ERGODICITY AND MIXING PROPERTIES OF THE NORTHEAST MODEL

GEORGE KORDZAKHIA,* *University of California, Berkeley*

STEVEN P. LALLEY,** *University of Chicago*

Abstract

The *northeast model* is a spin system on the two-dimensional integer lattice that evolves according to the following rule: whenever a site's southerly and westerly nearest neighbors have spin 1, it may reset its own spin by tossing a p-coin; at all other times, its spin remains frozen. It is proved that the northeast model has a phase transition at $p_c = 1 - \beta_c$, where β_c is the critical parameter for oriented percolation. For $p < p_c$, the trivial measure, δ_0, that puts mass one on the configuration with all spins set at 0 is the unique ergodic, translation-invariant, stationary measure. For $p \geq p_c$, the product Bernoulli-p measure on configuration space is the unique nontrivial, ergodic, translation-invariant, stationary measure for the system, and it is mixing. For $p > \frac{2}{3}$, it is shown that there is exponential decay of correlations.

Keywords: Northeast model; facilitated spin-flip system; oriented percolation; exponential mixing

2000 Mathematics Subject Classification: Primary 60J25
Secondary 60K35

1. Introduction

The *northeast model* is the simplest nontrivial facilitated spin-flip system on the two-dimensional integer lattice $\mathbb{Z}^2$. Interest in such systems dates to [3]; for a recent review, see [6]. The northeast model is specified by the following rules: spins take values in the two-element set $\{0, 1\}$; the spin at site $(x, y) \in \mathbb{Z}^2$ may flip only at times when the spins at sites $(x - 1, y)$ and $(x, y - 1)$ are both 1; and at such times the flip rates are

$$
\begin{aligned}
0 \to 1 \quad &\text{at rate } p, \\
1 \to 0 \quad &\text{at rate } 1 - p.
\end{aligned}
\tag{1}
$$

A rigorous construction of the process is outlined in Section 2. It will also be shown that the product Bernoulli-p measure ν_p on the configuration space is invariant and reversible. Note that, unlike many growth models (e.g. the contact process), the northeast model is *not* additive (in the sense of [5, Definition III.6.1]) or even monotone (in the sense that stochastic monotonicity is preserved by the evolution).

The northeast model is the natural two-dimensional analogue of the one-dimensional *east model* studied by Aldous and Diaconis [1]. The east model is the spin-flip model on the integer

lattice $\mathbb{Z}$ in which the spin at site x may flip only at times when the spin at site $x - 1$ is 1; the flip rates at such times are given by (1). The main results of [1] are rigorous bounds on relaxation rates to the product Bernoulli-p measure that are sharp in the limit $p \to 0$. The northeast model differs fundamentally from the east model in that the product Bernoulli-p measure ν_p is not even ergodic for $p < 1 - \beta_c$, where β_c is the critical parameter for oriented site percolation (for a general reference on percolation, see [4]). The reason for this is apparent: if $p < 1 - \beta_c$ then ν_p-almost every spin configuration must contain infinite southwest clusters of zero spins; these clusters must remain frozen at spin 0 forever.

We shall establish in this paper that the northeast model exhibits a phase transition at $p_c := 1 - \beta_c$, by showing that for $p > p_c$ the product Bernoulli-p measure is ergodic and mixing. We state this as follows.

Theorem 1. *The only ergodic, translation-invariant, stationary measures for the northeast model with flip parameter $p \geq 1 - \beta_c$ are ν_p and ν_0. The only ergodic, translation-invariant, stationary measure for the northeast model with flip parameter $p < 1 - \beta_c$ is ν_0. For $p > p_c$, the product Bernoulli-p measure ν_p is not only ergodic but also mixing.*

Here *mixing* means that correlations decay to zero as $t \to \infty$. Theorem 2, below, is a more precise statement of what we shall prove. For any spin configuration ζ, denote by ζ^Λ its restriction to the set $\Lambda \subset \mathbb{Z}^2$. Similarly, denote by ν_p^Λ the product Bernoulli-p measure on the restricted configuration space $\{0, 1\}^\Lambda$, by P_p^ζ the probability measure of the northeast model with rate p and initial configuration ζ, and by ξ_t the (random) spin configuration at time t in a realization of the northeast model.

Theorem 2. *If $p > 1 - \beta_c$ then ν_p is mixing for the northeast model in the following sense: for ν_p-almost every initial configuration σ, and for every finite set $\Lambda \subset \mathbb{Z}^2$ and every configuration $\zeta^\Lambda \in \{0, 1\}^\Lambda$,*

$$\lim_{t \to \infty} P_p^\sigma \{\xi_t^\Lambda = \zeta^\Lambda\} = \nu_p^\Lambda(\zeta^\Lambda).$$

In light of the results of [1], it is natural to ask if correlations decay exponentially for $p > p_c$. We conjecture that they do. We have been able to prove this only for $p > \frac{2}{3}$.

