean.† This certificate gives $\kappa = 14$ and $\rho = 12$, hence $\log_2(1/w_2) \approx 9.1{\cdot}10^8$.†
3. *End-to-end tests.* The construction is also run on explicit configurations with $\mathbb{Z}/9$ coefficients, on ${T^4}_{52}$ and ${T^4}_{64}$, with networks of types 0 and 1 placed in a capsule. An exact computation confirms that ω' carries no Θ-odd cycle and that $\omega' \cup x$ does;† these tests are not used in the proofs.
4. *What Theorem W excludes.* A natural route to Theorem A at $q = 1$ is to show that every mod-2 class lifts to a giant $\mathbb{Z}/4$-cycle above 1/2, that is, $L_4(P) = {\mathbb{F}_2}^6$. Indeed, by Lemma 19 the set $\{b \in {\mathbb{F}_2}^6 : 2\tilde{b} \in V_{\mathbb{Z}/4}(P\bullet)\}$ is $L_4(P)^\perp$, and it contains $V(P\bullet)$. Theorem W(iii) shows that this route fails near 1/2, since there $L_4(P) \subseteq C^\perp$.

9. Dualities and locks at q = 1

In this section $d = 4$ and plaquettes are 2-cells. By Proposition 8, Theorem A at $q = 1$ fails exactly when, for some $p_0 < 1/2$, $P_{p_0}(V \cap C^\perp \neq 0)$ does not tend to 0. Fix a subsequence of N and a class $z_0 \in C^\perp \setminus 0$ such that $P_{p_0}(z_0 \in V)$ stays bounded away from 0 along it. For $p \in (0, 1)$ let F_p be the set of classes z with $P_p(z \in V) \to 1$ along the subsequence. By sharp thresholds, $z_0 \in F_p$ for every $p \in (p_0, 1/2)$. F_p is a subspace invariant under coordinate permutations, and for $p < 1/2$ it is isotropic: $F_p \subseteq F_{1-p}$ by monotonicity, while $F_{1-p} \subseteq {F_p}^\perp$ since $V(P) = V(P\bullet)^\perp$ and $P\bullet$ has law P_p when P has law P_{1-p}. By Lemma 6′, F_p contains u or C. We call this a *lock* onto F_p. This section records three facts about locks (§§9.1–9.3).

For $z \in H$ write

$\mathrm{Pf}(z) := z_{12}\,z_{34} + z_{13}\,z_{24} + z_{14}\,z_{23} \in F_2$.

Since $e_I \cdot e_J = 1\{I \cap J = \emptyset\}$, Pf is a quadratic form whose polar form is the intersection form: $\mathrm{Pf}(z + z') = \mathrm{Pf}(z) + \mathrm{Pf}(z') + \langle z, z'\rangle$. A subspace on which Pf vanishes identically is *totally singular*; it is isotropic. Put $W_0 := \{w \in {F_2}^4 : \Sigma\, w_i = 0\}$ and $1 := (1, 1, 1, 1)$, and let

$\mathrm{Cyc} := \Lambda^2(W_0)$, $\mathrm{Cut} := 1 \wedge {F_2}^4$, $L_0 := \mathrm{span}\{e_{12} + e_{34},\, e_{13} + e_{24},\, e_{14} + e_{23}\}$, $E := u^{\perp}$.

Cyc is spanned by triangles and 4-cycles, Cut by the four stars, and L_0 is the oblique Lagrangian of §7.5; $L_0 = C + \langle u\rangle$. The permutation-invariant subspaces of H are 0, $\langle u\rangle$, C, L_0, Cyc, Cut, $C^{\perp}$, E and H, and Pf vanishes identically exactly on 0, C, Cyc and Cut.† In particular Cyc and Cut are totally singular Lagrangians, and $\mathrm{Pf}(u) = 1$. The nonzero invariant isotropic subspaces are $\langle u\rangle$, C, L_0, Cyc and Cut, so a lock is onto one of these.

9.1 The Pfaffian counts dual pairs

For a plaquette σ let σ^* be its dual plaquette (same centre, complementary directions), and h := (½, ½, ½, ½). A *duality* is a map $D(\sigma) = g(\sigma^*)$, where g is an isometry of R^4 with diagonal linear part (so that g descends to ${T^4}_N$ and acts trivially on $H_2(X; F_2)$) that maps the dual lattice $h + Z^4$ onto Z^4. Two families are relevant:

- translation dualities, $g(y) = y + h - t$ with $t \in Z^4$; D^2 is the translation by $1 - 2t$, so D is an involution iff $1 - 2t \in NZ^4$, which forces N odd;
- point-reflection dualities, $g(y) = m + h - y$ with $m \in Z^4$, so $D(c(x, I)) = c(m - x, I^c)$; these are involutions on every ${T^4}_N$.

D maps a plaquette of type I to one of type I^c, so it has no fixed plaquette, and an involutive D splits the plaquettes into *D-pairs* $\{\sigma, D\sigma\}$. Since g maps the dual complex onto the primal one and acts trivially on $H_2(X; F_2)$, Fact D(a) and the orthogonality $V(P\,\bullet) = V(P)^{\perp}$ noted after Proposition 6 give $V(D(\omega^c)) = V(\omega)^{\perp}$ for every configuration ω.

Theorem 32 (dual pairs). Let D be an involutive duality. For every mod-2 2-cycle S,

$\mathrm{Pf}([S]) \equiv \#\{\text{D-pairs } \{\sigma, D\sigma\} \subseteq S\} \pmod 2$.

Proof. For $x \in {F_2}^{X^{(2)}}$ put $Q(x) := \Sigma_{\text{pairs}}\, x_\sigma\, x_{D\sigma}$. Its polar form is $B(x, y) := Q(x + y) - Q(x) - Q(y) = \Sigma_\sigma\, x_\sigma\, y_{D\sigma} = |x \cap D(y)| \bmod 2$.

- (a) *On cycles, B is the intersection form.* A primal plaquette fixes its two non-type coordinates to integers and a dual plaquette fixes its two non-type coordinates to half-integers, so they meet only if their types are complementary, and then only if they have the same centre, i.e. are dual to each other; they then cross transversally in one point. A primal edge and a dual plaquette (or a dual edge and a primal plaquette) never meet: together they fix five coordinates, so some coordinate is fixed to an integer by one and to a half-integer by the other. Since g commutes with $\sigma \mapsto \sigma^*$, $D^2 = g^2$ on plaquettes, so g^2 is the identity on X and g also maps primal cells to dual cells. Hence σ meets $g(\sigma')$ iff $g(\sigma') = \sigma^*$, i.e. $\sigma' = g(\sigma^*) = D(\sigma)$. So for cycles S, S′ the set $|S| \cap g|S'|$ consists of transverse crossings, one for each $\sigma \in S \cap D(S')$, and $B(S, S') = \langle [S], [gS']\rangle = \langle [S], [S']\rangle$.
- (b) *Q vanishes on boundaries.* The faces of a 3-cube have types inside the cube's three directions, and D maps each type to its complement, which is not among them; so $Q(\partial c) = 0$. For a sum of cube boundaries, $Q(\Sigma\, \partial c_i) = \Sigma\, Q(\partial c_i) + \Sigma_{i<j}\, B(\partial c_i, \partial c_j) = 0$ by (a).
- (c) Hence $Q(S + \partial b) = Q(S) + Q(\partial b) + B(S, \partial b) = Q(S)$, so Q induces $\hat{q}\colon H \to F_2$ with $\hat{q}(z + z') = \hat{q}(z) + \hat{q}(z') + \langle z, z'\rangle$.

