Two-sided bounds for site percolation on the square lattice

Submitting contributor: Su XU

Preprint v1 — November 27, 2026

hexagon:2611.01212v1

Abstract

For independent site percolation on the infinite square lattice with nearest-neighbor adjacency, we prove $0.55855 \le p_c \le 0.63345$. For the lower bound we thin the children of a two-layer exploration according to their joint outcomes. A missing-neighbor constraint and bounds on the posterior of each explored column give stochastic domination at every depth. For the upper bound we couple boundary partitions on two offset tilings by 4×4 blocks. An intermediate occupation field with independent factors of sizes $9, 3, 3, 1$ respects both tilings. Integer transports certify both couplings. Relative to the endpoints 0.556921 and 0.6463 in Niedbala Giraudin's 2026 preprint, the lower bound increases by 0.001629 and the upper bound decreases by 0.01285. The finite certificates and independent checking programs accompany the paper.

Authorship and provenance. ChatGPT is the primary author of this work. Using the 6.2 Astra and 6.2 Sol models in Codex, ChatGPT independently developed the proofs, designed and ran the computations, verified the results, and wrote the entire manuscript. Su XU's contribution was limited to selecting the topic and providing the computing environment.

1 The interval and its comparison points

Let $S = (\mathbb{Z}^2, E_4)$, where $xy \in E_4$ if $\|x - y\|_1 = 1$. At density s, sites are open independently with probability s. Write $\theta_S(s) = \mathbb{P}_s(0 \leftrightarrow \infty)$ and $p_c(S) = \inf\{s : \theta_S(s) > 0\}$.

Theorem 1. *For Bernoulli site percolation on S,*

$$\frac{11171}{20000} \;\le\; p_c(S) \;\le\; \frac{12669}{20000}.$$

The endpoints give the closed interval $[0.55855, 0.63345]$ of width 0.07490. Theorem 1.1 of Niedbala Giraudin, version 1 [3], gives an interval of width 0.089379. The reduction is 0.014479, or about 16.2%. The lower endpoint here also exceeds the approximate search value 0.5569210257 reported in that work by approximately 0.0016289743. All terminating decimals used as parameters or theorem endpoints here denote exact rational numbers.

The two comparisons have different origins. Van den Berg and Ermakov [2] proved $p_c > 0.556$ by comparing a growth process with independent sites. Ashton [1] used conditional laws on an exploration tree to reach 0.5567. Wierman [9] developed the substitution method for this model. Wierman and Oberly [10] obtained $p_c \le 0.666894$, and Oberly and Wierman [6] reported $p_c < 0.662869$ in a poster. Niedbala Giraudin supplied the direct starting points for both sides: scalar messages on the exploration tree and two substitutions through an intermediate lattice.

In Section 2, the exploration sometimes discards a successful connection. This changes what a pair of recorded children reveals about their parent's hidden state. The resulting likelihood factors can be bounded even when the root is arbitrarily far away. In Section 3,

the intermediate field has a transpose-invariant law on nine sites, two independent three-site vectors with a common law, and an independent Bernoulli site. Each group lies entirely inside one block of each tiling, so dependence within a group is compatible with independent block couplings at both steps. The quoted numerical estimate $p_c \approx 0.592746$ from finite-width extrapolation [4] locates these improvements on the scale of the problem. The theorem determines an interval crossing 0.6.

2 The lower bound by a thinned exploration

We prove $q \le p_c(S)$ for $q = 11171/20000$. Put $p = 1 - q = 8829/20000$. The comparison via a two-layer graph is due to van den Berg and Ermakov [2]. We use the conditional tree-message approach of Niedbala Giraudin [3, Sec. 2]. The changes used here are a randomized rule for keeping children, a geometric restriction on available neighbors, and an invariant for the resulting conditional probabilities. We give the rules in detail because deleting one of two successful edges changes the posterior law of their common parent.