Theorem 3. *For each $p > \frac{2}{3}$ and each finite set Λ of sites, there exist constants $C_\Lambda, \alpha_\Lambda > 0$ such that, for any two configurations ζ^Λ and η^Λ,*

$$| P_p^{\nu_p} \{\xi_t^\Lambda = \zeta^\Lambda; \ \xi_0^\Lambda = \eta^\Lambda\} - \nu_p^\Lambda(\zeta^\Lambda)\nu_p(\eta^\Lambda)| \leq C_\Lambda \exp\{-\alpha_\Lambda t\} \quad \textit{for all } t > 0. \tag{2}$$

We further conjecture that the exponential decay parameters $\alpha_\Lambda \equiv \alpha_\Lambda(p)$ can be chosen so as not to depend on the sets Λ. If this is true then, for $p > p_c$, the northeast model has a positive *spectral gap* (see [1]), and the nature of the phase transition is reflected in the behavior of the spectral gap as $p \to p_c$.

Sections 2–4 are devoted to the proofs of Theorems 1–3. Theorems 1 and 2 will be proved in Sections 2 and 3, and Theorem 3 in Section 4. In the final section, Section 5, we state a conjecture about the propagation of influence in the northeast model.

2. Construction

Two constructions of the northeast model are possible, one using the Hille–Yosida theorem (see [5, Chapter 1]) and the other using a system of independent marked Poisson processes. The first has the advantage that it gives simple characterizations of invariant and reversible measures. The second yields detailed information about the time evolution of the process.

2.1. Generator and semigroup

Let X be the space of spin configurations on the lattice $\mathbb{Z}^2$ and, for any subset $\Lambda \subset \mathbb{Z}^2$, let X^Λ be the set of configurations on Λ. Denote by $C(X)$ the space of continuous, real-valued functions on X (relative to the product topology), and by $C_*(X)$ the subset consisting of those functions that depend only on finitely many coordinates. Note that $C_*(X)$ may be naturally identified with $\bigcup_\Lambda C(X^\Lambda)$, where the union is over all finite subsets $\Lambda \subset \mathbb{Z}^2$. For any configuration σ, define $A(\sigma)$ to be the set of sites that are flip-eligible in configuration σ, that is, those sites x whose nearest neighbors to the south and west both have spin 1 in configuration σ. Also, for any site x, denote by σ_x^+ and σ_x^- the configurations that agree with σ at all sites $y \neq x$ and have respective spins 1 and 0 at site x.

For any function $f \in C(X)$ and each site $x \in \mathbb{Z}^2$, define

$$
\nabla_x f(\sigma) = \begin{cases} (f(\sigma_x^+) - f(\sigma_x^-))p & \text{if } \sigma = \sigma_x^-, \\ (f(\sigma_x^-) - f(\sigma_x^+))(1-p) & \text{if } \sigma = \sigma_x^+. \end{cases} \tag{3}
$$

For any $f \in C_*(X)$, define

$$
\mathcal{L}f(\sigma) = \sum_{x \in \mathbb{Z}^2} \nabla_x f(\sigma)\, \mathbf{1}_{\{x \in A(\sigma)\}}. \tag{4}
$$

Observe that, because $f \in C_*(X)$ depends on only finitely many coordinates, the sum is finite and, hence, the definition is valid. It follows from [5, Section 1.3], that the operator $\mathcal{L}$ on $C_*(X)$ extends uniquely to a Markov generator and that $C_*(X)$ is a core for the generator. (Note: in [5] a larger core was used, but the graph of $\mathcal{L}$ restricted to $C_*(X)$ is easily seen to have closure equal to that of $\mathcal{L}$ restricted to that larger core.) The Hille–Yosida theorem therefore implies that the closure of $\mathcal{L}$ generates a unique Feller Markov semigroup on $C(X)$. This implies the existence of a Markov process with flip rates (1).

2.2. Stationary and reversible measures

A Borel probability measure μ on X is stationary for the Markov semigroup generated by $\mathcal{L}$ if and only if, for every $f \in C_*(X)$,

$$
\int_X \mathcal{L}f \, \mathrm{d}\mu = 0. \tag{5}
$$

A stationary measure μ is reversible if and only if, for all pairs $f, g \in C_*(X)$,

$$
\int_X f \mathcal{L}g \, \mathrm{d}\mu = \int_X g \mathcal{L}f \, \mathrm{d}\mu.
$$