- (d) $\hat{q} - Pf$ is therefore linear. A coordinate torus of type I contains no D-pair, since D changes the type, so $\hat{q}(e_I) = 0 = Pf(e_I)$. Hence $\hat{q} = Pf$. ■

Corollary 33. Let D be an involutive duality. For a set Y of plaquettes put $V(Y) := V(X^{(1)} \cup Y)$.

- (a) If $Y \cap D(Y) = \emptyset$, then V(Y) is totally singular.
- (b) (*Frozen structures.*) Let ω be random with $\omega \supseteq Y$ and $\omega \cap D(Y) = \emptyset$ almost surely, where $Y \cap D(Y) = \emptyset$. Then $V(Y) \subseteq V(\omega) \subseteq V(Y)^{\perp}$. If Y is invariant under coordinate permutations, V(Y) is one of 0, C, Cyc, Cut; in particular a frozen permutation-invariant structure never forces $u \in V$.
- (c) (*Self-dual product families.*) Let (ν_s) be a monotone family of product measures, $P_s(\sigma \text{ open}) = \pi_\sigma(s)$, with $\pi_{D\sigma}(1 - s) = 1 - \pi_\sigma(s)$ (so that $D(\omega_s^c)$ has law ν_{1-s}). For $s \le 1/2$ the majority set $\{\sigma : \pi_\sigma(s) > 1/2\}$ is D-free, and every cycle of ω_s with Pf = 1 contains a doubly open pair $\{\sigma, D\sigma\}$. Since $\pi_\sigma(s) + \pi_{D\sigma}(s) \le 1$, it contains an open plaquette σ with $\pi_\sigma(s) \le 1/2$.

Proof. (a) is Theorem 32. (b) $V(Y) \subseteq V(\omega)$ since $Y \subseteq \omega$, and $V(\omega)^{\perp} = V(D(\omega^c)) \supseteq V(D(D(Y))) = V(Y)$ since $D(Y) \subseteq \omega^c$. The list of permutation-invariant totally singular subspaces gives the rest, as $u \notin Cyc \cup Cut$. (c) For $s \le 1/2$, $\pi_\sigma(s) + \pi_{D\sigma}(s) \le \pi_\sigma(1 - s) + \pi_{D\sigma}(s) = 1$ by monotonicity and the hypothesis for $D\sigma$, so the majority set is D-free; the rest is Theorem 32. ■

So under an involutive duality a lock onto a subspace on which Pf does not vanish, such as $\langle u \rangle$ or L_0, cannot be produced by plaquettes that are open with probability one: every cycle whose class has Pf = 1 contains a doubly open pair. For even N the translation duality is not an involution, and the analogous count of ordered pairs is even on every cycle, so it carries no information. Point-reflection dualities are involutions on every T^4_N, so Theorem 32 still applies on even tori.

9.2 A symmetric, self-dual product family that locks

The following family shows which properties of Bernoulli percolation a proof of Theorem A at q = 1 must use.

Let $8 \mid N$. Call $\Sigma_j x_j \in Z/N$ the *exact level* of a cell $c = c(x, I)$, and let $\lambda(c)$ be its residue mod 8. Split the plaquettes into the classes $\mathcal{A} := \{\lambda \in \{0, 1\}\}$, $\mathcal{B} := \{\lambda \in \{4, 5\}\}$ and $\mathcal{F} := \{\lambda \in \{2, 3, 6, 7\}\}$. Fix $\varepsilon \in (0, 1/5)$, $\delta \in (0, 1/2)$ and a continuous, piecewise C^1 map φ: $[0, 1] \to [0, 1]$ with $\inf \varphi' > 0$, $\varphi(0) = 0$, $\varphi(1) = 1$ and $1 - \varepsilon \le \varphi \le 1 - \varepsilon/2$ on $[\delta, 1 - \delta]$. Let ν_s be the product measure with

$P(\sigma \text{ open}) = \varphi(s)$ on $\mathcal{A}$, $1 - \varphi(1 - s)$ on $\mathcal{B}$, s on $\mathcal{F}$,

and let $\omega_s := \{\sigma : U_\sigma \le P_s(\sigma \text{ open})\}$ for independent uniform labels U_σ.

Theorem 34 (a robust lock onto Cyc). For every $N \equiv 0 \bmod 8$,

$P(V(\omega_s) = Cyc \text{ for all } s \in [\delta, 1 - \delta]) \ge 1 - 48\, N^3\, (5\varepsilon)^{2N} / (5(1 - 5\varepsilon))$.

The family (ν_s) is a monotone family of product measures, invariant under coordinate permutations and under translations by v with $\Sigma v_i \equiv 0 \bmod 8$, it has finite energy, and it is exactly self-dual: for $\tau(c(x, I)) := c(x + 1_I, I^c)$ and $t := (6, 0, 0, 0)$, the configuration $\omega\widehat{\ }\sharp := \{\tau\sigma - t : \sigma \notin \omega_s\}$ has law ν_{1-s}, and $V(\omega\widehat{\ }\sharp) = V(\omega_s)^{\perp}$. It is also self-dual under the involutive point-reflection duality $D(c(x, I)) = c(m - x, I^c)$ with $\Sigma m_i \equiv 5 \bmod 8$.

Proof.

- *Self-duality.* $\lambda(\tau\sigma - t) = \lambda(\sigma) + |I| - 6 = \lambda(\sigma) - 4$, so $\sigma \mapsto \tau\sigma - t$ exchanges $\mathcal{A}$ and $\mathcal{B}$ and preserves $\mathcal{F}$, and the densities match: $P(\tau\sigma - t \in \omega\widehat{\ }\sharp) = 1 - \varphi(s) = 1 - \varphi(1 - (1 - s))$ for $\sigma \in \mathcal{A}$, $1 - (1 - \varphi(1 - s)) = \varphi(1 - s)$ for $\sigma \in \mathcal{B}$, and $1 - s$ for $\sigma \in \mathcal{F}$. The set $\tau(\omega^c)$ is the dual

complex translated by h, so $V(\omega^\frown\sharp) = V(\omega)^\perp$ by Fact D. The point reflection maps λ to $5 - \lambda$, so it too exchanges $\mathcal{A}$ and $\mathcal{B}$ and preserves $\mathcal{F}$, and the same densities match.

