Critical hard-core influence bounds via nonbacktracking collisions

With a Lean 4 formalization

Dev D. Goyal

30 October 2026; revised 1 November 2026

Abstract

For each fixed integer $\Delta \geq 3$, we establish a uniform $O_\Delta(\sqrt{n}\log(n+1))$ bound on the absolute row sums of the hard-core influence matrix on every finite n-vertex simple graph of maximum degree at most Δ, at all homogeneous fugacities $0 < \lambda \leq \lambda_c(\Delta) = (\Delta-1)^{\Delta-1}/(\Delta-2)^\Delta$. The bound holds under every feasible pinning. An explicit finite constant is given below. The proof uses exact vertex-deletion identities, a known local squared-message contraction, a self-avoiding-walk transfer of the type introduced by Sinclair, Srivastava, Štefankovič and Yin [9], and a finite reconstruction of the recent nonbacktracking collision argument of Bencs et al. [1]. The result addresses the hard-core case of the critical two-spin question of Bencs et al. [1, Section 5]. The actual Gibbs distribution, conditional influences, graph completion, path-count transfer, and collision estimate are formalized in Lean 4.

The contribution is the resulting general critical hard-core bound; its ingredients are credited prior mathematics. It recovers, up to a logarithm, an $O(\sqrt{n})$ bound that was stated in 2025 [3] and withdrawn in January 2026 [2]. Through a published proposition that we have not formalized, it gives a polynomial bound on the mixing time of Glauber dynamics at criticality. The explicit constant is large.

1 Result and scope

Let $G = (V, E)$ be a finite simple graph and $n = |V| \geq 1$. Its hard-core distribution at activity $\lambda > 0$ is

$$\mu_{G,\lambda}(I) = \frac{\lambda^{|I|}}{Z_G(\lambda)}, \qquad Z_G(\lambda) = \sum_{I \subseteq V \text{ independent}} \lambda^{|I|}.$$

We use the occupation variables $X_v = \mathbf{1}_{v \in I}$. A pinning consists of disjoint sets $P, Q \subseteq V$, whose vertices are required to be occupied and unoccupied respectively. It is feasible exactly when P is independent. Write $\mu^{P,Q}$ for the resulting conditional law. For a coordinate u with $0 < \mu^{P,Q}(X_u = 1) < 1$, define

$$\Psi^{P,Q}(u, v) = \mu^{P,Q}(X_v = 1 \mid X_u = 1) - \mu^{P,Q}(X_v = 1 \mid X_u = 0).$$

This convention has diagonal 1. Deterministic target coordinates contribute zero. A uniform absolute-row bound is the ℓ_∞ form of spectral independence [13]; it also controls the usual matrix eigenvalue notion. The main formal theorem states the absolute-row inequality directly.

Theorem 1.1 (Explicit critical hard-core bound)**.** *Fix an integer $\Delta \geq 3$, put $d = \Delta - 1$, and let G have maximum degree at most Δ. Suppose $0 < \lambda \leq d^d/(d-1)^{d+1}$. Set*

$$N = 2\Delta n, \quad r = \lfloor \sqrt{N} \rfloor + 1, \quad L = \lceil \log_d r \rceil, \quad M = 1 + 12\Delta d \frac{\lambda}{1+\lambda} rL.$$

For every feasible pinning (P, Q) and every nondegenerate source u,

$$\sum_{v \in V} \left| \Psi^{P,Q}(u, v) \right| \leq M. \tag{1}$$

In particular the bound is $O_\Delta(\sqrt{n} \log(n+1))$, uniformly in $0 < \lambda \leq \lambda_c(\Delta)$ and in the pinning.

The degree bound is fixed in the asymptotic notation. The displayed constant remains explicit when the degree varies. The variable n in the formal theorem is the original ambient vertex count. No expansion, girth, bipartiteness, connectivity, or minimum-degree assumption is imposed on G. Activity is homogeneous; arbitrary coordinate-dependent field tilts are not part of the theorem. The exact $O(\sqrt{n})$ hard-core upper bound and the general critical two-spin question remain unresolved by this result. The established hard-core lower bound is $\Omega(n^{1/3})$ for spectral independence [4]; accordingly, we make no assertion of optimality up to logarithms.

2 What is established by the formal development

All 32 Lean modules compile with the pinned toolchain, and the central declarations depend only on `propext`, `Classical.choice` and `Quot.sound`. A fresh replay on 30 October 2026, compiling all 32 modules from source against a Mathlib checkout built at the pinned revision, passed; its record is in `verification/`.