These criteria yield easy proofs that the product Bernoulli-p measure ν_p is stationary and reversible for the northeast model with flip parameter p. By (4) and (5), to prove that ν_p is stationary it suffices to show that, for any function $f : X \to \mathbb{R}$ that depends only on the spins in Λ and, for each $x \in \Lambda$,

$$
\int_{F(x)} \nabla_x f(\sigma) \, \mathrm{d}\nu_p(\sigma) = 0,
$$

where $F(x) = \{\sigma : x \in A(\sigma)\}$ is the set of configurations σ for which site x is flip-eligible. The event $\{x \in A(\sigma)\} = \{\sigma \in F(x)\}$ does not depend on the spin at x, so under ν_p this event is independent of $\sigma(x)$. Moreover, the event $\sigma = \sigma_x^+$ occurs if and only if $\sigma(x) = 1$, and is

hence independent of $\{\sigma \in F(x)\}$. Since $\sigma(x) = 1$ with probability p under ν_p, (3) implies that (with $q = 1 - p$)

$$\int_{F(x)} \nabla_x f(\sigma) \, d\nu_p(\sigma) = \int_{F(x)} ((f(\sigma_x^+) - f(\sigma_x^-))pq + (f(\sigma_x^-) - f(\sigma_x^+))qp) \, d\nu_p(\sigma)$$
$$= 0.$$

Similarly, to prove that ν_p is reversible, it suffices to show that, for any pair of functions $f, g \in C(X)$ that depend only on the spins in a finite set Λ,

$$\int_{F(x)} g(\sigma) \nabla_x f(\sigma) \, d\nu_p(\sigma) = \int_{F(x)} f(\sigma) \nabla_x g(\sigma) \, d\nu_p(\sigma).$$

Since the spin at x is independent of the event $\sigma \in F(x)$, this identity follows by a simple calculation similar to that above.

2.3. Construction via marked Poisson processes

Let $(\Omega, \mathcal{F}, \mathrm{P})$ be a probability space on which are defined countably many independent rate-one Poisson processes and countably many independent uniform(0,1) random variables. Assign one Poisson process to each site $x \in \mathbb{Z}^2$ and one uniform random variable to each occurrence in each Poisson process. The occurrences in the Poisson process attached to site x mark the times of flip opportunities for site x: at each such time, site x queries its neighbors to the south and west about their current spins; if both of these are 1 then x resets its spin according to the value, U, of the uniform random variable attached to the occurrence. The reset rule is

$$\begin{aligned} U \le p &\implies \text{reset spin to 1}, \\ U > p &\implies \text{reset spin to 0}. \end{aligned} \tag{6}$$

Note that such a reset does not necessarily change the spin at x: if the reset occurs at a time when the spin at x is 0 or 1, then the chance that the spin changes is p or, respectively, q. Thus, the flip rates at site x will agree with the specification (1).

In order that this construction uniquely specify the spin at site x at each finite time, it must be the case that the spins of the southwest neighbors of x at query times can be determined. For this to be the case, the initial configuration, ξ_0, must be specified. Throughout the remainder of this subsection and Subsection 2.4, we assume that $\xi_0 = \zeta$ for some fixed, nonrandom configuration ζ. (We suppress the dependence of the process ξ_t on ζ, but the reader should note that the algorithm specified below will give a different time evolution ξ_t for every different choice of ζ.) To determine the spin at site y at time t, we must search backward in time for occurrences in the Poisson process attached to site y. If there are no occurrences in this process between times zero and t, then the spin at site y at time t is set to equal its value at time zero. Otherwise, at s, the last occurrence time before t, the backward search moves to the south and west neighbors, y' and y'', of y and proceeds recursively. If it can then be determined that the spins at y' and y'' at time s are both 1, then the spin at y at time t is determined by the attached uniform random variable according to the rule (6). Otherwise, the query continues to the last occurrence in the Poisson process at y before time s. To show that this algorithm terminates, we must show that the backward tree of queries initiated at (y, t) is almost surely finite. This follows because the backward query tree is stochastically dominated by a simple binary fission process with fission rate one. (At each time a query is made, two new query processes are engendered. The additional queries in these offspring processes are mutually independent, *except* when they coincide by virtue of a merger, e.g. SW and WS).

2.4. Regeneration at reset times

Assume henceforth that the northeast process ξ_t has been constructed as in Subsection 2.3, using auxiliary marked Poisson processes and a specified initial configuration ξ_0. For each site x, define τ_x to be the time of the first reset at x. Note that τ_x may take the value ∞. Note also that τ_x need not be (and generally will not be) measurable with respect to the σ-algebra generated by $\{\xi_t\}_{0 \leq t < \infty}$.