- *Staircases.* For an integer ℓ let S_ℓ be the union of the 3-cells $y + [0, 1)^{[4]\setminus\{k\}}$, $k \in [4]$, with $\Sigma_j\, y_j \equiv \ell \bmod N$, and their faces. A plaquette $c(z, J)$ with $J^c = \{a, b\}$ lies in the four 3-cells $z + [0,1]^{J\cup\{a\}}$, $z + [0,1]^{J\cup\{b\}}$, $(z - e_a) + [0,1]^{J\cup\{a\}}$, $(z - e_b) + [0,1]^{J\cup\{b\}}$, of base levels Σz, Σz, $\Sigma z - 1$, $\Sigma z - 1$. So the plaquettes of S_ℓ are those of exact level ℓ or $\ell + 1$, each lies in exactly two 3-cells of S_ℓ, and for $\ell \equiv 0 \bmod 8$ they all lie in $\mathcal{A}$.
- *S_ℓ is a 3-torus carrying Cyc.* In R^4 let $R_\ell := \cup\{x + [0,1]^4 : \Sigma x \le \ell - 1\}$. A point w lies in R_ℓ iff $\Phi(w) := \Sigma_i \lceil w_i \rceil \le \ell + 3$, and Φ is nondecreasing along 1. So ∂R_ℓ is the graph of a continuous function over $\{\Sigma = 0\}$, and it is the lift of S_ℓ. Its stabilizer in NZ^4 is $\Gamma := NZ^4 \cap \{\Sigma = 0\} \cong Z^3$, so $S_\ell \cong R^3/\Gamma$. Γ is saturated in NZ^4, so $H_1(S_\ell; F_2) \to H_1(X; F_2)$ is injective with image W_0; since the inclusion is homotopic to a linear map of tori, the map on H_2 is Λ^2 of the map on H_1, hence injective with image $\Lambda^2(W_0) = \mathrm{Cyc}$.
- *Robustness.* Fix $\ell \equiv 0 \bmod 8$. For $s \in [\delta, 1 - \delta]$ the plaquettes of S_ℓ are open independently with probability $\ge 1 - \varepsilon$. Let G_ℓ be the 6-regular graph whose vertices are the $4N^3$ 3-cells of S_ℓ and whose edges are its plaquettes. By Poincaré–Lefschetz duality on the closed 3-manifold S_ℓ, if the closed plaquettes contain no cycle of G_ℓ that is essential in $H_1(S_\ell; F_2)$, the open plaquettes carry $H_2(S_\ell; F_2)$, hence $\mathrm{Cyc} \subseteq V$. An essential cycle lifts to a path between translates by some $v \in \Gamma \setminus \{0\}$, and $|v|_\infty \ge N$; adjacent 3-cells have centres differing by $\frac{1}{2}(e_a - e_b)$, so the cycle has length $\ge 2N$. Counting non-backtracking closed walks (at most $6\cdot 5^{k-1}$ of length k from a vertex) gives $P(\mathrm{Cyc} \not\subset V(\omega_\delta)) \le 24\, N^3\, (5\varepsilon)^{2N} / (5(1 - 5\varepsilon))$.
- *The lock.* The same bound applied to $\omega^\frown\sharp_{1-\delta}$, which has law ν_δ, gives $\mathrm{Cyc} \subseteq V(\omega^\frown\sharp_{1-\delta}) = V(\omega_{1-\delta})^\perp$, that is $V(\omega_{1-\delta}) \subseteq \mathrm{Cyc}^\perp = \mathrm{Cyc}$, off an event of the same probability. By monotonicity, $\mathrm{Cyc} \subseteq V(\omega_\delta) \subseteq V(\omega_s) \subseteq V(\omega_{1-\delta}) \subseteq \mathrm{Cyc}$ for $s \in [\delta, 1 - \delta]$. ■

In particular $\nu_{1/2}(V = \mathrm{Cyc}) \to 1$. Theorem A fails for this family, since $V = \mathrm{Cyc} \ne 0$ with high probability for every $s \in [\delta, 1/2)$, and the hypothesis of Proposition 7 fails too, since $e_{12} \notin \mathrm{Cyc}$. So does Lemma 4: for $n = e_{12}^* + e_{14}^* + e_{23}^* + e_{34}^*$, PD(n) is a 4-cycle, so it lies in Cyc, while $PD(s_1(n)) = e_{12} + e_{14} \notin \mathrm{Cyc}$; since D_n means $PD(n) \in V(P\,\bullet) = \mathrm{Cyc}$ on the lock, at $s = 1/2$ we get $R(n) \to 1$ and $R(s_1(n)) \to 0$. The point-reflection duality maps $\mathcal{A}$ onto $\mathcal{B}$, so $\mathcal{A}$, which is the majority set $\{\pi_\sigma(s) > 1/2\}$ for $s \in [\delta, 1/2]$, is D-free. By Corollary 33(a) and the list above, $V(\mathcal{A})$ is a permutation-invariant totally singular subspace. It contains the image Cyc of the staircases, so $V(\mathcal{A}) = \mathrm{Cyc}$: the lock is onto the subspace carried by the D-free majority set.

Corollary 35. The conclusion of Theorem A at $q = 1$ does not follow from the following properties of a family of measures on configurations: product structure (hence the Harris–FKG, van den Berg–Kesten and Reimer inequalities [vdBK85, Rei00], the OSSS inequality [OSSS05] and Russo's formula), finite energy, exchangeability of the labels within each orbit of the symmetry group, invariance under coordinate permutations, exact self-duality $V(\omega^\frown\sharp) = V(\omega)^\perp$, also under an involutive duality, and sharp thresholds with polynomial rates for invariant increasing events [BKKKL92]. A proof must therefore use a property that this family lacks, such as homogeneity (invariance under unit translations) or the sign changes of the hyperoctahedral group B_4 other than $x \mapsto -x$. The family is invariant under $x \mapsto m - x$ for $\Sigma\, m_i \equiv 3 \bmod 8$, which acts trivially on H_2.