The central declarations in `CriticalHardcore/Main.lean` are:

- `critical_hardcore_influence_rows`;
- `critical_hardcore_spectral_independence`;
- `critical_hardcore_pinned_rows`.

The last declaration exposes the original pinning sets, conditional probabilities, and nondegeneracy premise. Its assumptions are only a finite simple graph, positive activity at or below the stated threshold, the degree bound, and feasibility/nondegeneracy of the pinning. There is no assumed tree comparison or collision bound.

The finite Gibbs probabilities are defined as ratios of sums of $\lambda^{|I|}$ over actual independent sets satisfying the relevant events. All needed normalizers are proved positive. The formalization does not define the desired influence in terms of a surrogate recurrence and then assume that it represents the graph distribution.

3 Exact finite graph identities

For a set $S \subseteq V$, use $Z(S)$ and $\mathbb{E}_S$ for the partition function and expectation of the induced hard-core model on S. If $u \in S$, let $F(S, u) = S \setminus N_G[u]$. Splitting independent sets according to u

gives

$$Z(S) = Z(S \setminus \{u\}) + \lambda Z(F(S,u)), \qquad p(S,u) = \frac{\lambda Z(F(S,u))}{Z(S)}. \tag{2}$$

Both terms are positive. Also $0 < p(S,u) < 1$ and $p(S,u) \le p_{\max} := \lambda/(1+\lambda)$: removing additional vertices can only remove nonnegative summands from the partition function.

For $f : V \to \mathbb{R}$ define the additive observable $F_f(I) = \sum_{v \in I} f(v)$ and its conditional difference

$$D(S,u;f) = \mathbb{E}_S[F_f \mid X_u = 1] - \mathbb{E}_S[F_f \mid X_u = 0] = f(u) + \mathbb{E}_{F(S,u)}F_f - \mathbb{E}_{S\setminus\{u\}}F_f.$$

Choose any ordering $v_1, \dots, v_k$ of the neighbors of u lying in S. Put

$$S_i = S \setminus \{u, v_1, \dots, v_{i-1}\}, \qquad p_i = p(S_i, v_i).$$

Each v_i belongs to S_i. Successively conditioning v_i to be absent in these contexts gives the two exact identities

$$p(S,u) = \frac{\lambda \prod_{i=1}^{k}(1-p_i)}{1 + \lambda \prod_{i=1}^{k}(1-p_i)}, \tag{3}$$

$$D(S,u;f) = f(u) - \sum_{i=1}^{k} p_i D(S_i, v_i; f). \tag{4}$$

For completeness, $Z(S_i \setminus \{v_i\})/Z(S_i) = 1 - p_i$ by (2). Multiplying these ratios telescopes from $S \setminus \{u\}$ to $F(S,u)$ and proves (3). The mixture identity

$$\mathbb{E}_{S_i}F_f - \mathbb{E}_{S_i\setminus\{v_i\}}F_f = p_i D(S_i, v_i; f)$$

telescopes over the same contexts and proves (4). The neighbors need not be independent.

Self-avoiding graph reductions are established methodology [12]. Iterating these exact identities builds a finite deletion computation tree. A node stores (S,u) and its actual message $p(S,u)$; its children store the pairs (S_i, v_i). The free set strictly decreases at each edge, so fuel $|S|$ exhausts the recursion exactly. Child labels are distinct, and every descendant free set excludes each of its ancestor labels. Consequently each tree path labels a self-avoiding graph path, and two distinct tree paths with the same depth have different full label sequences.

For $|f(v)| \le 1$, taking absolute values in (4) bounds $|D(S,u;f)|$ by one plus the sum, over all nonempty downward tree paths, of the product of their messages, excluding the source root's own message. This is a finite identity and bound. The computation tree can have far more than n nodes; no bound on its size by n is used.

Finally,

$$D(S,u;f) = \sum_{v \in S} f(v)\Psi_S(u,v). \tag{5}$$

Taking $f(v)$ to be the sign of $\Psi_S(u,v)$ makes the right side its absolute row sum. Thus it suffices to bound the weighted tree expansion uniformly over all bounded observables. This also fixes the row orientation and the normalization of the influence matrix.

4 The known local inequality and the path estimate

The following local inequality is prior art; the stronger squared-message form is stated in the 2025 Chen–Jiang proceedings and traced there to Li–Lu–Yin's contraction argument [3, 5]. The

subsequent correction to the graph comparison does not invalidate this scalar inequality, which is proved directly here and in Lean.