Let '$\leq$' and '$<$' denote the natural weak and strong partial orders on $\mathbb{Z}^2$; that is, $x \leq y$ or $x < y$ if each coordinate of x is less than or equal to or, respectively, strictly less than the corresponding coordinate of y. For each site x, define

$$\mathcal{A}_x = \{y \colon y \leq x\}, \qquad \mathcal{B}_x = \{y \colon x \not\leq y\}.$$

Sites in $\mathcal{A}_x$ are those that may influence the times at which spin resets at x may occur, and sites in $\mathcal{B}_x$ are those whose reset times are not influenced by site x. Let N_t^x be the number of occurrences up to time t in the Poisson process attached to site x, and let $\{U_j^x\}_{j \geq 1}$ be the uniform$(0, 1)$ random variables attached to the occurrences of this Poisson process, listed in chronological order. Define the σ-algebras

$$\mathcal{G}_t^x = \sigma(N_s^x)_{s \leq t}, \qquad \mathcal{H}_t^x = \sigma(U_i^x)_{i \leq N_t^x}, \qquad \mathcal{F}_t^x = \sigma(\mathcal{G}_t^x \cup \mathcal{H}_t^x), \qquad \mathcal{I}_t^x = \sigma\left(\bigcup_{y \in \mathcal{A}_x} \mathcal{F}_t^y\right),$$

$$\mathcal{I}_t^\infty = \sigma\left(\bigcup_x \mathcal{I}_t^x\right), \qquad \mathcal{J}_t^x = \sigma\left(\bigcup_{y \in \mathcal{B}_x} \mathcal{F}_t^y\right), \qquad \mathcal{K}_t^x = \sigma(\mathcal{J}_t^x \cup \mathcal{G}_t^x).$$

Note that $\mathcal{I}_t^\infty$ is generated by everything that happens in the entire construction up to time t.

Proposition 1. *For each site x and each time $t > 0$, the conditional distribution of the spin $\xi_t(x)$ given the σ-algebra $\mathcal{K}_\infty^x$ on the event $\{\tau_x \leq t\}$ is Bernoulli-p, that is,*

$$\mathrm{P}\{\xi_t(x) = 1 \mid \mathcal{K}_\infty^x\} = p \quad \text{on } \{\tau_x \leq t\}. \tag{7}$$

Proof. Observe that the reset time τ_x, and in fact the entire string of flip opportunity times at site x up to time t, are measurable with respect to the σ-algebra $\mathcal{K}_t^x$. On the event $\{\tau_x \leq t\}$, there is a last flip opportunity at x before time t, and this occurs at the kth of the N_s^x occurrences, for some $k \geq 1$. The spin at x at time t is determined by the uniform random variable U_k^x. However, this random variable is independent of $\mathcal{K}_\infty^x$. Assertion (7) follows.

Corollary 1. *If $\mathrm{P}\{\tau_x < \infty\} = 1$ for every site x in a finite set Λ, then the joint distribution of the restricted configuration ξ_t^Λ converges weakly to the product Bernoulli-p measure ν_p^Λ as $t \to \infty$.*

Proof. Consider first, for illustration, the case $\Lambda = \{x, y\}$, where x is the westerly nearest neighbor of y. Observe that $\mathcal{K}_t^x \subset \mathcal{K}_t^y$ and recall that $\xi_t(x)$ is measurable with respect to $\mathcal{K}_t^y$. By hypothesis, $\mathrm{P}\{\tau_x \leq t\} \approx 1$ and $\mathrm{P}\{\tau_y \leq t\} \approx 1$ for large t. Consequently, by Proposition 1 (first for y, then x), for large t,

$$\begin{aligned}
\mathrm{P}\{\xi_t(x) = \xi_t(y) = 1\} &= \mathrm{E}\,\mathrm{E}(\mathrm{P}\{\xi_t(y) = 1 \mid \mathcal{K}_t^y\}\xi_t(x) \mid \mathcal{K}_t^x) \\
&\approx p\,\mathrm{E}\,\mathrm{P}\{\xi_t(x) = 1 \mid \mathcal{K}_t^x\} \\
&\approx p^2.
\end{aligned}$$

The general case is proved by induction on the size of Λ. Fix a site $y \in \Lambda$ such that $\Lambda' = \Lambda \setminus \{y\} \subset \mathcal{B}_y$; since Λ is finite, it must contain such a site. Observe that, for each $x \in \Lambda'$, the spin $\xi_t(x)$ is measurable with respect to $\mathcal{K}_t^y$. Let ζ^Λ be any spin configuration on Λ and let ζ^R be its restriction to a subset R of Λ. Then, by Proposition 1 and the induction hypothesis, for large t,

$$
\begin{aligned}
\mathrm{P}\{\xi_t^\Lambda = \zeta^\Lambda\} &= \mathrm{E}\,\mathrm{P}\{\xi_t(y) = \zeta^y \mid \mathcal{K}_t^y\}\mathbf{1}_{\{\xi_t^{\Lambda'}=\zeta^{\Lambda'}\}} \\
&\approx \nu_p^y(\zeta^y)\,\mathrm{P}\{\xi_t^{\Lambda'} = \zeta^{\Lambda'}\} \\
&\approx \nu_p^y(\zeta^y)\nu_p^{\Lambda'}(\zeta^{\Lambda'}) \\
&= \nu_p^\Lambda(\zeta^\Lambda).
\end{aligned}
$$

3. Stationary measures, ergodicity, and mixing

3.1. Examples of stationary measures

We have shown that the product Bernoulli-p measure ν_p is stationary for the northeast process. This is not, however, the only stationary measure; for instance, the measure ν_0 that assigns probability one to the configuration with all spins 0 is also stationary, and is ergodic. In fact, there are infinitely many distinct (and mutually singular) ergodic stationary distributions. A denumerable family may be built as follows.