Proof. The family of Theorem 34 has all the listed properties. For the sharp thresholds, let $\boldsymbol{\mathcal{E}}$ be an invariant increasing event with $\nu_s(\boldsymbol{\mathcal{E}}) = t$, and let $J(\sigma) := \nu_s(\boldsymbol{\mathcal{E}} \mid \sigma \text{ open}) - \nu_s(\boldsymbol{\mathcal{E}} \mid \sigma \text{ closed})$. J is constant on each symmetry orbit, and each orbit has at least $6N^4/8$ plaquettes. So the Bourgain–Kahn–Kalai–Katznelson–Linial inequality on the uniform labels gives $\Sigma_\sigma\, J(\sigma) \ge c\, t(1 - t) \log N$. Every $P_s(\sigma \text{ open})$ has derivative $\ge \min(1, \inf \varphi') > 0$, so Russo's formula gives $d\nu_s(\boldsymbol{\mathcal{E}})/ds \ge c'\, t(1 -$

t) log N. This yields windows of width O(1/log N) and polynomial rates. ■

The staircases are orientable, and the proof of Theorem 34 applies verbatim with coefficients in Q, Z/4 and $F_{\ell'}$ (ℓ' an odd prime). The robustness bound excludes every closed dual cycle whose lift is not closed, and Poincaré–Lefschetz duality on the orientable S_ℓ holds with any coefficients. The upper bound uses $V_A(P\,\bullet) \perp V_A(P)$ for the coefficient ring A (Lemma 19, or [DKS20, Lemma 9] over Q), and $\Lambda^2(\ker \Sigma)$ is a Lagrangian direct summand of $H_2(T^4; Z)$. So the family locks onto $\Lambda^2(\ker \Sigma)$ with these coefficients as well. This does not contradict [DKS20] or Proposition 10. [DKS20] uses the irreducibility of $H_2(T^4; F)$ under the point group, and Proposition 10 uses partners in B_4. Both need sign changes other than $x \mapsto -x$, and this family has only the central one, which acts trivially on H_2. In a B_4-invariant self-dual family with sharp thresholds, the argument of Proposition 10 excludes, below the self-dual point, every lock in which $V_{Z/4}$ contains an element of order 4, in particular every lock carried by orientable cycles. So a B_4-invariant counterexample would have to lock onto non-orientable cycles. Whether one exists is open.

9.3 An anisotropic window

Call a plaquette *vertical* if its type contains 1 and *horizontal* otherwise. Consider the product measure in which vertical plaquettes are open with probability $1 - a$ and horizontal ones with probability b, $0 < a, b \le 1$. Every duality maps type I to I^c and so exchanges vertical and horizontal plaquettes. Hence $D(\omega^c)$ has the law of the model with a and b exchanged, and $V(D(\omega^c)) = V(\omega)^\perp$. On the self-dual line $a = b$, $P(V = 0) = P(V = H) \le 1/2$, as in step 1 of the proof of Theorem A. The isotropic model is $a = 1 - p$, $b = p$.

Lemma 36 (slices). Let S be a mod-2 2-cycle, $i \in [4]$ and $r \in Z/N$. Send each plaquette $c(x, \{i, j\})$ of S with $x_i = r$ to the edge of T^3_N (coordinate i deleted) from $\hat{x}$ in direction j. The result $\sigma_{i,r}(S)$ is a mod-2 1-cycle, and for $j \ne i$ its winding parity in direction j is the coefficient z_{ij} of [S].

Proof. The edge of X from y to $y + e_i$ ($y_i = r$) lies in the six plaquettes $c(y, \{i, j\})$ and $c(y - e_j, \{i, j\})$, $j \ne i$, which are sent to the six edges of T^3_N at $\hat{y}$; since $\partial S = 0$, every vertex of $\sigma_{i,r}(S)$ has even degree. Its edges in direction j crossing $\{x_j = 1/2\}$ are the plaquettes of S of type $\{i, j\}$ with $x_i = r$ and $x_j = 0$; their number is the value on S of a cocycle Poincaré dual to the coordinate torus of type $\{i, j\}^c$, so its parity is z_{ij}. ■

Theorem 37. In the anisotropic model put $\kappa := 4b(2 - a)/a$. If $\kappa < 1$, then $P(V \ne 0) \to 0$ exponentially fast in N. More precisely, for $\theta \in (0, 1)$ with $4bA_\theta < 1$, where $A_\theta := (1 + (1 - a)^\theta)/(1 - (1 - a)^\theta)$,

$$P(V \ne 0) \le 3N^3 \sum_{n \ge N} \kappa^n + 3N^2(1 - a)^N + 3N^3\,(1 - a)^{(1-\theta)N} \cdot 4bA_\theta/(1 - 4bA_\theta),$$

and such a θ exists exactly when $\kappa < 1$, because A_θ decreases to $(2 - a)/a$ as $\theta \uparrow 1$ (strictly when $a < 1$).

Proof. Let $S \subseteq \omega$ be a cycle.

- *Horizontal coordinates.* Suppose $z_{34}([S]) = 1$. By Lemma 36 ($i = 3$, $r = 0$), $\sigma_{3,0}(S)$ is a 1-cycle of the torus with coordinates x_1, x_2, x_4 with odd winding in direction 4, so it contains a simple cycle γ with odd winding in direction 4, hence with at least N edges in direction 4. Its edges are distinct open plaquettes: direction $4 \leftrightarrow$ type $\{3, 4\}$ and direction $2 \leftrightarrow$ type $\{2, 3\}$ (probability b), direction $1 \leftrightarrow$ type $\{1, 3\}$ (probability $1 - a$). Write γ as a closed walk $(s_1, R_1, \dots, s_n, R_n)$ starting at a non-time step, where each s_m is one of the 4 steps $\pm e_2$, $\pm e_4$ and each R_m is a maximal run of time steps; a simple cycle does not reverse, so R_m is monotone, with a length $t_m \ge 0$ and one of 2 directions if $t_m \ge 1$; and $n \ge N$. Summing $\Pi_m\, b(1 - a)^{t_m}$ over the walks from a fixed vertex gives at most $\sum_{n \ge N} (4b(1 + 2\sum_{t \ge 1}(1 - a)^t))^n = \sum_{n \ge N} \kappa^n$, and there are N^3 starting vertices. The coefficients z_{23}, z_{24} are handled with the slice $i = 2$, $r = 0$ in the same way. This gives the first term.

