AVALANCHE SHARPNESS AND THE STATIONARY STATE OF THE BAK–SNEPPEN MODEL ON $\mathbb{Z}^d$

OLIVIER HÉNARD

ABSTRACT. We establish the stationary picture for the nearest-neighbour Bak–Sneppen model on $\mathbb{Z}^d$, for every $d \geq 1$. There exists a critical level t_c such that, as $N \to \infty$, every fixed family of stationary fitnesses on the torus $(\mathbb{Z}/N\mathbb{Z})^d$ converges, in exponential coordinates, to independent random variables with law $t_c + \mathrm{Exp}(1)$.

The proof relies on three estimates for avalanches. For $d \geq 2$, these replace the interval structure available in dimension one. We show that the threshold at which the mean avalanche range diverges coincides with the threshold for the appearance of an infinite avalanche, using the hierarchical structure of avalanches and an estimate on the creation of new sites under translation. We then prove that an avalanche reaching a macroscopic distance on the torus can, after any fixed positive increase of the level, be amplified into one covering the whole torus. Finally, below criticality, finite mean range yields exponential integrability of the avalanche range, giving a finite–infinite volume comparison.

These three ingredients allow us to implement the Meester and Znamenski program that connects these avalanche estimates to the stationary law of the finite-volume process.

1. Introduction and main results

The Bak–Sneppen model assigns a fitness to each vertex of a finite graph. At every step, the vertex of smallest fitness and its neighbours receive fresh independent values. We use a fixed deterministic ordering of the vertices to break ties. Introduced as a model of self-organized criticality [2], it is expected to develop a stationary fitness threshold, with excursions below that threshold occurring in avalanches on many scales. Fixing a level isolates these avalanches, whose hierarchical structure has guided the analysis of the model [8, 12].

Meester and Znamenski related avalanche criticality to the stationary large-volume limit [9, 10]; see also [4, 5, 13, 15]. Their program reduces the stationary product limit to the coincidence of three avalanche thresholds—finite mean duration, finite mean range, and the appearance of an infinite avalanche—together with suitable finite-volume estimates. We give this reduction in Section 6.

For the one-dimensional one-sided model, the range size admits a closed scalar evolution [6]. In the isotropic one-dimensional model, the range remains an interval, but the scalar description based on the graft length alone no longer closes, since the locations of the two endpoints matter. For $d \geq 2$,

Key words and phrases. Bak–Sneppen model, avalanche ranges, self-organized criticality.

avalanche geometry is more complex and largely unknown, and no comparable scalar description is available. Our aim is to replace these one-dimensional reductions by estimates that remain effective in arbitrary dimension. A key quantity is the overlap of independent avalanche ranges; the main point is that an upper bound on its expectation already suffices for sharpness.

We consider the model on the discrete torus

$$\Lambda_N = (\mathbb{Z}/N\mathbb{Z})^d, \qquad d \geq 1,$$

and use exponential coordinates for the fitnesses. Thus every refreshed fitness has law Exp(1), and an update at x refreshes the sites in

$$x + \mathcal{N}_d, \qquad \mathcal{N}_d := \{0, \pm e_1, \dots, \pm e_d\}.$$

Let π_N be the stationary law of this finite-volume Markov chain, and, for $x \in \mathbb{Z}^d$, let $F_{N,x}$ denote the stationary fitness at the image of x in Λ_N.

Our main result identifies the large-volume stationary state.

Theorem 1.1 (Stationary limit)**.** *There exists $t_c \in (0, \infty)$ such that, for every fixed collection of distinct sites $x_1, \dots, x_k \in \mathbb{Z}^d$,*

$$(F_{N,x_1}, \dots, F_{N,x_k}) \xrightarrow[N\to\infty]{\mathrm{d}} (t_c + E_1, \dots, t_c + E_k),$$

where $E_1, \dots, E_k$ are independent Exp(1) *random variables.*

For $y \geq 0$, let ν_y denote the law of $y + E$, where $E \sim \mathrm{Exp}(1)$. Under the increasing change of variables

$$u = 1 - e^{-f},$$

Theorem 1.1 is equivalent to the usual uniform representation: the limiting fitnesses are independent and uniform on

$$[1 - e^{-t_c}, 1].$$

The critical level t_c is characterized through infinite-volume avalanches. Fix $t \geq 0$, start from a configuration in which all fitnesses are larger than t, and force an update at the origin. Thereafter, as long as some current fitness is at most t, update the site carrying the smallest fitness and its neighbours. The procedure stops when all current fitnesses are again larger than t. The set of sites refreshed during this procedure (including the forced initial update) is the range of the t-avalanche.

Let $\xi(t)$ denote this range on $\mathbb{Z}^d$, let $\eta(t)$ be its duration, including the forced update and with value ∞ if the avalanche does not terminate, and put

$$r(t) := |\xi(t)|, \qquad R(t) := \mathbb{E}r(t), \qquad D(t) := \mathbb{E}\eta(t), \qquad \theta(t) := \mathbb{P}(r(t) = \infty).$$

We use throughout the natural monotone coupling in the level parameter, under which $\xi(s) \subseteq \xi(t)$ for $s \leq t$. Avalanche ranges are connected, since every newly refreshed site is a nearest neighbour of a site already reached by the avalanche.

Following Meester and Znamenski, introduce the three critical levels

$$t_{\mathrm{d}} := \sup\{t : D(t) < \infty\}, \quad t_{\mathrm{sus}} := \sup\{t : R(t) < \infty\}, \quad t_\infty := \inf\{t : \theta(t) > 0\}.$$

Every update refreshes at most $2d+1$ sites, hence $r(t) \le (2d+1)\eta(t)$, while positive mass at infinity forces infinite mean range. Therefore

$$t_{\mathrm{d}} \le t_{\mathrm{sus}} \le t_\infty.$$

Both inequalities could a priori be strict. Our first result identifies all three levels.

Theorem 1.2 (Avalanche sharpness)**.** *For the nearest-neighbour Bak–Sneppen model on $\mathbb{Z}^d$, $d \ge 1$,*

$$t_{\mathrm{d}} = t_{\mathrm{sus}} = t_\infty =: t_c.$$

Moreover, for every $t < t_c$,

$$t_c \le t + \frac{10}{R(t)}. \tag{1}$$

Remark 1.3. The equality

$$t_{\mathrm{d}} = t_{\mathrm{sus}}$$

was already derived by Meester and Znamenski [10, Theorem 3.3 and Section 6], through the range–duration relation, stated here as Lemma 2.2. The novelty here is the identification $t_{\mathrm{sus}} = t_\infty$, together with the quantitative bound (1).

Remark 1.4. Finiteness of t_c follows directly from (1), taking $t = 0$, since $R(0) = 2d + 1$. Positivity of t_c is elementary: in the exploration of a t-avalanche, each update produces $2d + 1$ fresh fitnesses, each of which is at most t with probability $1 - e^{-t}$. Ignoring overlaps therefore gives a domination by a Galton–Watson process with offspring distribution

$$\mathrm{Bin}(2d + 1, 1 - e^{-t}).$$

This process is subcritical whenever $(2d + 1)(1 - e^{-t}) < 1$, which implies finite mean avalanche duration for all sufficiently small $t > 0$, and hence $t_c > 0$.

For $d = 1$, Theorem 1.2 already suffices to complete the Meester–Znamenski program [9], recalled in Section 6, and derive the stationary-limit result. In dimensions $d \ge 2$, two additional finite-volume ingredients are required.

The first ingredient concerns avalanches on the finite torus. Let $\xi_N(t)$ denote the range of the t-avalanche on Λ_N, and set

$$r_N(t) := |\xi_N(t)|, \qquad R_N(t) := \mathbb{E} r_N(t), \qquad P_N(t) := \mathbb{P}(\xi_N(t) = \Lambda_N).$$

The following result says that once a supercritical avalanche has reached a macroscopic distance, any further fixed increase of the level amplifies it into complete coverage of the torus.

Theorem 1.5 (Macroscopic amplification)**.** *Fix*

$$t_c < s < t \qquad \text{and} \qquad 0 < \rho < \frac{1}{2}.$$

Let

$$E_N(s) := \{\xi_N(s) \cap \{x : d_{\Lambda_N}(0, x) = \lfloor \rho N \rfloor\} \ne \varnothing\}.$$

Then

$$\mathbb{P}\big(\xi_N(t) = \Lambda_N \mid E_N(s)\big) \longrightarrow 1.$$

Consequently,

$$\liminf_{N\to\infty} P_N(t) \geq \theta(s) > 0.$$

The complementary statement below shows that, in the subcritical regime, an avalanche is typically too small to detect the periodic geometry of the torus.

Theorem 1.6 (Subcritical exponential moments and locality)**.** *For every $t < t_c$, there exists $\lambda_t > 0$ such that*

$$\mathbb{E}[e^{\lambda_t r(t)}] < \infty.$$

In particular, all polynomial moments of $r(t)$ are finite. Moreover,

$$\sup_{0\leq u\leq t} |R_N(u) - R(u)| \longrightarrow 0.$$

1.1. **Proof strategy.** The proof for general $d \geq 1$ is organized around the following three mechanisms.

1. Infinite-volume sharpness: creation under translation. The level parameter gives avalanches a hierarchical structure. Conditionally on a finite current range A, each site $x \in A$ produces marks at rate one, each grafting an independent avalanche translated by x. The relevant geometric quantity is the part of this translated range lying outside A.

The deterministic input is

$$\sum_{x\in A} |(x + B) \setminus A| \geq c|A||B|, \qquad 0 \in B, \quad |B| \geq |A| \geq 3,$$

which follows from a translation-invariant total order on $\mathbb{Z}^d$. Applied to

$$\psi(t, \lambda) := \mathbb{E}[1 - e^{-\lambda r(t)}],$$

the hierarchical decomposition gives the Burgers-type inequality

$$\partial_t\psi(t, \lambda) \geq \frac{1}{10}\psi(t, \lambda)\partial_\lambda\psi(t, \lambda).$$

Since $\psi(t, 0+) = \theta(t)$, while $\partial_\lambda\psi(t, 0+) = R(t)$ whenever $r(t) < \infty$ almost surely, the characteristic flow prevents the boundary slope from becoming infinite strictly before positive boundary mass appears. This yields Theorem 1.2.