Let Γ be an infinite, connected subset of $\mathbb{Z}^2$ such that

(a) for every site $y \in \Gamma$, at least one of the southerly or westerly nearest neighbors of y is also an element of Γ, and

(b) every connected component of $\mathbb{Z}^2 \setminus \Gamma$ is finite.

Consider the initial configuration ξ_0 in which every site in Γ is assigned spin 0 and every other site spin 1. If the northeast process starts in this configuration then every site in Γ will remain frozen at spin 0, because any query tree (see Subsection 2.3) that begins at such a site will have leaves in Γ. Similarly, every site x one of whose southerly or westerly nearest neighbors is in Γ will remain frozen at spin 1 forever. Now consider the evolution ξ_t^Λ in a connected component, Λ, of $\mathbb{Z}^2 \setminus \Gamma$: because Λ is finite, and because its southwest border consists of sites that must remain frozen at spin 1, the process ξ_t^Λ is an ergodic, finite-state Markov process and, so, the distribution of ξ_t^Λ must converge to a stationary distribution μ^Λ as $t \to \infty$. The product, λ_Γ, of these measures μ^Λ over all components Λ with the point mass on the zero configuration in Γ is a stationary distribution for the northeast process. Since each of the factors μ^Λ is stationary and ergodic for the restricted process ξ_t^Λ, the product measure λ_Γ will be ergodic for the northeast process. Observe that the measures λ_Γ and $\lambda_{\Gamma'}$ will be mutually singular if $\Gamma \neq \Gamma'$, because different sets of sites (Γ and Γ') remain frozen forever.

The preceding construction shows that there are infinitely many ergodic stationary distributions for the northeast model. None of these distributions is translation invariant. It is natural to ask if there exist translation-invariant stationary distributions other than the product Bernoulli-p measures ν_p and ν_0. The answer is 'yes', as the following argument shows. Fix an integer $m \geq 2$ and let Γ_0 be the set of sites with at least one coordinate divisible by m. For each $x = (x_1, x_2) \in \mathbb{Z}^2$ with coordinates satisfying $0 \leq x_i < m$, let $\Gamma_x = \Gamma_0 + x$ be the translation of Γ_0 by x. Note that each of the sets Γ_x is invariant under translations by $(m, 0)$ and $(0, m)$,

and that $\Gamma_x + y = \Gamma_{x+y}$, where the addition in the subscript is done mod m. Define

$$\mu = \frac{1}{m^2} \sum_{x \in \mathbb{Z}_m^2} \lambda_{\Gamma_x}.$$

This measure is certainly translation invariant, because a translate of Γ_x is another, say Γ_y, and μ is stationary for the northeast process because it is a mixture of stationary distributions. Note that m is not ergodic, because the initially frozen set, Γ_x, must remain frozen forever.

3.2. Characterization of ergodic, translation-invariant measures

We now show that if μ is an ergodic, translation-invariant stationary distribution for the northeast model with flip parameter p, then either $\mu = \nu_p$ or $\mu = \nu_0$. Denote by $P_\mu = P \times \mu$ the probability measure on $\Omega \times \{0, 1\}^{\mathbb{Z}^2}$ according to which the initial configuration ξ_0 is chosen randomly from μ and the marked Poisson processes used in the construction of Subsection 2.3 are built on $(\Omega, \mathcal{F}, P)$. Under P_μ, the distribution of ξ_0 is translation invariant, and so there exists a constant, r, $0 \le r \le 1$, such that $P_\mu\{\xi_0(x) = 1\} = r$ and, hence (by stationarity),

$$P_\mu\{\xi_t(x) = 1\} = r \quad \text{for all } x \in \mathbb{Z}^2 \text{ and all } t \ge 0.$$

If $r = 0$ then $\mu = \nu_0$. If $r = 1$ then at every rational time every site has spin 1; however, this is impossible unless $p = 1$, because it would imply that no site ever flips to spin 0, as spin values are held for time intervals of positive duration. Thus, if $0 < p < 1$ then either $r = 0$ or $0 < r < 1$.