- *Vertical coordinates.* Suppose $z_{1j}([S]) = 1$ for some $j \in \{2, 3, 4\}$. By Lemma 36 with $i = j$ and $r = 0$, $\sigma_{j,0}(S)$ is a 1-cycle of the torus with coordinates x_1, x_k, x_l, where $\{k, l\} = [4] \setminus \{1, j\}$, with odd winding in direction 1. It contains a simple cycle γ with odd winding in direction 1. Every hyperplane $\{x_1 = r' + 1/2\}$, $r' \in Z/N$, is crossed by an odd number of its time edges, so γ has at least N of them, and each is a vertical plaquette (type $\{1, j\}$), open with probability $1 - a$. Its other edges are horizontal plaquettes (types $\{j, k\}$, $\{j, l\}$), open with probability b. If γ has no spatial edge it is a straight time circle: N^2 choices, probability $(1 - a)^N$. Otherwise write it as $(s_1, R_1, \dots, s_n, R_n)$ with $n \geq 1$ and monotone time runs of lengths t_m, $\Sigma\, t_m \geq N$. Using $(1 - a)^{\Sigma t} \leq (1 - a)^{(1-\theta)N}(1 - a)^{\theta\Sigma t}$ and $\Sigma_{t\geq 0}(1 - a)^{\theta t}\, 2^{[t\geq 1]} = A_\theta$, the expected number of open such γ is at most $N^3(1 - a)^{(1-\theta)N}\, \Sigma_{n\geq 1}(4bA_\theta)^n$. Sum over j. ■

Corollary 38 (anisotropic window). In the anisotropic model, $V = 0$ w.h.p. if $4b(2 - a) < a$, and $V = H$ w.h.p. if $4a(2 - b) < b$. So for each a every threshold of V, in particular a sharp transition curve if there is one, lies in the closed wedge $a/(4(2 - a)) \leq b \leq 8a/(1 + 4a)$. The self-dual line $b = a$ lies in its interior for every $a \in (0, 1]$; for small a and b the wedge is $a/8 \lesssim b \lesssim 8a$.

Proof. The first statement is Theorem 37. For the second, apply Theorem 37 to $D(\omega^c)$, whose parameters are (b, a), and use $V(\omega) = V(D(\omega^c))^\perp$. ■

For the isotropic model Theorem 37 gives only $p < (\sqrt{41} - 5)/8 \approx 0.175$, weaker than Proposition 6. Corollary 38 is useful in the anisotropic regime, where monotone comparison as in Proposition 6 gives nothing: it needs $\max(1 - a, b) < \sqrt{2} - 1$. The bounds are asymptotic. At $a = 0.2$, for instance, the bound of Theorem 37 drops below 1 only from $N = 93$, 134 and 384 for $b = 0.005$, 0.01 and 0.02 (and from $N = 38$ as $b \to 0$). The method uses one slice, so it cannot reach the self-dual line: Theorem 37 itself stops at $b = a/(4(2 - a)) \approx 0.028$ for $a = 0.2$, and at $N = 20$ the slices already contain cycles winding in a spatial direction at $b \approx 0.07$ (at $N = 10$ from $b \approx 0.09$), against the line $b = 0.2$.† Between that point and the line, the statement is again an F_2 sharpness problem for surfaces in four dimensions, now without a symmetry exchanging vertical and horizontal classes; its two-dimensional analogue, the critical line $p_h + p_v = 1$ of inhomogeneous bond percolation, rests on planar duality and crossing estimates [Kes82, §3.4], [Gri99, §11.9].

Appendix: verification table

Every statement marked † rests on the computation listed below. The supplementary material contains each program and each log, together with a README that gives the exact command for every row, and SHA-256 hashes of all files. The C and C++ programs are compiled with gcc or g++ and OpenMP (`capgen.cpp` also needs GMP); the Python programs need numpy, scipy, sympy, mpmath and, for one row, cvxpy; the Lean projects need Lean 4 (v4.35.0-rc3) and Mathlib (commit c55e6e78). Exact computations are deterministic. The multithreaded Monte Carlo rows use one random stream per thread, so a rerun reproduces their statistics but not their digits. Rows marked *cross-check* re-verify a statement that is proved by hand, or verify a computation independently; no proof depends on them.

Statement	**Programs**	**Logs**
§7, after Prop. 7: $\mu(e_I \in V) = 0.424$, 0.436, 0.452	`bern4d_h.c`, `vhist_summary.py`	`vhist_N*_p0.500.txt`, `vhist_summary.txt`
Prop. 8: the partnerless classes are the 15 nonzero classes of $C^\perp$	`z4partners.py`	`z4partners.txt`
After Prop. 9: S_5 fixes only u_Q, the only partnerless class of $A_4{}^*$	`perm_fixed.py`	`perm_fixed.txt`
Prop. 10(i): all 4032 elements of order 4 are partnered	`z4partners.py`	`z4partners.txt`