2. Supercritical amplification: translation expansion. An avalanche reaching distance of order N need not resemble a ball or be well mixed on the torus. We retain only one nearest-neighbour edge crossing each successive radius.

For $A \subset \Lambda_N$, let

$$b_A(z) := |\{x \in A : x + z \notin A\}|$$

be its translation boundary. It satisfies

$$b_A(z + w) \leq b_A(z) + b_A(w), \qquad \sum_{z\in\Lambda_N} b_A(z) = |A|(N^d - |A|).$$

Fix $0 < \rho < 1/2$, and let E be the event that an independent avalanche reaches distance ρN. On E, select uniformly one crossing edge at each radius $1 \leq j \leq L := \lfloor \rho N \rfloor$. Let κ be their expected endpoint counting

measure, with zero contribution off E. The selected edges need not form a connected path.

At a fixed supercritical level, $\mathbb{P}(E)$ is bounded away from zero uniformly in N, so $\kappa(\Lambda_N) = 2L\mathbb{P}(E)$ is of order N. Hyperoctahedral symmetry, subadditivity of b_A, and the crossing edges then imply

$$\sum_z \kappa(z) b_A(z) \geq cN \, \frac{|A|(N^d - |A|)}{N^d}.$$

Thus the averaged endpoint measure supplies enough translations to expand any nontrivial set A.

To turn expansion into covering, we use the Lyapunov function

$$\Phi(A) := \frac{N^d}{|A|} - 1.$$

Conditionally on macroscopic reach at the initial level, the preceding estimate implies exponential decay of the mean of Φ at rate of order N. Since every noncovering range satisfies $\Phi \geq 1/(N^d - 1)$, any fixed positive sprinkling forces complete coverage with high (conditional) probability.

3. Subcritical locality: exponential moments replace exact geometry. Below t_c, the hierarchy gives a transport inequality for the truncated exponential transform of the avalanche range. Finite susceptibility can thus be bootstrapped to exponential integrability of $r(t)$, making ranges of order N exponentially unlikely.

Every connected subset of $\mathbb{Z}^d$ with fewer than N vertices projects injectively onto Λ_N. A subcritical avalanche is therefore overwhelmingly unlikely to detect the periodic geometry. This gives the finite–infinite volume comparison of Theorem 1.6.

These mechanisms provide the inputs for the locking-threshold representation; see (25) and [9, Section 2]. Sharpness identifies the candidate threshold, while amplification and subcritical locality control the locking thresholds from above and below. The stationary limit then follows from a dimension-free finite-volume argument.

Scope. Several parts of the argument extend beyond $\mathbb{Z}^d$, for instance to suitable Cayley graphs and their finite quotients. One could formulate the proof in such an abstract setting, replacing translations and the amplification argument by appropriate group-theoretic and spectral assumptions. We have chosen not to do so: the lattice $\mathbb{Z}^d$ already contains the essential difficulties, and its geometry makes the amplification mechanism particularly transparent.

1.2. **Methodological perspective.** Two analogies help place the proof. The first is with coagulation and gelation. For solvable Smoluchowski coagulation equations, Laplace transforms lead to first-order nonlinear equations related to inviscid Burgers, while gelation corresponds to mass escaping from the finite-cluster distribution; see [1, 11]. Spatial overlaps prevent a closed equation for the avalanche law, but the translation estimate still yields

$$\partial_t \psi \geq c \, \psi \, \partial_\lambda \psi.$$

This one-sided inequality is enough for the characteristic argument: divergence of the susceptibility is transported into positive boundary mass at a higher level, signalling an infinite avalanche.

The second analogy concerns amplification. The measure κ may be viewed as the jump measure of a translation walk on Λ_N, and the expansion estimate is equivalent, up to constants, to a spectral-gap bound for this auxiliary walk; see Remark 4.4 and compare [3]. Here the estimate follows directly from subadditivity of the translation boundary and the selected crossing edges, without any shape or mixing assumption on the macroscopic range.

Disclosure on the use of AI. Human was used at the level of an annoyingly slow co-author (see Appendix A for a documentation of our precise interaction with AI). Joking aside, in view of the extent of the AI contribution, we do not intend to submit this work for publication.

Relation with [6]. The present paper supersedes the first version of [6], whose main goal was to establish sharpness in dimension 1. Accordingly, that work has been refocused on quantitative bounds for the critical tail of the avalanche range for the anisotropic model. This is a genuinely different question, on which the present method, despite its strength for proving sharpness, has nothing to say.

1.3. **Organization of the paper.** Section 2 gives the avalanche hierarchy and its basic identities. Section 3 proves infinite-volume sharpness. Section 4 establishes the finite-volume expansion inequality and the amplification theorem. Section 5 proves exponential moments of the subcritical range and the finite–infinite volume comparison. The stationary limit is deduced in Section 6.

2. Avalanche hierarchy and basic identities

We collect here the structural facts about avalanches which are common to all three parts of the proof.

Recall that $\xi(t)$ and $\xi_N(t)$ denote the ranges of the t-avalanches rooted at the origin in $\mathbb{Z}^d$ and Λ_N, respectively. We use the standard monotone graphical construction in the level parameter, so that

$$\xi(s) \subseteq \xi(t), \qquad s \le t.$$

The key property is regeneration in the level variable. We use the censored exploration filtration $(\mathcal{H}_u)_{u\ge 0}$: during the exploration of a u-avalanche, refreshed values at most u are revealed exactly, whereas values larger than u are revealed only through the event that they exceed u. The σ-field $\mathcal{H}_u$ records the explored genealogy, the identities of the refresh variables encountered, and this partially revealed information. The monotone graphical construction makes this filtration increasing.

Proposition 2.1 (Hierarchical avalanche decomposition)**.** *Consider either the infinite lattice or the finite torus. On the event that the range A at level u is finite, conditionally on $\mathcal{H}_u$, the terminal fitnesses at its sites are independent with law ν_u. Each site $x \in A$ therefore carries an independent*

rate-one clock in the level parameter. At a mark rooted at x and level u, an independent u-avalanche $\xi'(u)$ is grafted onto the current range:

$$A \longmapsto A \cup (x + \xi'(u)).$$

The graft is independent of the censored past. In infinite volume this description is used on finite-range states, up to explosion, or with the process stopped when the range reaches a prescribed finite size.

Proof. An avalanche with finite range has finite duration almost surely. Indeed, while the range has at most K sites, there are at most K active fitnesses at or below u. A block of K updates in which all fresh values exceed u removes these active fitnesses without creating new ones, and therefore forces termination within the block. Such a block has conditional probability at least

$$e^{-(2d+1)uK} > 0.$$

The probability of completing m blocks without either terminating or reaching more than K sites is consequently at most $(1 - e^{-(2d+1)uK})^m$. Letting $m \to \infty$, and then taking the union over K, proves the assertion.

The censored exploration of this finite avalanche identifies the last refresh variable assigned to each reached site. These variables are distinct, and each has been observed only to exceed u. Independence and the lack-of-memory property give independent terminal values $u + \mathrm{Exp}(1)$, conditionally on $\mathcal{H}_u$. Their excesses supply the independent level clocks. At the next triggering level, fresh refresh variables generate an independent avalanche at that level, translated to the triggering site. Iterating this argument gives the stated grafting description.

Equivalently, on the torus over any finite level interval, and in infinite volume over $[0, T]$ with $T < t_\infty$, independently for each site x, take Poisson marks (u, B) with intensity $du\,\mathbb{P}(\xi_N(u) \in dB)$ on the torus and $du\,\mathbb{P}(\xi(u) \in dB)$ in infinite volume. Starting from $\Xi_0 = \mathcal{N}_d$, retain a mark (u, B) whenever x belongs to the current range, in which case

$$\Xi_u = \Xi_{u-} \cup (x + B).$$

This realizes the preceding level-by-level description. On the torus it is directly non-explosive. In infinite volume, monotonicity gives

$$r(u) \le r(T) < \infty \qquad \text{a.s.,} \qquad u \le T,$$

and one may equivalently stop when the range reaches K sites and then let $K \to \infty$. □

Thus, conditionally on a current range A, marks occur at total rate $|A|$. This observation is the common source of the compensation identity of Section 3, the growth estimates of Section 4, and the transport inequality of Section 5. We shall also use the same construction with each graft decorated by the duration of the corresponding independent avalanche; this gives the range–duration identity below.

Let $\eta_N(t)$ denote the duration of the finite-volume t-avalanche. With $\eta(t)$ and $D(t)$ as above, set

$$D_N(t) := \mathbb{E}\eta_N(t).$$

Lemma 2.2 (Range–duration identity)**.** *For every $t < t_{\mathrm{sus}}$, $D(t) < \infty$, and on the torus $D_N(t) < \infty$ for every finite t. Moreover, R is continuous on $[0, t_{\mathrm{sus}})$, R_N is continuous on $[0, \infty)$, and D and D_N are C^1 on these respective domains, with*

$$D'(t) = R(t)D(t), \qquad D'_N(t) = R_N(t)D_N(t),$$

and

$$D(0) = D_N(0) = 1.$$

Consequently,

$$D(t) = \exp\left\{\int_0^t R(u)\,du\right\}, \qquad D_N(t) = \exp\left\{\int_0^t R_N(u)\,du\right\}. \tag{2}$$

Proof. Fix $T < t_{\mathrm{sus}}$, so that $R(u) \le R(T) < \infty$ for $u \le T$. In the decorated hierarchical construction a mark at level u adds exactly the duration $\eta'(u)$ of an independent u-avalanche. For $L \ge 1$, put $D_L(u) := \mathbb{E}[\eta(u) \wedge L]$. Since

$$(\eta + \eta') \wedge L - \eta \wedge L \le \eta' \wedge L,$$

compensation, first stopped when the range reaches K sites, gives after letting $K \to \infty$

$$D_L(t) \le 1 + \int_0^t R(u)D_L(u)\,du, \qquad 0 \le t \le T.$$

Here the rate of marks is the current range size, and the graft is independent of the censored past. The stopping can be removed because $r(T) < \infty$ almost surely, the observable is bounded by L, and its expected compensator is bounded by $LR(u)$. Gronwall's inequality yields

$$D_L(t) \le \exp\left\{\int_0^t R(u)\,du\right\}.$$

Monotone convergence as $L \to \infty$ proves $D(t) < \infty$, uniformly on $[0, T]$. The same argument on the torus uses $R_N(u) \le N^d$ and proves $D_N(t) \le e^{N^d t} < \infty$ on every bounded level interval.