2.1 The hidden columns and the observed tree

Over each $x \in \mathbb{Z}^2$ place a column with upper and lower sites, at heights 1 and 0. In the graph L, join the sites of each column, include all four-neighbor edges within each layer, and join the upper site over x to the lower site over $x + e_i$, $i = 1, 2$. All sites are independently open with probability p. Write a column state as $s = (u_s, l_s)$, with the upper bit first. A forward edge (direction $+e_i$) from state s to state t has a natural open connection precisely when

$$E(s,t) = (u_s u_t) \vee (l_s l_t) \vee (u_s l_t) = 1.$$

For a backward edge the indicator is $E(t, s)$. Take the column over the origin as root and condition on its state being 11, an event of probability p^2.

Assign observed value 1 to the root. Expand a previously unexpanded 1-vertex with at least one unassigned neighbor, assigning all its new neighbors in one batch. These neighbors are its tree children. Select an eligible vertex by increasing depth, then by the time its parent was expanded, then by the sibling order specified below. A nonroot vertex has type + or − according to the direction of its incoming edge. An outgoing direction of its own type is called majority, and one of the other type minority. After removing the incoming neighbor there are at most two majority and one minority slots. Among siblings expand majority children first, breaking ties (and ordering the root's children) by $+e_1, +e_2, -e_1, -e_2$. The next batch is determined by the observed bits and positions alone.

At the root, publish the natural connections. At any other vertex use the following rules. All extra draws are independent of the hidden column states and of one another, both within and across batches. Set

$$d = 337/5000, \qquad \gamma = 1459/2500.$$

For a parent of type +, call 01 weak and 10, 11 strong. For a parent of type −, interchange the layers. If two actual majority children both connect naturally and the parent is strong, publish just the first, just the second, or both as successes with respective probabilities $d, d, 1 - 2d$. In all other cases publish their natural bits. Use $h = \gamma$ if there are two actual majority children and exactly one published majority success, and $h = 1$ otherwise. For the minority child of a + parent retain a natural success of state 01 with probability h. For the minority child of a − parent do this to state 10. Publish every other natural success and publish all natural failures as zeros. Thus each successful tree edge has an open connection in L, and subsequent selections use only the published results.

2.2 Conditional probabilities at arbitrary tree depth

A likelihood ratio summarizes what the branch below a vertex reveals about its parent. A conditional probability carries evidence from the root side of the tree to the vertex about to be expanded.

Lemma 2.1 (Factorization of branch likelihoods)**.** *Fix a finite public history of positive probability. Conditional on the column states, its likelihood is a product of factors on the edges of the discovery tree and factors at their parents. Unassigned columns retain their independent priors. On deleting a success edge factor, the weights on the two resulting components can be summed separately.*

Proof. Replay the history. Its published bits and their positions determine the next batch at each step, so the choice of batch adds no likelihood factor. Given the states, the fresh draws in that batch give its record probability. Multiplying these conditional probabilities gives the likelihood of the history. Only assigned columns occur in it.

For the only joint draw, that on a pair of actual majority children, let e'_1, e'_2 denote the two natural majority indicators and $d_s = d$ for a strong parent, $d_s = 0$ for a weak one. In order $00, 10, 01, 11$ the record likelihoods are

$$(1-e'_1)(1-e'_2), \qquad e'_1(1-e'_2+d_s e'_2),$$
$$(1-e'_1+d_s e'_1)e'_2, \qquad (1-2d_s)e'_1 e'_2.$$

For record 10, for instance, the first edge must connect naturally. Given that connection, the second edge contributes 1 for a natural failure and d_s for a natural success. For record 11, keep $1-2d_s$ at the parent. With fewer than two majority children, each actual majority edge reports its natural bit. The majority record fixes h, leaving a single edge factor for the minority bit. The root uses only natural bits. Multiply by the column priors (with the root fixed at 11) to obtain the posterior weights. No factor crosses a cut in the tree except the edge factor itself. Assigning 11 to every success and 00 to every failure realizes the history with positive probability, since $1-2d > 0$. This also gives a positive normalizer on each side of a cut success edge. □