Statement	**Programs**	**Logs**
After Prop. 11: spectral radii $\geq \sqrt{2}$; every class from one or two hardness-free planes	`balanced.py`	`balanced.log`
Lemma 12′: the determinant formulas (cross-check)	`lemma12_dets.py`	`lemma12_dets.txt`
After Lemma 12′: $c_0 \approx 0.62$	`good_involutions.py`	`good_involutions.log`
§7: area $\approx 1.06 \cdot N^{3.35}$; 36 of 40 exact; flatter below 1/2	`minrep.c`, `run_probe*.sh`, `ilp_check.py`, `fit_area.py`, `crumple_analyze.py`	`fit_area.txt`, `crumple_tables_regen.txt`, `crumple_data/`
§7: fingered sheets, adversarial search at the largest scales	`fingers.py`, `fingers_opt.py`, `fingers_opt2.py`	`slabgm_data/fingers_*.txt`, `fingers_run1_regen.txt`
§7: unique crossing at $p \geq 0.54$ and at $p = 1/2$	`pairs.c`, `run_pairs.sh`	`slabgm_data/EP*.txt`, `pairs_p0.5_c*.txt`
Prop. 13: the sum −9 711 396; brute-force agreement	`prp.c`, `brute2.c`, `run_prp_checks.sh`	`prp_checks.txt`
§7: per-H inequality in $d = 2$	`prp.c`	`prp_2d.txt`
§7: twist-gap recurrences of orders 16, 69 and 17	`tmrel.c`, `bm3.py`	`bm3_w3.txt`, `bm3_w4.txt`, `w3_*.txt`
§7: Conjecture E on small tori	`cell3d.c`, `bond2d.c`, `bond2d_x.c`	`cell3d.txt`, `bond2d.txt`, `bond2d_4x5_record.txt`
§7: $R(s) - R(s+e) > 0$ on $Z_2 \times Z_2 \times Z_N$; ratio 0.2396	`tm3.c`, `bf2d.py`	`long224_*.out`, `tm3_vs_brute.txt`
§7: Monte Carlo in $d = 4$	`bern4d.c`, `bern4d_h.c`, `vhist_summary.py`	`bern4d_scan.txt`, `vhist_summary.txt`
§7: class-wise and cell-wise measures; $q = 2$ controls	`tm4.c`, `prof_scan2.sh`, `q2time3.py`, `q2cells.py`	`scan2_N*.txt`, `profs2_*.txt`, `tm4_percell_q1.txt`, `q2_time_inhom.txt`, `q2_cells.txt`
Prop. 14, $d = 3$: −243 (486/729), by three programs	`orbd.c`, `verify3d.py`, `v3d.py`	`orbd_prop14_d3.txt`, `verify3d.txt`, `v3d.txt`
Prop. 14, $d = 4$: −7776	`t_4d_embed.py`, `orbp.c`	`t_4d_embed.txt`
Prop. 14, $d = 2$: $-6.98 \cdot 10^{-9}$, two exact programs and one floating-point program	`blocknu.c`, `blockdisp.py`, `exact_mstar.py`; `exact_reldp.py`, `reldp.py`; `reldp_odd.c`, `t_rdp_W6.py`, `rdp.py`	`exact_mstar_Z6xZ39.txt`, `blockdisp_Z6xZ39.txt`, `exact_reldp_Z6xZ39.txt`, `t_rdp_W6.txt`
§7, mechanism: $\alpha = 0.1286$, first negative at $m = 13$; $W \leq 5$ achiral	`blocknu.c`, `exact_mstar.py`	`alpha_Z6xZ39.txt`, `exact_mstar_Z6xZ39.txt`, `chirality_small_W.txt`
§7, mechanism: $q = 2$ on $Z_7 \times Z_{232}$, $q = 4$ on $Z_7 \times Z_{378}$	`blocknu_q.c`, `t_q2exact.py`	`t_q2exact.txt`
§7: table of translation-orbit sums	`orbtr_q.c`, `orbtr.c`	`orbtr_q_small.txt`, `q44.out`, `q223_ie5.out`, `t223_ie6.out`, `q232_ie5.out`, `q232_ie6.out`
§7: $12q^3 - 4q$ on $Z_6 \times Z_2$; threshold → 1; thickening to $d = 3, 4$	`verify2d_orbit.py`, `orb2d_rk.c`, `orb2d_rkq.c`, `altfam.py`, `altfam_formula.py`, `verify_embed.py`	`verify2d_orbit.txt`, `orb2d_rk_6x2.txt`, `orb2d_rkq_6x2.txt`, `altfam.txt`, `altfam_formula_M3.txt`, `verify_embed.txt`
§7: 714 136 R-orbits, 20 negative rank parts, all achiral; rank parts for e_{13}*	`orbtr_R.c`, `verify_orbtr_witness.py`	`orbtrR_*_rank.txt`, `R232_ie6_wit.out`, `verify_orbtr_witness_R232_ie6.txt`
Prop. 17′(iii): for every block with at most n_W rows, $h_j = 0$ or $h_j \in K_W$; the extreme ratios	`spec_exh.c`	`spec_W*_n*.log`, `chain_spec.log`
§7.6: 17 two-row block types on Z_{11} whose antiperiodic part is not nonincreasing, all at $d = 5$	`uminus_exh.c`, `ufour_exh.c`, `indep_block.py`	`W10_n2.log`, `W11_n2.log`, `W12_n2.log`, `ufour_W*_n*.log`, `indep_block_W11.txt`

Statement	Programs	Logs
§7.6: $\Delta^{E}{}_{7} < 0$ on $Z_9 \times Z_2$ at $p = 1/2$ for E = all gaps thick, closing gap thick, and the total (two exact programs; interval arithmetic); Σ_k $\Delta^{E}{}_{k} > 0$ in the computed cases; for these events and a thin closing gap, $\Delta^{E}{}_{k} > 0$ on (0, 1) on every other torus with $W \geq 3$, $N \geq 2$, $WN \leq 18$ (exact counts; the sign on (0, 1) certified in interval arithmetic) and at every computed value with $W \leq 8$, $N \geq 2$ (floating point for $W \leq 7$); the failures on $Z_{10} \times Z_2$ and at $N = 1$	`bf2_helical.c`, `bf_helical.c`, `tgm3.c` (with `-DEXACT`), `witness_check.py`, `tgx_sectors.py`, `sign_cert.py`, `poly_positive.py`, `check_twist_positive.py`, `tgm.c`, `tgm2.c`, `make_jobs.py`, `analyse_tgm.py`, `make_tables.py`	`bf2_W*_N*.txt`, `bf_W*_N1.txt`, `W9x/`, `witness_W9_N2.log`, `sign_cert_WN18.log`, `sign_cert_N1_W10.log`, `poly_positive_small.log`, `poly_test.log`, `poly_W10_N2.log`, `check_twist_positive.log`, `tgx_W8_N1to4_p05.log`, `tableA.md`, `full_W*.tsv`
§7.6: $\Delta^{E}{}_{7} < 0$ at $p = 1/2$ on $Z_{11} \times Z_3$, $Z_{12} \times Z_3$, $Z_{13} \times Z_4$ and $Z_{14} \times Z_5$ for the three events, and the other sectors there (exact for the whole space on $Z_{11} \times Z_3$ and $Z_{12} \times Z_3$, otherwise floating point with a rounding bound; Monte Carlo agrees); the smallest failing widths for $N = 1, \ldots, 5$; $Z_{13} \times Z_5$ positive by Monte Carlo; the square tori with $3 \leq N \leq 11$	`tlt2.c`, `tlt4.c`, `tls.c`, `tls2.c`, `tls3.c`, `tls2x.c`, `tlt_thin.c`, `combine.py`, `exact_tls.py`, `merge_chunks.py`, `w14n5_analyse2.py`, `map_table.py`, `square_table.py`, `mc_check.c`, `mc_table2.py`	`p05_*.json`, `exact_*.txt`, `tables_final.md`, `mc_table.md`, `validation.log`, `w14n5_analyse2.log`, `mc_*.log`
§7.6, scaling limit (under the stated assumption): $\Delta_7 < 0$ on the grid $0.03 \leq \rho \leq 0.97$, Δ_{11} and $\Delta_{13} < 0$ at $\rho = 0.02$, $\Delta_5 < 0$ on the grid $2.82 \leq \rho \leq 12$, the sign changes; the negative tails at $\rho = 0.98, \ldots, 4$ and on the square from $k = 987$; the limit at $\rho = 1/3$ against the lattice	`cont_limit2.py`, `cont_bisect.py`, `cont_bisect2.py`, `cont_dense_scan.sh`, `dense_summary.py`, `cont_hik.py`, `cont_real_hik.py`, `cont_real_combine.py`, `cont_check.py`	`cont2/*.log`, `hik/*.log`, `cont_bisect_k*.log`, `cont_check.log`, `continuum.md`
§7: the orbit of 0x3f719edb6	`orbtr_r.c`, `vgfail.py`, `rorbit_rank.py`, `torbit_rank.py`, `verify_orbtr_witness.py`	`g232_ie5_wit.out`, `vgfail.txt`, `rorbit_rank.txt`, `orbtrT_232_ie5_rank.txt`, `verify_orbtr_witness_3f719edb6.txt`
§7: reflection averages +69 984 and +2 239 488	`v3dref.c`	`v3dref.txt`
§7: validation of the exact program in $d = 4$	`orbhom2.c`, `orbhom3.c`, `orb4mc.c`, `orb4ex.c`, `validate_orbhom.sh`	`orbhom_validation.txt`, `orb4ex_2x2x2x2_s7.out`
§7: exact and Monte Carlo orbit sums in $d = 4$	`orbhom2.c`, `orbhom3.c`, `orb4mc.c`, `orb4scan.c`, `psweep*.sh`, `susp_exact.sh`, `summarize_4d.py`	`summary_4d.txt`, `exact_*.out`, `mc_*.out`, `s1019_*.out`, `scan4d_*.out`, `exact_borderline.out`
§7: coarser groupings fail; subsets of classes on 3×4	`orbtr.c`, `orb4.c`, `orbhom2.c`, `orbtr_v2.c`, `run_subsets_3x4.sh`	`orbtr_grouping.txt`, `orb4_mode2.txt`, `exact_mode2_negatives_perclass.txt`, `subsets_3x4.txt`
After Prop. 17: values 64, 96; antiperiodic coefficients	`blocknu.c`, `blockrk.c`, `urk_exh.c`, `spec_search.py`	`blocknu_W8_reflsum.txt`, `blockrk_W8_33_33_15.txt`, `antiperiodic_W8_block.txt`, `urk_exh_W8_n2_dump.txt`, `spec_search_W8.txt`
Lemmas 3, 4 and 5(iv) on small tori (cross-check)	`check3d.c`, `checkq.c`, `check2d_q4.c`	`check3d.txt`, `checkq.txt`, `check2d_q4.txt`
Lemma 18(a): identities, and $\ell = 3, 5, 7$ exhaustively (cross-check); Smith form of the span of the forms $\langle z, gz \rangle$	`partner_rule.py`, `qspan_ref.py`	`partner_rule.txt`, `qspan_ref.log`
Lemma 29: (S0)–(S5) for $L_c \leq 7$, $2 \leq L_g \leq 7$; (S0)–(S5) and (U) for 255 pairs up to 32	`netstruct.py`, `net_ref.py`	`netstruct.log`, `net_ref.log`
After Lemma 30: modules computed directly; placements at the minimal L	`netcert.py`, `net_ref.py`, `placement_ref.py`	`netcert.log`, `net_ref.log`, `placement_ref.log`