We next record the continuity of the mean ranges. Couple the avalanches $\xi(u)$, $0 \le u \le T$, through the same hierarchical construction. Since $r(T) < \infty$ almost surely, only finitely many relevant marks occur below level T, and no deterministic level carries a mark. Hence, for any deterministic $t_n \to t$ in $[0, T]$, one has $\xi(t_n) = \xi(t)$, and therefore $r(t_n) = r(t)$, eventually almost surely. Since $r(t_n) \le r(T)$ and $\mathbb{E}r(T) = R(T) < \infty$, dominated convergence yields $R(t_n) \to R(t)$. Thus R is continuous on $[0, t_{\mathrm{sus}})$. The same argument on the torus is immediate from $r_N(T) \le N^d$, and shows that R_N is continuous on every bounded level interval.

We may now apply compensation to the full duration. Tonelli's theorem gives

$$D(t) = 1 + \int_0^t R(u)D(u)\,du, \qquad D_N(t) = 1 + \int_0^t R_N(u)D_N(u)\,du.$$

In infinite volume the stopping by range size is again removed by monotone convergence of the accumulated duration and compensator. The integral identities first give continuity of D and D_N. Together with the continuity of R and R_N, the fundamental theorem of calculus then yields the stated C^1

identities. Since $D, D_N \geq 1$, integrating $(\log D)' = R$ and $(\log D_N)' = R_N$ gives (2). □

3. Infinite-volume sharpness

We prove Theorem 1.2. The hierarchical avalanche decomposition provides an exact compensation formula for a bounded transform of the range, while a deterministic translation estimate ensures that an independent graft creates sufficiently many new sites on average. Their combination yields a Burgers-type differential inequality.

For $\lambda > 0$, put

$$\psi(t, \lambda) := \mathbb{E}[1 - e^{-\lambda r(t)}]$$

with the convention

$$e^{-\lambda \infty} = 0.$$

Lemma 3.1 (Avalanche compensation)**.** *For every $\lambda > 0$ and $0 \leq s < t < t_\infty$,*

$$\psi(t, \lambda) - \psi(s, \lambda) = \int_s^t H(u, \lambda)\, du, \tag{3}$$

where, if $\xi'(u)$ is an independent copy of $\xi(u)$,

$$H(u, \lambda) := \mathbb{E}\left[e^{-\lambda r(u)} \sum_{x \in \xi(u)} \left(1 - e^{-\lambda |(x + \xi'(u)) \setminus \xi(u)|}\right) \right]. \tag{4}$$

In particular, $t \mapsto \psi(t, \lambda)$ is locally Lipschitz on $[0, t_\infty)$. The functions H and $\partial_\lambda \psi$ are jointly continuous on $[0, t_\infty) \times (0, \infty)$, and

$$\partial_t \psi(t, \lambda) = H(t, \lambda), \qquad \partial_\lambda \psi(t, \lambda) = \mathbb{E}[r(t) e^{-\lambda r(t)}].$$

The t-derivative at $t = 0$ is understood from the right.

Proof. Fix $T < t_\infty$, and stop the hierarchical construction when the range reaches K sites. Before this stopping time the total rate of relevant marks is at most K, so compensation applies without any integrability assumption.

For

$$F_\lambda(A) := 1 - e^{-\lambda |A|},$$

a graft B rooted at $x \in A$ produces the increment

$$F_\lambda(A \cup (x + B)) - F_\lambda(A) = e^{-\lambda |A|} \left(1 - e^{-\lambda |(x+B) \setminus A|}\right).$$

Compensation for the stopped construction therefore gives the stopped version of (3).

The compensator is uniformly bounded in the stopping level, since

$$|A| e^{-\lambda |A|} \leq \frac{1}{e\lambda}.$$

Moreover,

$$\mathbb{P}(r(T) \geq K) \longrightarrow 0.$$

Dominated convergence as $K \to \infty$ proves (3) and (4). The same bound gives

$$|\psi(t, \lambda) - \psi(s, \lambda)| \leq \frac{|t - s|}{e\lambda},$$

hence local Lipschitz continuity.

For joint continuity, couple two independent range processes on $[0,T]$, with $T < t_\infty$. The pathwise local-constancy argument established in the proof of Lemma 2.2 applies verbatim: since $T < t_\infty$, the terminal ranges are finite almost surely, and for any deterministic $t_n \to t$ in $[0,T]$, both ranges are eventually unchanged almost surely. No moment assumption is needed here. If also $\lambda_n \to \lambda > 0$, the integrand defining H converges almost surely and is bounded by $1/(e\lambda_*)$, where $\lambda_* > 0$ is a common lower bound for the exponents. Dominated convergence gives joint continuity of H. The same bound justifies differentiation of ψ under the expectation and joint continuity of $\partial_\lambda \psi$; the integrand of ψ itself is bounded by one. Finally, (3) and continuity of H yield $\partial_t \psi = H$ at every level in the stated domain. □

The geometric input is the following elementary inequality.

Lemma 3.2 (Average creation under translation)**.** *Let $A, B \subset \mathbb{Z}^d$ be finite, with $0 \in B$, and write*

$$a := |A|, \qquad b := |B|.$$

If $b \geq a \geq 3$, then

$$\sum_{x \in A} |(x + B) \setminus A| \geq \frac{1}{5} ab. \tag{5}$$

Proof. Write $C(A,B) := \sum_{x\in A} |(x+B)\setminus A|$, and first assume $b = a$. Equip $\mathbb{Z}^d$ with the translation-invariant lexicographic order, and write

$$A = \{a_1 < \cdots < a_a\}, \qquad p := |\{z \in B : z > 0\}|, \qquad q := |\{z \in B : z < 0\}|.$$

Then $p + q = a - 1$. From a_i, at most $a - i$ positive displacements and $i - 1$ negative displacements can land in A. Hence

$$C(A,B) \geq \sum_{i=1}^{a} \big((p - a + i)_+ + (q - i + 1)_+\big) = \frac{p(p+1) + q(q+1)}{2} \geq \left\lfloor \frac{a^2}{4} \right\rfloor.$$

Here $v_+ := \max\{v, 0\}$.

For general $b \geq a$, choose uniformly a subset $B_0 \subset B$ of size a containing zero. The zero displacement contributes nothing, while every other displacement is retained with probability $(a-1)/(b-1)$. Applying the equal-cardinality case gives

$$\frac{a-1}{b-1} C(A,B) = \mathbb{E}[C(A,B_0)] \geq \left\lfloor \frac{a^2}{4} \right\rfloor.$$

Since $(b-1)/(a-1) \geq b/a$ and $a \geq 3$,

$$C(A,B) \geq \frac{b}{a} \left\lfloor \frac{a^2}{4} \right\rfloor \geq \frac{1}{5} ab,$$

as claimed. □

Remark 3.3. The proof of Lemma 3.2 uses only the existence of a translation-invariant total order and the cardinalities of A and B. In particular, neither isoperimetry nor any metric property of $\mathbb{Z}^d$ enters the infinite-volume sharpness argument.

We now combine the two lemmas.

Proposition 3.4 (Burgers comparison)**.** *For every $0 \le s < t < t_\infty$ and $\lambda > 0$,*

$$\psi(t,\lambda) - \psi(s,\lambda) \ge \frac{1}{10}\int_s^t \psi(u,\lambda)\,\partial_\lambda\psi(u,\lambda)\,du. \tag{6}$$

Equivalently,

$$\partial_t\psi(t,\lambda) \ge \frac{1}{10}\psi(t,\lambda)\partial_\lambda\psi(t,\lambda)$$

for every $0 \le t < t_\infty$.