Suppose that $0 < r < 1$. Since μ is assumed to be ergodic, the Birkhoff ergodic theorem implies that the long-time average spin value at any site x converges to r almost surely. Thus, with P_μ-probability one, site x flips its spin infinitely often and, so, $P_\mu\{\tau_x < \infty \text{ for all } x\} = 1$. Hence, Corollary 1 implies that, for μ-almost every initial configuration ζ, the conditional joint distribution of any finite spin block ξ_t^Λ given $\xi_0 = \zeta$ converges weakly to the product Bernoulli-p measure ν_p^Λ as $t \to \infty$. It now follows that $r = p$ and $\mu = \nu_p$.

3.3. Ergodicity and mixing of ν_p for $p > 1 - \beta_c$

To complete the proof of Theorem 1 and the first assertion of Theorem 2, we must show that if $p \ge 1 - \beta_c$ then ν_p is ergodic and mixing for the northeast model. The key is that if $p \ge 1 - \beta_c$ then, for ν_p-almost every spin configuration ζ, there are no infinite southwest clusters of zero spins (see [4, Chapter 12.8]). In particular, the size of the cluster containing the origin is almost surely finite.

Lemma 1. *Define $M^x(t)$ to be the number of flip opportunities at site x up to time t, and denote by x_s and x_w the southerly and westerly nearest neighbors of x, respectively. Then, for every initial configuration ζ,*

$$P\{M^x(\infty) < \infty \text{ and } M^{x_s}(\infty) = M^{x_w}(\infty) = \infty \mid \xi_0 = \zeta\} = 0.$$

Proof. It is enough to show that, on the event $G := \{M^{x_s}(\infty) = M^{x_w}(\infty) = \infty\}$, the spins at sites x_s and x_w will both almost surely be 1 at arbitrarily large times t. Denote by τ_y^k the time of the kth flip opportunity at site y. On the event G it must be that $\tau_y^k < \infty$ for all $k \ge 1$ and for both $y = x_s$ and $y = x_w$. At each time τ_y^k, the spin at site y is reset by a toss of a p-coin. Since the spin values at $y = x_s$ and $y = x_w$ play no role in determining the reset times at either x_s or x_w, it follows that the *pair* of spins at x_s and x_w will be reset by independent p-coin tosses infinitely many times. Hence, the spins will both be 1 at arbitrarily large times, with probability one.

Now assume that the initial configuration ξ_0 is chosen at random according to the product Bernoulli-p measure ν_p, independently of the marked Poisson processes used to determine the time evolution, and assume that $p \geq 1 - \beta_c$. To show that the process ξ_t is mixing (and, therefore, also ergodic), it suffices, by Corollary 1, to show that $\tau_x < \infty$ almost surely for each site x. We will show that in fact $M^x(\infty) = \infty$ almost surely.

Suppose, on the contrary, that for some site x there is positive probability, say $\varrho > 0$, that $M^x(\infty) < \infty$. Then, by Lemma 1, for at least one of the sites $y = x_s, x_w$ it must be the case that $M^y(\infty) < \infty$. It then follows that, for at least one of these two sites, the spin at this site stabilizes at 0, because if $\xi_t(y) = 1$ eventually and the other southwesterly nearest neighbor, y', has infinitely many flip opportunities, then $\xi_t(y') = 1$ for arbitrarily large times t and, so, x must have infinitely many flip opportunities, contradicting the hypothesis that $M^x(\infty) < \infty$. By induction, it follows that there is an infinite sequence, y_n, of sites such that (i) each y_{n+1} is either the southerly or westerly nearest neighbor of y_n, and (ii) the spin at each y_n eventually stabilizes at 0.

This contradicts the hypothesis that $p \geq 1 - \beta_c$; in particular, it implies (by translation invariance) that, for each $K < \infty$, there is probability at least ϱ that at large times t the origin will belong to a southwest cluster of zero-spin sites that has size greater than or equal to K. Since, by stationarity of ν_p, the size of the southwest cluster of zero spins containing the origin has the same distribution at time t as at time zero, this is impossible.

4. Exponential decay of correlations

In this section we shall establish the exponential decay of correlations given in (2) for all parameter values $p > \frac{2}{3}$. Without loss of generality, we may assume that the finite set Λ in (2) is a square, since every finite set is contained in a square; furthermore, we may assume that the northeast corner of the square Λ is the origin, $(0, 0)$.

Denote by $\partial_- \Lambda$ the exterior southwest boundary of Λ, that is, the set of all sites not in Λ that border Λ to the south or west. At any (Markov) time when the spins at all sites in $\partial_- \Lambda$ are 1, it is possible for the sites in Λ to begin to reset, starting from the southwest corner and proceeding north and east. In particular, it is possible that all spins in Λ will flip to 1 *before* any of the spins in $\partial_- \Lambda$ reset. Moreover, at any (Markov) time when all of the spins in Λ are 1, it is possible that the spins will then reset one at a time starting from the northeast corner and proceeding to the southwest, in the following order (for a 3×3 square):

$$
\begin{array}{ccc}
4 & 2 & 1 \\
7 & 5 & 3 \\
9 & 8 & 6
\end{array}
$$

We call such an occurrence a (total) Λ-*reset*. Define τ_Λ to be the time of the first Λ-reset. Observe that, at time τ_Λ, the spin at the southwest corner of Λ has just reset. At this time all spins in Λ have been reset at least once, in order, starting from the northeast corner, so the conditional distribution of the configuration $\xi_{\tau_\Lambda}^\Lambda$, given the evolution to time τ_Λ of the process in the region southwest of Λ, is the Bernoulli-p product measure ν_p^Λ.