Statement	Programs	Logs
§8.7, orders: the floating-point NS order of $\mathfrak{D}_{24}$ is the exact order (284,752 edges, no equal keys); gaps of $\langle \nu_0, \cdot \rangle$ up to l∞ norm 17 and of $\langle$G, $\cdot\rangle$ up to l∞ norm 2 (with a cross-check)	`nsorder_exact.py`, `keygap.py`, `vkey_gap.py`, `key_gaps.py`	`nsorder_exact_L24.log`, `keygap_R17.log`, `vkey_gap.log`, `key_gaps.log`
Prop. 31, base case L = 24: (C1) with locality, (C2), (C3) at all 168,712 edges	`capcert.py`, `caplib.py`, `intlink.py`, `funnel.py`, `cfg.py`, `runcap.sh`	`cap_19_10_24_*of12.log`, `SCRIPT_MD5_L24int.txt`
Prop. 31, base case: route structure and window types at L = 24; (F) at L = 24, 25, 30, 60	`analyze_routes.py`, `routes_ref.py`, `transport_reach.py`, `check_far_windows.py`	`analyze_routes_L24.log`, `routes_ref.log`, `transport_reach.log`, `check_far_windows.log`
Prop. 31: the other translate need not correspond (17,316 and 17,356 edges of $\mathfrak{D}_{25}$)	`make_maxable.py`, `preamble_check.py`	`make_maxable_24.log`, `make_maxable_25.log`, `preamble_check.log`
Prop. 31: integral link steps re-proved by explicit chains; exact integral test against a certified solver; independent replay of (C1) at L = 24, 25; an independent implementation of the search reproduces the base-case statistics chunk by chunk (cross-checks)	`certify_int_steps.py`, `test_bounds_torsion.py`, `test_bounds_real.py`, `zcert.py`, `c1_verify.py`, `refgeom.py`, `capgen.cpp`	`certify_int_steps.log`, `test_bounds_torsion.log`, `tbr_*.log`, `c1v_L24.log`, `c1v_L25.log`, `c1_L25.log`, `chain_capgen_19_10_24.log`
After Prop. 31: every edge of $\mathfrak{D}_{25}$, $\mathfrak{D}_{26}$ as in (C2) reached by the rule of the proof	`make_maxable.py`, `transport_reach_simple.py`, `transport_reach.py`	`make_maxable_*.log`, `transport_reach_simple.log`, `transport_reach.log`
After Prop. 31: full transport $\mathfrak{D}_{24} \to \mathfrak{D}_{25} \to \mathfrak{D}_{26}$ passes an independent checker	`capgen.cpp`, `transport_cert.py`, `capcheck.cpp`, `run_chain.sh`	`chain_capgen_19_10_2*.log`, `chain_tr_19_10_2*.log`, `chain_check_19_10_2*.log`, `chain_19_10_24_to34.out`
After Prop. 31: 1,416 designs at L = 24 and their transports to L = 25, 26, 30, verified by a second implementation	`verify_designs.py`, `verify_int_edges.py`, `refgeom.py`	`vd_*.log`, `vi_*.log`, `agg_vd.log`, `agg_vi.log`
After Prop. 31: coverage fails from L = 16 to 17 (56 edges)	`capgen.cpp`, `transport_cert.py`, `capcheck.cpp`, `run_chain.sh`	`chain_check_19_10_17.log`, `chain_batch.out`
§8.8, Remarks 1, 2: radii, ball sizes and $\log_2(1/w_2)$	`constants.py`	`constants.txt`
§8.8, Remark 2: second certificate for M = 2	`verify_funnel.py`, `funnel.py`, `colsearch.py`, `devices.py`, `netcheck.py`, `chunkcount.py`	`vf_*.log`, `vf_totals.log`, `SCRIPT_MD5.txt`, `netcheck_ctrunc25.log`, `chunkcount.log`
§8.8, Remark 3: end-to-end tests with Z/9	`e2e_cap.py`, `e2e_cap_cc.py`, `vimg_huge.c`, `vimg_h27.c`	`z9_t0.log`, `z9_cc_t1.log`
Prop. 31, base case L = 24: (C1) in local form, (C2), (C3) at all 168,712 edges, (H_{24}), (F); Lemma 26 and the 2-chain of Prop. 28 for ($\mathfrak{D}_{24}$, $\mathfrak{B}_{24}$, π_{24}) over every commutative ring; Lean kernel only, standard axioms	Lean 4 project `lean_kernel/` (`KernelV3/Main.lean`; leaves `KernelV3/Leaf/`, by `decide +kernel`)	`logs/axioms_v3.log`, build log
Lemma 30 for (ℓ, a, d) = (2,1,1), (2,2,1), (3,1,1), (5,1,1), all types (the third module statement with the chain anywhere on $\mathcal{N}$)	Lean 4 project `lean/` (`Char2Proof/NetsMain.lean`, `native_decide`)	`logs/axioms.log`
The same statements of Lemma 30 by the Lean kernel alone, on an explicit copy of the networks and certificates; the copy matched to the original data by a generator and a compiled comparison	`lean_kernel_nets/` (`KernelNets/`, by `decide +kernel`; `checks/`)	`logs/Axioms.log`, `logs/EvalEq.log`, `logs/dump_compare.log`, `logs/checks.log`