Proof. Fix $u < t_\infty$, and let

$$A := \xi(u), \qquad B := \xi'(u), \qquad X := |A|, \qquad Y := |B|.$$

Then X and Y are independent and identically distributed, finite almost surely, and

$$X, Y \ge 2d + 1 \ge 3.$$

For $x \in A$, put

$$d_x := |(x + B) \setminus A|.$$

On $\{Y \ge X\}$, Lemma 3.2 gives

$$\sum_{x\in A} d_x \ge \frac{1}{5}XY.$$

Since $r \mapsto 1 - e^{-\lambda r}$ is concave, vanishes at zero, and $0 \le d_x \le Y$,

$$1 - e^{-\lambda d_x} \ge \frac{d_x}{Y}(1 - e^{-\lambda Y}).$$

Consequently, on $\{Y \ge X\}$,

$$\sum_{x\in A}(1 - e^{-\lambda d_x}) \ge \frac{X}{5}(1 - e^{-\lambda Y}).$$

Lemma 3.1 therefore gives

$$H(u,\lambda) \ge \frac{1}{5}\mathbb{E}\left[Xe^{-\lambda X}(1 - e^{-\lambda Y})\mathbf{1}_{\{Y\ge X\}}\right]. \tag{7}$$

Set

$$f(r) := re^{-\lambda r}, \qquad g(r) := 1 - e^{-\lambda r}.$$

The ratio

$$\frac{f(r)}{g(r)} = \frac{r}{e^{\lambda r} - 1}$$

is decreasing on $(0,\infty)$. Hence, on $\{X < Y\}$,

$$f(Y)g(X) \le f(X)g(Y).$$

By exchangeability,

$$\begin{aligned}\mathbb{E}[f(X)g(Y)\mathbf{1}_{\{Y<X\}}] &= \mathbb{E}[f(Y)g(X)\mathbf{1}_{\{X<Y\}}] \\ &\le \mathbb{E}[f(X)g(Y)\mathbf{1}_{\{X<Y\}}].\end{aligned}$$

Therefore

$$\mathbb{E}[f(X)g(Y)\mathbf{1}_{\{Y\ge X\}}] \ge \frac{1}{2}\mathbb{E}[f(X)g(Y)].$$

Independence gives

$$\mathbb{E}[f(X)g(Y)] = \mathbb{E}[Xe^{-\lambda X}]\mathbb{E}[1 - e^{-\lambda X}] = \partial_\lambda\psi(u,\lambda)\psi(u,\lambda),$$

where differentiation under the expectation is justified by

$$re^{-\lambda r} \leq \frac{1}{e\lambda}.$$

Combining this with (7),

$$H(u,\lambda) \geq \frac{1}{10}\psi(u,\lambda)\partial_\lambda\psi(u,\lambda).$$

Integrating in (3) proves (6). □

The next lemma records the characteristic consequence of the Burgers comparison.

Lemma 3.5 (Comparison of the characteristics)**.** *Fix $t_0 < t_\infty$ and $\lambda_0 > 0$, and put*

$$q := \psi(t_0, \lambda_0).$$

Define

$$\lambda(s) := \lambda_0 - \frac{q}{10}(s - t_0)$$

as long as $\lambda(s) > 0$. Then

$$\psi(s, \lambda(s)) \geq q$$

for every $s \geq t_0$ such that

$$s < t_\infty, \qquad \lambda(s) > 0.$$

Proof. Fix $t_1 > t_0$ with $t_1 < t_\infty$ and $\lambda(t_1) > 0$, and put $v(s) := \psi(s, \lambda(s))$ for $s \in [t_0, t_1]$. Lemma 3.1 permits the ordinary chain rule. Proposition 3.4 therefore gives

$$v'(s) = \partial_t\psi(s,\lambda(s)) - \frac{q}{10}\,\partial_\lambda\psi(s,\lambda(s)) \geq \frac{1}{10}\partial_\lambda\psi(s,\lambda(s))(v(s) - q).$$

Set $a(s) := \frac{1}{10}\partial_\lambda\psi(s,\lambda(s))$, which is continuous on this interval. Multiplication by the integrating factor gives

$$\frac{d}{ds}\left[e^{-\int_{t_0}^{s} a(u)\,du}(v(s) - q)\right] \geq 0.$$

Since $v(t_0) = q$, it follows that $v(s) \geq q$ on $[t_0, t_1]$. The choice of t_1 was arbitrary. □

Proof of Theorem 1.2. Fix $t < t_\infty$ and $\lambda > 0$, and set

$$q := \psi(t, \lambda) > 0.$$

Consider

$$\lambda(s) = \lambda - \frac{q}{10}(s - t).$$

Its hitting time of zero is

$$s_* = t + \frac{10\lambda}{q}.$$

We claim that

$$s_* \geq t_\infty.$$

Indeed, if $s_* < t_\infty$, then Lemma 3.5 gives

$$\psi(s, \lambda(s)) \geq q, \qquad t \leq s < s_*.$$

Since $r(s_*) < \infty$ almost surely and $r(s) \le r(s_*)$,

$$0 \le \psi(s, \lambda(s)) \le \mathbb{E}[1 - e^{-\lambda(s) r(s_*)}] \longrightarrow 0 \qquad (s \uparrow s_*),$$

by dominated convergence, a contradiction. Therefore

$$t + \frac{10\lambda}{\psi(t,\lambda)} \ge t_\infty,$$

or equivalently

$$\frac{\psi(t,\lambda)}{\lambda} \le \frac{10}{t_\infty - t}.$$

Since $r(t) < \infty$ almost surely,

$$\frac{1 - e^{-\lambda r(t)}}{\lambda} \uparrow r(t) \qquad (\lambda \downarrow 0).$$

Monotone convergence yields

$$R(t) \le \frac{10}{t_\infty - t}.$$

Thus $R(t) < \infty$ for every $t < t_\infty$, and hence

$$t_{\mathrm{sus}} \ge t_\infty.$$

The reverse inequality is immediate. Therefore

$$t_{\mathrm{sus}} = t_\infty =: t_c.$$

It remains to identify the duration threshold. Since $r(t) \le (2d+1)\eta(t)$, finite mean duration implies finite mean range, so $t_{\mathrm{d}} \le t_{\mathrm{sus}}$. Conversely, Lemma 2.2 gives $D(t) < \infty$ for every $t < t_{\mathrm{sus}}$, hence $t_{\mathrm{d}} \ge t_{\mathrm{sus}}$. Thus

$$t_{\mathrm{d}} = t_{\mathrm{sus}} = t_\infty = t_c,$$

and (1) follows. □

Remark 3.6. The quantity

$$\psi(t,\lambda) = \mathbb{E}[1 - e^{-\lambda r(t)}]$$

may be viewed as the probability that the avalanche range is hit by an independent Poisson field of intensity λ. Infinite susceptibility corresponds to an infinite right slope at $\lambda = 0$, whereas an infinite avalanche corresponds to positive mass at the boundary:

$$\psi(t, 0+) = \theta(t).$$

The Burgers comparison forces positive values of ψ to propagate towards this boundary. Estimate (1) expresses quantitatively that the boundary slope cannot diverge strictly before positive boundary mass appears.

4. Finite-volume amplification

We prove Theorem 1.5. Fix

$$d \geq 1, \qquad t_c < s < t, \qquad 0 < \rho < \frac{1}{2},$$

and put

$$\Lambda_N := (\mathbb{Z}/N\mathbb{Z})^d, \qquad V := N^d, \qquad L := \lfloor \rho N \rfloor.$$

For $j \leq L$, identify

$$S_j := \{x \in \mathbb{Z}^d : \|x\|_1 = j\}$$

with its canonical image in Λ_N. Since $2L < N$ for all sufficiently large N, this identification is unambiguous.

Let

$$E_N(u) := \{\xi_N(u) \cap S_L \neq \varnothing\}.$$

The finite- and infinite-volume graphical constructions agree until the range first reaches S_L. Therefore

$$\mathbb{P}(E_N(u)) = \mathbb{P}(\xi(u) \cap S_L \neq \varnothing) \geq \theta(u). \tag{8}$$

In particular,

$$\mathbb{P}(E_N(u)) \geq \theta(s) > 0, \qquad u \geq s.$$

For $A \subset \Lambda_N$, write $a := |A|$ and define its translation boundary:

$$b_A(z) := |\{x \in A : x + z \notin A\}|, \qquad z \in \Lambda_N.$$

Since A and $A - z$ have equal cardinality,

$$b_A(z) = \frac{1}{2}|A \triangle (A - z)|, \qquad b_A(-z) = b_A(z), \qquad b_A(z+w) \leq b_A(z) + b_A(w). \tag{9}$$

Moreover,

$$\sum_{z \in \Lambda_N} b_A(z) = a(V - a), \tag{10}$$

since both sides count pairs in $A \times A^c$ according to their displacement.

Lemma 4.1 (Short translations control all translations)**.** *Let*

$$f : \mathbb{Z}/N\mathbb{Z} \to [0, \infty)$$

satisfy $f(0) = 0$ *and*

$$f(r + r') \leq f(r) + f(r').$$

For $L = \lfloor \rho N \rfloor$*, with* $0 < \rho < 1/2$ *fixed, there exists* $C_\rho < \infty$ *such that, for all sufficiently large* N*,*

$$\sum_{r \in \mathbb{Z}/N\mathbb{Z}} f(r) \leq C_\rho \sum_{j=1}^{L} f(j).$$

Proof. Set

$$S := \sum_{j=1}^{L} f(j).$$

For $1 \leq q < L$, subadditivity gives

$$f(L) \leq f(q) + f(L - q).$$

Summing over q yields

$$(L-1)f(L) \le 2S.$$

Since $L-1 \ge \rho N/2$ for all sufficiently large N, it follows that $f(L) \le 4S/(\rho N)$.

For $0 \le r < N$, write $r = mL + q$, where $0 \le q < L$. Then

$$f(r) \le mf(L) + f(q).$$

Also $L \ge \rho N/2$, hence $m \le N/L \le 2/\rho$, and each value of q occurs for at most $2/\rho + 1$ values of r. Therefore

$$\sum_{r \in \mathbb{Z}/N\mathbb{Z}} f(r) \le \frac{2N}{\rho} f(L) + \left(\frac{2}{\rho} + 1\right) S \le \left(\frac{8}{\rho^2} + \frac{2}{\rho} + 1\right) S.$$

□

Let

$$H_d = (\mathbb{Z}/2\mathbb{Z})^d \rtimes \mathfrak{S}_d$$

be the hyperoctahedral group, acting on Λ_N by sign changes and coordinate permutations.