Lemma 4.1 (Local squared-message contraction)**.** *Let $d \geq 2$, $0 \leq \lambda \leq d^d/(d-1)^{d+1}$, and let $p_1, \ldots, p_k \in [0,1]$ with $k \leq d$. If $p = \lambda \prod_i (1-p_i)/(1+\lambda \prod_i (1-p_i))$, then*

$$dp \sum_i p_i \leq 1.$$

Proof. Pad the list with zeros to length d, without changing its sum or product. Write $s = \sum_i p_i$ and $A = \prod_i (1-p_i)$. If $ds \leq 1$, the conclusion follows from $p \leq 1$. Otherwise apply AM–GM to the $d+1$ nonnegative numbers

$$1-p_1, \ldots, 1-p_d, \quad (ds-1)/d.$$

Their average is $(d-1)/d$. Therefore

$$A(ds-1) \leq \frac{(d-1)^{d+1}}{d^d} = \frac{1}{\lambda_c(d+1)}.$$

It follows that $\lambda A(ds-1) \leq 1$, which rearranges to $d\lambda As/(1+\lambda A) \leq 1$. □

Every nonroot node in the deletion tree has at most $d = \Delta - 1$ children: its graph parent has already been deleted. Hence Lemma 4.1 applies at every such node. The original root may have $d+1$ children and is treated separately.

Lemma 4.2 (Weighted paths versus their count)**.** *Let $d \geq 2$, and let a finite rooted tree carry messages $0 \leq p_x \leq p_{\max}$ satisfying $dp_x \sum_{y \text{ child of } x} p_y \leq 1$ at every node. Let a_k be the number of root-to-depth-k paths and let*

$$M_k = \sum_{(x_0,\ldots,x_k)} \prod_{i=0}^{k} p_{x_i}.$$

Then $M_k \leq p_{\max}\sqrt{a_k/d^k}$.

Proof. Define

$$E_k = \sum_{(x_0,\ldots,x_k)} \left(\prod_{i=0}^{k-1} p_{x_i}^2 \right) p_{x_k}.$$

At a subtree rooted at x, $E_0(x) = p_x$ and $E_{k+1}(x) = p_x^2 \sum_y E_k(y)$. Induction and the local inequality give

$$d^{k+1} E_{k+1}(x) \leq dp_x^2 \sum_y p_y \leq p_x.$$

The sum of squared full path weights is at most $p_{\max} E_k$, because the last message is at most $p_{\max}$. Cauchy–Schwarz on the finite list of paths therefore gives

$$M_k^2 \leq a_k p_{\max} E_k \leq p_{\max} p_{x_0} \frac{a_k}{d^k} \leq p_{\max}^2 \frac{a_k}{d^k}.$$

All quantities are nonnegative, so taking square roots proves the claim. No probabilistic independence of tree branches is involved. □

5 Regular completion and exact walk normalization

Lemma 5.1 (Explicit regular completion)**.** *Every finite graph G of maximum degree at most $\Delta \geq 1$ embeds as an induced subgraph of a simple Δ-regular graph H with exactly $2\Delta|V(G)|$ vertices.*

Proof. Take vertices (v, a, b) with $v \in V(G)$, $a \in \mathbb{Z}/\Delta\mathbb{Z}$, and $b \in \{0, 1\}$. Retain a copy of every original edge for each fixed (a, b). At each v, let $\delta_v = \Delta - \deg_G(v)$. For each $j \in \{0, \ldots, \delta_v - 1\}$, join $(v, a, 0)$ to $(v, a + j, 1)$ for all a. These are δ_v disjoint perfect matchings on the copy coordinates. They have no loops and are disjoint from the original edges, which change v. Thus every new vertex has degree $\deg_G(v) + \delta_v = \Delta$. Fixing (a, b) gives an induced original copy. The stated vertex count is immediate. □

Let H now be Δ-regular, $d = \Delta - 1$, and fix an oriented edge $e = (a, b)$. A nonbacktracking walk starts with this exposed edge and then, at each step, chooses uniformly among the d neighbors different from its previous vertex. Let

$$s_e(k) = \mathbb{P}(\text{the initial edge and the next } k \text{ steps have no repeated vertex}).$$

With a fixed enumeration of each directed edge's available next neighbors, there is a bijection between length-k choice words over $\{1, \ldots, d\}$ and these nonbacktracking continuations. Hence

$$s_e(k) = \frac{\#\{\text{self-avoiding continuations of } e \text{ with } k \text{ further steps}\}}{d^k}. \tag{6}$$

The sequence is nonincreasing and vanishes once its required number of distinct vertices exceeds $|V(H)|$.