For a + edge, sum the likelihood of the observed branch below the edge against the column priors in that branch, keeping the parent state fixed. The edge factor depends on an upper-open parent in the same way for states 10 and 11. Divide by the common value at these states and call the normalized value at 01 the *branch ratio* r. For a − edge interchange the layers. Write

$$F(v,w) = \frac{p+qv}{p+qv+qw}, \quad B = \frac{1}{q}, \quad B_* = \frac{q}{q^2+d(1-q^2)}, \quad K = \frac{1}{1-2d}, \quad a_h = \frac{1}{1+hq}.$$

A failed child is never expanded, so its edge factor is averaged against its prior alone. A missing slot contributes 1. An ordinary majority failure has ratio B, since weak and strong failure probabilities are q and q^2. If the other majority child was published as a success, the failure factor for a strong parent integrates to $q^2+d(1-q^2)$, while that for a weak parent still gives q. Their likelihood ratio (weak over strong) is B_*. Record 11 multiplies the parent's weak state weight relative to 11 by K. A failed minority edge gives $1/(1-hp)$, the ratio of q to $q(1-hp)$. For a successful unexpanded child the branch ratio is a_h, where $h = 1$ also handles majority edges.

For an expanded successful + child let V be the product of its two majority branch ratios and W its minority branch ratio, including ones for missing slots. Before applying its incoming edge factor, the relative masses of its $01, 10, 11$ states are qAV, qW, p, where $A = K$ for its own double-success record and $A = 1$ otherwise. For this forward edge, a lower-only parent requires the child's lower site to be open. For an upper-open parent the state 10 is admitted with its

incoming retention probability h and each lower-open state is admitted with probability 1. The branch likelihoods at the two parent types are therefore proportional to $p + qAV$ and $p + qAV + qhW$, giving the ratio $F(AV, hW)$. Reflection gives the same formula for a $-$ child. In particular a success ratio is at most 1.

Lemma 2.2 (Missing neighbor)**.** *A minority child of a nonroot vertex has at most one actual majority child of its own.*

Proof. Suppose x was reached from $x - e_i$ and its minority child is $y = x - e_j$. Then $i \neq j$, since $x - e_i$ was already assigned. The site $x - e_i - e_j$, a neighbor of $x - e_i$, had been assigned by the end of the step which discovered x, since that expansion assigned every remaining neighbor of $x - e_i$. Since y has type $-$, this site at $y - e_i$ occupies one of its two majority slots. Reversing directions treats the other sign. □

Fix a nonroot successful vertex b immediately before expansion. For each edge $x \to y$ on its discovery path omit its success edge factor and integrate out the component on the root side, retaining every other factor there, including any local double-success factor at x. Define the *outgoing message* of x to be the probability in that normalized marginal that its upper site is open if y is $+$, or its lower site if y is $-$. Write z_y for this message and h_y for the retention parameter on the success edge to y (equal to 1 on a majority or root edge).

Lemma 2.3 (Probability bounds at every depth)**.** *Put $T = 33/100$, $Z = 243/625$, $M = 41009/100000$ and $N = 2367/5000$. In every finite history as above, each branch ratio lies in $[T, B]$, and each path message is at least Z. At the vertex b, the effective input $v = h_b z_b$ satisfies*

$$\begin{cases} M \le v \le 1, & \text{if both majority slots are actual,} \\ \min\{M, \gamma N\} \le v \le 1, & \text{otherwise.} \end{cases}$$

Proof. From leaves to the root. The function F increases in its first argument and decreases in its second. The exact inequalities for $h \in \{1, \gamma\}$ are

$$1 \le B_* \le B, \qquad 1 \le (1 - hp)^{-1} \le B, \qquad a_h \ge a_1 \ge T, \qquad F(T^2, B) \ge T.$$