Proposition 4.2 (Translation expansion)**.** *Let E be an event with positive probability. Suppose that on E, for every $1 \le j \le L$, we are given a nearest-neighbour edge*

$$(W_j, Z_j), \qquad W_j \in S_{j-1}, \quad Z_j \in S_j,$$

whose joint conditional law given E is H_d-invariant. Define

$$\kappa := \mathbb{E}\left[\mathbf{1}_E \sum_{j=1}^{L} (\delta_{W_j} + \delta_{Z_j})\right].$$

Then

$$\kappa(\Lambda_N) = 2L\mathbb{P}(E),$$

and there exists $c_{d,\rho} > 0$ such that, for every $A \subset \Lambda_N$, writing $a = |A|$, and all sufficiently large N,

$$\sum_{x \in A, z \in \Lambda_N} \kappa(z) \mathbf{1}_{\{x+z \notin A\}} \ge c_{d,\rho}\, \kappa(\Lambda_N) \frac{a(V-a)}{V}. \tag{11}$$

The constant is uniform in N, A, and in the law of the selected edges.

Proof. Symmetrize the translation boundary by putting

$$f(z) := \frac{1}{|H_d|} \sum_{g \in H_d} b_A(gz), \qquad I := \sum_z \kappa(z) b_A(z).$$

The function f is subadditive and H_d-invariant, with $f(0) = 0$ and $f(-z) = f(z)$. Invariance of κ and (10) give

$$I = \sum_z \kappa(z) f(z), \qquad \sum_{z \in \Lambda_N} f(z) = a(V-a).$$

Fix $z \in S_j$, and let $a_1, \dots, a_d$ be its absolute coordinates, so that $\sum_i a_i = j$. Permute the coordinates to put a_i first, then reflect that coordinate.

The difference of the two resulting vectors is $2a_i e_1$, and subadditivity and invariance give

$$f(2a_i e_1) \le 2f(z).$$

Consequently,

$$f(2je_1) \le \sum_{i=1}^{d} f(2a_i e_1) \le 2d\, f(z),$$

and the endpoints Z_j yield

$$I \ge \frac{\mathbb{P}(E)}{2d} \sum_{j=1}^{L} f(2je_1).$$

For each selected edge, $Z_j - W_j = \pm e_\ell$, so

$$f(e_1) = f(Z_j - W_j) \le f(Z_j) + f(W_j).$$

Summing gives $I \ge \mathbb{P}(E) L f(e_1)$. Together the two estimates imply

$$I \ge \frac{\mathbb{P}(E)}{4d} \left\{ \sum_{j=1}^{L} f(2je_1) + Lf(e_1) \right\}. \tag{12}$$

To control all translations, decompose each vector into its coordinates. By subadditivity and coordinate symmetry,

$$\sum_{z \in \Lambda_N} f(z) \le dN^{d-1} \sum_{r \in \mathbb{Z}/N\mathbb{Z}} f(re_1).$$

If N is odd, multiplication by two permutes $\mathbb{Z}/N\mathbb{Z}$, so the last sum equals $\sum_r f(2re_1)$. If N is even, writing each odd translation as an even translation followed by e_1 gives

$$\sum_{r \in \mathbb{Z}/N\mathbb{Z}} f(re_1) \le \sum_{r \in \mathbb{Z}/N\mathbb{Z}} f(2re_1) + \frac{N}{2} f(e_1).$$

Lemma 4.1, applied to $r \mapsto f(2re_1)$, gives

$$\sum_{r \in \mathbb{Z}/N\mathbb{Z}} f(2re_1) \le C_\rho \sum_{j=1}^{L} f(2je_1).$$

Since $L \ge \rho N/2$ for all sufficiently large N, the extra term in the even case satisfies

$$\frac{N}{2} f(e_1) \le \frac{1}{\rho} L f(e_1).$$

Thus in both parity cases we conclude that

$$a(V-a) \le C_{d,\rho} N^{d-1} \left\{ \sum_{j=1}^{L} f(2je_1) + Lf(e_1) \right\}. \tag{13}$$

Combining (12) and (13) yields

$$I \ge c_{d,\rho} \mathbb{P}(E) N \frac{a(V-a)}{V}.$$

Since

$$\kappa(\Lambda_N) = 2L\mathbb{P}(E) \le 2\rho N \mathbb{P}(E),$$

the preceding estimate, after changing the constant, gives

$$I \geq c_{d,\rho}\,\kappa(\Lambda_N)\frac{a(V-a)}{V}.$$

Finally,

$$I = \sum_{x\in A, z\in\Lambda_N} \kappa(z)\mathbf{1}_{\{x+z\notin A\}},$$

which proves (11). □

Remark 4.3. We need both endpoints of the selected edges to handle parity. The vertices Z_j alone control the even translations $2je_1$. When N is even, these do not control odd translations. Using also W_j, the nearest-neighbour relation $Z_j - W_j = \pm e_\ell$ yields $f(e_1) \leq f(Z_j) + f(W_j)$, which bridges the two parity classes. For odd N, multiplication by 2 is invertible modulo N, and this extra ingredient is unnecessary.

Remark 4.4. Proposition 4.2 also gives a spectral-gap bound for the auxiliary translation walk. Put $m := \kappa(\Lambda_N)$ and $\mu := \kappa/m$, and let $P_\mu(x,y) := \mu(y-x)$. By symmetry of κ, this kernel is reversible for the uniform law. The corresponding continuous-time generator is

$$L_\kappa f(x) := \sum_z \kappa(z)\big(f(x+z) - f(x)\big) = m(P_\mu - I)f(x).$$

The proposition gives the conductance bound

$$\inf_{0<|A|\leq V/2} \frac{\sum_z \mu(z) b_A(z)}{|A|} \geq \frac{c_{d,\rho}}{2}.$$

Hence Cheeger's inequality [7, Theorem 13.10] implies that the spectral gap λ_κ of $-L_\kappa$ satisfies

$$\lambda_\kappa = m\ \mathrm{gap}(I - P_\mu) \geq \frac{c_{d,\rho}^2}{8}\, m.$$

Conversely, a spectral-gap estimate of the form $\lambda_\kappa \geq c'_{d,\rho}\kappa(\Lambda_N)$ automatically provides an alternative proof of Proposition 4.2. Indeed, writing π for the uniform law, the Dirichlet form associated with L_κ is

$$\mathcal{E}_\kappa(f,f) := -\langle f, L_\kappa f\rangle_\pi = \frac{1}{2V}\sum_{x,z}\kappa(z)\big(f(x+z)-f(x)\big)^2.$$

Thus

$$\mathcal{E}_\kappa(\mathbf{1}_A,\mathbf{1}_A) = \frac{1}{V}\sum_z \kappa(z) b_A(z), \qquad \mathrm{Var}_\pi(\mathbf{1}_A) = \frac{|A|(V-|A|)}{V^2}.$$

Poincaré's inequality [7, Lemma 13.7 and Remark 13.8],

$$\lambda_\kappa\,\mathrm{Var}_\pi(f) \leq \mathcal{E}_\kappa(f,f),$$

then implies (11), up to constants.

We now apply Proposition 4.2 to the avalanche. On $E_N(u)$, connectedness of $\xi_N(u)$ implies that, for every $1 \leq j \leq L$, the range contains an edge

$$(W_j, Z_j), \qquad W_j \in S_{j-1}, \quad Z_j \in S_j.$$

Choose one uniformly among all such edges, independently for each j conditional on the range. This selection is equivariant under H_d, so its conditional law given $E_N(u)$ is H_d-invariant. Set

$$\kappa_u := \mathbb{E}\left[\mathbf{1}_{E_N(u)} \sum_{j=1}^{L} (\delta_{W_j} + \delta_{Z_j})\right].$$

Since

$$\kappa_u(\Lambda_N) = 2L\mathbb{P}(E_N(u)),$$

while (8) gives

$$\mathbb{P}(E_N(u)) \geq \theta(s), \qquad u \geq s,$$

and, since $L \geq \rho N/2$ for all sufficiently large N,

$$\kappa_u(\Lambda_N) \geq \rho\,\theta(s) N.$$

Hence Proposition 4.2 gives $c_{d,s,\rho} > 0$ such that, uniformly in $u \geq s$, $A \subset \Lambda_N$, and large N,

$$\sum_{x \in A, z} \kappa_u(z) \mathbf{1}_{\{x+z \notin A\}} \geq c_{d,s,\rho} N \frac{|A|(V - |A|)}{V}. \tag{14}$$

By Proposition 2.1, conditionally on the current range A at level u, every $x \in A$ carries marks at rate one, and a mark rooted at x grafts an independent u-avalanche translated by x.

For such a graft, the event $E_N(u)$ and the selected endpoints refer to the independent avalanche. Let $D_A(x)$ be the number of distinct sites outside A among the selected translated endpoints, with $D_A(x) = 0$ when $E_N(u)$ fails. A vertex on S_j can occur only as Z_j, and possibly also as W_{j+1}; hence every selected endpoint has multiplicity at most two. Therefore

$$0 \leq D_A(x) \leq 2L, \qquad D_A(x) \geq \frac{1}{2} \mathbf{1}_{E_N(u)} \sum_{j=1}^{L} \left(\mathbf{1}_{\{x+W_j \notin A\}} + \mathbf{1}_{\{x+Z_j \notin A\}}\right).$$

Summing over the possible roots, taking expectations, and using the definition of κ_u and (14) gives, after decreasing the constant,

$$\sum_{x \in A} \mathbb{E}[D_A(x)] \geq c_{d,s,\rho} N \frac{|A|(V - |A|)}{V}. \tag{15}$$

The expectation is over the independent graft and the auxiliary selection of its crossing edges, with A fixed. The actual number of new sites created by a graft rooted at x dominates $D_A(x)$.