Apply Lemma 4.2 to a child subtree of the original source u, whose first child label is v. Prepending u to a tree path yields an injective map into the self-avoiding paths counted by (6), after embedding G into H. Different tree paths remain different vertex sequences. Thus

$$M_k \leq p_{\max}\sqrt{s_{(u,v)}(k)}, \qquad \sum_w |\Psi_G(u, w)| \leq 1 + \sum_{v \text{ a child of } u} p_{\max} \sum_{k \geq 0} \sqrt{s_{(u,v)}(k)}. \tag{7}$$

The displayed infinite sum is only notation for an eventually zero sequence; the Lean proof uses uniform bounds on all finite partial sums. The number of children in the last expression is at most Δ.

6 A finite square-root survival bound

This section reconstructs the regular-graph collision method in Bencs et al. [1]. Its purpose is to discharge the complete walk prerequisite in Lean and obtain the square-root survival sum needed by (7). Neither the basic collision phenomenon nor the endpoint-packing idea is claimed as new here.

Proposition 6.1 (Fixed frontier length)**.** *Let $d \geq 2$. For every $(d + 1)$-regular finite simple graph H with N vertices, every initial oriented edge e, every integer $h \geq 0$, and every finite cutoff m,*

$$\sum_{k=0}^{m-1} \sqrt{s_e(k)} \leq (2h + 1)\left(\left\lfloor \frac{2N}{d^h} \right\rfloor + 4d^{h+1}\right). \tag{8}$$

6.1 Endpoint packing along a surviving history

Put $\ell = h+1$. At an oriented edge position i of a self-avoiding history, a frontier is a nonbacktracking path of ℓ edges beginning with that edge. There are d^h such frontiers. Call a frontier safe if, from every oriented edge along it, no nonbacktracking continuation adding between one and ℓ edges can end at an already visited vertex of the old history up to position $i+1$.

Safe frontiers at all positions of one self-avoiding history have distinct endpoints. Here is the full injection argument. Suppose two length-ℓ frontiers $a_0, \dots, a_\ell$ and $b_0, \dots, b_\ell$ have the same last vertex, with the second one starting no earlier in the history. Then a_0 is in the second frontier's old history. If $a_{k+1} = b_{k+1}$ but $a_k \neq b_k$, the walk

$$b_k, a_{k+1}, a_k, a_{k-1}, \dots, a_0$$

is nonbacktracking. Its first two vertices are the kth oriented edge of the second frontier, and it adds $k+1 \leq \ell$ edges to reach the old history. This contradicts safety. Starting at their equal endpoints and descending in k shows equality of every vertex of the two frontiers. Their initial vertices then identify the same history position, because the history is self-avoiding. The frontiers are identical.

Consequently the total number of safe frontiers over every history position is at most N. Call a position good when at least half of its d^h frontiers are safe. Along any surviving history the number of good positions is therefore at most

$$K = \left\lfloor \frac{2N}{d^h} \right\rfloor. \tag{9}$$

The same bound applies to any selected trial-start positions.

6.2 A bad start forces a uniform collision chance

Fix the past history. A partial choice word is dangerous if some continuation of length at most ℓ makes the complete graph walk repeat a vertex. If a length-h word has no dangerous prefix, its frontier is safe: a short return to an old vertex from any frontier edge can be realized as such a killing continuation from the corresponding prefix. At a bad start, at least half of the d^h words must therefore encounter danger.

For each dangerous prefix choose one killing continuation and pad it to exactly ℓ choices. Padding is harmless because a collision, once present, persists under extension. For any dangerous length-h word w, stop at its first dangerous prefix p and encode w as

$$p \parallel \text{chosen padded killing word for } p \parallel \text{the original suffix of } w.$$

This is a killed word of length $h + \ell = 2h+1$. The map is injective: the first dangerous prefix is still p in the encoded word, since all earlier prefixes are unchanged; after recovering p, delete its fixed length-ℓ killing word to recover the original suffix. More formally, two minimal dangerous prefixes of a common word are comparable in prefix order and must coincide.

Thus at least $d^h/2$ words of length $2h+1$ are killed. Conditional on any bad trial start, the survival probability through the next $B = 2h+1$ steps is at most

$$q = 1 - \frac{1}{2d^{h+1}}. \tag{10}$$

This argument uses counts of complete words, so it does not condition on a retrospectively selected frontier or assume independence of danger events.

6.3 Adaptive trials and summation

Break the continuation into blocks of length B. A trial start's good/bad status can depend on the whole history. Represent all surviving block extensions as a finite rooted tree with their conditional probabilities. At a good node the outgoing mass is at most 1; at a bad node it is at most q. Every live root-to-node path has at most K good starts, by (9).