Failures and missing slots have ratios in $[1, B]$, and successful leaves have ratio $a_h \le 1$. If all ratios below a successful child are in $[T, B]$, then $AV \ge T^2$ and $hW \le B$. Its ratio $F(AV, hW)$ is therefore in $[T, 1]$. Induction up the finite tree proves the first assertion. At a minority child, Lemma 2.2 strengthens $V \ge T^2$ to $V \ge T$, because a missing slot contributes 1. Thus a minority side branch with parameter h contributes at least $D_h = \min\{a_h, F(T, hB)\}$, including a failed or missing branch.

From the root to b. The root sends 1. If x is $+$ and receives message $z = z_x$, its incoming factor multiplies its 10 weight by $h_x z$, relative to the common multiplier of its $01, 11$ weights. Here h_x is the incoming retention probability. This follows because a child in state 10 requires an upper-open parent, whereas either lower-open state connects to every nonempty parent. The parent is nonempty on the root side because it is the root or was reached by a success edge. Interchanging the layers treats a $-$ vertex.

Suppose $z \ge Z$. If the path changes type at x, both majority slots are side branches. In the $+$ case the favorable states for the outgoing message are $01, 11$, of relative weights qAV, p, and the remaining state 10 has weight $qh_x z$. The message is $F(AV, h_x z) \ge F(T^2, 1) \ge N \ge Z$. If the path continues with the same type, let t be its minority side-branch ratio. The following table gives lower bounds for the outgoing message.

Other majority slot at x	All steps	Final step to b
Published failure	$F(D_\gamma Z, B_*)$	$F(a_\gamma Z, B_*)$
Published success	$F(D_1 Z, K)$	$F(a_1 Z, K)$
Absent	$F(D_1 \gamma Z, 1)$	$F(a_1 \gamma Z, 1)$

Indeed, for + continuation the favorable states are $10, 11$. If r' is the ratio of the other majority slot (1 when absent), the three weights in order $10, 11, 01$ are qth_xz, p, qAr', so the outgoing message is $F(th_xz, Ar')$. In the first row the state 01 has multiplier B_*, and the minority side branch has $h = \gamma$. In the second row its multiplier is at most K, since a successful sibling has ratio at most 1, and the minority branch has $h = 1$. In both rows x has two actual majority children, so by the missing-slot argument it could not have been born as a minority child of a nonroot vertex. Thus $h_x = 1$. In the last row the missing slot contributes 1 and $h_x \geq \gamma$. For the last step b has remained eligible from discovery onwards because the set of unassigned neighbors only shrinks. By majority-first sibling order, its minority sibling, if present, has not yet been expanded. A successful such sibling has ratio a_h, while a failed or absent one contributes at least $1 \geq a_h$. This gives the last column.

Direct rational comparison gives each entry in the middle column at least Z, and each in the last column at least M. For example the first row of the last column equals $57516274601589148515132 3/14025098178893763851513 23$. Induction along the path now gives $z_y \geq Z$ at every step. At its last step a same-type birth gives $z_b \geq M$ and $h_b = 1$. A type-changing birth gives $z_b \geq N$ and $h_b \geq \gamma$, and Lemma 2.2 rules out two actual majority slots at b in that case. A child of the root has $h_b z_b = 1$. These are all possibilities and $\min(M, \gamma N) = 3453453/12500000$. □

2.3 Finite comparison and passage to the infinite lattice

At expansion a + vertex has no descendant observations. Its nonempty states $01, 10, 11$ have prior weights pq, pq, p^2. Its incoming success edge rules out 00 and multiplies only the 10 weight by v, relative to the other two. Thus the posterior weights are proportional to q, qv, p, with v in the interval of Lemma 2.3. For actual majority and minority slot counts $(m, k) \in \{0, 1, 2\} \times \{0, 1\}$ with $m + k > 0$, compute its published-vector law as follows. Conditional on the parent being weak, the m natural majority successes are independent with probability p; conditional on either strong state use $1 - q^2$. If $m = 2$ and the parent is strong, move fraction d of the double-success mass to each of the two one-success masks. For each resulting mask set h by the rule above. If $k = 1$, append a success bit with probability p for parent state 10, and $p + hpq$ for $01, 11$. Use layer interchange for the other direction.

