hexagon:2611.00973v1

Columnar order in a system of $2 \times 2 \times 2$ hard cubes on $\mathbb{Z}^3$

Generated by GPT-6 Astra
Prompted by Qidong He
Results not human-verified

Abstract

We study cubes of side two with integer centres and disjoint interiors in three dimensions. At sufficiently large fugacity the periodic Gibbs measures are the mixtures of twelve extremal measures. In each, most cubes lie in parallel columns of cross section 2×2; unit translation symmetry is preserved along the columns. We first prove layering. Within every favoured layer the cube centres then select a transverse coordinate and its parity. To show that successive layers make the same choice, we count cubes that straddle their common boundary. Reflection gives two tests of the same count: one based on replacing an aligned pair by one cube, the other on the vacancies required by a cube at an incompatible parity. Disagreement percolation gives uniqueness within each phase and exponential decay of correlations. The axial correlation length is of order $\sqrt{z}$ at fugacity z, and the axial head covariance at this scale has the one-dimensional hard-dimer profile.

1 Introduction

A fully packed gas of lattice cubes need not resemble a crystal. Tile each slab of thickness two independently by cubes: for instance, in one slab their rows can run east–west and in the next north–south. Even if all the rows run in the same direction, successive slabs can give them different offsets. All these packings have the same density. Why should equilibrium at a large but finite fugacity favour parallel columns extending throughout space?

Vigneshwar, Mandal, Damle, Dhar and Rajesh [10] observed a columnar phase at the highest densities in their simulations of cubes of side two. They identified cubes crossing the boundaries of the slabs as a source of alignment: after giving these cubes a separate activity, their first-order calculation assigns a free-energy cost to a misaligned layer [10, Eqs. (28)–(35)]. We study the gas with a single activity. The issue is to detect the preference for alignment in the presence of all the other allowed cubes.

1.1 Mathematical background

For integer centres the exclusion here is also Euclidean exclusion at distance strictly less than two: a forbidden displacement has every component in $\{-1, 0, 1\}$. This is the sliding case $D^2 = 4$ in the classification of dense packings by Mazel, Stuhl and Suhov [7, Sec. 2.3]. The corresponding two-dimensional example is the gas of 2×2 squares. Ramola and Dhar [8] analysed its high-activity columnar phase by resolving the motion along a row. Hadas and Peled [5] proved that its periodic equilibria are the mixtures of four columnar phases. They conjectured dk^{d-1} extremal periodic phases for side-k cubes in dimension d [5, Sec. 10.2]. Their periodic Gibbs chessboard estimate and their use of disagreements to distinguish Gibbs states provide two tools for the present argument. In three dimensions an additional question arises even after the planar registry of every layer has been identified: two adjacent layers must be shown to agree.

1.2 Model and main results

A configuration is a set of occupied sites, or *heads*, in $\mathbb{Z}^3$. A head u represents a cube with centre u and support $u + [-1, 1]^3$. Two distinct heads are compatible when their Chebyshev distance is at least two. Write n_u for the occupancy indicator. Given the configuration outside a finite set of sites, the probability of a compatible inside pattern with n heads is proportional to z^n. A Gibbs measure has these conditional laws for every finite set. A measure is periodic if it is invariant under translations in a finite-index sublattice.

The eight possible parities of a head form $\mathcal{P} = \{0, 1\}^3$. A face of $\mathcal{P}$ fixes one coordinate parity; an edge E fixes two. Write w for the free coordinate of an edge. Its *favoured columns* are the lines of possible heads parallel to e_w at the two prescribed transverse parities. Each physical unit cell belongs to exactly one of their 2×2 footprints. Write $b_S(B)$ for the fraction of the unit cells of a box B covered by cubes with head parity in $S \subseteq \mathcal{P}$; a vacant cell contributes zero. For a fixed coordinate j of E, call the plane of centres $u_j = m$ a *favoured host* if m has the prescribed parity. For a square Q in that plane, let $a_S(j, m; Q)$ be the fraction of Q covered by the footprints of cubes with heads *in that plane* of parity in S. Footprints are clipped to Q. We write $Q_L = [0, L]^2$ in transverse coordinates and $B_L = [0, L]^3$.

Theorem 1.1 (Columnar phases)**.** *For all sufficiently large z there is a Gibbs measure μ_E for every one of the twelve parity edges E, with the following properties.*

1. *Each μ_E is periodic and extremal among all Gibbs measures. Every periodic Gibbs measure has a unique representation as a convex combination of the μ_E.*

2. *For each $\varepsilon > 0$ and all sufficiently large z, μ_E-almost surely,*

$$\liminf_{L\to\infty} b_E(B_L) \geq 1 - \varepsilon, \qquad \liminf_{L\to\infty} a_E(j, m; Q_L) \geq 1 - \varepsilon \quad \text{for every favoured host } (j, m).$$

3. *Lattice translations, coordinate permutations and coordinate reflections carry μ_E to the state labelled by the transformed edge. In particular the translation periods of μ_E are exactly the vectors even in the two fixed coordinates; a unit shift along w preserves it.*

Theorem 1.2 (Correlations)**.** *There exist $c, C, z_0 > 0$ such that for $z \geq z_0$, bounded local functions f, g with finite nonempty head supports A, B satisfy*

$$|\operatorname{Cov}_{\mu_E}(f, g)| \leq C\|f\|_\infty\|g\|_\infty \min(|A|, |B|) \exp[-c \operatorname{dist}_\infty(A, B)/\sqrt{z}].$$

For heads on a favoured column and integers $t \geq 0$ with $t/\sqrt{z} \to s \geq 0$ as $z \to \infty$,

$$(-1)^t \operatorname{Cov}_{\mu_E}(n_u, n_{u+te_w}) \longrightarrow \tfrac{1}{4}e^{-s}.$$

Fix a physical unit square transverse to the columns and let D_j indicate that a cube centred at axial height j covers this square in projection. At each sufficiently large fixed z, the limit

$$\lim_{m\to\infty} -\frac{1}{2m} \log \frac{\operatorname{Cov}_{\mu_E}(D_0, D_{2m})}{\operatorname{Var}_{\mu_E}(D_0)}$$

exists and is bounded above and below by positive constants times $z^{-1/2}$. Its reciprocal is the axial correlation length of this observable.

1.3 Proof strategy

1. Find the layers. In a physical slice the cubes cut out squares of side two. Along an entirely covered row their head coordinates have one parity in the direction of the row. Section 3 bounds the probability that two specified rows have opposite parities and two specified columns also have opposite parities in the same square. Section 4 uses the resulting plane constraints in three directions. A box without a favoured parity face would have too much of its volume pinned to a single head parity, losing the entropy of sliding along a column. The labels of overlapping boxes then propagate one common signed coordinate through space. For example, most cubes can be assigned to hosts $y = 2m$.

2. Choose a second parity in each layer. In any such host, dense squares typically have one dominant transverse parity, and each of the two remaining head parities occupies a positive fraction of their area. This last condition makes the label distinguishable on overlapping squares. The estimates on exceptional boxes give a connected family of squares with one label in each favoured plane. At this point a different transverse coordinate, or the opposite sign of the same coordinate, is still possible in the very next plane.

3. Compare adjacent layers. Suppose the favoured planes $y = 0, 2$ have different labels. Choose a transverse position used by the lower plane but disfavoured by the upper plane. A line through this position at $y = 1$ counts actual intervening cubes. Section 5.5 constructs two Gibbs measures with exactly the same head distribution on this line. In the continuation from the lower host, an aligned pair may be replaced by one intervening cube at relative weight $1/z$, giving head probability at least c/z. In the continuation from the upper host, the same intervening head violates two coordinate parities. The occupied runs on its four physical rows in each of these directions must end in sixteen distinct vacancies unless a run is long. After bounding the long runs the head probability is $o(1/z)$. Equality of the line marginals excludes a mismatch. The deduction from layering to aligned columns takes place at one fixed high fugacity.

4. Distinguish the Gibbs measures. Once the parity edge is fixed, two independent samples can disagree along a favoured column until a vacancy blocks the path. A cube outside the favoured columns can allow the path to change columns. Section 6 bounds both possibilities and uses finite-component interchange to prove uniqueness. Comparison with a one-dimensional dimer chain gives the correlation scale and the leading density of straddlers that admit the reverse pair replacement.

2 Preliminaries

2.1 Cells, Gibbs boundaries and pressure

Coordinates of sites are integer head positions. We count volume in *physical cells*, the open unit boxes between consecutive coordinate planes. In dimension $d = 2, 3$ a cell has 2^d possible owners, its corners; at most one is occupied. A block D is a rectangle of such cells, and a function of D may use all head variables in its closure. In particular, two neighbouring blocks have disjoint cells but share boundary sites. Write $|D|$ for its number of cells. If S is a set of parities, the ownership fraction b_S defined in the introduction satisfies

$$b_S(B_L) = \frac{8}{L^3} \sum_{u \in \{0,\dots,L-1\}^3} \mathbf{1}_{\{u \bmod 2 \in S\}} n_u + O(L^{-1}). \tag{2.1}$$

Only cubes meeting the boundary are counted with a different multiplicity on the two sides.

An even rectangular torus T is a product of cycles of even lengths. Write π_T for its Gibbs law,

$V = |T|$ for its volume, and H for its number of vacant cells. Define

$$q = z^{-1}, \qquad \lambda = \sqrt{q}, \qquad p_d(q) = \lim_T \frac{1}{V} \log \sum_{\omega \text{ legal}} q^{H(\omega)/2^d}.$$

The normalized weight $q^{H/2^d}$ equals $z^{-V/2^d} z^{|\omega|}$; p_d is the ordinary pressure minus $2^{-d} \log z$. For a periodic field ϕ add the factor $\exp(\sum_u \phi(u) n_u)$ and denote the resulting normalized pressure by $p_d[\phi]$, at the same z. In particular $p_S(q, g) = p_3[-g\mathbf{1}_{\{u \bmod 2 \notin S\}}]$.

Lemma 2.1 (Changing the boundary)**.** *Let B be a rectangular set of sites, let $\partial^+ B$ be its exterior one-site collar in Chebyshev distance, and bound all activities by M. For any nonnegative function of the sites in B, its expectations for any two Gibbs boundaries* outside the collar *differ by at most a multiplicative factor $(1 + M)^{2|\partial^+ B|}$.*

Proof. Given a legal pattern in B and any exterior beyond $\partial^+ B$, the sum of the allowed collar weights lies between 1 (empty collar) and $(1 + M)^{|\partial^+ B|}$. Apply these bounds to numerator and denominator, comparing each exterior with the free law on B. □

For period-two activities the limiting pressure exists and is the same for free and even periodic boxes. Indeed, tile a large box by even subboxes and delete sites along their internal boundaries. The logarithmic error in additivity is bounded by a constant times their total boundary size, by the empty-collar comparison above. Take the large-box limit and then increase all subbox sides. The same argument applies in a bounded period-two field.

Lemma 2.2 (Pressure bound on an empirical mean)**.** *For $d = 3$, any bounded period-two field ϕ, and a periodic Gibbs measure μ at zero field,*

$$\langle \phi n \rangle_\mu \le p_3[\phi] - p_3[0], \tag{2.2}$$

where the left side is the expectation averaged over a common period. In any Gibbs measure, the upper limit of the empirical mean of $\phi(u) n_u$ in B_L is almost surely at most the same pressure difference.

Proof. Express the exponential moment in a free box as the ratio of the partition functions with and without the field. Lemma 2.1 then gives

$$\log \mathbb{E}_\mu \exp\Big(\sum_{u \in B_L} \phi(u) n_u \Big) \le |B_L| (p_3[\phi] - p_3[0] + o(1)).$$

Here and in this estimate only, the sum uses the half-open set of sites in (2.1); odd side lengths change a boundary term. Jensen's inequality gives (2.2); the exponential Markov inequality and Borel–Cantelli give the samplewise assertion. □

2.2 The chessboard estimate

On a torus a coordinate reflection ϑ through integer sites fixes two antipodal planes. Given their complete head patterns, the strict halves are conditionally independent: opposite interior sites are at normal distance at least two. Their conditional laws are reflections of one another. Hence for a function f on a closed half, $\mathbb{E}_{\pi_T}(f\, \vartheta \overline{f}) \ge 0$. This is *reflection positivity*. Every such reflection preserves the parity of each head. Both assertions also hold for period-two parity fields.

Tile T by copies of D with an even number of tiles in each direction. Write $\vartheta_i f$ for the copy of f on tile i, reflecting at successive boundaries, and $k = |T|/|D|$. The chessboard seminorm of f and the chessboard estimate for distinct tile indices $i_1, \dots, i_n$ are

$$\|f\|_{D,T} = \Big(\mathbb{E}_{\pi_T} \prod_i \vartheta_i |f|\Big)^{1/k}, \qquad \mathbb{E}_{\pi_T} \prod_{\ell=1}^{n} |\vartheta_{i_\ell} f_\ell| \le \prod_{\ell=1}^{n} \|f_\ell\|_{D,T}. \tag{2.3}$$

This is the site-reflection estimate [3, Theorem 5.8]; successive coordinate reflections give the same inequality for rectangles. Write $\|f\|_D = \limsup_T \|f\|_{D,T}$ along admissible dyadic tori with all lengths tending to infinity. An event in a seminorm denotes its indicator. The seminorm is monotone on nonnegative functions and subadditive: expand the product of the reflected sums in its definition and use (2.3) term by term. Grids may be translated. We use only dyadic side lengths for grids that have to fit the same torus.

Two consequences will be used with different block sizes. First, if one cell v incurs a factor $f_v = e^{-g/8}$ when its owner lies outside $S_v \subseteq \mathcal{P}$, and $f_v = 1$ otherwise, then

$$\log \Big\| \prod_{v \in D} f_v \Big\|_D \le \sum_{v \in D} \big(p_{S_v}(q, g) - p_3(q)\big). \tag{2.4}$$

To see this, disseminate (place reflected copies of) the block product and apply (2.3) on the unit-cell grid. A single factor disseminated to every cell acts on each off-S_v head eight times, producing precisely the stated pressure. Second, a bound on a block seminorm gives a probability bound even in a periodic Gibbs law without reflection positivity. This extension is due to Hadas and Peled for squares [5, Proposition 3.8]. The following proof includes families of events constructed using several grids.

Lemma 2.3 (Transfer to periodic Gibbs measures)**.** *For a finite observation G, suppose that any m distinct translated copies on a fixed rectangular grid of periods have joint torus probability at most v^m in the thermodynamic upper limit. A Gibbs law invariant under those periods satisfies $\mu(G) \le v$. In particular, the right side of (2.3), with thermodynamic seminorms, bounds the probability of any specified finite collection of tile events in every periodic Gibbs law.*

Proof. For the last assertion choose grid spacings that preserve the law, are multiples of twice the tile widths, and separate the finite tile families. Equation (2.3) then gives the hypothesis with v equal to the product of their seminorms. For the general assertion place n^3 translates G_i in a box of side $O(n)$ and expand $\prod_i [1 + (e^t - 1)\mathbf{1}_{G_i}]$, $t > 0$. The torus upper limit of its expectation is at most $[1 + (e^t - 1)v]^{n^3}$. Lemma 2.1 compares this expectation with that under μ at logarithmic cost $O(n^2)$, with z and the grid fixed. Jensen gives $tn^3\mu(G) \le O(n^2) + n^3 \log[1 + (e^t - 1)v]$. Send $n \to \infty$ and then $t \downarrow 0$. □

Constants denoted by c, C are positive, may change between inequalities, and are independent of z once the stated tolerances and grid ratios are fixed.

2.3 Exceptional sets and planar connectivity

Hereafter an index set D is *jointly β-rare* if

$$\Pr(F \subset D) \le \beta^{|F|} \quad (F \subset \mathbb{Z}^3 \text{ finite}). \tag{2.5}$$

Use nearest neighbours for paths unless stated otherwise; $*$-neighbours have distance one in the sup norm. We need the exceptional set to be small on lines as well as in volume. The following form of

the connected-set count is useful several times; compare the quantitative Peierls estimate of Hadas and Peled [5, Lemma 6.4]. For a set of sites its *box* is its coordinate bounding box, enlarged by a fixed integer in each direction. At level one take the boxes of the $*$-components of D. At each successive level join boxes that intersect or are $*$-adjacent, then take one box of the union for each connected family. Write R_i for the union at level i. The enlargement may be different at different levels.

Lemma 2.4 (Size and line occupancy of the boxes)**.** *Start from a set satisfying (2.5), and fix a finite number of levels and their enlargements. If β is small enough all the families joined at these levels are finite almost surely. At every such level i, for any site v and $n \geq 1$,*

$$\Pr\{v \text{ lies in a level-}i \text{ box of diameter plus one at least } n\} \leq C(C\beta^c)^n.$$

For any $\theta > 0$ and any coordinate line segment I of n sites,

$$\Pr\{|I \cap R_i| \geq \theta n\} \leq C_\theta (C_\theta \beta^{c_\theta})^n. \tag{2.6}$$

The constants depend only on the number of levels, enlargements and, in (2.6), on θ.

Proof. There are at most 26^{2a} $*$-connected sets of size a containing a given site, by a traversal of a spanning tree. The probability that one contains at least αa marks is at most $(2 \cdot 26^2 \beta^\alpha)^a$, by summing over subsets in (2.5).

We first locate the marks that account for the size of a box. Each box Q at a fixed level admits a $*$-connected set inside Q of size at most $C_i m_Q$ containing m_Q selected original marks, where $1 + \operatorname{diam}_\infty Q \leq C_i m_Q$. Different boxes of the same level use disjoint selected marks, even if the boxes intersect: their constituent families at the previous level are disjoint. At level one take the marked component itself. At the next level take a simple chain of constituent boxes between opposite extremes of a longest coordinate of the family. Set m_Q to the sum of their selected counts. Join their connected sets along this chain, at cost at most a constant times the sum of the constituent diameters plus one per box; the final enlargement costs only a fixed constant. The resulting connected set witnesses both assertions. It may also be joined to any specified point of Q at cost $O(m_Q)$.

It remains to justify that these successive boxes are finite. The rooted count first excludes infinite marked components. At any level already constructed it excludes arbitrarily large boxes meeting a fixed finite region, by connecting a witness to that region. Only finitely many smaller boxes can meet it, since their selected marks are disjoint and lie in a bounded region. Thus the graph of touching boxes is locally finite. An infinite family starting at a given box would contain arbitrarily long simple chains; the chain construction would give arbitrarily large connected sets with a fixed positive fraction of marks rooted at that box. Summing the rooted count rules this out, at all integer roots simultaneously. Induction therefore applies through the prescribed levels. The same count with a root at v and size at least n/C_i proves the diameter bound (using that the size is at least the number of selected marks).