Lemma 2. *To prove (2) it suffices to prove that, for all $t \geq 0$,*

$$
\mathrm{P}_p^{\nu_p}\{\tau_\Lambda \geq t\} \leq C_\Lambda \exp\{-\alpha_\Lambda t\}. \tag{8}
$$

Proof. This follows from Proposition 1 by an argument similar to that in the proof of Corollary 1.

Whether or not a Λ-reset will occur quickly depends on the current values of the spins at the sites to the southwest of Λ, in particular on those that directly border Λ. Since any such site with spin 0 will freeze the spins of its northern and eastern neighbors, it is impossible for Λ to reset until all the southwest boundary sites have assumed spin value 1, unless all spins in Λ initially have spin 1. Similarly, any of the boundary sites that initially have spin 0 cannot flip until their southwest neighbors have assumed spin value 1, and so on. Thus, at any time t the obstruction to a Λ-reset is the union of the oriented zero-clusters attached to the southwest boundary of Λ. Denote this union by K_t^{Λ}; that is, K_t^{Λ} is the set of all sites y for which there exists a (northeast-oriented) path of zero spins in the configuration ξ_t, beginning at y and ending at the southwest boundary of Λ. Fix $t_* > 0$ and define stopping times $T_0 \leq S_0 \leq T_1 \leq S_1 < \cdots$ by setting

$$S_0 = \min\{t \geq 0 \colon K_t^{\Lambda} = \varnothing\},$$
$$T_0 = 0,$$
$$S_n = \min\{t \geq T_n \colon K_t^{\Lambda} = \varnothing\}, \qquad\qquad n \geq 1,$$
$$T_n = \min\{t > S_{n-1} \colon K_t^{\Lambda} \neq \varnothing\} \wedge (S_{n-1} + t_*), \qquad n \geq 1,$$

where t_* is a fixed, nonrandom time.

Lemma 3. *Assume that $p > \frac{2}{3}$. Then there exist constants $C, \alpha > 0$, possibly depending on the size of the square Λ and on the choice of the constant t_*, such that, for all n and all $t > 0$,*

$$P_p^{\nu_p}\{S_n - T_n \geq t \mid \mathcal{I}_{T_n}^{\infty}\} \leq C \exp\{-\alpha t\}.$$

Proof. Let $X_t := |K_t^{\Lambda}|$ denote the cardinality of the zero-cluster attached to Λ at time t; then S_n is the first time after T_n when $X_t = 0$. At any time when the cluster K_t^{Λ} is nonempty, its exterior southwest boundary consists of $B \geq 1$ sites, all with spin 1, and its interior southwest boundary (the set of sites that are flip-eligible) consists of $A \geq 1$ sites, all with spin 0; since each interior boundary site is bordered by at most two exterior boundary sites, $B \leq 2A$.

The cluster can grow or shrink by only one site at a time, either by addition of a site in the exterior southwest boundary or deletion of a site in the interior southwest boundary. Since there is always at least one interior boundary site that is flip-eligible, the rate at which jumps in X_t occur is at least p. Moreover, at any addition/deletion event, the conditional probability that the event is a deletion, given the current configuration, is at least $pA/(pA + (1-p)B)$. This is because the instantaneous rate of deletions is pA, whereas the instantaneous rate of additions is *at most* $(1-p)B$. (All of the spins in the exterior southwest boundary of K_t^{Λ} are 1, but some of these may be frozen.)

Since $B \leq 2A$ and $p > \frac{2}{3}$, we have $(pA)/(pA + (1-p)B) > \frac{1}{2}$. Thus, the jump process X_t is stochastically dominated by a reflecting random walk on the nonnegative integers with negative drift. It follows routinely that, for ν_p-almost all initial configurations, the times it takes to reduce the size of the zero-cluster by one have uniformly bounded, exponentially decaying tails. Let N be the size of the zero-cluster with a random initial configuration (N can be arbitrarily large), and let $Y_1, \ldots, Y_N$ be the consecutive random times required to reduce the size of the cluster by one. Then $S_0 - T_0 = \sum_{j=1}^{N} Y_j$. For all sufficiently small $\theta > 0$, the moment generating functions of Y_j are finite and arbitrarily close to one. Since, under the product Bernoulli-p measure, N has an exponentially decaying tail in n, the moment generating function of $S_0 - T_0$ is finite for all sufficiently small $\theta > 0$ and, hence, $S_0 - T_0$ has an exponentially decaying tail. Note that $X_{T_n} \leq 1$ for all $n > 0$. It follows immediately that, for all $n > 0$, $S_n - T_n$ also has an exponentially decaying tail.