Proposition 4.5 (Supercritical covering)**.** *For every $t_c < s < t$ and every fixed $0 < \rho < 1/2$,*

$$\mathbb{P}(\xi_N(t) = \Lambda_N \mid E_N(s)) \longrightarrow 1.$$

Consequently,

$$\liminf_{N \to \infty} P_N(t) \geq \theta(s) > 0.$$

Proof. Condition throughout on $E_N(s)$. Connectedness gives $r_N(s) \geq L+1$, and the range can only increase. For a nonempty $A \subset \Lambda_N$, put

$$\Phi(A) := \frac{V}{|A|} - 1.$$

Let $a := |A| \geq L$. A graft rooted at $x \in A$ adds $\Delta := |(x + \xi'_N(u)) \setminus A|$ new sites, where $\xi'_N(u)$ is an independent u-avalanche. Since $\Delta \geq D_A(x)$, $D_A(x) \leq 2L \leq 2a$, and $v \mapsto v/(a+v)$ is increasing,

$$\Phi(A) - \Phi(A \cup (x + \xi'_N(u))) = \frac{V\Delta}{a(a+\Delta)} \geq \frac{VD_A(x)}{a(a+D_A(x))} \geq \frac{VD_A(x)}{3a^2}.$$

Let $\mathcal{L}_u$ denote the generator of the range process. Each root has rate one, so (15) gives, uniformly for $u \geq s$ and $a \geq L$,

$$-\mathcal{L}_u\Phi(A) \geq \frac{V}{3a^2}\sum_{x\in A}\mathbb{E}[D_A(x)] \geq \frac{c_{d,s,\rho}N}{3}\frac{V-a}{a} = \frac{c_{d,s,\rho}N}{3}\Phi(A).$$

The conditioning event belongs to $\mathcal{H}_s$, and all subsequent grafts retain their independent laws. Dynkin's formula and Gronwall's inequality therefore yield

$$\mathbb{E}[\Phi(\xi_N(t)) \mid E_N(s)] \leq \frac{V}{L}\exp\left\{-\frac{c_{d,s,\rho}N}{3}(t-s)\right\}. \tag{16}$$

If $A \neq \Lambda_N$, then $\Phi(A) \geq 1/(V-1)$. Consequently,

$$\mathbb{P}(\xi_N(t) \neq \Lambda_N \mid E_N(s)) \leq \frac{V(V-1)}{L}\exp\left\{-\frac{c_{d,s,\rho}N}{3}(t-s)\right\} \longrightarrow 0. \tag{17}$$

Finally,

$$P_N(t) \geq \mathbb{P}(E_N(s))\mathbb{P}(\xi_N(t) = \Lambda_N \mid E_N(s)),$$

and (8) gives $\liminf_{N\to\infty} P_N(t) \geq \theta(s) > 0$. □

Proof of Theorem 1.5. This is Proposition 4.5. □

Remark 4.6. The estimate (17) controls covering throughout the growth from a macroscopic arm to the whole torus. Since $V = N^d$ and $L \geq \rho N/2$ for all sufficiently large N, its prefactor satisfies

$$\frac{V(V-1)}{L} \leq \frac{2}{\rho}N^{2d-1}.$$

Thus the comparison also supplies a covering bound after a sprinkling $C(\log N)/N$, with C sufficiently large depending on d, s, ρ.

5. Subcritical locality

We prove Theorem 1.6. Recall that

$$R(t) = \mathbb{E}r(t), \qquad t_c = t_{\mathrm{sus}},$$

so that $R(t) < \infty$ for every $t < t_c$. The hierarchical decomposition first upgrades this finite susceptibility to exponential integrability of the avalanche range. Local injectivity then compares infinite- and finite-volume ranges, while Lemma 2.2 transfers the resulting convergence of mean ranges to mean durations. The following proof is inspired by [14].

Proposition 5.1 (Subcritical exponential moments)**.** *For every $t < t_c$, there exists $\lambda_t > 0$ such that*

$$\mathbb{E}[e^{\lambda_t r(t)}] < \infty.$$

In particular, $\mathbb{E}[r(t)^p] < \infty$ for every integer $p \geq 1$.

Proof. Fix $L \geq 1$ and, for $\lambda \geq 0$, set

$$F_L(u, \lambda) := \mathbb{E}[e^{\lambda(r(u)\wedge L)}].$$

By Proposition 2.1, conditionally on a current range A of size x, marks occur at total rate x, and each mark grafts an independent avalanche B with $|B| \overset{\mathrm{d}}{=} r(u)$. Since

$$|A \cup (z + B)| \leq x + |B|,$$

if $x < L$, then

$$e^{\lambda((x+|B|)\wedge L)} - e^{\lambda x} \leq e^{\lambda x}\big(e^{\lambda(|B|\wedge L)} - 1\big),$$

whereas the truncated observable is constant when $x \geq L$. Conditioning on the current range and using the independence of the graft therefore gives, for almost every $u < t_c$,

$$\partial_u F_L(u, \lambda) \leq \big(F_L(u, \lambda) - 1\big)\partial_\lambda F_L(u, \lambda). \tag{18}$$

Indeed, the first factor produced by the rate is bounded by

$$\mathbb{E}\Big[r(u)e^{\lambda r(u)}\mathbf{1}_{\{r(u)<L\}}\Big] \leq \partial_\lambda F_L(u, \lambda).$$

We claim that exponential integrability propagates over every interval $[s, s + h]$ such that $hR(s) < 1$. Suppose $\mathbb{E}[e^{\alpha_0 r(s)}] < \infty$ for some $\alpha_0 > 0$. Since

$$\frac{\mathbb{E}[e^{\alpha r(s)}] - 1}{\alpha} \longrightarrow R(s) \qquad (\alpha \downarrow 0),$$

we may choose $0 < \alpha < \alpha_0$ such that, writing $q := \mathbb{E}[e^{\alpha r(s)}]$,

$$h(q - 1) < \alpha.$$

For $s \leq u \leq s + h$, set

$$\lambda(u) := \alpha - (q - 1)(u - s), \qquad G_L(u) := F_L(u, \lambda(u)).$$

Then $\lambda(u) > 0$, and (18) yields

$$G_L'(u) \leq \partial_\lambda F_L(u, \lambda(u))\big(G_L(u) - q\big), \qquad G_L(s) \leq q.$$

Hence $G_L(u) \leq q$ on $[s, s + h]$. In particular,

$$F_L(s + h, \alpha - h(q - 1)) \leq q, \qquad \alpha - h(q - 1) > 0.$$

Letting $L \to \infty$ gives

$$\mathbb{E}[e^{(\alpha - h(q-1))r(s+h)}] \leq q < \infty.$$

Now fix $t < t_c$. Since $R(t) < \infty$, choose a partition $0 = s_0 < \cdots < s_m = t$ with mesh smaller than $R(t)^{-1}$. The initial range $r(0) = |\mathcal{N}_d| = 2d + 1$ is deterministic. If $r(s_j)$ has an exponential moment, then, by monotonicity of R,

$$(s_{j+1} - s_j)R(s_j) \leq (s_{j+1} - s_j)R(t) < 1,$$

so the preceding propagation argument gives an exponential moment for $r(s_{j+1})$. Iterating finitely many times yields

$$\mathbb{E}[e^{\lambda_t r(t)}] < \infty$$

for some $\lambda_t > 0$. □

The preceding argument is independent of the geometry of $\mathbb{Z}^d$. Geometry enters only when these range moments are used to compare the infinite lattice with the torus.

Lemma 5.2 (Local injectivity)**.** *Let*

$$\varpi_N : \mathbb{Z}^d \to \Lambda_N$$

be the canonical projection. If $A \subset \mathbb{Z}^d$ is connected and

$$|A| < N,$$

then ϖ_N is injective on A.

Consequently, for every K with

$$K + 2d < N,$$

the infinite- and finite-volume avalanche constructions can be coupled in physical update time, using the same refresh variables, so as to agree, including their durations, until their common range first contains at least K sites.

Proof. A connected set with k vertices has graph diameter at most $k - 1$. Thus, if $|A| < N$,

$$\operatorname{diam}_{\ell^1}(A) \le |A| - 1 < N.$$

On the other hand, distinct $x, y \in \mathbb{Z}^d$ satisfying

$$\varpi_N(x) = \varpi_N(y)$$

have

$$\|x - y\|_1 \ge N.$$

Hence ϖ_N is injective on A.

For the second assertion, couple the elementary Bak–Sneppen updates in $\mathbb{Z}^d$ and on the torus using the same refresh variables, as long as the explored range projects injectively. Avalanche ranges are connected, and one elementary Bak–Sneppen update can add at most $2d$ new sites to the range. Hence, while the current range contains fewer than K sites, after the next update it contains fewer than $K + 2d < N$ sites and still projects injectively. The two explorations therefore remain identical until the claimed stopping time. □

Proposition 5.3 (Subcritical finite-volume approximation)**.** *For every $t < t_c$,*

$$\sup_{0 \le u \le t} |R_N(u) - R(u)| \longrightarrow 0. \tag{19}$$

Moreover,

$$D_N(t) \longrightarrow D(t). \tag{20}$$

Finally, for every fixed K, for all sufficiently large N and every $1 \le k \le K$,

$$\mathbb{E}\left[\eta_N(t)\mathbf{1}_{\{r_N(t)=k\}}\right] = \mathbb{E}\left[\eta(t)\mathbf{1}_{\{r(t)=k\}}\right]. \tag{21}$$

Proof. Take

$$K_N := \lfloor N/3 \rfloor.$$

For large N,

$$K_N + 2d < N,$$

so Lemma 5.2 couples the two avalanches identically until their common range reaches K_N. In particular, on $\{r(u) < K_N\}$,

$$r_N(u) = r(u).$$

Since

$$r_N(u) \le N^d,$$

we obtain

$$|R_N(u) - R(u)| \le N^d \mathbb{P}(r(u) \ge K_N) + \mathbb{E}[r(u)\mathbf{1}_{\{r(u)\ge K_N\}}].$$

Fix an integer $p > d$. By monotonicity in the level parameter and Proposition 5.1,

$$\mathbb{P}(r(u) \ge K_N) \le \frac{\mathbb{E}[r(t)^p]}{K_N^p},$$

and

$$\mathbb{E}[r(u)\mathbf{1}_{\{r(u)\ge K_N\}}] \le \frac{\mathbb{E}[r(t)^p]}{K_N^{p-1}},$$

uniformly for $u \le t$. Hence

$$\sup_{u\le t} |R_N(u) - R(u)| \le C_{p,t}\left(N^{d-p} + N^{1-p}\right) \longrightarrow 0,$$

which proves (19).