For the line estimate, intersect the boxes with I. Greedily choose a longest remaining interval and discard all intervals it meets. The discarded intervals lie in its threefold enlargement. The disjoint intervals chosen therefore have total length at least $\theta n/3$. Join the witness for each of their boxes to its interval and adjoin I. The resulting connected set contains an endpoint of I and $M = \sum_Q m_Q \geq \theta n/(3C_i)$ distinct marks. It has at most $n + C_i' M \leq C_{i,\theta} M$ sites and at least n sites. Sum the rooted count over its possible sizes. □

Lemma 2.5 (Connection around planar defects)**.** *Let $A \subset \mathbb{Z}^2$ have finite $*$-components. If infinitely many sites are outside the bounding rectangles of these components, all such sites belong to the unique infinite nearest-neighbour component of A^c.*

Proof. For any finite $*$-component K of A, the sites at sup distance one from K which can reach infinity in $\mathbb{Z}^2 \setminus K$ by nearest-neighbour paths are nearest-neighbour connected [9, Theorem 4]. All these sites avoid A, by maximality of K. A point outside the bounding rectangle of K can reach infinity avoiding K on a coordinate ray. A point in an infinite component of A^c can also reach infinity avoiding K. Now join any two points of the kinds specified in the statement, or in an infinite component of A^c, by a finite unrestricted path. If it hits K, the site before its first visit and that after its last visit lie on the connected outer boundary just described. Replace this part by a path on that boundary. Each replacement removes visits to at least one component of A and introduces none. The resulting path avoids A. Infinitely many sites are thus joined, and this component contains every infinite component of A^c. □

3 Close packings and the cost of defects

3.1 Sliding cubes and a reversible replacement

A perfect packing on an even torus has one head for every eight cells. It may consist of independently tiled slabs as in the introduction, or of parallel columns in which the axial parity is chosen independently for each column. The following elementary part of the close-packing geometry will be useful at positive activity; compare the coordinate-line argument in Mazel, Stuhl and Suhov [6, proof of Theorem 3.4].

Lemma 3.1 (Encoding a perfect packing by line parities)**.** *On a torus of sides $n_1, n_2, n_3 \geq 4$ the number of perfect packings is at most $2^{n_1 n_2 + n_1 n_3 + n_2 n_3}$.*

Proof. On a physical coordinate line of a perfect packing, consecutive owners cover consecutive intervals of length two. Their centre coordinates along the line therefore have one common parity. List this bit for all lines in all three directions. The three lines through any cell give the three bits of its owner, choosing exactly one of its eight corners. This proves the count. □

Corollary 3.2 (Cost of covering prescribed cells)**.** *In any periodic Gibbs measure at fugacity z, any m prescribed distinct physical cells are all covered with probability at most $e^{-m p_3(q)}$.*

Proof. As in the chessboard estimate for full coverage by squares [5, Corollary 4.5], the unit-cell event of coverage, repeated on every cell, imposes a perfect packing of normalized weight one per configuration. Its log partition function per cell tends to zero by Lemma 3.1. Its chessboard seminorm is therefore $e^{-p_3(q)}$. Apply (2.3) on unit cells and Lemma 2.3. □

Lemma 3.3 (Replacing an aligned pair)**.** *Fix a possible centre u and a coordinate j. In any Gibbs measure ν, the probability that u is occupied and its cube can be legally split into the pair centred at $u - e_j, u + e_j$ is exactly*

$$q\,\nu(n_{u-e_j} = n_{u+e_j} = 1). \tag{3.1}$$

Proof. Condition on all sites other than these three. Whenever the pair is allowed, it excludes u. Its union contains the cube at u, so replacing the pair by that cube is always legal. The image consists precisely of the singly occupied patterns admitting the inverse split; the two weights are z^2 and z. Integrate over the fixed exterior. □

To use (3.1) one must know both where a layer has many candidate pairs and where a cube of an incompatible parity is rarer than q. We begin with finite-volume estimates for the projected squares in a slice between two consecutive planes of centres.

3.2 Transfer between two adjacent planes

A physical slice between centre heights j and $j+1$ receives squares from both planes. After projection their union is a legal square configuration. To estimate events on this union, let G be the exclusion graph within one plane: its vertices are potential centres and its edges join incompatible pairs. Write $\mathcal{I}(G)$ for its independent sets and $Z_G(\boldsymbol{a}) = \sum_{S\in\mathcal{I}(G)} \prod_{v\in S} a_v$. The sets occupied in the two planes must be disjoint and their union independent. For activities $\boldsymbol{a}, \boldsymbol{b}$ and $f : \mathcal{I}(G) \to [0,\infty)$ put

$$M_f(S,T) = \prod_{v\in S} \sqrt{a_v} \prod_{v\in T} \sqrt{b_v}\, f(S\cup T) \quad \text{for compatible } S,T, \qquad M_f(S,T) = 0 \text{ otherwise.}$$

The square roots assign half of each plane's weight to each of its two transfers. This is the usual trace construction for periodic lattice systems [1, Ch. 2]; the following bound controls a test on the *union*.

Write $\|\cdot\|$ for the Euclidean operator norm. The one-vertex matrix $\left(\begin{smallmatrix} 1 & \sqrt{\beta} \\ \sqrt{\alpha} & 0 \end{smallmatrix}\right)$ has squared norm $1+t$, where

$$t = \tau(\alpha,\beta) = \tfrac{1}{2}\bigl(\alpha+\beta-1+\sqrt{(1+\alpha+\beta)^2-4\alpha\beta}\bigr), \qquad F(z) = \tau(z,z) = z + \tfrac{1}{2}(\sqrt{1+4z}-1).$$

Increasing either argument cannot decrease the norm of its entrywise nonnegative matrix. Its quadratic equation gives

$$t \ge \max(\alpha,\beta), \qquad (t-\alpha)(t-\beta) = \frac{\alpha\beta}{1+t} = \alpha+\beta-t. \tag{3.2}$$

Lemma 3.4 (Deleting one site)**.** *Let A, B be nonnegative square matrices and I a coordinate inclusion with $I^*AI \le B$ entrywise. For $\alpha,\beta \ge 0$,*

$$C = \begin{pmatrix} A & \sqrt{\beta} IB \\ \sqrt{\alpha} BI^* & 0 \end{pmatrix} \quad \textit{satisfies} \quad \|C\|^2 \le \|A\|^2 + \tau(\alpha,\beta)\|B\|^2.$$

The block-matrix calculation is given in Appendix B.

Lemma 3.5 (An increasing test on the union)**.** *If $S \subset T$ implies $f(S) \le f(T)$, then*

$$\|M_f\|^2 \le \sum_{U\in\mathcal{I}(G)} f(U)^2 \prod_{v\in U} \tau(a_v,b_v). \tag{3.3}$$

Proof. Split the rows and columns according to presence of a vertex x. The resulting matrix has the block form of Lemma 3.4. Here A uses f on $G-x$, and B uses $U \mapsto f(U\cup\{x\})$ on the graph obtained by deleting x and its neighbours. The inclusion I identifies the independent sets on the smaller graph with those independent sets of $G-x$ avoiding the neighbours of x. Monotonicity gives $I^*AI \le B$. Apply the block lemma with $\alpha = a_x, \beta = b_x$ and induct on $|G|$: the two terms on the right count respectively $x \notin U$ and $x \in U$ in (3.3). For the empty graph the bound is equality. □

3.3 Pressure bounds for specified parities

Write P_d for ordinary pressure and $p_d(q) = P_d(z) - 2^{-d}\log z$, $q = z^{-1}$, for normalized homogeneous pressure. With parity-dependent activities the same reference term is subtracted. Let $p_S(q,g)$ be the normalized pressure of cubes with activity z on centre parities in S and ze^{-g} elsewhere.

For an even normal circumference $2m$ and A transverse sites, the partition sum with alternating plane activities is $\mathrm{Tr}[(M_1M_1^*)^m] \le 2^A Z_G(\boldsymbol{\tau})^m$ by (3.3) with $f=1$. Here $\tau_v = \tau(a_v,b_v)$. Taking

logarithms per cell and sending all sides to infinity removes an error at most $(\log 2)/(2m)$. Repeat along each coordinate. For any activities of period two this gives

$$P_d \le 2^{-d} \log(1+T), \tag{3.4}$$

where T is the last entry after the two activities differing only in the first coordinate have been merged by τ, then the two differing only in the second, and so on. The coordinate order is free.

Lemma 3.6 (Pressure of independent columns)**.** *For small q,*

$$p_{\mathrm{col}}(q) := \tfrac{1}{4}\operatorname{arsinh}(\sqrt{q}/2) = \tfrac{1}{8}\log[(1+F(z))/z] \ \le p_3(q) \le C\sqrt{q}.$$

Proof. Fixing parities in two coordinates leaves one independent dimer chain per four transverse cells. The same factorization for hard squares appears in [8, Eqs. (10)–(12)]. At activity z the two characteristic roots of the dimer recurrence are $\Lambda_\pm = (1 \pm \sqrt{1+4z})/2$. On an even cycle of length n the normalized partition sum is $z^{-n/2}(\Lambda_+^n + \Lambda_-^n) = 2\cosh(n \operatorname{arsinh}(\sqrt{q}/2))$. This gives the lower bound. For the upper bound use (3.4) with $T = F^{\circ 3}(z)$ and $F(x) \le x + \sqrt{x}$. □

Lemma 3.7 (Pressure in a parity field)**.** *For $16\sqrt{q} \le g \le 1$ and small q,*

$$p_S(q,g) \le \begin{cases} Cq/g, & S \text{ contained in one vertex of the parity cube,} \\ p_{\mathrm{col}}(q) + Cq/g, & S \text{ contained in one edge.} \end{cases}$$

Proof. Increasing activities, enlarge S to the vertex or edge. Write $b = ze^{-g}$. In merging a favoured branch $A \ge z$ with an unfavoured branch $B = F^{\circ k}(b)$, $k \le 2$, one has $B \le (\sqrt{b} + k/2)^2 \le b + 2\sqrt{z} + 1$ and $A - B \ge zg/4$. Equation (3.2) hence gives

$$0 \le \tau(A,B) - A \le B/(A-B) \le C/g. \tag{3.5}$$

For a single vertex there are three such increments in (3.4); for an edge merge its free coordinate first, so the favoured branch starts at $F(z)$ and has only two increments. Use Lemma 3.6 and subtract $(\log z)/8$. □

3.4 Transferring a square event to a cube slice

Lemma 3.8 (A square test in a physical slice)**.** *Let $\mathcal{A}$ be an increasing hard-square event on an $a \times b$ rectangle. Test it on one physical slice of an $a \times b \times c$ cube block by projecting the squares owned by the two adjacent centre planes. Then*

$$\|\mathcal{A}_{\mathrm{slice}}\|_{a\times b\times c} \le \|\mathcal{A}\|^{1/2}_{a\times b, F(z)} (F(z)/z)^{ab/8}.$$

The second factor is at most $\exp(Cab\sqrt{q})$.

Proof. On the transverse torus of area A, let $\mathcal{A}_*$ be the joint event on all reflected copies of the rectangle. Write N for the even normal circumference and $M = M_1$ for the homogeneous transfer. Exactly N/c physical slices receive a test. Each has transfer $K = M_{\mathbf{1}_{\mathcal{A}_*}}$, since either order of its adjacent planes gives the same union. The numerator for the dissemination is therefore bounded by $2^A \|K\|^{N/c} \|M\|^{N-N/c}$, and the denominator $\operatorname{Tr} M^N$ is at least $\|M\|^N$. By (3.3), $\|K\|^2 \le Z_G(F(z); \mathcal{A}_*)$, whereas $\|M\|^2 \ge \|Me_\varnothing\|^2 = Z_G(z)$. Each transverse independent set has at most $A/4$ heads, so $Z_G(F(z))/Z_G(z) \le (F(z)/z)^{A/4}$. Taking the $abc/(AN)$ power of the probability gives the asserted finite transverse seminorm bound with an extra factor $2^{abc/N}$. This tends to one in the thermodynamic limit. □

3.5 The square transfer for prescribed rows

Consider the square gas at activity z and put $u = z^{-1/2}$. A *physical row* of $[0,a] \times [0,b]$ is a strip of unit cells of height one; columns are defined by interchanging the coordinates. A fully covered row has a single horizontal centre parity: consecutive length-two intervals along it meet end to end. The two outermost intervals may extend beyond the rectangle. We will impose two fully covered rows of opposite horizontal parities. We first describe the transfer they have in common, whether or not two vertical columns are also prescribed.

Reduction to one reflected strip. Take a, b even, $a, b \geq 4$, and disseminate the test on a $2ma \times 2nb$ torus. Let Z_* denote its partition sum normalized by $z^{-\mathrm{area}/4}$. Cut at horizontal centre lines b apart. Given the head positions on the cuts, interiors of consecutive strips do not interact. Assigning half of a cut head's activity exponent to each side gives alternating strip matrices D, D^{T}, related by reflection. This remains true with an extra activity factor depending on horizontal centre parity. Consequently

$$Z_*(2ma, 2nb) = \mathrm{Tr}[(DD^{\mathsf{T}})^n] \leq (\mathrm{Tr}\, DD^{\mathsf{T}})^n = Z_*(2ma, 2b)^n. \tag{3.6}$$

On this shorter torus let $N = 2b$. Retain every other vertical centre column and sum over the intervening columns. A configuration S on a retained column is a matching of adjacent physical cells on the vertical N-cycle. Denote its number of dimers by $n(S)$ and its set of covered cells by $U(S)$. Each prescribed horizontal row runs across the whole torus. One row therefore forces a specified cell to be in $U(S)$ on every retained column, while the opposite row forces a different specified cell to be outside $U(S)$ there. Keep these two constraints and discard the other copies of the row constraints. Cut the cycle at the required vacant cell. The matching is now uniquely determined by its covered set: every occupied run on the remaining path has even length.

For consecutive retained columns with states S, T, the intervening column admits dimers on the complement of $U(S) \cup U(T)$. Its normalized transfer is

$$K(S,T) = u^{N-n(S)-n(T)} Z_{\mathrm{dim}}\big((U(S) \cup U(T))^c; u^{-2}\big). \tag{3.7}$$

Here $Z_{\mathrm{dim}}(D; z)$ sums $z^{|M|}$ over all dimer matchings M in D; the factor before it includes two physical columns of normalization and half the activity of each adjacent state. The matrix K is indexed only by matchings obeying the two specified cell constraints.

Lemma 3.9 (A bound independent of the circumference)**.** *The nonnegative matrix polynomial $K(u)$ of (3.7) obeys $\|K(1/N)\| \leq C$ with an absolute constant.*

Proof. Let $U = U(S)$ and $V = U(T)$, with $n(S) \geq n(T)$ by symmetry. Insert S into any allowed matching on the intervening column. The resulting matching leaves $V \setminus U$ vacant. Under the unrestricted cycle law, the probability of f specified vacancies is at most $\Lambda_+^{-f} \leq u^f$, with Λ_+ as in Lemma 3.6 at $z = u^{-2}$. Indeed deleting the specified cells leaves paths, whose length-k partition sums are at most Λ_+^k by the dimer recurrence; the even cycle sum is at least Λ_+^N. Undoing the insertion of S gives

$$K(S,T) \leq 2\cosh(N\,\mathrm{arsinh}(u/2))u^{|U \triangle V|/2}, \tag{3.8}$$

since $n(S) - n(T) + |V \setminus U| = |U \triangle V|/2$.

We bound the matrix on the right over *all* covered sets on the path, temporarily dropping the required occupied cell. Join two sets in a graph with adjacency matrix J if their Hamming distance is two. The symmetric difference of their two matchings consists of $r = |U \triangle V|/2$ alternating paths:

the endpoints are exactly the cells covered in only one matching, and the underlying path has no cycle. Switching these paths in any order gives $r!$ walks of length r in J. Hence $u^r \le (e^{uJ})_{U,V}$.

To bound $\|J\|$, let L be the sum of the $N-1$ single-bit flips on all binary words of the path. On words with all runs of ones even, split $L = L_+ + L_-$ according as one adds or removes a cell. After either flip there is exactly one odd run. For addition, an output odd run of length $2h+1$ has $h+1$ predecessors (deleting alternating cells). Summed over additions from a fixed input, these multiplicities are at most N: each vacant cell contributes one plus half the lengths of the occupied runs on either side, and each run is counted at most twice. Cauchy–Schwarz at each output gives $\|L_+\|^2 \le N$. For removal an output has at most two predecessors, since the cell to be restored must touch an end of its odd run; hence $\|L_-\|^2 \le 2N$. The compression of L^2 to the valid words is $(N-1)\mathrm{Id} + 2J$. The preceding bounds give $\|J\| \le CN$. Set $u = 1/N$ in (3.8) and bound its right side in norm by $2\cosh(N \operatorname{arsinh}(u/2))e^{u\|J\|}$.

Finally K is a polynomial in u with nonnegative coefficients: the exponent of a matching of size k in (3.7) is $N - n(S) - n(T) - 2k \ge 0$, since the $2k$ cells avoid $U \cup V$. □

Lemma 3.10 (Dimers at orders zero and one)**.** *Let P_i project onto the retained states whose dimer centres all have parity $i = 0, 1$, $P = P_0 + P_1$. These projections are disjoint. For $0 \le uN \le 1$,*

$$K(u) = P + uK_1 + E(u), \qquad \|K_1\| \le CN, \qquad \|E(u)\| \le C(uN)^2. \tag{3.9}$$

Within either class P_i, K_1 joins states differing by the addition or removal of one dimer, every such entry equal to one. Between the two classes the linear term has norm at most two.

Proof. Every state contains the prescribed covered cell, so $P_0P_1 = 0$. In (3.7) maximize the number of dimers in each complementary interval. The smallest exponent of u is

$$\tfrac{1}{2}|U(S) \triangle U(T)| + \#\{\text{odd-length gaps in the complement of } U(S) \cup U(T)\}.$$

It is zero precisely when $S = T$ and all gaps are even, that is, the nonempty matching has one centre parity. Each even gap has a unique perfect matching, so $K(0) = P$. Lemma 3.9 and nonnegativity of the coefficients bound the linear coefficient by CN and the sum of orders ≥ 2 at u by $C(uN)^2$.