Induction on the remaining number of trials and good-start budget proves

$$s_e(Br) \le q^{(r-K)_+} \qquad (r \ge 0). \tag{11}$$

At a good node one consumes one unit of budget and at most unit mass; at a bad node one preserves the budget and multiplies by at most q. An empty extension list contributes zero. These statements apply to every child separately, so the induction remains valid for history-dependent statuses. It does not replace adaptive conditional laws by a product of unconditional probabilities.

By monotonicity of s_e and a geometric sum,

$$\begin{aligned}\sum_{k=0}^{m-1} \sqrt{s_e(k)} &\le B \sum_{r\ge 0} (\sqrt{q})^{(r-K)_+} \\ &= B\left(K + \frac{1}{1-\sqrt{q}}\right) \le B\left(K + \frac{2}{1-q}\right) = B\left(K + 4d^{h+1}\right).\end{aligned}$$

In the Lean development the same computation is carried out with finite sums, avoiding any convergence premise. This proves Proposition 6.1.

6.4 Choosing the square-root scale

Set $r = \lfloor \sqrt{N} \rfloor + 1$, $L = \lceil \log_d r \rceil \ge 1$, and $h = L - 1$. Then

$$r \le d^L \le dr, \qquad N < r^2, \qquad \left\lfloor \frac{2N}{d^{L-1}} \right\rfloor \le 2dr.$$

Substituting into (8) gives

$$\sum_{k=0}^{m-1} \sqrt{s_e(k)} \le 12drL. \tag{12}$$

The formal natural-number definitions are

$$\begin{aligned}\mathrm{frontierLength}(d, N) &= \mathrm{clog}_d(\lfloor \sqrt{N} \rfloor + 1), \\ \mathrm{collisionScale}(d, N) &= 12d(\lfloor \sqrt{N} \rfloor + 1)\,\mathrm{frontierLength}(d, N).\end{aligned}$$

Lean also proves the real-variable comparison

$$\mathrm{collisionScale}(d, N) \le 12d(\sqrt{N} + 1)(1 + \log_d(\sqrt{N} + 1)).$$

This establishes the stated asymptotic order without suppressing a graph-dependent constant.

7 Completing the theorem, including pinnings

Combine the induced embedding from Lemma 5.1, the path transfer (7), and (12), with $N = 2\Delta n$. There are at most Δ source children, each contributing at most $p_{\max}\,\mathrm{collisionScale}(d, N)$. The source's diagonal term is 1. This proves (1) for every induced free-vertex set $S \subseteq V$ and every $u \in S$.

Now fix feasible occupied/unoccupied pins P, Q. The residual free set is

$$R = V \setminus (P \cup Q \cup N_G(P)).$$

The map $J \mapsto P \cup J$ is a bijection from independent sets of $G[R]$ to independent sets satisfying the original pins. Their weights differ by the common positive factor $\lambda^{|P|}$, so this map identifies the residual law exactly with the homogeneous hard-core law on R. The same factor cancels from both additional conditional distributions used to define an influence. Contributions of the occupied pins to an additive observable cancel from their difference.

Every vertex in R has a nondegenerate marginal: the empty residual configuration and its singleton both have positive weight. Vertices outside R are forced, with marginal either zero or one. Thus the nondegenerate source coordinates are exactly R, and fixed target coordinates contribute zero. The already proved free-set row inequality on R gives the original pinned row inequality. All bounds retain the original n, so they are uniform over the pinning. This completes Theorem 1.1.

8 A conditional consequence for critical Glauber dynamics

This section derives a mixing bound from a published proposition. The Markov-chain mixing theorem is not formalized in this package, and we could not re-examine the current text of that proposition while preparing this release, so the bound below is conditional on it as quoted.

The corrected Chen–Jiang theorem gives the subcritical spectral-independence bound ρ/δ, where

$$\rho = 2 + \frac{4}{\Delta - 2}, \qquad \lambda = (1 - \delta)\lambda_c(\Delta).$$

Their Proposition A.1 [2], based on the homogeneous reweighting localization framework of [4], requires this property for all pinnings. It does not require arbitrary site-dependent fugacities. Theorem 1.1 supplies an additional uniform cap $M_n = O_\Delta(\sqrt{n}\log n)$ for this entire homogeneous family. Use $K(\delta) = \min\{M_n, \rho/\delta\}$, with $K(0) = M_n$.