For every upward-closed $U \subseteq \{0, 1\}^{m+k}$, subtract the independent Bernoulli(q) probability of U from the probability under this mixture. Multiplying by $1 + qv$ gives an affine function of v, so the two endpoints of its proved interval suffice. All input is rational. To enumerate U, run over all subsets of masks, retaining a subset precisely when adding any bit to any member yields another member. The five pairs $(0, 1), (1, 0), (1, 1), (2, 0), (2, 1)$ have respectively $1, 1, 4, 4, 18$ nontrivial upward-closed subsets. With two endpoints and two signs this gives 112 comparisons. The minimum is attained at $v = M$ for the event that at least one of the two majority children succeeds. Thinning a double success always preserves this event. Its probability margin is the difference between the two probabilities of a double failure, which gives

$$p^2 - \frac{q^3 + q^4(qM + p)}{1 + qM} = \frac{14394176127638528447}{131099282080000000000000000} > 0.$$

At the root the four success bits are independent with common parameter $1 - q^2 > q$. The accompanying file `stage3/lower/refined-certificate.json` records the scalar inequalities and comparisons. The programs `stage3/lower/certificate.py` and `stage3/lower/audit.py` both take `--refined`. The latter enumerates the raw child states and checks 3936 identities for conditional record likelihoods.

Lemma 2.4 (Global domination along the exploration)**.** *Conditional on the root column being* 11*, one can couple the exploration to a field ξ of independent Bernoulli(q) sites on all of S. At every assigned site its comparison bit is at most its published bit. If the root lies in an infinite ξ-open cluster, the root column lies in an infinite open cluster of L.*

Proof. Let H be the public history just before a batch, $A = A(H)$ its set of new sites and Y the vector about to be published there. Write ν_A for the Bernoulli(q) product law on A. The upward-event comparisons above and the finite monotone coupling theorem [8] give a joint law Q_H for (Y, C) with first marginal the published law given H, second marginal ν_A, and $C \le Y$ coordinatewise. Fix such a law for each finite history. Run the original exploration. After it publishes Y, draw the comparison vector C from $Q_H(C \mid Y)$ using fresh randomness. At the root use an independent Bernoulli(q) comparison bit below its published value 1. Only published bits enter the exploration rule.

To identify the law on the entire lattice, condition on H and on all comparison bits assigned before this batch. For a fixed H, the conditional probability of any particular past comparison record, even given the hidden field and the exploration draws, is the product of its successive comparison kernels. Each used only the corresponding public prefix, published vector and a fresh independent draw. This product is a function of the public and comparison records alone. Conditioning on past comparisons therefore leaves the law of the next Y given H intact. Averaging $Q_H(C \mid Y)$ against this law gives exactly ν_A conditional on the past, with A now fixed.

Initially give every site an independent Bernoulli(q) reserve bit, independent of the entire construction. At each assignment replace the reserves on A by C, and keep all other coordinates. Given the past just described, the assigned coordinates are fixed, and the unused reserves are independent of both the record and the new draws on the now specified sites. On A the old and new vectors both have law ν_A, independently of the remaining reserves. The conditional law of the whole field is thus the same before and after replacement. Integrating over the past proves that every finite step preserves its product marginal, also when exploration has stopped. Each coordinate changes at most once. On any fixed finite set the values eventually stabilize, so their limiting joint law is again the product law.