For two states in the same parity class, each dimer has one of $N/2$ disjoint positions. An entry of order one requires a change of one dimer, and conversely this gives coefficient one. For opposite pure parities, suppose the first nonzero exponent is one. Their covered sets differ in two cells and the gaps of their union are even. If they had equal size their union would have odd size, impossible on the even cycle. Hence one set contains the other. Each run of the larger set must lose both endpoints: an endpoint remaining in the smaller set would force a dimer of the larger set's parity there. Thus the larger set is one run, and the smaller is its interior. From any state there is at most one contraction and one extension; the coefficient is one in either case. The resulting symmetric adjacency has norm at most two. □

Lemma 3.11 (Four prescribed lines)**.** *On $[0, a] \times [0, b]$, let W require full coverage of two prescribed rows of opposite horizontal centre parities and two prescribed columns of opposite vertical centre parities. If $bu \le c$ for a sufficiently small absolute c, then*

$$\|W\|_{a\times b} \le [Ca(u + u^2b^2)]^{1/2} \exp(Cabu).$$

Proof. Use (3.6) and the transfer (3.7). Each prescribed vertical column has one retained centre column as a boundary. Every dimer there projects onto the prescribed column, so all such dimers must have its vertical parity. Keep one requirement of each parity per horizontal period $2a$. They

force visits to P_0 and P_1, defined in Lemma 3.10. If the same retained column is forced to have both parities, the required occupied cell makes the event empty.

Put $Q = \mathrm{Id} - P$ and $\rho = \|K\|$. When both pure classes exist, (3.9) gives $1 \le \rho \le e^{CuN}$, $\|QKP\| \le CuN$ and $\|QKQ\| \le CuN \le 1/2$. Every path from P_0 to P_1 has a segment going directly from P_0 to P_1, of norm at most $2u + C(uN)^2$, or such a segment with one or more intermediate Q states. Summing the latter norms over all lengths costs at most $\|PKQ\|^2/(1 - \|QKQ\|) \le C(uN)^2$. Allow at most s choices for the first step of the segment and bound the remaining transfers by powers of ρ. Entrywise nonnegativity then gives

$$\|P_1 K^s P_0\| \le Cs[u + (uN)^2]\rho^s.$$

In one horizontal period the two gaps between the required visits have lengths $s_1 + s_2 = a$. Apply this bound in each gap. The transfer over m periods thus has trace at most $2^N[Ca(u+u^2N^2)]^{2m}\rho^{ma}$, where 2^N bounds its dimension. The full normalized torus partition function is at least one (use a perfect packing). By (3.6), take the $1/(4mn)$ power of the constrained partition sum, then let $m, n \to \infty$ at fixed a, b. Since $N = 2b$, this gives the stated seminorm. □

4 Layering

We first identify a preferred coordinate parity near a given point. We then compare the labels of overlapping cubes. Corollary 4.16 shows that one parity accounts for all but a small fraction of the physical volume. The second transverse parity is still free to vary from one layer to the next.

4.1 Finding a preferred parity in a cube

Write $b_S(B)$ for the fraction of the physical cells of B whose owner has parity in $S \subseteq \mathcal{P} = \{0,1\}^3$, and $V(B)$ for its number of vacant cells. Thus b_S is always divided by *all* cells of B. The face $F(j,s) = \{r : r_j = s\}$ fixes one coordinate parity. We will find such a face in a cube of side N; later we compare the faces found in overlapping cubes. Throughout this subsection $q = 1/z$, $\lambda = \sqrt{q}$. Choose dyadic lengths

$$h \asymp q^{-13/50}, \qquad a \asymp q^{-1/5}, \qquad a \mid h. \tag{4.1}$$

Here N may be h or any fixed dyadic multiple of h. Constants in block estimates may depend on that multiple. The reason for using these sizes is that $\lambda h^2 \to \infty$, whereas $\lambda a h \to 0$. Thus a pressure difference of order λ is visible over a whole square of side h, while a thinner test rectangle still has negligible pressure error. Compare the use of rectangles on opposite sides of this area scale for hard squares in [5, Sec. 4].

We use two estimates from the preceding sections. First, Lemma 3.6 gives $0 \le p_3(q) \le C\lambda$ for the pressure of the weights $q^{V/8}$. Multiplication by e^V replaces q by $e^8 q$. Disseminating the test on a rectangle D and using Markov's inequality therefore gives

$$\|\{V(D) > s\}\|_D \le \exp[-s + C_0\lambda|D|]. \tag{4.2}$$

This also bounds the seminorm of the event of more than s vacancies in one specified physical slice of D.

Second, consider a single physical slice, the unit interval between two consecutive centre heights. Its covering squares come from the two bordering centre planes. A row is *clean* if all its cells in the observation square are covered. Consecutive owners in that row have centres two apart along it; thus the centres covering a clean row have one common parity in the row direction, called its registry. Use the same definition for a column. Call an N-square slice *ambiguous at* r if each registry occurs on at least rN clean rows and on at least rN clean columns.

Lemma 4.1 (Two registries in both directions)**.** *For $0 < r \le 1/10$, in a specified N-square physical slice,*

$$\|\{\text{ambiguous at } r,\ V(\text{slice}) \le a/10\}\|_D \le Cr^{-2}q^{1/100}.$$

Here D has that square as its footprint and has depth two or N.

Proof. In a rectangle of sides a, N, prescribe two clean rows of opposite registries and two clean columns of opposite registries. The square estimate of Lemma 3.11, at the transfer activity $F(z)$, bounds the corresponding square seminorm by $[Ca(\lambda' + (\lambda' N)^2)]^{1/2}\exp(CaN\lambda')$, where $\lambda' = F(z)^{-1/2} \le \lambda$. By (4.1) this is $O(q^{7/50})$. Adding a compatible square cannot alter any prescribed owner. Lemma 3.8 therefore bounds the three-dimensional seminorm for these four prescribed lines by $Cq^{7/100}$, since $aN\lambda = O(q^{1/25})$.

Take bands of a rows starting at $0, a/2, \dots, N-a$. If the event in the statement occurs, at most $a/10$ rows are not clean. Some band contains at least $ra/4$ clean rows of each registry. Otherwise every band would have more than $0.8a$ rows of one registry; successive bands must then favour the same registry since their union has only $1.5a$ rows. The disjoint bands starting at multiples of a would give fewer than rN clean rows of the other registry. In the band just found the more numerous registry has at least $0.45a$ clean rows. Two uniform ordered row positions thus select clean registries $0, 1$ with probability at least cr. Two uniform column positions do so with probability at least cr as well: at most $a/10$ entire columns are not clean, and a clean column retains its registry on restriction to the band. The average of the four-line indicators in this band is consequently at least c^2r^2.

For each summand its seminorm on D has the same upper bound $Cq^{7/100}$ as on the band. Indeed in the full dissemination of D the reflected bands occupy distinct tiles of the grid whose origins are $0 \bmod a$ or $a/2 \bmod a$, according to the original band. Reflection in a side of D preserves either grid. Pullbacks can change the test positions within a tile; the four-line bound is uniform in those positions. Apply the chessboard estimate (2.3) on that grid and take one root per D-tile. This use of a finer grid is the recursive chessboard argument of [5, Lemmas 3.4 and 3.7]. Summing the averages over at most $2N/a = O(q^{-3/50})$ bands and using the triangle inequality proves the result. □

Lemma 4.2 (Few ambiguous slices in one cube)**.** *Take $r = q^{1/1000}$, $\eta = q^{1/500}$. For each $P > 0$ and small enough q, the cube seminorm of the event that some direction has at least ηN ambiguous slices with at most $a/10$ vacancies is at most q^P.*

Proof. Fix the direction and group consecutive slices into $N/2$ disjoint slabs of depth two, starting at a cube boundary. A slab with either of its slices exceptional has seminorm at most $2Cq^{1/125}$ by Lemma 4.1. At least $m = \lfloor \eta N/2 \rfloor$ distinct slabs must be exceptional. For any specified m slabs, their reflected copies inside all the disseminated cubes are distinct tiles of one depth-two grid. Equation (2.3) thus bounds the cube seminorm by the product of the m slab bounds. Summing over choices and directions gives

$$3\binom{N/2}{m}(2Cq^{1/125})^m \le 3(C'q^{1/125}/\eta)^m.$$

Here $m \to \infty$ and $q^{1/125}/\eta = q^{3/500}$. □

We now need a deterministic observation. At each physical cell three slices meet, one in each normal direction. A slice label specifies a parity in *one of its two transverse directions*. For example, a slice normal to x can prescribe y or z, never x.

Lemma 4.3 (Three families of slice labels)**.** *Suppose in each direction all but a fraction σ of the slices are labelled as above. At most a fraction v of all cells are vacant or have an owner that violates*

a label. If each bit in each direction occurs in the ownership of at least a fraction δ of all cells and $3\sigma + 12\sqrt{v} + v \le \delta/2$, then at least $c\delta^{3/2}$ of the volume lies in cells at which the three labels specify all three bits consistently.

Proof. In each of the three families discard the unlabelled slices and each of the four possible label classes whose fraction is at most $\sqrt{v}$. The remaining cells form a Cartesian product; the discarded fraction is at most $s = 3\sigma + 12\sqrt{v}$. If two surviving classes from different families imposed contrary bits on the same coordinate, their intersection would have fraction greater than v. Its owners could never satisfy both; thus all surviving labels agree whenever they specify the same coordinate. Give each surviving cell an owner parity satisfying its labels, keeping the actual parity wherever it already satisfies them. In the conditional uniform distribution on the surviving product, each bit of each coordinate then still has probability at least $\delta - s - v \ge \delta/2$.

For a uniform cell in this product let A, B, C be the probabilities that the slices normal to x, y, z, respectively, prescribe the next coordinate in the cycle x, y, z, x. These three choices are independent because each depends only on its own normal position. The probability d of three distinct prescriptions and the probabilities k_i of leaving coordinate i unspecified are

$$\begin{aligned} d &= ABC + (1-A)(1-B)(1-C), \\ k_x &= B(1-C), \qquad k_y = (1-A)C, \qquad k_z = A(1-B). \end{aligned}$$

If $0 < A, B, C < 1$, each coordinate receives prescriptions from both possible families. Their consistency forces them all to prescribe the same bit in that coordinate; its other bit can occur only where it is unspecified. Thus $k_i \ge \delta/2$ and $d^2 \ge 4k_xk_yk_z$. For a boundary value, relabel axes so that $A = 0$. All x-slices now prescribe z. If $B > 0$, one y-slice also prescribes z, forcing all x-slices to give that same bit throughout the product, a contradiction. Hence $B = 0$, and the same reasoning gives $C = 0$, whence $d = 1$. Permuting the names of the axes covers either orientation of the cycle. Multiply d by $1 - s$ to recover the fraction of the original volume. □

Proposition 4.4 (One preferred face)**.** *Take $\delta = q^{1/5000}$. For an absolute K_v and every $P > 0$, at all sufficiently small q the cube B of side N satisfies*

$$\Big\| \{V(B) > K_v\lambda N^3\} \ \cup \ \{\ b_{F^c}(B) > \delta \text{ for every parity face } F\ \} \Big\|_B \le q^P.$$

Proof. By (4.2) the vacancy exception costs less than any fixed power of q if $K_v > 2C_0$. With fewer vacancies, in each direction at most $10K_v\lambda N^2/a = O(q^{9/50})$ of the slices have more than $a/10$ vacancies. By Lemma 4.2, outside another event smaller than any power, the high-vacancy and ambiguous slices together have fraction at most 2η per direction. In each remaining slice one family of lines has fewer than rN clean representatives of a registry. Label that slice by the opposite bit along those lines. If the slice has V_i vacancies, at most $rN^2 + NV_i$ of its cells can then violate its label. Summing in all three directions, including all vacant cells, gives a fraction at most $v_0 = 3r + (3N+1)K_v\lambda = O(r)$.

Suppose every face has outside ownership greater than δ. As $6\eta + 12\sqrt{v_0} + v_0 = o(\delta)$, Lemma 4.3 gives at least $c\delta^{3/2}N^3$ cells where the prescribed parity is a singleton. For a fixed set of slice labels let S_u be the allowed owner parities at cell u. Except on cells meeting an unlabelled slice, the three prescriptions have at least two distinct target coordinates, so S_u lies in an edge or is empty. Penalize each covered cell violating its prescriptions by $e^{-g/8}$. To bound the seminorm of the event for this set of labels we may undo at most v_0N^3 such penalties. Equation (2.4) then bounds its logarithm by

$$\frac{gv_0N^3}{8} + \sum_{u\in B} \big(p_{S_u}(q,g) - p_3(q)\big) \le N^3\big[-c'\lambda\delta^{3/2} + Cq/g + gv_0/8\big].$$

Indeed Lemma 3.7 bounds the pressure of a singleton by Cq/g, that of an edge by $p_{\mathrm{col}} + Cq/g$, and $p_3 \geq p_{\mathrm{col}} \geq c\lambda$. At all other cells the penalty only decreases the pressure. Set $g = \lambda/\sqrt{r}$; this is between 16λ and one for small q. The two error terms are $O(\lambda\sqrt{r}) = o(\lambda\delta^{3/2})$. There are at most 5^{3N} choices of labels, including the unlabelled option. Their log count is $O(N)$, whereas $\lambda\delta^{3/2}N^2 \asymp q^{-197/10000} \to \infty$. The union therefore has seminorm smaller than every fixed power of q. Combining the three bounds (with a larger power if necessary) proves the proposition. □

4.2 The transverse parity in one centre plane

Fix $F = F(j, s)$. A *favoured plane* is a centre plane $x_j = m$ with $m \equiv s \bmod 2$. For a transverse square Q, let $a_r(j, m; Q)$ be the fraction of Q covered by footprints of heads of parity r *in this plane*, clipping each footprint at the boundary. Write $a_S = \sum_{r\in S} a_r$. These are area fractions; $b_S(B)$ continues to denote a volume fraction. The following identity converts volume control in a cube into area control at every favoured height.

Lemma 4.5 (Coverage at each favoured height)**.** *Let $B = Q \times [t, t + N]$ be a cube of even side $N \geq 4$ with integer boundaries and normal coordinate j. If $b_{F^c}(B) \leq u$, the fraction of Q missed by its favoured plane at any $m \in [t, t + N]$, including a boundary plane, is at most $Cu + 2V(B)/N^2$.*

Proof. Let c_i be the covered area from centres at height i, and let v_i count vacancies between heights i and $i+1$ over Q. Disjointness and the fact that a cube extends one unit above and below its centre give

$$c_i + c_{i+1} = N^2 - v_i, \qquad c_{i+2} - c_i = v_i - v_{i+1}. \tag{4.3}$$

Any two plane areas of the same parity in the closed interval differ by at most $\sum_i v_i = V(B)$. At least $N/2 - 1 \geq N/4$ interior planes have the unfavoured parity. Their total area, counted twice in volume, is at most uN^3. One has area at most $2uN^2$; every unfavoured plane therefore has area at most $2uN^2 + V(B)$. Every favoured plane has an unfavoured neighbour within the closed interval. The first identity in (4.3) now bounds its deficit by $2uN^2 + 2V(B)$, as claimed. □

An h-square Q in a centre plane is *dense* if $a_{\mathcal{P}}(j, m; Q) \geq 1 - d_0$. Call it *usable* if it is dense and some edge $E \subset F(j, m \bmod 2)$ satisfies

$$a_{E^c}(j, m; Q) \leq e, \qquad a_r(j, m; Q) \geq \kappa \quad (r \in E). \tag{4.4}$$

Thus both endpoint parities must be represented. If $e < \kappa$ there can be only one such edge, since two distinct edges have at most one endpoint in common.

Proposition 4.6 (The edge of a dense square)**.** *There is an absolute $\varepsilon_v > 0$ such that if $0 < e < \kappa$ and $d_0 + e + \kappa \leq \varepsilon_v$, then for the slab $D = Q \times [m - 1, m + 1]$ of side lengths $h, h, 2$,*

$$\|\{Q \text{ is dense but not usable}\}\|_D \leq C_e q^{1/100} + 4e^{-c\lambda h^2}.$$

The same assertion holds at any fixed dyadic multiple of h with constants allowed to depend on the multiple.

Proof. If every edge has more than e area outside it, then both bits of each transverse coordinate have area greater than e in this centre plane. They retain that lower bound in the physical slice above it: the heads at height $m + 1$ add disjoint footprints. Except on an event of seminorm at most $\exp(-a/10 + 2C_0\lambda h^2)$, by (4.2), this slice has at most $a/10$ vacancies. In particular at most $a/10 = o(h)$ rows or columns in each direction contain a vacancy. A bit with area $> e$ must thus

occur on at least $eh/2$ clean lines of its direction. Lemma 4.1 at $r = \min(e/2, 1/10)$ gives the first term.

Otherwise take an edge with $a_{E^c} \le e$. If an endpoint has mass less than κ on a dense square, the other endpoint owns at least $1 - d_0 - e - \kappa$ of Q; all those footprints cover both physical slices of D. The event $b_{\{r\}}(D) \ge 1 - \varepsilon_v$ for this one endpoint r has seminorm at most

$$\exp\{|D|[Cq/g + g\varepsilon_v/8 - p_{\mathrm{col}}(q)]\} \le e^{-c\lambda|D|}.$$

Here (2.4) applies the field for a singleton and Lemma 3.7 bounds its pressure. Choose $g = K_1\lambda$ with K_1 a sufficiently large constant, then ε_v sufficiently small. Sum over the four possible endpoints in this centre plane. Finally $\lambda h^2 = o(a)$, so the preceding vacancy tail can be absorbed in $C_e q^{1/100}$. □

4.3 Labels on overlapping cubes

Fix $\kappa < \varepsilon_v/4$ and small $d_0, e > 0$ as in Proposition 4.6, with $e < \kappa/16$. Choose a dyadic $K > 16/\kappa$; fine squares have side h and origins on the grid of spacing $d = h/K$. Write $H = Mh$ for a larger side, where $M \ge 4$ is a fixed dyadic integer. The cube $B_v = Hv + [0, H]^3$ is indexed by $v \in \mathbb{Z}^3$. Its *parent* $P_v = Hv + [0, 2H]^3$ contains the eight children B_{v+w}, $w \in \{0, 1\}^3$. A child belongs to eight parents; the label we eventually assign to it must be determined by the child alone.

Call P_v *good* if the following two conditions hold:

1. $V(P_v) \le K_v\lambda|P_v|$ and $b_{F^c}(P_v) \le \delta = q^{1/5000}$ for at least one face F;
2. for each normal direction and each *interior* integer height at most b fine squares in the parent are dense but not usable. We inspect all fine squares contained in the transverse square of the parent with origins on the d-grid.