The published proposition bounds the mixing time by

$$Cn^2 \log\Delta \, \exp\left(\int_0^\theta \frac{K(\delta)}{1-\delta}\, d\delta\right),$$

when $(1 - \theta)\lambda_c(\Delta) \le 1/(2\Delta)$. Since $\lambda_c(\Delta) \le 12/\Delta$, the choice $\theta = 23/24$ meets this condition. For large enough n, $a = \rho/M_n \le \theta$ and

$$\begin{aligned}\int_0^\theta \frac{K(\delta)}{1-\delta}\, d\delta &= M_n \log\frac{M_n}{M_n - \rho} + \rho \log\frac{\theta(M_n - \rho)}{\rho(1-\theta)} \\ &\le 24\rho + \rho\log\frac{23M_n}{\rho}.\end{aligned}$$

The finitely many smaller n are absorbed into a degree-dependent constant. Therefore

$$T_{\mathrm{mix}} = O_\Delta\left(n^{3+2/(\Delta-2)}(\log n)^{2+4/(\Delta-2)}\right). \tag{13}$$

The corrected comparison bound in [2] is $O(n^{4+4/(\Delta-2)}\log\Delta)$. Equation (13) improves the polynomial exponent, conditionally on the quoted proposition. It is not a claim of optimal critical mixing.

9 Novelty assessment and significance

9.1 The version issue that must not be missed

The APPROX/RANDOM 2025 proceedings version of Chen–Jiang states a stronger $O(\sqrt{n})$ critical hard-core influence bound [3]. However, the authors' 7 January 2026 revision explicitly withdraws that claim in Remark 1.2: the required comparison between general graphs and a truncated tree quantity was incorrect [2]. The corrected paper retains the optimal subcritical result and an improved polynomial mixing bound. Search engines still prominently return the obsolete proceedings statement.

Our proof makes no such truncated-tree comparison. It retains the entire finite deletion tree, which may be exponentially large, and controls its weighted paths by counts of self-avoiding graph paths. Every use of the original graph size comes from actual vertices in a regular completion and the endpoint-packing argument. This is the specific logical difference from the withdrawn route.

9.2 Closest valid sources

Bencs et al. prove a nonbacktracking self-intersection bound and derive an Ising application [1]. Their Section 5 asks whether critical two-spin systems on bounded-degree regular graphs always have $\widetilde{O}(\sqrt{n})$ spectral independence. Theorem 1.1 addresses the hard-core case of that question, and covers maximum-degree graphs through regular completion. It does not settle the entire two-spin question.

The step that bounds the weighted tree paths by self-avoiding walk counts through a potential and the Cauchy–Schwarz inequality follows the connective-constant method of Sinclair, Srivastava, Štefankovič and Yin [9]. Pei [10] proves an $O(\sqrt{n})$ bound without the logarithm for ferromagnetic Ising models at the uniqueness threshold, uniformly in fields and pinnings; that result concerns a different model. Dozier [11] gives an independent birthday-paradox bound for nonbacktracking walks on regular graphs.

Efthymiou's work on the local connective constant and nonbacktracking matrices is close methodological prior art [6]. Its checked hard-core statements retain strict subcritical slack. It is important background for the potential/path-count viewpoint. We make no claim that the general use of self-avoiding expansions, potentials, or nonbacktracking counts to bound influence is new.

The 2026 edge-tilting field-dynamics paper [7] gives other critical mixing results and a spectral-independence estimate with a positive slack parameter; it does not supply the universal critical finite-size bound proved here. Peters–Regts–Reppekus [8] relate very strong spatial mixing, zero-freeness, and spectral independence. Their hypothesis is a uniform decay property of a graph family; it is not an unconditional critical square-root bound on all bounded-degree graphs.

9.3 What is and is not proposed as original

The proposed original mathematical contribution is the universal critical hard-core ℓ_∞ influence bound and its resulting quantitative consequences. Its proof aligns an exact Gibbs deletion expansion with the finite square-root survival estimate. The local contraction and the collision input are prior results, rederived rather than postulated in the formal development. Formalizing known ingredients is part of verification, not evidence of their novelty.

The result gives a finite-size critical correlation bound for the hard-core model on all graphs of maximum degree Δ, under every feasible pinning, and through existing machinery a better mixing exponent. Its ingredients are known; the main new observations are that the hard-core argument needs the sum of square roots of survival probabilities rather than an expected self-intersection time, and a tail form of the collision bound. The constant is large: the bound improves on the trivial bound n only for $n \geq 3.4 \times 10^6$ when $\Delta = 3$ and $n \geq 4.2 \times 10^6$ when $\Delta = 4$. Monte Carlo estimates on random cubic graphs with up to 10^5 vertices suggest that the $\sqrt{n}$ order of the survival sum is real and that the proved constant is roughly 40–70 times too large; these are numerical observations, not theorems.