If exploration stops, every neighbor of every published 1 has been assigned. Starting from an open comparison root, induction along any comparison-open path then keeps it in the finite assigned set. Its next vertex is assigned and its comparison 1 forces a published 1 there. Thus an infinite comparison cluster at the root forces infinitely many expansions. The expanded columns are distinct, and each nonroot one is joined to its parent by a natural open connection. The open sites within each such column are connected, so their union contains an infinite open cluster at the root column. □

The matching-lattice comparison. Removing the conditioning in Lemma 2.4 gives probability at least $p^2\theta_S(q)$ that the bottom site of the root column lies in an infinite open cluster of L. To compare L with the eight-neighbor matching lattice S^*, put $\varphi(x, y, j) = (x-y, x+y+j)$. Given $(a, b) \in \mathbb{Z}^2$, choose the unique $j \in \{0, 1\}$ congruent to $b - a$ modulo 2. The unique preimage has $x = (a+b-j)/2$ and $y = (b-a-j)/2$. The four same-layer edges at a site give all four diagonal neighbors of its image. The two intercolumn upper-to-lower edges incident to it give both horizontal neighbors. Its column edge gives the upward vertical neighbor for $j = 0$ and the downward one for $j = 1$. Hence the image omits precisely the vertical edges from (a, b) to $(a, b+1)$ with $b - a$ odd. Each image vertex has degree seven, compared with eight in S^*. Thus the image is a proper spanning subgraph, and

$$\varphi(L) \subsetneq S^*, \qquad p_c(S^*) \le p_c(L).$$

Indeed, use the same independent site bits on the common vertex set. Then every open path of $\varphi(L)$ is also open in S^*. If $\theta_S(q) > 0$, the root estimate gives $p_c(L) \le p$, and Russo's matching identity $p_c(S) + p_c(S^*) = 1$ [7, Thm. 1] gives $p_c(S) \ge 1 - p = q$. If $\theta_S(q) = 0$, monotonicity gives $p_c(S) \ge q$ directly.

3 Upper bound by a correlated intermediate field

Write S for the nearest-neighbour square lattice. We use the two substitutions of Niedbala Giraudin [3, Sec. 3]. Here the occupation law on the intermediate graph factors over the intersections of the two tilings. This permits exact local comparison on both tilings with the same infinite-volume marginal.

3.1 The graphs and the common field

Adjoin both diagonals of every unit square whose lower left corner (x, y) satisfies $x \equiv y \pmod 2$ to obtain a graph J on $V(S)$. Retaining only the diagonals of the odd–odd squares gives an intermediate graph H. The centres of the selected squares of J, joined when the squares share a vertex, form a rotated square lattice. Each site is the common corner of exactly two selected squares and hence corresponds to an edge of this centre lattice. Two sites are adjacent in J precisely when their edges share a centre, since the four corners of each selected square form a clique. Thus independent site occupation on J is independent bond occupation on a square lattice. By Kesten's theorem [5], every parameter greater than $1/2$ gives an infinite cluster with positive probability.

For $z = (0, 0)$ or $(1, 1)$ tile the vertices by $z + 4k + \{0, 1, 2, 3\}^2$, $k \in \mathbb{Z}^2$. Use the first tiling for $J \to H$ and the second for $H \to S$. The deleted diagonals lie in even–even squares at the first step and odd–odd squares at the second. A unit square crossing a tile boundary has a lower left coordinate congruent to 3 or 0 modulo 4, respectively, in at least one direction. It follows that source and target have identical cross-tile edges at each step.

Choose independent occupation vectors on all $4\mathbb{Z}^2$-translates of the following four sets:

$$\{1, 2, 3\}^2, \qquad \{1, 2, 3\} \times \{0\}, \qquad \{0\} \times \{1, 2, 3\}, \qquad \{(0, 0)\}.$$

They partition the vertices. They are precisely the nonempty intersections of tiles of the first tiling with tiles of the second. In particular each tile in *either* tiling contains four entire factors, and different tiles in that tiling have independent occupation vectors.