Here b is a fixed integer chosen before M. A parent that is not good is called bad.

Lemma 4.7 (Joint bound for bad parents)**.** *For every $P > 0$, one can choose $b = b(P, K)$ so that, for any fixed M and all sufficiently small q, every periodic Gibbs measure satisfies*

$$\Pr(P_v \text{ is bad for every } v \in A) \le q^{P|A|} \tag{4.5}$$

for every finite $A \subset \mathbb{Z}^3$. The same bound holds in the thermodynamic limsup of symmetric tori.

Proof. Choose at least $|A|/8$ indices of one parity modulo two in all three coordinates; their parents have disjoint interiors. If all fail, at least half fail the same numbered condition. Choosing that subset costs at most $2^{|A|+1}$. For condition 1 the chessboard estimate on the parent grid and Proposition 4.4 give an arbitrary power of q per selected parent.

For condition 2 choose in each of k selected parents one direction, one height and b distinct offending fine squares. There are at most $C_{M,K,b}H$ choices per parent, since only the number of possible heights grows with q^{-1}. Sort the bk depth-two slabs by normal, centre height modulo two and transverse origin modulo h in both coordinates. In the largest of these $6K^2$ classes at least $bk/(6K^2)$ distinct slabs belong to one chessboard. Indeed every slab is contained in its parent and the chosen parents have disjoint interiors. For any $c_h < 1/100$, Proposition 4.6 bounds the seminorm per slab by q^{c_h} for small enough q. The bound for a prescribed set of k parents is therefore

$$\left[C_{M,K,b}Hq^{c_h b/(6K^2)}\right]^k. \tag{4.6}$$

Since $H = O_M(q^{-13/50})$, take $c_h b/(6K^2) > 13/50 + 16(P + 2)$ to absorb the two subset losses. These arguments use (2.3) on symmetric tori or Lemma 2.3 to bound the same fixed tile intersections in a periodic Gibbs law. □

The geometry of the good-parent condition has two simple consequences. First, bounds known on most overlapping windows pass to a whole square. Secondly, unusable windows cannot separate the four transverse children into regions with different labels.

Lemma 4.8 (Averaging the windows of a square)**.** *In an H-square Q of the d-grid there are $n_c^2 = (MK - K + 1)^2$ fine squares. Suppose at least a fraction $1 - \eta$ have edge E in (4.4), with $\eta \le 1/4$. Then*

$$a_{E^c}(j, m; Q) \le 4/M + e + \eta, \qquad a_r(j, m; Q) \ge \kappa/4 \quad (r \in E). \tag{4.7}$$

Proof. Each physical cell is in at most K^2 of the windows, and exactly K^2 if it lies at least h from the boundary. The boundary strip has area at most $4H^2/M$; the summed outside areas are at most $(e + \eta)n_c^2 h^2 \le (e + \eta)K^2 H^2$. For either endpoint the summed areas are at least $(1 - \eta)n_c^2 h^2 \kappa$. Dividing by $K^2 H^2$ proves the lower bound since $M \ge 4$ and $(1 - \eta)(1 - 1/M)^2 \ge 1/4$. □

Proposition 4.9 (One edge in all four transverse children)**.** *Fix $\eta > 0$ and choose M sufficiently large in terms of b, K, η. Let P_v be good with face $F(j, s)$. At every interior favoured height m there is one edge $E_m \subset F$ such that, in each of the four transverse H-squares, at least a fraction $1 - \eta$ of all fine windows have that edge. Each of the four H-squares has total covered fraction at least $1 - C(\delta + \lambda H)$ from this centre plane.*

Proof. Lemma 4.5 bounds the deficit on the parent footprint by $C(\delta + \lambda H)(2H)^2$. Its deficit on any fine square is no greater, so all fine windows are dense when q is small at fixed M; dividing instead by H^2 proves the last assertion. At most b fine origins are unusable. Two usable windows at neighbouring origins have the same edge: shifting a window removes and adds strips of relative area $1/K$. Thus every endpoint mass changes by at most $1/K < \kappa - e$.

There are $2MK - K + 1$ rows and as many columns of origins. All rows and columns containing no unusable origin form a connected grid: each such row meets each such column, and both families are nonempty once $2MK - K + 1 > b$. At most b^2 origins lie off this grid. Hence all but b^2 windows share one edge. In any child square there are n_c^2 windows; choose M with $b^2/n_c^2 \le \eta$. □

We can now label a child B using only its own occupancies. Let $Q_j(B)$ be its transverse square and $J_F(B)$ the strictly interior heights of parity s for $F = F(j, s)$. In particular $H/2 - 1 \le |J_F(B)| \le H/2$. Set $\kappa_0 = \kappa/4$ and choose

$$0 < \epsilon_0 < \kappa_0, \quad 0 < \gamma < 1/2, \quad 0 < \delta_0 < 1/4, \qquad 2\delta_0/\kappa_0 < \gamma. \tag{4.8}$$

Call F *admissible for* B if $b_{F^c}(B) \le \delta_0$ and each $m \in J_F(B)$ admits an edge $E_m \subset F$ such that

$$a_{E_m^c}(j, m; Q_j(B)) \le \epsilon_0, \qquad a_r(j, m; Q_j(B)) \ge \kappa_0 \quad (r \in E_m). \tag{4.9}$$

The edge at each height is unique since $\epsilon_0 < \kappa_0$. If one edge occurs at a fraction at least $1 - \gamma$ of these heights, label B by that edge; otherwise label it by F. If no face is admissible leave B unlabelled. If there is a choice of admissible faces use the first in a fixed order, and denote the resulting label by $S(B)$.

Lemma 4.10 (Agreement of admissible faces)**.** *If $b_{\mathcal{P}}(B) > 2\delta_0$, every admissible face gives the same label $S(B)$.*

Proof. Two opposite faces cannot both be admissible: their outside fractions sum to $b_{\mathcal{P}}(B)$. If two faces F, G with different normals are admissible, put $E = F \cap G$. Each $E_m \ne E$ in the list for F has an endpoint outside G. Equation (4.9) gives at least $\kappa_0 H^2$ footprint area at that endpoint.

It occupies two physical slices in B, and different heights in this list occupy disjoint slices. Since $b_{G^c}(B) \le \delta_0$, the fraction of such heights is at most $\delta_0 H/(2\kappa_0|J_F(B)|) \le 2\delta_0/\kappa_0 < \gamma$. Thus F gives E; reversing the faces shows that G does also. □

Lemma 4.11 (Ownership implied by a label)**.** *Whenever B receives a label,*

$$b_{S(B)^c}(B) \le t_0, \qquad t_0 = \delta_0 + \epsilon_0 + \gamma + 2/H. \tag{4.10}$$

If this label is an edge E, each endpoint $r \in E$ satisfies

$$b_{\{r\}}(B) \ge (1-\gamma)(1-2/H)\kappa_0. \tag{4.11}$$

Proof. A face label has outside mass at most δ_0 by admissibility. For an edge label E, let F be the face that supplied it. The mass outside F is at most δ_0, and the two boundary centre planes normal to F cover at most $2/H$ of the cube. Among the interior favoured heights, those with $E_m \ne E$ cover at most $2\gamma|J_F(B)|/H \le \gamma$. The rest contribute at most ϵ_0 outside E by (4.9); each contributes at least $2\kappa_0/H$ to either endpoint. This gives both estimates. □

Proposition 4.12 (Labels of the eight children)**.** *On every good parent with face $F(j,s)$, all eight children have labels contained in $F(j,s)$. On each of the two layers of four children perpendicular to j, the four labels are identical. The tolerance t_0 in (4.10) can be made arbitrarily small, keeping $\kappa_0 > 0$ fixed and achieving any prescribed power P in (4.5).*

Proof. A child's outside-F fraction and its vacancy fraction are at most 8δ and $8K_v\lambda$, by the first good-parent condition. For the interior favoured heights of each child, the common list of edges from Proposition 4.9 satisfies (4.9) by (4.7) whenever $4/M + e + \eta \le \epsilon_0$. Thus F is admissible once $8\delta \le \delta_0$. The respective four children with the same j-interval read the very same list at the very same heights; their labels agree. Lemma 4.10 permits any choice of admissible face since $1 - 8K_v\lambda > 2\delta_0$.

For the parameter order, first choose κ, κ_0; next choose $\epsilon_0, \gamma, \delta_0$ satisfying (4.8) and below the desired error. Choose d_0, e, η with $d_0 + e + \kappa < \varepsilon_v$, $e < \kappa/16$, $e + \eta < \epsilon_0/2$. Fix K, then b for the desired P by Lemma 4.7, then M large enough for Proposition 4.9 and $4/M < \epsilon_0/2$. Finally decrease q. All estimates now hold simultaneously. □

4.4 A common layering parity

Index the cubes B_v of side H by $v \in \mathbb{Z}^3$. For this subsection record the label of B_v by the set $S(v)$ of coordinate parities it fixes. Thus a face fixes one pair (j,s) and an edge fixes two. We identify such a set of fixed coordinates with its corresponding face or edge when discussing ownership. The eight indices of the children of P_v form an elementary cube in this index lattice. Proposition 4.12 says that on each good elementary cube some (j,s) belongs to all eight labels, and the labels are constant on each of the two layers normal to j. We show that at least one such coordinate parity persists across the whole lattice, apart from a small exceptional set.

By Lemmas 4.7 and 2.3 the bad elementary cubes are jointly ϵ-rare in every periodic Gibbs measure, with $\epsilon \to 0$ as $q \downarrow 0$. Mark the set D of all vertices of these cubes. Each of n specified marked vertices can be assigned one of eight incident bad cubes, and at least $n/8$ of the assigned cubes are distinct. Hence

$$D \text{ is jointly } \beta\text{-rare}, \qquad \beta = 8\epsilon^{1/8}. \tag{4.12}$$

Lemma 4.13 (A change of label propagates across a square)**.** *Let the four vertices of a coordinate unit square lie outside D. If the labels on one i-bond differ, the opposite i-bond has the same ordered pair of labels, when both bonds are oriented alike.*

Proof. A good elementary cube containing this square obeys the rule above. Since one i-bond changes label, its two constant layers must be normal to i. Moving either end across the square stays within its layer. □

Construct R_1, R_2 from this D as in Section 2.3, enlarging by two index steps at each level.

Lemma 4.14 (At most one direction of extended changes)**.** *For each pair of adjacent coordinate planes of indices, call a transverse position good if both endpoints of the bond between the planes are outside D. There is a unique infinite component of good positions. On it either all these bonds preserve their label or every bond has the same ordered pair of distinct labels. The latter alternative occurs in at most one coordinate direction.*

Proof. A planar $*$-component of bad positions lifts into one $*$-component of D: successive bad endpoints are separated by at most one in each coordinate, including the normal coordinate. Its rectangle on either plane therefore lies in the padded level-one box of that component. By (2.6) and Borel–Cantelli, on every fixed coordinate half-line the upper density of R_1 and of R_2 is less than $1/8$ (use threshold $1/9$ and sufficiently small β). Lemma 2.5 gives the claimed component, containing each position with an endpoint outside R_1. On it a change at any bond propagates with the same ordered labels by Lemma 4.13. For level pairs in different directions, their intersections give the four parallel lines through the corners of a unit square. The density bound supplies a translate of that square with all four vertices outside R_2. A change on the infinite component for either level pair would force equality along the two bonds in the other direction. Thus changes on both components are impossible. □

Proposition 4.15 (A common signed coordinate)**.** *There is a nonempty random set of signed coordinates S_∞ such that $S_\infty \subset S(v)$ for every $v \notin R_2$. More precisely, one can choose an axis k such that in each plane $v_k = t$ the infinite component of the complement of R_1 has constant label U_t and contains every vertex outside R_2. The intersection of all U_t is S_∞, and it contains a signed k-coordinate. If all U_t coincide, any of their signed coordinates may be chosen. The set S_∞ and the exceptional boxes commute with translations.*

Proof. Take for k the direction of changes in Lemma 4.14 if they occur. In a plane normal to k, a planar $*$-component of R_1 meets only one family of joined first boxes; its whole rectangle lies in a level-two box. By Lemma 2.5 and (2.6) the asserted infinite component exists. Any bond in it goes in a direction other than k and belongs to the infinite component for its level pair, hence preserves its label. Successive planes normal to k have corresponding transverse sites outside R_2 by the line bound. Their labels U_t, U_{t+1} are either equal or differ across a good k-bond and then share a signed k-coordinate. Starting at one change and proceeding in either direction, this sign survives every step: equality preserves it and a label cannot contain both signs. If there are no changes take any provisional axis for the argument; all labels outside R_2 then coincide. Taking their signed coordinate as k gives the stated description in that direction as well. The intrinsic intersection and the boxes are translation equivariant. □

We call any member (k, s) of S_∞ a *common parity* and refer to the choice of k in Proposition 4.15 as the first layering direction. By (4.10) all but a small fraction of cells in each cube outside R_2 are either vacant or covered by a head with this parity. Vacancies themselves occupy at most $C\sqrt{q}\,H^3$ cells there, since every such cube belongs to a good parent.

Corollary 4.16 (Layering)**.** *For every $\varepsilon > 0$ and all sufficiently large z, in every periodic Gibbs measure there is almost surely a signed coordinate (j, s) such that*

$$\liminf_{L\to\infty} b_{F(j,s)}(B_L) \geq 1 - \varepsilon.$$

In fact every member of S_∞ has this property.

Proof. Choose $t_0, \theta < \varepsilon/4$. In the index cube $[0, n)^3$, apply (2.6) to the n^2 lines parallel to one coordinate axis. For small q the sum over n of the resulting failure probabilities is finite. Almost surely R_2 therefore occupies at most θn^3 indices for all sufficiently large n. Each remaining cube has $F(j, s)$-ownership at least $1 - t_0 - C\sqrt{q}$ by (4.10) and the vacancy bound above. Sum over the cubes and discard a boundary fraction $O(H/L)$ for general L. Proposition 4.15 ensures $S_\infty \neq \varnothing$. □

5 Partial translational symmetry breaking

Fix a first layering direction as in Corollary 4.16. In each favoured plane we find a connected collection of squares whose origins cover all but a small fraction of the grid. These squares have the same transverse registry. To prove that this registry agrees in successive planes, we compare the probabilities of a cube on their common boundary.

5.1 A transverse registry in each layer

Fix a common parity (k, s) from Proposition 4.15. We now work in the *physical* plane of cube centres $x_k = m$, $m \equiv s \pmod 2$; recall that its area counts the square footprints of cubes centred in that plane. A registry in this plane is an edge of the full parity cube contained in the face $r_k = s$. Our first aim is to assign one edge E_m to the whole plane. At this point it may depend on m.

Fix $\kappa_h > 0$ small enough that Proposition 4.6 allows the following test with arbitrarily small fixed $d, e > 0$. An H-square in this plane is dense if its covered area fraction is at least $1 - d$. It is usable with edge G if in addition the footprints off G have area fraction at most e and each endpoint of G has area fraction at least κ_h. Fix a dyadic integer K_1 with $1/K_1 < \kappa_h - e$. We inspect overlapping H-squares with origins on a grid of spacing H/K_1; any fixed integer translate of this grid is allowed. Neighbouring usable squares have the same edge: a shift by one grid step can change the area fraction of an endpoint by at most $1/K_1$.

Lemma 5.1 (Density in windows meeting a coarse boundary)**.** *Let the origin of an H-square be in the transverse projection of a coarse cube B_v. Use the normal index $v_k = \lfloor m/H \rfloor$. If none of the at most four coarse cubes in this row touched in positive area by the square belongs to R_2, its missing footprint fraction at m is at most $4C(t_0 + \sqrt{q}\,H)$, where t_0 is the right-hand side of (4.10).*

Proof. In each touched cube the outside-(k, s) ownership fraction is at most t_0, and there are at most $C\sqrt{q}\,H^3$ vacant cells. Lemma 4.5, applied on its closed normal interval, bounds the missing area in its transverse square by $C(t_0 + \sqrt{q}\,H)H^2$, including when m is a boundary height. Sum over the touched cubes. Choose t_0 small compared with d and then q small; $\sqrt{q}\,H \to 0$. □

Proposition 5.2 (The registry of a physical plane)**.** *For every sufficiently small fixed $d, e, \theta > 0$ the tolerances in the block labels can be chosen so that the following holds at all sufficiently large z. On the event that (k, s) is common, in every favoured plane m there is a connected collection of usable H-squares with one edge E_m. Its origins contain all sites outside a set A_m of boxes on the grid of spacing H/K_1. These boxes have the diameter and line bounds of Lemma 2.4, with a rarity base tending to zero with q, uniformly in m and in the grid offset. Their bounds hold in the original periodic Gibbs law.*

On expanding transverse squares, the lower limits of the footprint fractions of E_m and of each of its endpoints are at least

$$(1 - \theta)(1 - d - e) \quad \text{and} \quad (1 - \theta)\kappa_h, \tag{5.1}$$

respectively. At side length L the same estimates with an $o(1)$ boundary loss hold eventually, simultaneously for all favoured $0 \le m \le L$.

Proof. Work on the original lattice of coarse indices v. Add a mark at v in the row $v_k = \lfloor m/H \rfloor$ if one of the inspected squares with origin in its half-open transverse square is dense but unusable. There are exactly K_1^2 such origins. A specified square has seminorm at most $\alpha(q) \to 0$ on its $H \times H \times 2$ slab by Proposition 4.6; H is a fixed multiple of the side used there. Include also the original marks D of (4.12), throughout three dimensions. The combined marks are jointly rare with a base tending to zero. Indeed for n indicated indices there are at most $(1 + K_1^2)^n$ assignments of witnesses (old mark or a specified window offset). At least $n/(1 + K_1^2)$ have one type. For an offset type the corresponding slabs are distinct tiles of one chessboard; use Lemma 2.3. For old marks use (4.12).

Form three levels of boxes from the combined marks as in Lemma 2.4, with enlargement eight in coarse units at every level; denote their unions by R_1', R_2', R_3'. The union R_2' contains the one-site neighbourhood of R_2. Indeed each old first box, with one extra site of margin, is inside a new first box (the old margin was two). Touching old boxes are therefore represented in one new cluster, and taking its box with the larger margin proves the assertion at level two. In a coarse square outside R_2', every inspected window anchored there is dense by Lemma 5.1 and usable by the additional mark test.