The post-result review checked primary papers, corrected versions, conference proceedings, and related publications, with broader discovery searches over preprints, institutional repositories, thesis records, bibliographic indexes, MathOverflow, and GitHub. Search terms included hard-core/hardcore, critical/uniqueness threshold, total influence, ℓ_∞ spectral independence, self-avoiding paths, nonbacktracking walks, and square-root finite-size bounds. No valid earlier theorem with the same general conclusion was located in the reviewed sources. Searches cannot exclude unpublished work or guarantee exhaustive indexing. The accessible bibliographic coverage included DBLP; direct access to zbMATH and MathSciNet did not succeed, and some thesis PDFs were unavailable. The GitHub search is likewise not exhaustive. No priority claim is made beyond this documented search.

10 Formal correspondence and reproducibility

10.1 Source map

Files under `lean/CriticalHardcore/`	Mathematical role
`GibbsModel.lean`, `GibbsInfluence.lean`	Actual weighted independent-set sums, positive normalizers, conditional probabilities, and row-sum duality.
`DeletionRecursion.lean`, `ComputationTree.lean`, `ComputationGeometry.lean`	Exact neighbor deletion, terminating computation tree, valid states, and self-avoiding labels.
`LocalContraction.lean`, `WeightedTree.lean`, `WeightedTotals.lean`, `WeightedPaths.lean`	AM–GM contraction, second-moment energy, weighted path sums, and sequence injectivity.
`PathCountBridge.lean`, `InfluenceTransfer.lean`	Injection into actual graph path counts and transfer back to actual Gibbs influences.

Files under `lean/CriticalHardcore/`	Mathematical role
`RegularCompletion.lean`	Explicit induced embedding into a Δ-regular graph on $2\Delta n$ vertices.
`Nonbacktracking*.lean`	Actual graph walks, uniform choice-word bijections, endpoint packing, local collision, normalized survival, and adaptive trials.
`FirstDanger.lean`, `AdaptiveTrials.lean`, `WordTrial*.lean`	Injective killing-word encoding and finite history-dependent budget induction.
`CollisionScale.lean`, `CollisionBound.lean`, `ScaleInterpretation.lean`	Explicit square-root/logarithm scale and the completed actual-graph walk bound.
`Pinning.lean`, `SpectralStatement.lean`, `Main.lean`, `Audit.lean`	Actual pinning laws, nondegenerate coordinates, central theorem, and executable statement/axiom audit.

10.2 Pinned environment and clean build

The compiler is Lean 4.19.0, commit `6caaee842e94`. Mathlib is pinned to

`c44e0c8ee63ca166450922a373c7409c5d26b00b`.

The included `lake-manifest.json` also pins all eight transitive packages. The verification script checks the compiler version, all package commits, and clean tracked dependency sources. It copies the project proofs into a fresh directory and rebuilds every local module in dependency order with `autoImplicit=false`. Its import path excludes all previously built project objects. Pinned compiled dependency caches are reused read-only; this is not a claim to rebuild the Lean compiler and all of Mathlib from source in the present session.

From a clean checkout of the delivered `lean/` directory on Linux or WSL, with Elan, Python 3, Git, curl, tar, and network access available, run:

```
lake exe cache get
python3 verify.py
```

The toolchain file selects Lean 4.19.0, and the manifest records exact dependency revisions. For an already downloaded pinned cache, run instead:

```
python3 verify.py --lean /path/to/lean-4.19.0/bin/lean \
  --mathlib /path/to/pinned/mathlib4
```

The verifier creates three audit files next to the project:

- `verification/build.log`;
- `verification/axioms.txt`;
- `verification/certificate.json`.

The certificate records every module result, source SHA-256 digests, dependency commits, and central axiom dependencies. The included certificate is evidence for the delivered source; rerunning the script produces a new certificate for the local copy.

No `sorry`, `admit`, custom `axiom`, unchecked native decision shortcut, or trust-level override is allowed. The only permitted axioms are the standard Lean foundations `propext`, `Classical.choice`, and `Quot.sound`. Their names are reported explicitly rather than hidden behind an assertion of being "axiom-free."

10.3 Independent checks and remaining uncertainty

The referee check described in Section 11 covered root versus nonroot degree, path multiplicity, recursion termination, positive conditioning events, forced coordinates, oriented-edge lengths, adaptive trial dependence, and the regular completion's exact size, and found no error. These are AI checks, not external peer review. The central inequality and its square-root/logarithm interpretation are the formal claims. The mixing consequence imports published analysis and is not independently Lean formalized. The originality assessment remains bounded by the literature coverage documented here. The proof does not establish optimal hard-core scaling, arbitrary inhomogeneous activity, or the full critical two-spin conjecture.