By Lemma 2.2,

$$D_N(t) = \exp\left\{\int_0^t R_N(u)\,du\right\}, \qquad D(t) = \exp\left\{\int_0^t R(u)\,du\right\}.$$

Therefore (19) implies (20).

Finally fix K, and take N large enough that

$$K + 1 + 2d < N.$$

Apply Lemma 5.2 with $K + 1$. The two avalanches then agree until their common range first contains at least $K + 1$ sites. On the event that the final range has size $k \le K$, this stopping level is never reached. Hence the entire finite-volume avalanche is identical to the infinite-volume avalanche, including its duration, and conversely. This proves (21). □

Proof of Theorem 1.6. The exponential-moment statement is Proposition 5.1, and the convergence of the mean ranges is (19). □

6. From avalanches to the stationary state

We now complete the passage from avalanche criticality to the stationary state. The argument is the finite-volume mechanism introduced by Meester and Znamenski [9, 10], but we give it in full because it is dimension-free and is an essential part of the present theorem. Write

$$V := |\Lambda_N| = N^d, \qquad \nu_h := \mathcal{L}(h + E), \quad E \sim \mathrm{Exp}(1).$$

All neighbourhoods on the torus are understood with repetitions removed.

Proposition 6.1 (Finite-volume ergodicity)**.** *For every N, the Bak–Sneppen chain on Λ_N has a unique stationary law π_N. From every deterministic initial configuration, its law converges to π_N in total variation. Moreover π_N is translation invariant.*

Proof. If one update refreshes the whole torus, then after one step the configuration consists of independent Exp(1) variables, independently of the initial state, and the conclusion is immediate. Assume henceforth that one update does not cover Λ_N.

Fix $h > 0$, write P for the transition kernel and $V = N^d$, and set

$$C_h := \{f : f_x > h \text{ for every } x \in \Lambda_N\}.$$

With probability at least $q_h := e^{-(2d+1)hV} > 0$, every fitness sampled in the next V updates exceeds h, uniformly in the current state. On this event no value at most h is created. As long as such values remain, the selected minimum is at most h and is refreshed, so their number decreases at each update. After at most V updates none remains. Hence

$$P^V(f, C_h) \geq q_h, \qquad f \in [0, \infty)^{\Lambda_N}. \tag{22}$$

Iteration gives a geometric tail for the hitting time of C_h, uniformly in the initial state.

For the minorization on C_h, choose a finite nearest-neighbour walk starting at the first updated site and visiting every vertex. At each intermediate update, require that the prescribed next vertex receive a fitness at most h, while every other refreshed vertex receives a fitness above h; at the final update require all fresh values to exceed h. At each intermediate step the prescribed next vertex is then the unique site with fitness at most h, so the avalanche follows the walk, covers the torus, and terminates. These finitely many requirements have positive probability, so $P_N(h) > 0$. Since the avalanche duration is integer-valued, there exists $m \geq 1$ such that

$$p_h := \mathbb{P}\big(r_N(h) = V,\ \eta_N(h) = m\big) > 0.$$

This range–duration event is measurable with respect to the censored level-h exploration, whose law, up to translation of its starting point, does not depend on $f \in C_h$. On this event every site has been refreshed, so Proposition 2.1 implies that the terminal fitnesses are independent with common law ν_h, independently of the censored exploration. Thus, setting

$$\mu_{N,h} := \nu_h^{\otimes \Lambda_N},$$

we obtain

$$P^m(f, \cdot) \geq p_h\, \mu_{N,h}(\cdot), \qquad f \in C_h. \tag{23}$$

Combining (22) and (23) by the Markov property gives

$$P^{V+m}(f,\cdot) \geq q_h p_h\, \mu_{N,h}(\cdot) \qquad \text{for every } f. \tag{24}$$

Hence $\lambda \mapsto \lambda P^{V+m}$ is a contraction in total variation with factor $1 - q_h p_h$. The space of probability measures is complete for this metric, so Banach's fixed-point theorem gives a unique fixed point π_N. Since $\pi_N P$ is also P^{V+m}-invariant, uniqueness gives $\pi_N P = \pi_N$, so π_N is the unique invariant law of P. Iterating the contraction and using that Markov kernels are nonexpansive in total variation yields

$$\|\lambda P^n - \pi_N\|_{\mathrm{TV}} \leq (1 - q_h p_h)^{\lfloor n/(V+m) \rfloor} \|\lambda - \pi_N\|_{\mathrm{TV}} \longrightarrow 0$$

for every initial law λ.

Finally, start the chain from independent Exp(1) fitnesses. Its law is translation invariant, ties have probability zero at every time, and the Bak–Sneppen dynamics preserves translation invariance. The total-variation convergence just proved therefore implies that its limit π_N is translation invariant. □

We next introduce the locking variables. Start the finite-volume chain from independent Exp(1) fitnesses and set $Y_{N,x}(0) = 0$. If $i_N(n)$ is the site of minimum fitness at time n and

$$M_N(n) := F_{N,i_N(n)}(n),$$

define

$$Y_{N,x}(n+1) = \begin{cases} 0, & x \in i_N(n) + \mathcal{N}_d, \\ \max\{Y_{N,x}(n), M_N(n)\}, & x \notin i_N(n) + \mathcal{N}_d. \end{cases} \tag{25}$$

For $y, z \geq 0$, set

$$H_y(z) := \begin{cases} 0, & z < y, \\ 1 - e^{-(z-y)}, & z \geq y. \end{cases}$$

Proposition 6.2 (Conditional exponential representation)**.** *Conditionally on $Y_N(n) = (Y_{N,x}(n))_{x\in\Lambda_N}$, the fitnesses $(F_{N,x}(n))_{x\in\Lambda_N}$ are independent and*

$$F_{N,x}(n) \sim \nu_{Y_{N,x}(n)}.$$

Consequently, for distinct $x_1, \ldots, x_k \in \Lambda_N$ and $z_1, \ldots, z_k \geq 0$,

$$\mathbb{P}\big(F_{N,x_j}(n) \leq z_j,\ 1 \leq j \leq k\big) = \mathbb{E}\left[\prod_{j=1}^{k} H_{Y_{N,x_j}(n)}(z_j)\right]. \tag{26}$$

Proof. The assertion holds at time zero. Assume it holds at time n, and condition on $Y_N(n)$, the location $i_N(n)$, and the value $M_N(n)$ of the minimum. By the lack-of-memory property, the unrefreshed coordinates remain independent and can be written as

$$\max\{Y_{N,x}(n), M_N(n)\} + E'_{N,x},$$

where the $E'_{N,x}$ are independent Exp(1) variables. The refreshed coordinates are fresh independent exponentials. Thus, conditionally on these variables, the fitnesses at time $n+1$ are independent exponentials above the thresholds in (25). Since this conditional law depends on the conditioning variables only

through $Y_N(n+1)$, the same statement holds conditionally on $Y_N(n+1)$ alone. This proves the claim by induction. □

Put

$$G_N(n,t) := \mathbb{P}(Y_{N,0}(n) \le t).$$

Translation invariance of the initial law and of the dynamics makes the choice of the origin immaterial.

Proposition 6.3 (Subcritical locking thresholds)**.** *For every $t < t_c$,*

$$\lim_{N\to\infty} \limsup_{n\to\infty} G_N(n,t) = 0. \tag{27}$$

Proof. Let

$$K_N(n,t) := \sum_{x\in\Lambda_N} \mathbf{1}_{\{Y_{N,x}(n)\le t\}}.$$

Let C_t be the set of configurations with all fitnesses above t. By (22), its hitting time has finite mean. Discard the initial segment up to the first visit to C_t, and from then on decompose the trajectory into successive t-avalanches. If T_j is the starting time of the j-th one and (r_j, η_j) its range–duration pair, then $T_{j+1} - T_j = \eta_j$, and these pairs are i.i.d. with the law of $(r_N(t), \eta_N(t))$. At the first update of the avalanche the current minimum exceeds t, so every site outside its eventual range has its locking threshold raised above t. During the remaining updates the selected minima are at most t, and only sites in the avalanche range are refreshed. Hence

$$K_N(n,t) \le r_j, \qquad T_j < n \le T_{j+1}.$$

For all sufficiently large N, a one-step noncovering t-avalanche has positive probability, so the cycle length has span one. The delayed renewal–reward theorem and translation invariance therefore give

$$\limsup_{n\to\infty} G_N(n,t) \le \frac{\mathbb{E}[\eta_N(t) r_N(t)]}{V D_N(t)}.$$

For every fixed $K \ge 1$, using $r_N(t) \le V$,

$$\limsup_{n\to\infty} G_N(n,t) \le \frac{K}{V} + \frac{\mathbb{E}[\eta_N(t)\mathbf{1}_{\{r_N(t)>K\}}]}{D_N(t)}. \tag{28}$$

Now

$$\mathbb{E}[\eta_N(t)\mathbf{1}_{\{r_N(t)>K\}}] = D_N(t) - \sum_{k=1}^{K} \mathbb{E}[\eta_N(t)\mathbf{1}_{\{r_N(t)=k\}}].$$

By Proposition 5.3, as $N \to \infty$ this converges to $\mathbb{E}[\eta(t)\mathbf{1}_{\{r(t)>K\}}]$, while $D_N(t) \to D(t) < \infty$. Letting first $N \to \infty$ and then $K \to \infty$ in (28) proves (27). □

Proposition 6.4 (Supercritical locking thresholds)**.** *For every $b > t_c$,*

$$\lim_{N\to\infty} \liminf_{n\to\infty} G_N(n,b) = 1. \tag{29}$$

Proof. Fix $t_c < t < b$. Starting from the process defining G_N, ignore the initial segment up to the first visit to C_t. Thereafter run successive t-avalanches until the first one covering Λ_N, and then continue until the next visit to C_b.