In the row $v_k = \lfloor m/H \rfloor$, each planar $*$-component of R_2' lies in one family of touching second boxes, hence its entire bounding rectangle lies in a third box. By Lemma 2.5 and (2.6), the complement of R_2' in the row has an infinite connected component containing all sites outside R_3'. All fine-grid origins anchored in this component form a connected set: within one coarse square they form a $K_1 \times K_1$ grid, and adjoining coarse squares join along a side. Neighbouring usable windows agree, so they carry one edge. Take A_m to consist of the fine origins anchored in R_3' on the row. Passing from coarse to fine grid multiplies diameters and line lengths by at most K_1, a fixed number, up to end intervals. The same rooted and line bounds follow with changed constants; for segments shorter than $2K_1$ use the one-site marginal. The only step requiring a common parity was the density implication, so these probability estimates are unconditional.

Choose one fixed offset within each coarse square. The corresponding H-windows tile the physical plane. On each row of these disjoint windows at most fraction θ of the origins is exceptional, except with probability exponentially small in the number of windows by (2.6) on the coarse row. Sum over $O(L/H)$ rows and $O(L)$ heights, then over integer L. At fixed z, H is fixed, so the sum converges. Discarding incomplete boundary windows costs $O(H/L)$. This proves (5.1) by Borel–Cantelli and gives the simultaneous statement. All fixed planes and offsets may be taken together. Different offsets give the same E_m: a private endpoint of either proposed edge would have limiting lower area $(1 - \theta)\kappa_h$, exceeding the upper area $\theta + d + e$ allowed off the other. We take the fixed errors small enough for this strict inequality. □

Lemma 5.3 (Reading aligned physical registries on the coarse grid)**.** *Use the first common normal k of Proposition 4.15. If all its favoured physical planes carry the same edge E in Proposition 5.2, every label outside R_2 is E.*

Proof. Fix a coarse plane $v_k = t$ with constant label U_t on the infinite component of the complement of R_1. Its complement has upper area density at most θ: it lies in R_2, and the line estimate (2.6) may be summed over the rows of an expanding square. Consider the corresponding slab of physical thickness H. If U_t is an edge $G \neq E$, choose $r \in E \setminus G$. By (4.10), the upper ownership of r on the slab is at most $\theta + t_0$: every nonexceptional cube with label G has outside-G ownership at most t_0. But each favoured centre plane strictly inside the slab supplies two physical layers on which these footprints have lower area $(1 - \theta)\kappa_h$ by (5.1). The layers are disjoint and cover at least $H - 2$ of its H levels. This contradicts the upper bound for small θ, t_0 and large H.

Otherwise U_t fixes just (k, s). On the connected component every incident good parent must certify this face: its signed bit must belong to each child label, and here there is only one bit. At any strictly internal favoured height m, adjacent children in the component thus read the same letter by Proposition 4.9 applied to a shared good parent. Denote the common letter by G_m. If $G_m \neq E$, an endpoint $r \in E \setminus G_m$ occupies at most $\theta + \epsilon_0$ of the whole plane in upper area density, by (4.9) on the component. This contradicts (5.1). Thus every letter in each child's word is E, which makes its label E by the reading rule. This excludes the one-bit alternative. □

Proposition 5.4 (Recognizing the registry from its columns)**.** *Let E be an edge with free coordinate w. Suppose that in a periodic Gibbs law almost surely every axial column of centres at the transverse parities of E has lower limiting covered length at least $1 - u$ from its own cubes. For sufficiently small fixed u and errors in the block labels, and sufficiently large z, every coarse label off R_2 is E. Moreover, in each favoured plane of* either *fixed coordinate the connected collection of windows in Proposition 5.2 carries E.*

Proof. Every physical unit cell in the transverse plane belongs to exactly one favoured 2×2 footprint. Thus on any fixed finite physical cross-section times $[0, L]_w$ the lower limiting E-ownership is at least $1 - u$, by summing the corresponding finitely many column estimates. We use this first on the $H \times H$ cross-section of a fixed half-line of coarse cubes parallel to w. For the first common face $F = \{r : r_k = s\}$, the upper ownership off F on this cylinder is at most $\theta + t_0$ by (4.10) and the line bound for R_2. Thus $E \cap F$ has lower ownership at least $1 - u - \theta - t_0$.

This intersection must be all of E. Indeed, a singleton r cannot have nearly full ownership on the cylinder: the singleton-field estimate in the proof of Proposition 4.6 bounds the seminorm of $b_{\{r\}}(B) \geq 1 - \varepsilon_v$ by $\exp(-c\sqrt{q}\,|B|)$ on any rectangular block B, since it uses a field on physical cells and the homogeneous singleton pressure. For B of dimensions H, H, L with dyadic L this is summable. Apply Lemma 2.3 and Borel–Cantelli. Choose $u + \theta + t_0 < \varepsilon_v$; the empty intersection is excluded as well. Consequently k is a fixed coordinate of E with the specified sign.

Fix a favoured plane of k. In a line of disjoint H-windows in that plane along w, all but upper limiting fraction θ carry the physical registry E_m, by the line bound in Proposition 5.2 (or its coarse-row proof). If $E_m \neq E$, a parity in $E_m \setminus E$ therefore covers lower fraction $(1 - \theta)\kappa_h$ along a strip of width H. These same cubes cover both physical slices adjacent to m, contradicting the $1 - u$ lower E-ownership on that fixed $2 \times H$ cross-section as soon as $u < (1 - \theta)\kappa_h$. Hence all these E_m equal E. Lemma 5.3 identifies every coarse label as E; both its coordinates are now common. Apply exactly the same strip comparison to the other fixed coordinate, using its own physical-plane estimates in Proposition 5.2. □

5.2 Two rows of opposite parity in a layer

We now keep the two prescribed rows of Section 3.5 and dispense with its two prescribed columns. An activity penalty for one horizontal centre parity will instead suppress the two-row event. First work in the square gas on an $a \times b$ block, with a, b even, $a, b \geq 4$. Let W require two specified physical rows to be fully covered with opposite prescribed horizontal centre parities. Heads have activities v on the favoured horizontal parity and ve^{-g} on the other. Write $P_{2,W}(v, g)$ for the limsup of the ordinary log partition sum per unit area under full dissemination of W on $2ma \times 2nb$ tori as $m, n \to \infty$, and $P_2(v, g)$ for pressure without W.

Lemma 5.5 (Square pressure with two prescribed rows)**.** *For $u = v^{-1/2}$, $K_* \geq 1$ fixed, $K_* u \leq g \leq 2K_* u$ and $bu \leq c_{K_*}$,*

$$P_{2,W}(v, g) - \tfrac{1}{4} \log v \leq u^2/(4g) + u/(2b) + C_{K_*} b u^2.$$

Proof. Reduce the vertical circumference to $N = 2b$ by (3.6). Take the retained centre columns of (3.7) to be the unfavoured ones. The two rows supply exactly the required and forbidden cells used there. If $D_g(S, S) = e^{-gn(S)/2}$, the transfer after summing the intervening favoured column is $K_g = D_g K(u) D_g$. Recall from (3.9) that $P = P_0 + P_1$ projects onto states of pure vertical parity. Since $n(S) \le N/2$, Taylor's formula gives

$$PK_gP = P + uPK_1P - gP\,\mathrm{diag}(n)P + O_{K_*}((uN)^2)$$

in operator norm. In either parity, the $N/2$ possible dimer positions are disjoint. The required and forbidden cells each fix one of their occupancy bits. By Lemma 3.10, within this parity the first-order correction is a sum, over the free bits, of $\left(\begin{smallmatrix} 0 & u \\ u & -g \end{smallmatrix}\right)$, plus nonpositive contributions from fixed occupied bits. Its largest eigenvalue is at most $(N/2)(\sqrt{u^2 + g^2/4} - g/2) \le Nu^2/(2g)$. Transitions between the two parities add at most $2u$ by the same lemma.

Thus $\|PK_gP\| \le U := 1 + Nu^2/(2g) + 2u + C_{K_*}(uN)^2$, using symmetry and entrywise nonnegativity. For $Q = \mathrm{Id} - P$, (3.9) and $0 < D_g \le 1$ give $\|QK_gP\| \le CuN$ and $\|QK_gQ\| \le 1/2$. Bounding the two blocks by their norms yields $\|K_g\| \le U + 2(CuN)^2$. There are ma transfers in $2ma$ columns. Use $\mathrm{Tr}\, K_g^{ma} \le 2^N \|K_g\|^{ma}$ in (3.6) and divide by area: the excess pressure is at most $(2N)^{-1} \log \|K_g\|$. Substitution gives the claim. □

In the cube gas use coordinates (x, y, w), fix desired parities in x, y, and set $\lambda = z^{-1/2}$. Consider a physical block D of dimensions $a \times 2 \times b$. In one of its two x, w slices impose W on two rows parallel to x. Let $\mathcal{A}$ require W and the following two volume bounds in D: cubes of incorrect x-parity own at most a fraction ε of the cells, and likewise for incorrect y-parity. An incorrect cube in both coordinates contributes to both bounds.

Lemma 5.6 (Reducing the cube test to two square pressures)**.** *For $g_x, g_y > 0$ put $v = z + e^{g_x}/(e^{g_y} - 1)$. Then*

$$\frac{1}{|D|} \log \|\mathbf{1}_{\mathcal{A}}\|_D \le \frac{1}{4}[P_{2,W}(v, g_x) + P_2(v, g_x)] - P_3(z) + \frac{\varepsilon(g_x + g_y)}{8}. \tag{5.2}$$

Proof. On the full dissemination insert activity factors e^{-g_x} and e^{-g_y} for incorrect x and y parities, respectively. Each affected head owns eight cells; summing over the disjoint physical blocks of a torus of volume V costs at most $e^{V\varepsilon(g_x+g_y)/8}$ on the constrained event. Drop the volume bounds and keep the row constraints.

Let A be the area transverse to y, N the normal circumference. The two neighbouring y planes have activity arrays a_x and $a_x e^{-g_y}$ in one order or the other, with $a_x = z e^{-g_x \mathbf{1}_{\{x \text{ incorrect}\}}}$. Exactly $N/2$ physical slices receive the reflected row constraints W_*. These are increasing: inserting a compatible square cannot change the owners of the already covered cells of the two rows. The trace of the product of transfers is bounded by $2^A \|M_{\mathbf{1}_{W_*}}\|^{N/2} \|M_1\|^{N/2}$, since transposition does not affect the norm. Apply (3.3) to bound the two squared norms by the constrained and ordinary square partition sums at activities $\tau(a_x, a_x e^{-g_y})$. By (3.2) these activities are at most $a_x + 1/(e^{g_y} - 1)$, hence at most $v e^{-g_x \mathbf{1}_{\{x \text{ incorrect}\}}}$. Raising the activities increases either partition sum. Take logs, divide by $V = AN$, subtract the original cube pressure and pass to the limit. The exponent $N/2$ and the squared norm in (3.3) give the factor $1/4$ for each square pressure in (5.2). □

Proposition 5.7 (Cost of opposite row parities)**.** *There are absolute $b_0, c_0, \varepsilon_0, c > 0$ such that, for small $q = z^{-1}$, $b \ge b_0$, $b\lambda \le c_0$ and $\varepsilon \le \varepsilon_0$,*

$$\|\mathbf{1}_{\mathcal{A}}\|_D \le e^{-cab\lambda}.$$

Proof. Use (5.2) with $g_x = K_x\lambda$, $g_y = K_y\lambda$ and large fixed K_x, K_y. Then $u = v^{-1/2} \le \lambda$, $u/\lambda \to 1$ and $\log(v/z) \le Cq/g_y$. Lemma 5.5 bounds $P_{2,W}$. For the unmarked P_2, first merge the free coordinate in (3.4), leaving $F(v), F(ve^{-g_x})$. Their gap is at least $vg_x/4$ for $g_x \ge 32u$, as in (3.5). Therefore, with $p_{\mathrm{line}}(u) = \frac{1}{2}\operatorname{arsinh}(u/2)$,

$$P_2(v, g_x) \le \tfrac{1}{4}\log v + p_{\mathrm{line}}(u) + Cu^2/g_x, \qquad P_3(z) \ge \tfrac{1}{8}\log z + \tfrac{1}{2}p_{\mathrm{line}}(\lambda)$$

by (3.4) and Lemma 3.6, respectively. The ordinary square pressure contributes at most $p_{\mathrm{line}}(\lambda)/4$ beyond its reference term in (5.2), while the cube denominator contributes $-p_{\mathrm{line}}(\lambda)/2$. The resulting upper bound per cell is

$$-\tfrac{1}{4}p_{\mathrm{line}}(\lambda) + C\lambda(K_x^{-1} + K_y^{-1} + 1/b) + C_{K_x}b\lambda^2 + (K_x + K_y)\varepsilon\lambda/8.$$

Since $p_{\mathrm{line}}(\lambda) \ge c'\lambda$, first choose K_x, K_y, b_0 large, then c_0, ε_0 small, then q small. The bound is $-c''\lambda$. Multiply by $|D| = 2ab$. □

5.3 The cost of violating two parities

Fix a parity edge E prescribing bits in coordinates j, k, with w the remaining coordinate. We first work in a periodic Gibbs law ν for which the conclusions of Proposition 5.2 hold with label E in *every* favoured j- and k-plane. Take its tolerances sufficiently small, as specified below. We also assume invariance under shifts by two and reflection positivity through all integer planes perpendicular to j or k. The reflection construction of Section 5.5 will produce such a law.

A physical row parallel to j has fixed unit intervals in k, w. Write $X_v = 1$ if a head in the plane $j = v$ covers these two intervals in projection. Its owner covers the two row cells on either side of v; the four possible owners at v are mutually exclusive. At a height v of the parity opposite to the j-bit of E, put

$$A_m = \{X_v = X_{v+2} = \cdots = X_{v+2m-2} = 1\}, \qquad a_m = \nu(A_m), \quad a_0 = 1.$$

Invariance under shifts by two makes a_m independent of the choice of v on this row. We need a short run to be unlikely even when $m \ll H$, where $H \asymp q^{-13/50}$ is the width in the plane estimate.

Lemma 5.8 (A run of wrong parity)**.** *Uniformly over physical rows parallel to either fixed coordinate,*

$$a_m \le \exp(-cq^{1/50}m), \qquad 1 \le m \le H/2.$$

Proof. First take $m = H/2$. Its dimer intervals then cover H cells. Use precisely this interval as the j-side of an $H \times 2 \times b$ block in coordinates j, k, w. Centre its k-side on the favoured plane adjacent to the row. Choose a dyadic multiple b of H with $c_0/2 \le b\sqrt{q} \le c_0$, where c_0 is from Proposition 5.7. Its j, w footprint consists of b/H successive H-squares. With exception probability $Ce^{-cb/H}$, cubes of parity in E from this k-plane cover at least $1 - \delta_1$ of the footprint. Indeed, apply the line estimate of Proposition 5.2 on its grid of spacing H/K_1: retaining every K_1th origin multiplies the maximum exceptional fraction by K_1. Choose that fraction and the errors on the remaining squares with sum (after multiplication) less than $\delta_1 < \min(1/8, \varepsilon_0)$, with ε_0 from Proposition 5.7. These cubes cover the same footprint in *both* physical slices of the block.

By the vacancy estimate (4.2) and Lemma 2.3, the chosen slice has at most $b/4$ vacancies except with probability $e^{-b/4 + C\sqrt{q}Hb}$. Thus at least $3b/4$ rows there are fully covered. Each has one j parity, since successive dimer centres are two apart. At least one of these rows has the prescribed j parity, because the fraction covered by the opposite parity is at most δ_1. On $A_{H/2}$ the specified row has the

opposite parity. Sum the two-row estimate of Proposition 5.7 over at most b choices of the first row, using Lemma 2.3. This gives

$$a_{H/2} \le Ce^{-cb/H} + e^{-b/4+C\sqrt{q}Hb} + be^{-cHb\sqrt{q}} \le C'e^{-c'b/H}.$$

Here $H\sqrt{q} = o(1)$ and $H^2\sqrt{q} \asymp q^{-1/50}$.

To pass to shorter runs, reflection Cauchy–Schwarz gives

$$a_m^2 \le a_{m-1}a_{m+1}. \tag{5.3}$$

For $m = 2s+1$ reflect halfway between consecutive tested heights and use the products of the nearest s and $s+1$ indicators on one side. Their reflected square expectations are a_{2s} and a_{2s+2}. For $m = 2s$, reflect about a tested height, including that height in both products: $X_v^2 = X_v$ gives a_{2s-1} and a_{2s+1}. Thus $a_m \le a_{H/2}^{2m/H}$; the statement also holds at a zero by (5.3). Since $b/H^2 \asymp q^{1/50}$, this proves the estimate. Interchange j, k for the other direction. □

Proposition 5.9 (A head outside its prescribed parity)**.** *Under the assumptions above, if a site u violates $s = 1$ or 2 of the prescribed bits of E, then*

$$\nu(n_u = 1) \le Cq^{19s/25}.$$

Proof. Follow each occupied row through the cube until its first vacancy on either side. This is the nearest-vacancy construction of Hadas and Peled for squares [5, proof of Theorem 9.1]. In three dimensions a violated direction gives four rows, hence eight vacancies. If two directions are violated, these sets are disjoint: the j-end vacancies lie beyond the cube's j-extent, whereas the k-end vacancies lie inside that extent. Opposite ends or different parallel rows also give different cells.

Put $W = \lceil q^{-3/100} \rceil$. If an end vacancy is farther than $2W + 2$ cells, or absent, there are W consecutive wrong-parity dimers on that side starting at a prescribed height. By Lemma 5.8, the union of these alternatives has probability at most $8se^{-cq^{1/50}W}$. Otherwise there are $O(W^{8s})$ choices for the $8s$ vacancies. For any finite set of cells S and any periodic Gibbs law ν,

$$\nu(S \text{ all vacant}) \le \left(q^{1/8}e^{-p_3(q)}\right)^{|S|} \le q^{|S|/8}. \tag{5.4}$$

Indeed, disseminating one vacant unit cell gives the empty torus, of normalized weight $q^{V/8}$ in volume V; apply the unit-cell chessboard inequality and Lemma 2.3. Compare [5, Corollary 4.5] for squares. The resulting bound $8se^{-cq^{1/50}W} + Cq^sW^{8s}$ proves the claim. □

5.4 Repeating line observations by reflection

The bound in Proposition 5.9 on a head violating two parities requires reflection positivity. A periodic Gibbs measure in which two layers disagree need not have this property. We now use reflection to construct a Gibbs measure that retains specified observations along a line.