Statement	Programs	Logs
§8.8, Remark 2: (C1), and (C2) and (C3) with the integral link criterion at all 77,712 relevant edges, for the hole of each of the five networks; Lemma 26 and the 2-chain of Prop. 28 for this region over every commutative ring; Lean kernel only, standard axioms	Lean 4 project `lean_kernel_ct25/` (`KernelCT25/MainCT25.lean`; leaves `KernelCT25/*/Leaf/`, by `decide +kernel`)	`logs/axioms.log`, build and replay logs
§8.7 and §8.8, Remark 2, orders: the integer key of the Lean projects induces the order of §8.7 on the NS edges of $\mathfrak{D}_{24}$, and the floating-point order of the certificate of Remark 2 on the NS edges of its region	`test_nskey24.py` (with `nsorder_exact.py`), or the direct 60-digit check `ns_order_check.py`; `test_nskey_ct25.py`	`test_nskey24_D24.log`, `ns_order_D24.log`, `test_nskey_ct25_CT25.log`
§9, before §9.1: the nine permutation-invariant subspaces of H, which of them are isotropic, and which are spanned by vectors with Pf = 0 (an isotropic subspace spanned by such vectors is totally singular); the log calls Cyc "Z"	`dichotomy_check.py`	`dichotomy_check.txt`
After Cor. 38: at a = 0.2 a slice contains a cycle with odd spatial winding with probability $\geq 1/2$ from b ≈ 0.07 at N = 20 and from b ≈ 0.09 at N = 10 (40 samples per point)	`slice_wrap.py`	`slice_a0.2.txt`

Data and code availability

The programs, logs and certificates of the Appendix and the Lean projects, with the README and the hash list, are archived at https://doi.org/10.5281/zenodo.23631702, and https://doi.org/10.5281/zenodo.23477694 points to the latest version of that archive. This paper is archived at https://doi.org/10.5281/zenodo.23716219. What the Lean projects check, and what is left outside Lean, is stated in §8.6, §8.7 and Remark 2 of §8.8; the README gives the commands.

Use of AI tools

Generative AI tools were used to find and prove the results, to write the programs, to review the proofs and to draft the text. The author is responsible for all statements in the paper.

References

- [AG91] M. Aizenman, G. Grimmett, Strict monotonicity for critical points in percolation and ferromagnetic models, J. Stat. Phys. 63 (1991) 817–835.
- [BBR14] P. Balister, B. Bollobás, O. Riordan, Essential enhancements revisited, arXiv:1402.0834.
- [BDC12] V. Beffara, H. Duminil-Copin, The self-dual point of the two-dimensional random-cluster model is critical for q $\geq$ 1, Probab. Theory Related Fields 153 (2012) 511–542.
- [BKKKL92] J. Bourgain, J. Kahn, G. Kalai, Y. Katznelson, N. Linial, The influence of variables in product spaces, Israel J. Math. 77 (1992) 55–64.
- [BR06] B. Bollobás, O. Riordan, Sharp thresholds and percolation in the plane, Random Structures Algorithms 29 (2006) 524–548.
- [BS83] C. Borgs, E. Seiler, Lattice Yang–Mills theory at nonzero temperature and the confinement problem, Comm. Math. Phys. 91 (1983) 329.

- [DKS20] P. Duncan, M. Kahle, B. Schweinhart, Homological percolation on a torus: plaquettes and permutohedra, Ann. Inst. Henri Poincaré Probab. Stat. 61 (2025) 2235–2261; arXiv:2011.11903.
- [dMU21] L. de Moura, S. Ullrich, The Lean 4 theorem prover and programming language, in: Automated Deduction – CADE 28, Lecture Notes in Comput. Sci. 12699, Springer, 2021, 625–635.
- [DS22] P. Duncan, B. Schweinhart, Topological phases in the plaquette random-cluster model and Potts lattice gauge theory, Comm. Math. Phys. 406 (2025), Paper No. 145; arXiv:2207.08339.
- [DS26] P. Duncan, B. Schweinhart, A topological formula for Potts lattice gauge theory correlations, arXiv:2607.02434.
- [DSZ87] P. di Francesco, H. Saleur, J.-B. Zuber, Relations between the Coulomb gas picture and conformal invariance of two-dimensional critical models, J. Stat. Phys. 49 (1987) 57–79.
- [GG06] B. T. Graham, G. R. Grimmett, Influence and sharp-threshold theorems for monotonic measures, Ann. Probab. 34 (2006) 1726–1745.
- [Gri99] G. Grimmett, Percolation, 2nd ed., Grundlehren der mathematischen Wissenschaften 321, Springer, 1999.
- [HS16] Y. Hiraoka, T. Shirai, Tutte polynomials and random-cluster models in Bernoulli cell complexes, RIMS Kôkyûroku Bessatsu B59 (2016) 289–304.
- [Kes82] H. Kesten, Percolation Theory for Mathematicians, Progress in Probability and Statistics 2, Birkhäuser, 1982.
- [Mat20] The mathlib Community, The Lean mathematical library, in: Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020), ACM, 2020, 367–381.
- [OSSS05] R. O'Donnell, M. Saks, O. Schramm, R. A. Servedio, Every decision tree has an influential variable, in: Proc. 46th IEEE Symposium on Foundations of Computer Science (FOCS 2005), 31–39.
- [Pin94] H. T. Pinson, Critical percolation on the torus, J. Stat. Phys. 75 (1994) 1167–1177.
- [Rei00] D. Reimer, Proof of the van den Berg–Kesten conjecture, Combin. Probab. Comput. 9 (2000) 27–32.
- ['tH79] G. 't Hooft, A property of electric and magnetic flux in non-abelian gauge theories, Nucl. Phys. B 153 (1979) 141.
- [TY85] E. T. Tomboulis, L. G. Yaffe, Finite temperature SU(2) lattice gauge theory, Comm. Math. Phys. 100 (1985) 313.
- [vdBK85] J. van den Berg, H. Kesten, Inequalities with applications to percolation and reliability, J. Appl. Probab. 22 (1985) 556–569.