The file `stage4/alternative/u63345/parameters.json` gives the rational vector h specifying the law. Index the nine sites by increasing y and then increasing x, using the first site as the least significant bit. Among the 512 masks, 64 are symmetric 3×3 binary matrices. Transposition therefore has 288 orbits. Sort these orbits by least mask and divide the orbit mass equally among its members. The first 288 entries of h are their masses. The following eight entries specify the three-site mask masses in binary order, with the coordinate increasing from the least significant bit. Use that law for each of the two three-site vectors. The last two entries give the closed and open masses of the corner. Each of these three probability lists is nonnegative and sums to one. This prescription defines a single field on H, including on tiles based at $(1, 1) + 4k$.

3.2 What the integer transports certify

A port is the incidence of a cross-tile edge with its vertex inside the tile. Incidences from the same vertex have distinct labels. For a given occupation mask, put two distinct ports in the same class precisely when their vertices are open and in the same open component within the tile. Every port at a closed vertex forms its own singleton class. Write $\pi \preceq \rho$ if each class of π is contained in a class of ρ. Increasing the set of open vertices can only coarsen the partition.

Let $P = 50000001/10^8$ and $u = 12669/20000$. For each of the two steps the finite claim is the existence of a coupling with $\pi_{\mathrm{source}} \preceq \pi_{\mathrm{target}}$. Each block law is invariant under the two-element group generated by coordinate exchange. For source and target partition orbits i, j, join i to j when they contain representatives $\pi_i \preceq \rho_j$. All equalities of π_i must hold in the *same* ρ_j. Scale the two orbit mass vectors to integers a_i, b_j with common total mass D.

The supplementary lists of flows f_{ij} have positive integers on the listed edges, zero elsewhere, and satisfy

$$\sum_j f_{ij} = a_i, \qquad \sum_i f_{ij} = b_j.$$

For each positive edge, fix one comparable pair and apply a uniformly chosen group element to that pair. This gives uniform marginals on the two orbits, including when one orbit is a singleton. Mixing with weights f_{ij}/D gives the required coupling on the original labelled partitions.

Step	Ports	Orbits (source, target)	Comparable pairs	Positive flows
$J(P) \to H$	28	$(8660, 7052)$	1 569 136	17 526
$H \to S(u)$	16	$(6171, 4135)$	690 957	11 143

For each graph and offset the checker reconstructs the ports and internal edges from the coordinates and computes the partitions of all 2^{16} masks by depth-first search. At independent density s the weight of a mask with k occupied sites is $s^k(1-s)^{16-k}$. On H it is the product of the four specified mask probabilities. The checker forms all orbit sums, verifies equal mass within each orbit, reconstructs the comparability relation and checks every flow and marginal. It also compares the two copies of the intermediate parameter vector with h. These operations use exact integers and rationals. The certificate files `step1.json` and `step2.json` are in the same directory as the parameters.

3.3 Passage to the infinite lattice

On each tile of one tiling draw an independent copy of its partition coupling. Conditional on the two partitions, sample each side's occupation mask with its specified conditional law. The marginals on the whole plane are the required ones by tile independence. For either marginal form an auxiliary graph with a vertex for every tile partition class and an edge between the two incident classes of each cross-tile edge. A closed port contributes a degree-one vertex.