Lemma 5.10 (Conditioning at negligible entropy cost)**.** *On tori T_ℓ with all sides tending to infinity let $\mathcal{A}_\ell$ be events with $-|T_\ell|^{-1} \log \pi_{T_\ell}(\mathcal{A}_\ell) \to 0$. Condition on $\mathcal{A}_\ell$ and average over translations in a fixed rectangular period lattice, with torus sides divisible by its periods. Every local limit of these averaged laws is a Gibbs measure for the original activity.*

Proof. This is the local form of the Gibbs variational principle (see [4, Chap. 15]); we give the argument for the hard constraint. If ν_ℓ is the averaged law, its relative entropy $D(\nu_\ell \| \pi_{T_\ell})$ is at most $-\log \pi_{T_\ell}(\mathcal{A}_\ell)$, by convexity and translation invariance of π_{T_ℓ}. Fix a finite set of head sites K and

a finite $Y \subset K^c$ containing all its interaction neighbours outside K. Place $N \geq c_{K,Y}|T_\ell|$ disjoint copies of $K \cup Y$ on a sufficiently widely spaced grid of periods, with no interaction between different copies of K. Let U be the union of the K's. Under π_{T_ℓ} their conditional law given U^c is the product of the Gibbs kernels γ_K determined by the respective Y's. The chain rule for relative entropy and subadditivity of conditional entropy give a lower bound on the total relative entropy by the sum of the N corresponding conditional relative entropies. Replacing conditioning on U^c in each term by conditioning on its own Y decreases the bound by convexity; the N terms now coincide. Thus

$$\mathbb{E}_{\nu_\ell} D\big(\nu_\ell(n_K \in \cdot \mid n_Y) \,\|\, \gamma_K(\cdot \mid n_Y)\big) \leq N^{-1} D(\nu_\ell \| \pi_{T_\ell}) \;\longrightarrow 0.$$

The left side is continuous in the finite marginal on $K \cup Y$: use the difference of the two finite Shannon entropies for the conditional entropy, and restrict the log kernel to compatible patterns. In the limit the two conditional laws coincide. Let Y increase to K^c and apply martingale convergence. □

Corollary 5.11 (Repeating a frequently observed block)**.** *Suppose a block test A_ℓ on D_ℓ has probability at least $a_\ell > 0$ in a periodic Gibbs law. If $|D_\ell| \to \infty$ and $|\log a_\ell| = o(|D_\ell|)$, there are arbitrarily large compatible tori on which the full dissemination of A_ℓ has log probability $-o(|T_\ell|)$.*

Proof. Use Lemma 2.3 on tori with the prescribed periods: the thermodynamic seminorm is at least a_ℓ. Choose arbitrarily large tori on which it is at least $a_\ell/2$. By (2.3) the log probability per site of the disseminated event is then at least $\log(a_\ell/2)/|D_\ell|$. □

Fix an axial coordinate w. A line observation is made on the sequence $h = (n_{(x,y,w)})_{w \in \mathbb{Z}}$ at fixed x, y. Its *empirical law* records the frequencies of finite patterns under shifts by two along w (in particular it distinguishes the even and odd sites). In a law invariant under this shift these frequencies converge almost surely for all patterns simultaneously. The limiting law is an ergodic component for the shift by two [4, Chap. 14]. Let ϑ now denote reversal of a sequence through the integer site $w = 0$.

Lemma 5.12 (Preserving line observations)**.** *Take a fixed closed transverse rectangle and finitely many marked transverse positions. At each prescribe a shift-by-two ergodic law m_i for its axial sequence. For successively finer lists of cylinder frequencies, suppose that the tests of closeness to the respective m_i, imposed separately on these lines over $[0, R]$, hold jointly with probability bounded below for all sufficiently large dyadic R at each fixed accuracy, in a periodic Gibbs law. Repeat the rectangle by reflection in each transverse direction, with consistent tests at shared positions. There is a Gibbs law in which each copy of line i has marginal*

$$\tfrac{1}{2}(m_i + \vartheta m_i).$$

It is stationary under axial shifts by two and under transverse periods of the pattern preserving the line tests. It is reflection positive through every transverse integer plane that pairs the off-plane tests and leaves only single-line tests on its fixed planes.

Proof. At accuracy δ keep finitely many words and count their occurrences at even offsets whose supports stay in the interval; denote the probability lower bound by a_δ. Choose R tending to infinity so that $|\log a_\delta|/R \to 0$ and the ratio of the longest word diameter to R tends to zero as the lists increase. Apply Corollary 5.11 to the rectangular block with this axial length. Lemma 5.10 makes the ensuing limit a Gibbs measure. Average over shifts by two along w and over fixed transverse periods preserving the tests. On each indicated line, half of the axial tiles read the words forwards and half backwards; positions within a word diameter of an end contribute a vanishing fraction. Each integer reflection along w differs from ϑ by an even shift. Hence the averaged one-line marginals on

every fixed cylinder converge to the displayed mixture. In particular, any parity head frequency or mean covered fraction invariant under reversal has its prescribed value almost surely on that line, since both laws in the mixture are ergodic and have that value. Countably many lines may be used at once.

For a transverse reflection as in the statement, condition on the two fixed planes of the torus. The indicator of the disseminated constraint is a product of reflected pairs of half-torus factors and factors measurable on the fixed planes. Thus the two strict halves are still conditionally independent reflected copies. Translation averages over the specified periods preserve this assertion of positivity (each translated cut is also a cut of the pattern), and local limits preserve the resulting quadratic inequalities. □

Here is how to choose a finite list of m_i that meets the hypothesis. Start in a Gibbs law invariant under the even axial shift, and take the random vector of the separate empirical laws of the indicated lines. On a positive-probability event specified by their frequencies, choose a point of that event in the support of the restricted vector distribution, with all coordinates ergodic. Such points have full conditional measure. Each neighbourhood of the chosen vector then has positive probability. Almost-sure convergence of the finite-pattern frequencies gives the required lower bound at every fixed accuracy (use a smaller neighbourhood first). This selects the line tests simultaneously while leaving their joint law unspecified.

5.5 Straddling cubes between adjacent layers

Use coordinates (x, y, w), with favoured hosts at $y = 0, 2$. A *straddler* has centre at $y = 1$ and occupies cells in both layers. Vigneshwar et al. [10, eqs. (28)–(35)] computed the first-order cost of misaligning one layer by giving these cubes an auxiliary activity. Here we compare two Gibbs measures of the pure gas that give an identical marginal to one line of straddlers.

Let $L_{rt} = \{(r, t, w) : w \in \mathbb{Z}\}$, and fix an axial residue $a \in \{0, 1\}$. On such a line the *density at parity a* counts heads with $w \equiv a$ per unit length of the whole line. Under a law invariant by two axial steps, its expectation is one half the head probability at $w = a$. The line's *own coverage* is the fraction of axial cells covered by dimers centred on that line; up to $O(1/R)$ on length R, it is twice its total head density.

Lemma 5.13 (Selecting the two host lines)**.** *Suppose that a periodic Gibbs measure gives positive probability to both of the following conditions holding together in the limit of transverse R-squares with fixed origin:*

1. *heads in $y = 2$ with x odd cover a fraction at least $1 - u$;*
2. *heads in $y = 0$ with x even and $w \equiv a$ have density at least $\kappa > 0$ per unit transverse area.*

For sufficiently small $u = u(\kappa)$ there is a Gibbs law μ_*, invariant under periods $(2, 4, 2)$ and reflection positive at $y = 0, 2$, in which almost surely L_{00} has density at parity a at least $b = b(\kappa) > 0$ and L_{12} has own coverage at least $1 - C\sqrt{u}$.

Proof. Average the original law over its finitely many translates by $2\mathbb{Z}^3$. The stated joint event retains positive probability (the identity translate has positive weight). In large even squares the simultaneous finite tests with thresholds $1 - 2u, \kappa/2$ therefore have probability bounded below. The odd x-lines in the upper host occupy disjoint strips of width two. By averaging, all but an $O(\sqrt{u})$ fraction cover at least $1 - C\sqrt{u}$ of their axial intervals. In the lower host a fraction at least $c\kappa$ of the even x-lines have at least $c\kappa R$ tested heads: each line contains at most R heads. Discarding the two boundary strips costs $O(1/R)$. Match each even line to its right odd neighbour. For $C\sqrt{u} < c\kappa/2$ a

positive fraction of these pairs pass both tests. By transverse stationarity their expected fraction equals the probability for the fixed pair L_{00}, L_{12}. Passing to the almost-sure axial empirical limits, with slack in both thresholds, gives positive probability for the required pair of densities.

Apply Lemma 5.12 to separate empirical observations on these two lines in $[0,1]_x \times [0,2]_y \times [0,R]_w$. Under reflection in the tile faces, L_{00} repeats at even $x,\ y = 0 \bmod 4$, and L_{12} at odd $x,\ y = 2 \bmod 4$. The lemma gives an unrestricted Gibbs measure with these line frequencies. About $y = 0$ or 2, the factors strictly on opposite sides are reflected copies and factors on the cut involve only its own heads. Thus the two stated cuts retain reflection positivity. Average by shifts with periods $(2, 4, 2)$, which preserve these families of cuts. □

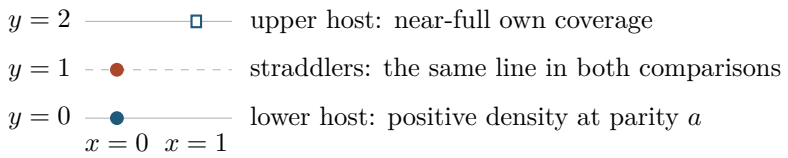


Figure 1: Each marked point represents an axial line parallel to w. The first reflection uses only the two host lines. The next two constructions retain the same observation from the middle line.

Lemma 5.14 (Two continuations with an identical straddler marginal)**.** *Under the conclusion of Lemma 5.13 there are two periodic Gibbs laws ν_0, ν_2, each invariant under all even translations and reflection positive through all integer x, y planes. Their head sequences on the separating line L_{01} have the same law. Under ν_0 the density at parity a on L_{00} is at least b almost surely; under ν_2 every line with x odd and y even has own coverage at least $1 - C\sqrt{u}$ almost surely.*

Proof. In μ_* select a jointly attainable triple of ergodic empirical laws for L_{00}, L_{01}, L_{12} by Lemma 5.12. Write m for the selected line law on L_{01}. The outer two laws have the asserted frequencies almost surely by Lemma 5.13. For either $i = 0, 2$, let H_i be the finite empirical test for that host line and A the test for m on L_{01}. Fix the list of patterns and its tolerance. For all sufficiently large dyadic R, $\mu_*(H_i A)$ is bounded below by a positive constant for that list.

Reflection θ_i about $y = i$ fixes H_i. Reflection Cauchy–Schwarz with $H_i A$ and 1 gives

$$\mu_*(H_i A\, \theta_i A) \geq \mu_*(H_i A)^2.$$

Disseminate this joint event on $[0,1]_x \times [i-1, i+1]_y \times [0,R]_w$. It prescribes identical, separate copies of A on its two y-boundaries at $x = 0$; its one host test is halfway between them. As the pattern list increases and the tolerance decreases, choose R so the logarithm of the probability bound divided by R tends to zero. Lemmas 5.10 and 5.12 then give ν_i.

In both disseminations the middle line of Figure 1 has the same law $(m + \vartheta m)/2$, where ϑ reverses the axial coordinate. The average comes from the two axial reflection orientations. Integer reflections preserve parity and own coverage. The repeated tests consist of A at every even x, odd y, and H_i at even y, with x even for $i = 0$ or odd for $i = 2$. Thus they are paired or fixed at every integer x, y cut. Conditional independence given the cut planes gives reflection positivity for those cuts, also after averaging by all even shifts. □

Proposition 5.15 (Excluding an incompatible pair of hosts)**.** *For every $\kappa > 0$ there is $u > 0$ such that, at sufficiently large z, the two density conditions in Lemma 5.13 have joint probability zero under every periodic Gibbs measure.*

Proof. Otherwise let h be the common probability of a head at $(0,1,a)$ in the two continuations of Lemma 5.14. In ν_0 a head at $(0,0,a)$ has probability $p \geq 2b$. Reflection about $y = 1$ makes the probability of heads at both $(0,0,a)$ and $(0,2,a)$ at least p^2. By the pair replacement (3.1), the intervening head therefore satisfies $h \geq qp^2 \geq 4b^2q$.

In ν_2 all columns of the edge $F = \{x \text{ odd}, y \text{ even}\}$ have near-full own coverage. Choose u sufficiently small to apply Proposition 5.4. It identifies the plane labels in both fixed directions with F, including the line estimates of Proposition 5.2. The site $(0,1,a)$ violates both bits of F. The reflection positivities of ν_2 therefore give $h \leq Cq^{38/25}$ by Proposition 5.9. Since b depends only on κ, the two bounds contradict one another for small q. □

5.6 Agreement of the layer registries

Each favoured plane perpendicular to a common signed coordinate (k, s) carries a label from Proposition 5.2. We regard that label as a full parity edge E_m: its two endpoints both have k-bit s. Distinct edges in that face differ by at least one endpoint.

Lemma 5.16 (Agreement along a common coordinate)**.** *For every sufficiently small fixed $\varepsilon > 0$ and sufficiently large z, on the event that (k, s) is common, the edges E_m are almost surely equal at all favoured heights. For their common value E, both the cubic ownership and the ownership in each such plane have lower limit at least $1 - \varepsilon$; each endpoint has lower limiting footprint fraction at least $\kappa_h/2$ in these same averages.*

Proof. In the host-plane estimate fix the endpoint margin κ_h first. Apply Proposition 5.15 with $\kappa = \kappa_h/8$, obtaining its tolerance $u > 0$. Choose the errors d, e, θ of Proposition 5.2 so that $(1-\theta)(1-d-e) > 1 - \min(\varepsilon, u)$ and $\theta < 1/4$. Its area bound (5.1) yields the plane conclusions with endpoint margin $(1-\theta)\kappa_h$.

If $E_m \neq E_{m+2}$, take a vertex $r \in E_m \setminus E_{m+2}$. In plane m its density in *heads per unit area* is at least $(1-\theta)\kappa_h/4 > \kappa$: each head there covers four transverse cells, up to a boundary error of order $1/R$ on an R-square. Plane $m+2$ has coverage greater than $1-u$ in its own edge. Its fixed transverse bit excludes r. Translating and permuting coordinates puts these two tests in exactly the form of Proposition 5.15. For each fixed m and unequal pair of labels this event has probability zero in the original periodic law. The countable union proves agreement at every height.

For a cube of side R, the simultaneous-height assertion of Proposition 5.2 gives (5.1), up to $o(1)$, in *every* favoured plane in $[0, R]$. Each of its footprints covers the same transverse area in its two adjacent physical slices. Sum over the disjoint pairs of slices, discarding at most the two end slices, to obtain the cubic bounds. □

Proposition 5.17 (One columnar edge)**.** *For every sufficiently small $\varepsilon > 0$, at sufficiently large z every periodic Gibbs measure has almost surely a parity edge E of lower limiting cubic ownership at least $1 - \varepsilon$. In both its fixed directions every favoured plane carries this same edge, with the plane estimates of Proposition 5.2. Every coarse cube B_v with $v \notin R_2$ has label $S(B_v) = E$.*

Proof. Take the preferred common coordinate of Proposition 4.15 and apply Lemma 5.16. By Lemma 5.3, agreement in those planes gives $S(B_v) = E$ for every $v \notin R_2$. The other prescribed bit of E is therefore common as well. Apply Lemma 5.16 in this second direction. Its agreed edge must also be E, since each endpoint of E has cubic ownership at least $\kappa_h/2$ in lower limit, and the complement of the second edge has upper limiting ownership at most $\varepsilon < \kappa_h/2$. □

In particular E is determined by the configuration irrespective of the observation grid: it is the unique edge with lower limiting cubic ownership at least $1 - \kappa_h/4$. Every other edge omits

an endpoint of cubic mass at least $\kappa_h/2$. This description uses a tail event invariant under even period translations (changing finitely many heads, or shifting a box by a fixed even vector, changes its ownership by $o(1)$). Conditioning on any of the twelve events of positive probability therefore preserves the Gibbs specification and periodicity. A limit of symmetric torus measures gives at least one such sector, and lattice symmetries carry it to all twelve.

5.7 Area cost of an incompatible pair of finite squares

In a block Q_R with transverse sides R, R and normal interval $[0, 2]$, let $M_R(F, r; d, k)$ denote the following event. In plane 0, heads with transverse parity on a given parity-square edge F cover at least $1 - d$ of their square. In plane 2 there are at least kR^2 heads with a given transverse parity r; count centres in the interior, so differences from clipped area divided by four are $O(R)$. These tests leave all other heads unrestricted.

Proposition 5.18 (Area cost for an incompatible pair)**.** *If $r \notin F$, then for each $k > 0$, all sufficiently small fixed $d > 0$, and all sufficiently large z, there is $c = c(z, d, k) > 0$ such that, for all sufficiently large dyadic R,*

$$\|\mathbf{1}_{M_R(F,r;d,k)}\|_{Q_R} \le e^{-cR^2}.$$

Proof. For $r \notin F$ suppose that along a sequence $R \to \infty$ the negative logarithm of the seminorm is $o(R^2)$. Along the torus limsup defining the norm choose tori much larger than R on which the full dissemination has logarithmic probability $-o(V)$, since there are $V/(2R^2)$ tiles. The first test then holds on all squares in planes 0 mod 4, and the second in planes 2 mod 4. Condition on the dissemination, average under even transverse shifts and normal shifts by four, and take a subsequential limit ν. By Lemma 5.10, ν is a periodic Gibbs measure of the original gas.

Let U be the limiting F-coverage in plane 0 and V_2 the limiting r-head density in plane 2, on transverse squares. Both exist by the periodic ergodic theorem. Summing the constraints over the disjoint square interiors in each tested plane and then averaging over translates gives $\mathbb{E}_\nu U \ge 1 - d$, $\mathbb{E}_\nu V_2 \ge k$: after averaging these are expectations of bounded local observables over a transverse period cell, so they pass to the limit. Since both averages lie in $[0, 1]$,

$$\nu(U \ge 1 - \sqrt{d},\ V_2 \ge k/2) \ge k/2 - \sqrt{d}.$$

For $0 < k < 1$, take $\sqrt{d} < k/2$ and smaller than the tolerance in Proposition 5.15 at threshold $k/2$. The intersection would have positive probability, contrary to that proposition with the roles of the two heights interchanged. Thresholds $k \ge 1$ describe an empty event. This proves the area bound; Lemma 2.3 gives it also for the probability of one such block in every periodic Gibbs measure. □

Corollary 5.19 (The test for a compatible pair)**.** *If $r \in F$ and $0 < k < \kappa_h/8$, the seminorm of $M_R(F, r; d, k)$ on Q_R tends to one for each sufficiently small fixed $d > 0$ and sufficiently large z.*

Proof. Take a periodic state in the sector with transverse label F and even normal bit. Proposition 5.17 and Lemma 5.16 give the first inequality with slack and limiting r-head density at least $\kappa_h/8 > k$ in the second plane. Thus its probability tends to one; it is bounded above by its seminorm by Lemma 2.3. □

6 Columnar order

Proposition 5.17 gives twelve disjoint events, one for each choice of two transverse parities. We show that each event supports exactly one periodic Gibbs law and determine its correlations along the remaining axis.