The initial delay and the final recovery have finite mean by (22). Moreover $P_N(t) > 0$ by Proposition 6.1, and $D_N(t) < \infty$ by Lemma 2.2, so the waiting time for the first covering t-avalanche also has finite mean.

At the end of this covering avalanche, every site has been refreshed. Conditionally on its t-censored exploration, the terminal fitnesses are independent with law ν_t, by Proposition 2.1. Censor the subsequent recovery at level b: reveal fitnesses at most b exactly and reveal larger fitnesses only through their exceedance of b. At the designated return to C_b, each terminal value has been observed only to exceed b. Thus, conditionally on the censored cycle history, the terminal vector has law

$$\mu_{N,b} = \nu_b^{\otimes \Lambda_N}.$$

Treat the trajectory up to the first such return as the initial delay, and repeat the construction from that point. Fresh future refresh variables and this censored regeneration make the successive pairs of waiting and recovery lengths independent and identically distributed. Let A_N be the time spent in successive t-avalanches up to and including the first covering one in a cycle started with law $\mu_{N,b}$, and let B_N be the subsequent recovery time to C_b. For large N, the cycle length has span one, since a one-step noncovering t-avalanche may be inserted before the first covering one with positive probability.

At the end of the covering avalanche all locking thresholds are at most t, and throughout the recovery every selected minimum is at most b. Thus

$$K_N(n, b) = V$$

during the whole recovery phase. By the delayed renewal–reward theorem,

$$\liminf_{n\to\infty} G_N(n, b) \geq \frac{\mathbb{E}B_N}{\mathbb{E}A_N + \mathbb{E}B_N}. \tag{30}$$

Since the successive t-avalanches are i.i.d. until the first covering one, Wald's identity gives

$$\mathbb{E}A_N = \frac{D_N(t)}{P_N(t)}. \tag{31}$$

To bound B_N from below, use the decorated construction of Proposition 2.1. Since the covering t-avalanche has range Λ_N, every site is active in the level extension from t to b. For each $x \in \Lambda_N$, let τ_x be the first mark above t on the vertical line through x. Then $\tau_x - t$ are i.i.d. Exp(1), and, conditionally on $\tau_x = u \leq b$, the corresponding graft contributes a duration of mean $D_N(u)$. Hence, by Tonelli and monotonicity of D_N,

$$\mathbb{E}B_N \geq V \int_t^b e^{-(u-t)} D_N(u)\, du \geq V(1 - e^{-(b-t)}) D_N(t). \tag{32}$$

Combining (30)–(32),

$$1 - \liminf_{n\to\infty} G_N(n, b) \leq \frac{1}{V(1 - e^{-(b-t)}) P_N(t)}. \tag{33}$$

Choose $s \in (t_c, t)$. By Theorem 1.5,

$$\liminf_{N\to\infty} P_N(t) \geq \theta(s) > 0.$$

Since $V = N^d \to \infty$, the right-hand side of (33) tends to zero, proving (29). □

Proposition 6.5 (Locking-threshold concentration and product limit)**.** *For every $\varepsilon > 0$,*

$$\lim_{N\to\infty} \limsup_{n\to\infty} \mathbb{P}\big(|Y_{N,0}(n) - t_c| > \varepsilon\big) = 0. \tag{34}$$

Consequently, for every fixed collection of distinct sites $x_1, \dots, x_k \in \mathbb{Z}^d$, the corresponding marginal of π_N converges to $\nu_{t_c}^{\otimes k}$.

Proof. For $0 < \varepsilon < t_c$, Propositions 6.3 and 6.4, applied at $t_c - \varepsilon$ and $t_c + \varepsilon$, give (34); for $\varepsilon \geq t_c$, the lower-tail event is empty and the same conclusion follows from the supercritical bound. For fixed $x_1, \dots, x_k$, and N large enough that their images in Λ_N are distinct, translation invariance and the union bound imply

$$\lim_{N\to\infty} \limsup_{n\to\infty} \mathbb{P}\left(\max_{1\leq j\leq k} |Y_{N,x_j}(n) - t_c| > \varepsilon\right) = 0.$$

Fix $z_1, \dots, z_k \geq 0$ and define

$$\Psi(y_1, \dots, y_k) := \prod_{j=1}^{k} H_{y_j}(z_j).$$

The function Ψ is bounded and continuous. If

$$\omega_\varepsilon := \sup_{\max_j |y_j - t_c| \leq \varepsilon} |\Psi(y_1, \dots, y_k) - \Psi(t_c, \dots, t_c)|,$$

then $\omega_\varepsilon \to 0$ as $\varepsilon \downarrow 0$. By Proposition 6.2,

$$\begin{aligned} &\left|\mathbb{P}(F_{N,x_j}(n) \leq z_j,\ 1 \leq j \leq k) - \prod_{j=1}^{k} H_{t_c}(z_j)\right| \\ &\qquad \leq \omega_\varepsilon + \mathbb{P}\left(\max_{1\leq j\leq k} |Y_{N,x_j}(n) - t_c| > \varepsilon\right). \end{aligned}$$

For fixed N, Proposition 6.1 lets $n \to \infty$ and replaces the law of the chain by π_N. Letting next $N \to \infty$ and finally $\varepsilon \downarrow 0$ proves convergence of the joint distribution functions to $\prod_j H_{t_c}(z_j)$, the distribution function of $\nu_{t_c}^{\otimes k}$. □

Proof of Theorem 1.1. Proposition 6.5 gives the asserted convergence in exponential coordinates. Under the inverse transformation $u = 1 - e^{-f}$, a random variable with law $t_c + \mathrm{Exp}(1)$ becomes uniform on $[1 - e^{-t_c}, 1]$. This proves the theorem. □

Appendix A. Use of AI

The Laplace transform of the additive process in which X_t is replaced by $X_t + X'_t$ an independent copy of X_t at rate X_t is known to satisfy a Burgers-type equation, and the phase transition is then visible using the characteristics. For two years, we had been looking for a way to relate this mechanism to the evolution of the avalanche range. During that period, we repeatedly prompted earlier AI models to find a suitable comparison or

domination by solutions of this equation. These attempts were unsuccessful and, in particular, repeatedly led to incorrect boundary-layer arguments.

The connection with the Smoluchowski–Burgers picture was therefore part of our input. In August 2026, ChatGPT 5.6 Sol produced the argument that made this connection effective and that is used in the present paper. The key idea found by the model was to restrict attention to grafts involving large avalanches, derive a first-moment lower bound on the number of newly created vertices, and combine this estimate with concavity to obtain the required Burgers-type comparison.

The amplification argument was subsequently obtained by the same model through a spectral-gap approach. This spectral formulation was again found by the model; we then suggested looking for a purely combinatorial proof, which led to the formulation ultimately retained in the paper.

Finally, finiteness of all polynomial moments in the subcritical regime was first obtained by a recursive argument. We suggested seeking a recursive proof of exponential moments inspired by [14]; the resulting argument was then developed with AI assistance and led to the stronger estimate used in the paper.

References

1. David J. Aldous, *Deterministic and stochastic models for coalescence (aggregation and coagulation): a review of the mean-field theory for probabilists*, Bernoulli **5** (1999), no. 1, 3–48.
2. Per Bak and Kim Sneppen, *Punctuated equilibrium and criticality in a simple model of evolution*, Physical Review Letters **71** (1993), 4083–4086.
3. Persi Diaconis and Laurent Saloff-Coste, *Comparison techniques for random walk on finite groups*, Annals of Probability **21** (1993), no. 4, 2131–2156.
4. Andrew Gillett, Ronald Meester, and Misja Nuyens, *Bounds for avalanche critical values of the Bak–Sneppen model*, Markov Processes and Related Fields **12** (2006), no. 4, 679–694.
5. Andrew J. Gillett, Ronald W. J. Meester, and Peter van der Wal, *Maximal avalanches in the Bak–Sneppen model*, Journal of Applied Probability **43** (2006), no. 3, 840–851.
6. Olivier Hénard, *Sharpness of the avalanche phase transition in the Bak–Sneppen model*, arXiv preprint, arXiv:2608.10992 (2026).
7. David A. Levin and Yuval Peres, *Markov chains and mixing times*, 2 ed., American Mathematical Society, Providence, RI, 2017.
8. Sergei Maslov, *Infinite series of exact equations in the Bak–Sneppen model of biological evolution*, Physical Review Letters **77** (1996), 1182–1185.
9. Ronald Meester and Dmitri Znamenski, *Limit behavior of the Bak–Sneppen evolution model*, Annals of Probability **31** (2003), 1986–2002.
10. ———, *Critical thresholds and the limit distribution in the Bak–Sneppen model*, Communications in Mathematical Physics **246** (2004), 63–86.
11. Govind Menon and Robert L. Pego, *Approach to self-similarity in Smoluchowski's coagulation equations*, Communications on Pure and Applied Mathematics **57** (2004), no. 9, 1197–1232.
12. Maya Paczuski, Sergei Maslov, and Per Bak, *Avalanche dynamics in evolution, growth, and depinning models*, Physical Review E **53** (1996), 414–443.
13. Serguei Popov and Stanislav Volkov, *On extinction and survival of the Bak–Sneppen model on arbitrary graphs*, arXiv preprint arXiv:2506.15500 (2026).
14. Oliver Riordan and Lutz Warnke, *The evolution of subcritical achlioptas processes*, Random Structures & Algorithms **47** (2015), no. 1, 174–203.
15. Stanislav Volkov, *Rigorous upper bound for the discrete Bak–Sneppen model*, Journal of Statistical Physics **186** (2022), 8.

Institut de mathématique d'Orsay, Université Paris-Saclay, 91405 Orsay, France

Email address: olivier.henard@universite-paris-saclay.fr