An infinite simple open path crosses infinitely many finite tiles and gives a connected set meeting infinitely many tiles in the auxiliary graph. Coarsening the partitions maps this set into a target component meeting all those tiles, because cross-tile edges agree. The target auxiliary graph is locally finite, so this component contains an infinite simple path. Every internal vertex of that path joins two distinct ports in one class. Their incident sites are open and connected within the tile, including the case of two ports at one open site. Dropping the initial edge, we can therefore join these connections by genuine open cross-tile edges. Each tile has finitely many classes, so the connected set thus obtained meets infinitely many tiles. Applying this implication twice transfers positive probability of an infinite cluster to $S(u)$. The two couplings may be drawn separately because the intermediate field has the same global law at both steps. On the translation-invariant target, the event that an infinite cluster exists is the countable union of the events that a fixed site belongs to one. Positive probability for the union therefore gives positive probability at the origin. By the definition of the site threshold,

$$p_c \le u = \frac{12669}{20000} = 0.63345.$$

4 Reproducing the finite claims

Download and extract this version's TeX source archive. From the directory containing `bounds.tex`, run

```
python -B anc/verify.py
```

with Python 3.10 or later. The lower checks pair `anc/lower_certificate.py` and `anc/lower_audit.py` with `anc/refined-certificate.json`. The upper checks use `anc/parameters.json` and `anc/step2.json`. The command joins `anc/step1_partitions.json` and `anc/step1_flows.json` to recover the first transport byte for byte, verifies its SHA-256 digest, and restores the paths cited in Sections 2 and 3. The test reconstructs the lower certificate, recomputes the local likelihoods from hidden spins and checks the two transports at the densities in Theorem 1. One upper checker uses depth-first search and builds the partition order. A second grows the components of a mask by adding one vertex and checks containment on the listed positive arcs. Both recalculate probabilities from site coordinates. All acceptance comparisons use integers or fractions. The independent checks operate on the finite inputs; the coverage of all exploration depths and the transfer of infinite clusters are proved in Sections 2 and 3.

The upper search started from a witness at $u = 0.6335$. Its intermediate law was adjusted using violated transport cuts as constraints and converted to rational weights. At iteration 54 (numbered from zero) both integer transports at $u = 0.63345$ carried the full marginal mass. Rebuilding the two transports from the frozen law produced byte-identical files. On the recorded Windows run with Python 3.13.7 the search used 52.1 seconds and a peak process working set of 210 MiB, while reconstruction from the final parameters used 21.6 seconds and 175 MiB. This deposit includes the accepted parameters, exact witnesses, checking programs, file digests and the recorded transport reconstruction result.

Acknowledgment. Comments on a prepublication draft led us to organize the exploration proof into four lemmas and to identify explicitly the edges omitted by its embedding in the matching lattice.

References

[1] C. Ashton, $p_c^{\mathrm{site}}(\mathbb{Z}^2) \geq 0.5567$, computer-assisted proof, October 2026.

[2] J. van den Berg and A. Ermakov, A new lower bound for the critical probability of site percolation on the square lattice, Random Structures Algorithms **8** (1996), 199–212.

[3] D. Niedbala Giraudin, Computer-assisted bounds $0.556921 \leq p_c \leq 0.6463$ for site percolation on the square lattice, Hexagon 2611.00065v1, November 2026.

[4] J. L. Jacobsen, Critical points of Potts and $O(N)$ models from eigenvalue identities in periodic Temperley–Lieb algebras, arXiv:1507.03027 (2015).

[5] H. Kesten, The critical probability of bond percolation on the square lattice equals 1/2, Commun. Math. Phys. **74** (1980), 41–59.

[6] S. P. Oberly and J. C. Wierman, A new upper bound for the site percolation threshold of the square lattice, Johns Hopkins University Design Day poster, 2024.

[7] L. Russo, On the critical percolation probabilities, Z. Wahrscheinlichkeitstheorie verw. Gebiete **56** (1981), 229–237.

[8] V. Strassen, The existence of probability measures with given marginals, Ann. Math. Statist. **36** (1965), 423–439.

[9] J. C. Wierman, Substitution method critical probability bounds for the square lattice site percolation model, Combin. Probab. Comput. **4** (1995), 181–188.

[10] J. C. Wierman and S. P. Oberly, A new upper bound for the site percolation threshold of the square lattice, in *Combinatorics, Graph Theory and Computing*, Springer Proc. Math. Stat. **462** (2024), 129–138.