6.1 From a registered layer to individual axial lines

Fix an edge E and two independent configurations sampled from periodic Gibbs measures in its sector, as defined in Proposition 5.17. They need not have the same law. Write x, y for the fixed coordinates and w for the free coordinate of E. A head is *wrong in* x if its x-parity is opposite to the one fixed by E, and likewise for y. Join two centres of the disagreement set (present in exactly one sample) when their Chebyshev distance is one. We will bound the radius of a component in this graph. On a region without wrong heads a path of disagreements stays on one favoured column; a vacancy can then stop it. The first task is to bound where a wrong head could allow the path to change columns.

Use $\lambda = \sqrt{q}$ and the cubes of side $H \asymp q^{-13/50}$ of Lemma 4.7. On any rectangular grid let $N_a(i)$ be the indices at Chebyshev distance at most a from i. We also use $N_a(i)$ for the union of their blocks when discussing physical regions. As in (2.5), a set of marks has base β if m prescribed distinct indices are all marked with probability at most β^m. Exceptional *regions* will be made by repeatedly enclosing the marked components in boxes, as in Lemma 2.4; it is the initial marks, before enclosure, that have this joint probability bound.

Lemma 6.1 (Control in every favoured plane of a cube)**.** *Given $\delta > 0$ and $P < \infty$, one can choose the fixed parameters of the H-grid and marks of base q^P, for all sufficiently small q, with the following property. Outside the third enclosure of these marks, in every favoured x- or y-plane meeting a closed H-cube, its H-square has fraction at least $1 - \delta$ covered by footprints of heads from that plane with parity in E.*

Proof. Choose the tolerances d, e, κ for the square test (4.4): a usable square has missing fraction at most d, fraction off its edge at most e, and mass at least κ on each endpoint. Keep $\kappa > 0$ fixed and take d, e small compared with δ, κ. Fix dyadic K' with $1/K' < \kappa - e$. Choose an integer B sufficiently large as below, then a dyadic M' large enough that $1/M'$ and $B^2/(M'K' - K' + 1)^2$ are small compared with δ, κ. We inspect windows of side $h' = H/M'$ on mesh h'/K'. Finally choose the original layer test in (4.10) with error $t_0 \ll d/(M')^2$. Enlarge the original dyadic multiple M to at least M'; thus h' is a fixed multiple of the original fine side h. The edge reported by the new choice of parameters is still E by the intrinsic sector criterion after Proposition 5.17. The earlier dense-plane estimate (Proposition 4.6) applies to h' with the same exponent.

For each H-cube inspect its concentric box of side $5H$. Mark the cube if, at some strictly internal height in one direction, more than B of the windows in that plane are dense but unusable. To bound joint occurrences, retain one congruence class of indices with disjoint expanded boxes. There are $O(H)$ choices of direction and height in each chosen box and only a constant number of possible lists of B window origins, since $H/h' = M'$ is fixed. Of all retained witnesses at least the fraction $1/(6K'^2)$ belong to one $h' \times h' \times 2$ chessboard (choose normal, centre parity and transverse mesh residues). They are distinct. The bound per chosen box, before the fixed colouring loss, is consequently

$$CH\,(Cq^{c_h})^{B/(6K'^2)},$$

where $c_h > 0$ is the fixed exponent in Proposition 4.6. The transfer to periodic laws uses Lemma 2.3. Choose B large enough to give any required power of q. Include cubes incident to failed original parents, also with arbitrarily high power by Lemma 4.7. The union still has base q^P: on a specified set at least half have the same type, and there are at most 2^m assignments of types to m indices.

Take enclosures with fixed padding, say 20 index steps at each of the first three levels. Outside the second enclosure the patch of side $5H$ and its neighbours avoid the original second exceptional region: the starting marks include the old ones, and at each enclosing step the larger padding includes their fixed neighbourhood. Proposition 5.17 therefore gives report E on these neighbouring H-cubes,

in particular wrong ownership at most t_0 for each fixed coordinate. Plane conservation (Lemma 4.5, including boundary heights) bounds the missing area of each favoured plane by $C(t_0 + \lambda H)H^2$ in each such cube. A window of side h' meets at most four transverse H-squares. Its missing fraction is at most $4C(M')^2(t_0 + \lambda H) < d$.

Now fix one favoured plane. Work at the coarse normal index containing its height, using either index at a boundary. All its windows in a patch outside the second enclosure are dense, so at most B are unusable. Rows and columns of window origins free of unusable sites form a connected set omitting at most B^2 origins. Overlapping usable windows one mesh step apart have the same edge (the change of any one parity mass is at most $1/K'$). By Lemma 4.8, in every aligned inscribed H-square this edge has both endpoint fractions at least $\kappa/4$ and off-edge fraction at most $4/M' + e + B^2/(M'K' - K' + 1)^2$. Adjacent patches share an entire such square; these bounds give them the same edge. By Lemmas 2.5 and 2.4, the planar complement of the second enclosure has a connected component containing everything outside the third. The line bound of the latter lemma and a union bound over the lines of a growing planar square show that this component has lower area density arbitrarily near one for small q. Its common edge is E: if it differed, one of its endpoints outside E would occupy a positive fraction of this plane, contrary to the near-total E-coverage in Proposition 5.17 with error smaller than $\kappa/8$. The missing and off-edge fractions sum to at most δ by the choices above. This argument applies on a single probability-one event to all the countably many planes. □

Lemma 6.2 (An estimate on each axial line)**.** *For any fixed $0 < \eta < 1/8$ the marks in Lemma 6.1 can be enlarged, still with arbitrary prescribed base q^P, so that the following holds off their third enclosure. On* each *physical line parallel to w through the H-cube, at most ηH cells have owners wrong in x, and at most ηH have owners wrong in y.*

Proof. For a slab of dimensions $a \times 2 \times b$ centred at a favoured y-height (and with x, y interchanged), consider either of its physical $a \times b$ slices. Mark a failure if its vacancy count exceeds s, or if the whole slab has at most the tolerance ε_0 of Proposition 5.7 in each wrong parity and this slice has two fully covered x-rows of opposite parity. By the vacancy tilt (4.2) and that proposition the sum of the two seminorms is at most

$$e^{-s+C\lambda ab} + 2b^2 e^{-c\lambda ab}. \tag{6.1}$$

For $a = b = H$, $s = \eta H$, this is smaller than every fixed power of q: $\lambda H \to 0$ and $\lambda H^2 \asymp q^{-1/50}$. Add a mark on each H-cube containing a slice of this test that fails, including the adjoining slab at boundary heights. There are $O(H)$ choices per cube and a fixed number of grids (orientation and centre parity); colouring and Lemma 2.3 give arbitrary polynomial joint bounds as above. Enlarge the marks *before* taking their enclosures.

Apply Lemma 6.1 with $\delta < \min(\varepsilon_0, 1/4)$. The registered host covers both slices of its slab, so each has wrong fraction at most δ. At most ηH of the H rows parallel to x in one slice have a vacancy. At least one of the remaining rows must have correct x-parity, since otherwise the wrong fraction is at least $1 - \eta$. By the second test in (6.1) every fully covered row has that same parity. At a *fixed* transverse location, wrong x-ownership along w can thus occur only at the at most ηH row heights containing a vacancy. Exchange x, y for the second bound. □

6.2 Confinement of disagreement paths

Put $t = \log(1/q)$ and $\xi = q^{-1/2}$. Besides the H-cubes, use the following dyadic lengths:

$$W \asymp t^4, \qquad b \asymp \xi t^{-5/2}, \qquad T \asymp C_0 \xi t, \qquad W \mid H \mid b \mid T.$$

We choose C_0 below and then fix a large dyadic constant D. Partition the centres into smaller boxes $W \times W \times T$ and larger boxes $H \times H \times DT$, using half-open integer intervals and common grid origins. Distances between indices always refer to the grid being discussed. The cost $\lambda W b \asymp t^{3/2}$ makes the strip failure in (6.1) very unlikely. The length T allows a vacancy to interrupt paths along a column. This is the rectangle criterion of Hadas and Peled [5, Sec. 8.3] applied to columns with two fixed coordinates: the following estimates exclude side exits in both transverse directions.

For the second test below we use Corollary 3.2 and Lemma 3.6: m prescribed cells are all covered with probability at most $e^{-mp_3(q)}$, and $p_3(q) \geq c/\xi$.

For a smaller box J set $J^+ = N_2(J)$ on its own grid. For one configuration mark J in either of the following circumstances.

1. For a centre u in J^+ and one fixed coordinate $i = x, y$, each of the four physical i-rows through its prospective cube has a vacancy on each side of that cube, within $3W + 2$ cells. All eight vacancies must be present.

2. For some favoured column through J one of the two axial intervals of length T immediately preceding and following J has no vacancy along a prescribed physical line in its footprint (for a column centred at (x_0, y_0) use the cell with lower corner $(x_0 - 1, y_0 - 1)$).

Let F be the set of smaller boxes so marked in either sample.

Lemma 6.3 (Rarity on the smaller grid)**.** *The set F has base q^{c_1} for some fixed $c_1 > 0$.*

Proof. For the first test there are $O(W^2T)$ candidate centres, two directions and $O(W^8)$ choices for the eight distinct cells. Equation (5.4) gives $CqW^{10}T = q^{1/2-o(1)}$. For the second test Corollary 3.2 gives $CW^2e^{-cT/\xi}$; choose C_0 large. For a prescribed collection of boxes, colour their indices modulo a fixed integer in each direction so the witnesses of one colour use disjoint cells. In the largest class retain at least a quarter with the same test and sample, summing over the at most 4^m assignments. For separated boxes the sums over witnesses multiply because the corresponding unit-cell bounds hold for their union. These fixed losses give the stated base. □

Apply (6.1) also with $a = W$, $s = b/4$ to every $W \times 2 \times b$ slab; use the aligned W, b grids and every favoured normal height. Call a smaller box J *regular* if every such slab meeting $N_6(J)$ passes both tests in both samples and every H-cube meeting any of these slabs is outside the third enclosure of Lemma 6.2 in both samples. In a regular region a wrong head requires the first mark: by $H \mid b$ the per-line bound gives wrong fractions at most η separately in every slab and slice. Choose $\eta < \varepsilon_0$. At least $3b/4$ rows in the slice are fully covered, and all of these have the correct parity by the same reasoning as in Lemma 6.2. An unbroken run of wrong dimers along a row from the head extending beyond $3W + 2$ would therefore contain a forbidden fully covered aligned W-interval. This holds for all eight directions of travel in its four rows.

Consequently, if J is regular and $J \notin F$, every head in J^+ in both samples belongs to a favoured column. A disagreement path through J must follow a single such column in J^+, since distinct favoured columns have transverse distance at least two. In each of the two adjacent intervals the second test supplies a vacancy in the first sample. Neither centre bordering that cell belongs to the first sample, and the second sample cannot contain both. Thus there is a centre absent in both on each side of J. The entire disagreement component meeting J stays on its one column and has axial span at most $3T + 2$.

Lemma 6.4 (Confinement on the larger grid)**.** *There is a set of larger-box indices of base q^{c_2}, $c_2 > 0$ fixed, with the following property. If i is outside its third enclosure, every disagreement component meeting i stays in $N_1(i)$.*

Proof. Mark a larger box i on any of the following observations: (a) it contains an initially marked H-cube from Lemma 6.2 in either sample; (b) some $W \times 2 \times b$ slab meeting $N_4(i)$ fails a strip test in either sample; (c) inside $N_2(i)$ there is a connected set of $n = \lfloor D/2 \rfloor$ smaller boxes, all in F. Here and below smaller boxes are connected by Chebyshev steps of their indices.

For (a) there are $DT/H \le q^{-1}$ candidate indices; the exponent P in Lemma 6.2 is arbitrary. For (b) each large box has polynomially many tests, each with seminorm bounded by (6.1), smaller than every power of q because $\lambda W b \asymp t^{3/2}$ and $\lambda W = o(1)$. On separated N_4-neighbourhoods the witnesses use distinct slabs. They lie on a fixed number of grids (choose normal and its parity); their locations on these grids are counted in the polynomial number of choices. For (c), a connected set of n indices containing a given one has at most 26^{2n} possibilities, by following a tour of a spanning tree. The F bound of Lemma 6.3 therefore gives

$$CD(H/W)^2\, 26^{2n} q^{c_1 n}.$$

For separated N_2-neighbourhoods the witness indices are distinct, so the same bound per large box holds jointly. Since $(H/W)^2 \le q^{-1}$, take D large depending only on c_1, then q small. Colouring overlapping observations and summing over the three types gives base q^{c_2} for their union.

We specify how the enclosures are compared across grids. In coarse H-cube indices the projection onto larger-box indices is $(i_x, i_y, i_w) \mapsto (i_x, i_y, \lfloor i_w/(DT/H) \rfloor)$. It preserves connectedness and commutes with taking a coordinate bounding box. Each step of padding by k coarse units projects to at most k large-grid units. At each of three enclosing steps take the large-grid padding at least the coarse padding plus ten. Induction then shows that the projected coarse exceptional boxes, with their four-step large-grid neighbourhoods, lie inside the corresponding large-grid enclosure: images of touching boxes still touch, and a bounding box contains all the boxes of its cluster. Thus if i is outside the third large enclosure, every smaller box of $N_2(i)$ is regular. Here the slabs and coarse cubes needed for regularity lie in $N_4(i)$ because the additional distances beyond $N_2(i)$ are at most $7W + H + 4 < 2H$ transversely and $6T + b + H + 4 < DT$ axially.

Follow a path from i until it first leaves $N_1(i)$. If it meets a box outside F, regularity confines the whole component to one column with axial span at most $3T + 2$. It cannot cross a transverse margin of H or an axial margin of DT. Otherwise the visited smaller boxes form a connected subset of F in $N_2(i)$, containing at least $\min(H/W, D) - 1 \ge n$ distinct indices: a path leaving $N_1(i)$ must cross an entire larger box in some coordinate. This gives mark (c) at i, also impossible. □

Proposition 6.5 (Radius of a disagreement component)**.** *Let $D(u)$ be the component of the centre u in the disagreement set, empty when the two samples agree there. Uniformly over the two periodic laws in sector E, for all $r \ge 0$ and all sufficiently large z,*

$$\Pr\{D(u) \text{ contains } v \text{ with } \|v - u\|_\infty \ge r\} \le Ce^{-cr/\sqrt{z}}.$$

Proof. If $r > 6DT$, a component reaching this far has diameter greater than that of any three-by-three-by-three union of larger boxes. Lemma 6.4 puts *every* box it visits inside the third enclosure. These visited indices form one connected set, hence lie in a single fourth enclosure box containing the index of u. Its diameter in index units is at least $cr/(DT)$. The diameter bound of Lemma 2.4 gives $C(Cq^{c_3})^{cr/(DT)} \le C'e^{-c'r/\xi}$, since $DT \asymp \xi \log(1/q)$ up to a fixed constant.

For $1 \le r \le 6DT$ let J be the smaller box containing u. With probability at least $1 - Cq^{c_4}$ its larger box lies outside the third enclosure and $J \notin F$, by the marginal bounds of Lemmas 6.3 and 2.4. Then a component meeting u follows one favoured column. If it reaches distance r, along that column it contains a segment of $\lfloor r \rfloor$ consecutive steps from u in one direction. The two samples must alternate their heads on this segment by exclusion. Each physical axial cell between consecutive centres, at a

fixed transverse cell of the column, is therefore covered in the first sample. For a favoured u this is one of two *prescribed* segments; for other u the event is impossible on the indicated regular boxes. By Corollary 3.2 the desired probability is at most $Cq^{c_4} + 2e^{-p_3(q)\lfloor r \rfloor}$. Since $r/\xi \le CDC_0 \log(1/q)$, both terms are at most $C'e^{-c'r/\xi}$ after decreasing c'. For $0 \le r < 1$ enlarge C. □

6.3 Uniqueness, symmetry and exponential decay

The passage from disagreement paths to Gibbs measures uses van den Berg's criterion [2, Theorem 1]. We state the form for local correlations as well, following Hadas and Peled [5, Theorem 8.2 and Lemma 8.8]. Its short proof for our exclusion graph is in Appendix A.

Lemma 6.6 (Interchange of finite disagreements)**.** *If independent samples of two Gibbs measures μ, ν have only finite disagreement components almost surely, then $\mu = \nu$ and this measure is extremal among all Gibbs measures. For two independent samples of this measure and bounded observables f, g depending respectively on finite centre sets A, B,*

$$|\operatorname{Cov}_\mu(f,g)| \le 2\|f\|_\infty \|g\|_\infty \Pr(\textit{one disagreement component meets both } A \textit{ and } B).$$

Proof of the phase classification and covariance bound. Proposition 6.5 excludes an infinite component at each centre and hence anywhere. Lemma 6.6 implies that each sector has at most one periodic Gibbs law and that this law is extremal. Proposition 5.17 decomposes every periodic Gibbs measure into its disjoint sectors; a symmetric torus limit and cubic symmetries give a law in each of the twelve sectors. Write μ_E for the unique one in sector E. Its weight in any periodic mixture is the probability of the intrinsic tail event of that sector. A lattice isometry g permutes the sector events, including when it reflects a coordinate. Indeed choose an even common period a; the ergodic theorem for $a\mathbb{Z}^3$ gives, in each coordinate orthant, the same almost-sure limit of the ownership averages: the conditional expectation of the ownership in a period cell given the invariant sigma-field. Changing the origin by a bounded amount changes the average by $O(1/L)$. Thus g sends the sector of E to that of gE, and $g_*\mu_E = \mu_{gE}$ by uniqueness. Since distinct edges give distinct laws, the stabilizer of μ_E is exactly the stabilizer of E; it contains the unit translation along w.