11 Research process and AI disclosure

This is fully autonomous agentic research. Dev D. Goyal is responsible for publication. The work has not been externally peer reviewed, no expert validation is claimed, and the novelty assessment is provisional.

The research, the proof, the Lean formalization and an earlier version of this note were produced by autonomous AI research agents in an earlier session. The records of that session state that the note was prepared with ChatGPT; the model version is not recorded.

Opus 5.5 by Anthropic used with Claude Code assessed the package for this release. It coordinated an independent AI referee check, which recomputed exact influence row sums by brute force on 57 small graphs at the critical activity, with and without pinnings, checked the collision ingredients on small regular graphs, and compared the Lean statements with the claims. It ran the fresh Lean replay described in Section 2 and revised this manuscript: it added the credit to Sinclair, Srivastava, Štefankovič and Yin, the comparison with recent Ising results, the discussion of the constant, and the conditional status of the mixing corollary. On 1 November 2026 it also revised the front matter and moved this section to the end, without changing the mathematics. AI referee checks are not human peer review.

Lean verification establishes only the formal statements described in Section 2; the remaining uncertainty and limits are stated in Section 10.3.

References

[1] F. Bencs, L. A. Goldberg, M. Jenssen, M. Jerrum, G. Pete, G. Regts, and Y. Yin. *On the self-intersection time of non-backtracking random walks.* arXiv:2608.09729v1, 10 August 2026. https://arxiv.org/abs/2608.09729.

[2] Z. Chen and T. Jiang. *Improved Mixing of Critical Hardcore Model*. arXiv:2505.07515v2, 7 January 2026. In particular Remark 1.2, the corrected subcritical theorem, and Proposition A.1. https://arxiv.org/html/2505.07515v2.

[3] Z. Chen and T. Jiang. *Improved Mixing of Critical Hardcore Model*. APPROX/RANDOM 2025, LIPIcs 353, Article 51. https://doi.org/10.4230/LIPIcs.APPROX/RANDOM.2025.51. The critical $O(\sqrt{n})$ conclusion in this version was withdrawn in [2]; the local contraction is independently checked in this note.

[4] X. Chen, Z. Chen, Y. Yin, and X. Zhang. *Rapid Mixing at the Uniqueness Threshold*. arXiv:2411.03413v3, 7 January 2026; STOC 2025 proceedings version. Use the corrected arXiv version for current claims. https://arxiv.org/abs/2411.03413.

[5] L. Li, P. Lu, and Y. Yin. *Correlation Decay up to Uniqueness in Spin Systems*. SODA 2013, pp. 67–84; revised full version arXiv:1111.7064v3. https://arxiv.org/abs/1111.7064.

[6] C. Efthymiou. *On sampling two spin models using the local connective constant*. arXiv:2411.08179v2; SODA 2026. https://arxiv.org/html/2411.08179v2. https://doi.org/10.1137/1.9781611978971.39.

[7] X. Chen, Z. Ju, T. Miao, Y. Yin, and X. Zhang. *Edge-Tilting Field Dynamics: Rapid Mixing at the Uniqueness Threshold and Optimal Mixing for Swendsen-Wang Dynamics*. arXiv:2604.10525v2, May 2026. https://arxiv.org/abs/2604.10525.

[8] H. Peters, G. Regts, and J. Reppekus. *Decay of correlations and zeros for the hard-core model*. arXiv:2603.17858v2, 15 June 2026. https://arxiv.org/html/2603.17858v2.

[9] A. Sinclair, P. Srivastava, D. Štefankovič, and Y. Yin. *Spatial mixing and the connective constant: optimal bounds*. Probability Theory and Related Fields 168 (2017), 153–197.

[10] Y. R. Pei. *Critical and near-critical influence bounds for ferromagnetic Ising models*. arXiv:2609.20866, September 2026. https://arxiv.org/abs/2609.20866.

[11] B. Dozier. *The Birthday Paradox for non-backtracking walks on regular graphs*. arXiv:2608.26321, August 2026. https://arxiv.org/abs/2608.26321.

[12] D. Weitz. *Counting independent sets up to the tree threshold*. STOC 2006, pp. 140–149. https://doi.org/10.1145/1132516.1132538.

[13] N. Anari, K. Liu, and S. Oveis Gharan. *Spectral Independence in High-Dimensional Expanders and Applications to the Hardcore Model*. FOCS 2020; SIAM Journal on Computing. https://doi.org/10.1137/20M1367696.