For the covariance bound in Theorem 1.2 use $\mu = \nu = \mu_E$. A component meeting A and B reaches distance at least $\operatorname{dist}_\infty(A,B)$ from one centre of the smaller support. Sum Proposition 6.5 over that support and apply Lemma 6.6. □

6.4 Pressure and reflection symmetry within a phase

Let $p(q) = p_3(q)$, $\lambda = \sqrt{q}$, $\xi = 1/\lambda$, and let μ_E denote the unique periodic state with edge E. We first compute the density of vacancies. This fixes the scale of the axial lower bound below.

Lemma 6.7 (Leading pressure)**.** $p(q) = \sqrt{q}/8 + o(\sqrt{q})$ *as* $q \downarrow 0$.

Proof. By Proposition 5.17 and (2.1), for any fixed sufficiently small $\varepsilon > 0$ the mean density of off-E heads in μ_E is at most $\varepsilon/8$ for q small, averaging over one period. Apply (2.2) with a penalty g on those heads and use Lemmas 3.6 and 3.7:

$$p_{\mathrm{col}}(q) \le p(q) \le p_E(q,g) + g\varepsilon/8 \le p_{\mathrm{col}}(q) + Cq/g + g\varepsilon/8.$$

With $g = \sqrt{q/\varepsilon}$, first send $q \downarrow 0$, then $\varepsilon \downarrow 0$. The independent-column pressure is $p_{\mathrm{col}}(q) = \frac{1}{4}\operatorname{arsinh}(\sqrt{q}/2) \sim \sqrt{q}/8$. □

Corollary 6.8 (Vacancy density)**.** *Uniformly over periodic Gibbs measures the vacancy probability averaged over their periods is* $v_{\mathrm{av}} = (1/2 + o(1))\sqrt{q}$.

Proof. Set $s = \log z$ and $f(s) = p(e^{-s})$, a convex function. If d is the period-averaged head density, (2.2) for constant positive and negative fields says that $d - 1/8$ is a subgradient of f. Hence for every $a > 0$,

$$\frac{f(s) - f(s+a)}{a} \le 1/8 - d \le \frac{f(s-a) - f(s)}{a}.$$

Divide by $e^{-s/2}$, take $s \to \infty$ using Lemma 6.7, and let $a \downarrow 0$. Both secant bounds tend to $1/16$. Now $v_{\mathrm{av}} = 1 - 8d$. □

Lemma 6.9 (Reflection positivity in a specified phase)**.** *The state μ_E is reflection positive through every integer coordinate plane.*

Proof. On symmetric tori give the heads in E the additional factor e^h, $h > 0$. These activities preserve all integer reflections, so the conditional-half argument of Section 2 applies. Take a torus limit ν_h, then a subsequential local limit as $h \downarrow 0$. Continuity of the local Gibbs kernels gives a periodic Gibbs measure of the original gas; positivity survives both limits.

Write $d_E(\nu)$ for the period-averaged density of E-heads. Applying (2.2) in the two directions, first in any unperturbed periodic state ν and then in ν_h with the field removed, gives $h d_E(\nu) \le p_3[h\mathbf{1}_E] - p_3[0] \le h d_E(\nu_h)$. Consequently the limit maximizes d_E. By the periodic decomposition just proved, its unique maximizer is μ_E: that phase has E-ownership near one, whereas each other phase has near-full ownership on its own edge, with equal mean mass at its two endpoints by unit translation along its free axis. At most one of these endpoints belongs to E. This identifies the reflection-positive limit. □

6.5 Correlations along a column

Fix one physical unit square in the transverse plane. Let D_j be the indicator that a head in the plane $w = j$ covers that square in projection. It is the sum of the indicators of the four possible heads, which exclude one another. Write $C_m = \mathrm{Cov}_{\mu_E}(D_0, D_m)$. The positive inverse length in the following statement is defined by the large-distance limit at *fixed* z.

Proposition 6.10 (Even axial correlations)**.** *All C_{2m} are positive. Their exponential rate exists and satisfies*

$$c/\sqrt{z} \le \alpha(z) := \lim_{m\to\infty} -\frac{1}{2m} \log \frac{C_{2m}}{C_0} \le (1 + o(1))/\sqrt{z}.$$

Proof. Reflection in an integer head plane interchanges the two parities of the transverse cell *lower endpoint*. Thus vacancy probability v on our line equals v_{av} by the symmetries of μ_E. The mutually exclusive events D_j, D_{j+1} cover exactly the event that the intervening cell is occupied. Unit axial stationarity gives $p = \mu_E(D_j) = (1 - v)/2$, $C_0 = p(1 - p)$ and $C_1 = -p^2$.

For $F_j = D_j - p$ the Gram matrix of the form $(F_i, F_j) \mapsto \mu_E(F_{-i}F_j) = C_{i+j}$ for $i, j \ge 0$ is positive semidefinite by Lemma 6.9. Its minors give $C_1^2 \le C_0 C_2$ and $C_{2m}^2 \le C_{2m-2}C_{2m+2}$. Consequently

$$C_{2m}/C_0 \;\ge\; \left[(1 - v)/(1 + v)\right]^{2m} > 0.$$

The logarithms are convex and bounded above by zero, so their secant slopes from zero have a limit. Corollary 6.8 bounds $\alpha(z)$ above by $\log[(1 + v)/(1 - v)] \sim \sqrt{q}$. The covariance bound proved from Proposition 6.5 and Lemma 6.6 bounds it below by $c\sqrt{q}$. □

Lemma 6.11 (The one-dimensional comparison)**.** *For the nearest-neighbour exclusion chain of activity z, let $\Lambda_\pm = (1 \pm \sqrt{1 + 4z})/2$, $\rho = |\Lambda_-|/\Lambda_+$. An interval whose two ends are at distance at least d from the prescribed exterior of a longer interval has marginal probabilities within relative error*

$O(\rho^{d-1})$ *of the stationary chain, uniformly in legal boundary states, for* $z \geq 1$ *and* ρ^{d-1} *small. In the stationary chain the covariance of the head indicators at separation* t *is* $p_1(1-p_1)(-\rho)^t$*, with* $p_1 = (\Lambda_+ - 1)/(2\Lambda_+ - 1)$*.*

Proof. Use the symmetric transfer matrix with rows $(1, \sqrt{z})$ and $(\sqrt{z}, 0)$ as in the one-dimensional hard-square calculation [8, Eqs. (10)–(12)]. Its eigenvalues are $\Lambda_\pm$. Both coordinates of the positive unit eigenvector u_+ are bounded below uniformly for $z \geq 1$. Any nonzero nonnegative boundary vector b propagated by $d-1$ steps, divided by $\Lambda_+^{d-1}\langle b, u_+\rangle$, equals u_+ up to coordinatewise relative error $O(\rho^{d-1})$. Inserting this at both ends of the short interval proves the bound for each legal pattern after normalization; the forbidden patterns have zero mass under both laws. The stationary chain has one-site marginal u_+^2 coordinatewise and second transition eigenvalue $-\rho$; its centred one-site functions form a one-dimensional eigenspace. These facts give the stated mean and covariance. □

Proposition 6.12 (The dimer profile of the cube gas)**.** *On* N *distinct favoured head columns, observe the occupancy sequences on a common axial interval* I *of length at most* $C_1\xi\log(1/q)$*. Their joint marginal in* μ_E *is within* $C_{C_1}Nq^{13/50}\log(1/q)$ *in total variation of* N *independent stationary dimer chains of activity* z*. In particular the axial head covariance has the scaling limit of Theorem 1.2.*

Proof. By Lemma 6.9 and the host conclusion of Proposition 5.17, the estimate for one wrong parity in Proposition 5.9 applies: any head outside the favoured columns has probability at most $Cq^{19/25}$ per site. Extend I at each end by $\ell = \lceil 8\xi\log(1/q)\rceil$ sites to get I^+. Condition on all sites except the N selected columns over I^+. At most $9N(|I^+|+2)$ exterior sites are within interaction range of these columns. Except with probability $CN(|I^+|+2)q^{19/25}$ none of these sites carries a head of wrong transverse parity. On such an exterior the conditional law factorizes into N one-dimensional finite-interval kernels with their two prescribed ends: different favoured columns have transverse separation at least two. Lemma 6.11 now costs $O(N\rho^{\ell-1})$ in total variation, with $-\log\rho = 2\,\mathrm{arsinh}(\sqrt{q}/2) \sim \sqrt{q}$. Combine the bounds using $19/25 - 1/2 = 13/50$. For $N = 1$ and $t/\xi \to s$, the limiting signed covariance is $\frac{1}{4}e^{-s}$. □

Corollary 6.13 (Density of the reversible straddler)**.** *Suppose a site* u *violates just one fixed coordinate of* E*, say* j*. In* μ_E *the event* $\mathcal{R}_u$ *that a cube at* u *can be replaced by two cubes centred at* $u - e_j, u + e_j$ *has probability* $(1/4 + o(1))/z$*.*

Proof. By the exact replacement identity (3.1), $\mu_E(\mathcal{R}_u) = q\,\mu_E(n_{u-e_j} = n_{u+e_j} = 1)$. The two heads belong to different favoured columns at the same axial height, so Proposition 6.12 gives joint probability $1/4 + o(1)$. □

A Interchanging two configurations on finite components

Proof of Lemma 6.6. For a finite set A of centres and a finite box $V \supset A$, swap the two configurations on all disagreement components meeting A if they are all contained in V; otherwise leave the configurations unchanged. The rule is determined inside V and its one-site collar, by exploring from A up to the first exit. The disagreement set is unchanged by the swap, so the map is an involution. It cannot introduce an incompatible pair: an adjacent disagreeing centre is in the same component and swaps with it, while a head present in both configurations was already compatible. At fixed exteriors the sum of the two counts of heads in V is unchanged. The product of the two conditional Gibbs laws, whose weight is z to this sum, is therefore invariant. This is the finite-component interchange of [5, Lemma 8.8] on the cube exclusion graph.

If all components are finite, the components meeting A lie in V with probability tending to one as V exhausts the lattice. Thus the two marginals on A agree. This holds for every finite A, so $\mu = \nu$.

Suppose $\mu = s\rho + (1-s)\rho'$ with $0 < s < 1$ and ρ, ρ' Gibbs. Then $\mu \otimes \rho$ is absolutely continuous with respect to $\mu \otimes \mu$, so it too has only finite disagreements. The same interchange gives $\mu = \rho$ and hence $\mu = \rho'$, proving extremality.

Finally, for two independent copies the covariance is half the expectation of $(f(\sigma) - f(\sigma'))(g(\sigma) - g(\sigma'))$. On the event that no disagreement component meets both supports, swapping the components meeting A changes the sign of the first factor and leaves the second factor and the event unchanged. Use the finite-box involutions and bounded convergence to obtain zero contribution on that event. On its complement the absolute product is at most $4\|f\|_\infty \|g\|_\infty$. □

B A block-matrix inequality for adjacent planes

Proof of Lemma 3.4. Put $u = \|A\|$, $v = \|B\|$, $t = \tau(\alpha, \beta)$. The case $tv = 0$ is immediate. For $u \le v$ use the norms of the four blocks to obtain $\|C\| \le v\|D_s\|$, where $D_s = \left(\begin{smallmatrix} s & \sqrt{\beta} \\ \sqrt{\alpha} & 0 \end{smallmatrix}\right)$ and $s = u/v$. The diagonal entries of $(s^2 + t)\mathrm{Id} - D_s^* D_s$ are nonnegative; its determinant equals $(t-\alpha)(t-\beta) + s^2(t-\alpha-\beta) = \alpha\beta(1-s^2)/(1+t)$. This proves the claim when $u \le v$.

Suppose $u \ge v$, and put $w = \sqrt{u^2 + tv^2} > u$. For any matrix X write $\widehat{X} = \left(\begin{smallmatrix} 0 & X \\ X^* & 0 \end{smallmatrix}\right)$. After reordering coordinates, $\widehat{C}$ has diagonal blocks $\widehat{A}, 0$ and upper off-diagonal block $I_2 J \widehat{B}$, where $I_2 = \mathrm{diag}(I, I)$ and $J = \mathrm{diag}(\sqrt{\beta}, \sqrt{\alpha})$ on the two copies. To show $\|\widehat{C}\| \le w$ it suffices by spectral symmetry to prove $w\mathrm{Id} - \widehat{C} \succeq 0$. Take the Schur complement of $w\mathrm{Id} - \widehat{A}$. Its subtracted term is entrywise nonnegative (expand the resolvent in powers of $\widehat{A}$). It therefore suffices to bound that term on nonnegative test vectors.

The scalar inequality $(w-x)^{-1} \le (w+x)/(w^2-u^2)$ for $|x| \le u$ bounds the compressed resolvent in quadratic form by $(w\mathrm{Id} + I_2^* \widehat{A} I_2)/(w^2 - u^2)$. On nonnegative test vectors one may further replace $I_2^* \widehat{A} I_2$ by $\widehat{B}$, by the entrywise hypothesis. Thus the numerator of the subtracted term is bounded in quadratic form by $\widehat{B} J (w\mathrm{Id} + \widehat{B}) J \widehat{B}$. In left and right singular-vector bases its blocks are $r^2 \left(\begin{smallmatrix} w\alpha & r\sqrt{\alpha\beta} \\ r\sqrt{\alpha\beta} & w\beta \end{smallmatrix}\right)$, $0 \le r \le v$. Their largest eigenvalue is at most $wv^2 t$: after dividing the two-by-two matrix by w, use $r/w \le (1+t)^{-1/2}$ and $(t-\alpha)(t-\beta) = \alpha\beta/(1+t)$. Since $wv^2 t = w(w^2 - u^2)$, the Schur complement is positive semidefinite. □

C A bound on axial parity imbalance without periodicity

Theorem 1.1 identifies the two axial parities by symmetry in a periodic phase. A field estimate gives a samplewise bound that also applies in an arbitrary Gibbs measure on the event of large ownership by a specified edge. Index the endpoints r_0, r_1 of E by their bit in the free coordinate w and set $\Lambda_L = \{0, \dots, L-1\}^3$,

$$D_L = L^{-3} \sum_{x \in \Lambda_L} (-1)^{x_w} n_x, \qquad V_L = L^{-3} \sum_{x \in \Lambda_L} \mathbf{1}_{\{x \bmod 2 \notin E\}} n_x.$$

Lemma C.1 (A staggered field along a column)**.** *For $0 < t \le 1$, $|h| = t\lambda$ and $32\lambda \le g \le 1$, put $\phi(x) = h(-1)^{x_w} - g\mathbf{1}_{\{x \bmod 2 \notin E\}}$. Then*

$$p_3[\phi] - p_3[0] \le C(q/g + \lambda t^2). \tag{C.1}$$

Proof. Merge the activities in the two fixed directions first in (3.4). For each sign of the free bit there is one high value $a = ze^{\pm h}$ and three low values ae^{-g}. At each merge with a low branch the gap is at least $ag/4$: the low branch is at most $ae^{-g} + \sqrt{a}$. The identity (3.2) bounds the increment

by C/g as in (3.5). Thus the two remaining values are at most $ze^{\pm h}\exp(Cq/g)$. Since $1+\tau(A,B)$ is a squared matrix norm, scaling the two inputs by $c \geq 1$ increases it by at most the factor c. Before the common factor $\exp(Cq/g)$ is applied, the last merge at $ze^{\pm h}$, divided by z, gives

$$(1+\tau(ze^{h}, ze^{-h}))/z = \cosh h + q/2 + \sqrt{\sinh^2 h + q\cosh h + q^2/4}.$$

Its value at zero gives the independent-column pressure after taking one eighth of its logarithm. The square root at zero is at least λ; subtracting the two values of the displayed expression costs $O(h^2/\lambda) = O(\lambda t^2)$. Lemma 3.6 and the cost Cq/g of the first two merges imply (C.1). □

Proposition C.2 (Axial parity balance on a prescribed edge)**.** *For every Gibbs law, almost surely on* $\{\liminf_L b_E(B_L) \geq 1-\varepsilon\}$*, both* $\limsup_L |8D_L|$ *and* $\limsup_L \max_{i=0,1} |b_{\{r_i\}}(B_L) - 1/2|$ *are at most* $C(\sqrt{\varepsilon}+\sqrt{q})^{1/2}$ *for sufficiently small* ε, q*.*

Proof. By (2.1), $\limsup V_L \leq \varepsilon/8$ on the stated event. For the two signs of h apply the samplewise assertion of Lemma 2.2 and (C.1) to $hD_L - gV_L$. Set $\delta = \sqrt{\varepsilon}+\lambda$, $g = \lambda/\delta$, $t = \sqrt{\delta}$. Then $q/g = \lambda\delta$ and $g\varepsilon \leq \lambda\delta$, which gives $\limsup |D_L| \leq C\sqrt{\delta}$. The difference of the two endpoint ownerships differs from $8D_L$ by at most $8V_L + O(L^{-1})$; their sum has lower limit at least $1-\varepsilon$ and is at most one. □

References

[1] R. J. Baxter, *Exactly Solved Models in Statistical Mechanics*, Academic Press, 1982.

[2] J. van den Berg, A uniqueness condition for Gibbs measures, with application to the 2-dimensional Ising antiferromagnet, Comm. Math. Phys. **152** (1993), 161–166.

[3] M. Biskup, Reflection positivity and phase transitions in lattice spin models, in *Methods of Contemporary Mathematical Statistical Physics*, Lecture Notes in Mathematics **1970**, Springer, 2009.

[4] H.-O. Georgii, *Gibbs Measures and Phase Transitions*, 2nd ed., De Gruyter, 2011.

[5] D. Hadas and R. Peled, Columnar order in random packings of 2×2 squares on the square lattice, J. Eur. Math. Soc. **28** (2026), 5023–5106.

[6] A. Mazel, I. Stuhl and Y. Suhov, Kepler's conjecture and phase transitions in the high-density hard-core model on $\mathbb{Z}^3$, arXiv:2112.14250v2 (2023).

[7] A. Mazel, I. Stuhl and Y. Suhov, The hard-core model on $\mathbb{Z}^3$ and Kepler's conjecture, arXiv:2304.08642 (2023).

[8] K. Ramola and D. Dhar, High-activity perturbation expansion for the hard square lattice gas, Phys. Rev. E **86** (2012), 031135.

[9] Á. Timár, Boundary-connectivity via graph theory, Proc. Amer. Math. Soc. **141** (2013), 475–480.

[10] N. Vigneshwar, D. Mandal, K. Damle, D. Dhar and R. Rajesh, Phase diagram of a system of hard cubes on the cubic lattice, Phys. Rev. E **99** (2019), 052129.