Proof of (8). Let G_n be the event that there is a Λ-reset between times S_n and T_{n+1}. Since, at time S_n, the zero-cluster $K_{S_n}^\Lambda$ is empty, the conditional probability of G_n, given the history of the process ξ_t up to time S_n, is bounded below by a constant, $\delta > 0$. (Recall that the spins in $\partial_- \Lambda$ must stay fixed at 1 until there is an occurrence in one of the Poisson processes N_t^y attached to a site $y \in \partial_- \Lambda$, regardless of what happens elsewhere.) Hence, by Lemma 3,

$$\tau_\Lambda \leq S_0 + \sum_{j=1}^H R_j, \tag{9}$$

where S_0, H, and $R_1, R_2, \ldots$ are mutually independent, H is the number of tosses of a δ-coin until the first occurrence of a head, and

$$P\{R_j \geq t\} = C \exp\{-\alpha t\} \quad \text{for all } t \geq \log(C/\alpha).$$

Elementary arguments similar to those given in the proof of Lemma 3 show that the random variable on the right-hand side of (9) has a finite moment generating function and, hence, has an exponentially decaying tail.

5. Influence propagation

A key component of the argument of [1] (although not explicitly noted as such) is a proof that the influence of a spin-1 site is propagated forward at a definite linear rate. In two dimensions, propagation of influence is a much more subtle business, as there is no longer a unique path along which it occurs. A natural way to measure influence propagation in the northeast model is by use of *influence regions* R_t, defined as follows. Consider an initial configuration in which all sites in the first quadrant are initially set to have spin 0 and all other sites are initially set to have random spins (using independent p-coin tosses). For each time $t \geq 0$, define

$$R_t = \{\text{sites in quadrant one that are flipped at least once by time } t\}.$$

Figure 1 shows the results of a simulation (performed by Marc Coram of the University of Chicago) in which sites in quadrants two, three, and four are initially determined by independent p-coin tosses with $p = 0.8$. The figure shows the state of the system at time $t = 1044.65$; the size of the box shown is 500×500. The influence region R_t is shown in white, and the darker-gray region in quadrants two, three, and four, Q_t, consists of those sites that have been queried at least once by time t. See Section 2 for an explanation of the query process. The black specks indicate those sites in R_t or Q_t whose spin values are 0 at time t.

The results of this and other simulations suggest the following conjecture.

Conjecture 1. *For $p > p_c$, the influence region R_t has a definite, nonrandom limit shape S. That is, with $\hat{R}$ denoting the union of all unit squares centered at points of R,*

$$\hat{R}_t/t \to S$$

in the Hausdorff metric for some nonrandom set S contained in the first quadrant of $\mathbb{R}^2$.

Note that, because the northeast process is not additive (in the sense of [5]), the usual methods for proving shape theorems (see, e.g. [2] and [7]) are not applicable.
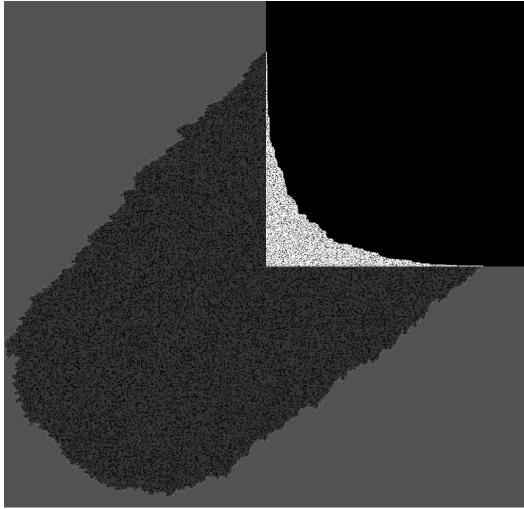
FIGURE 1: Influence propagation in the northeast model.

References

[1] ALDOUS, D. AND DIACONIS, P. (2002). The asymmetric one-dimensional constrained Ising model: rigorous results. *J. Statist. Phys.* **107,** 945–975.

[2] COX, J. T. AND DURRETT, R. (1981). Some limit theorems for percolation processes with necessary and sufficient conditions. *Ann. Prob.* **9,** 583–603.

[3] FREDRICKSON, G. H. AND ANDERSEN, H. C. (1984). Kinetic Ising models of the glass transition. *Phys. Rev. Lett.* **53,** 1244–1247.

[4] GRIMMETT, G. (1999). *Percolation*, 2nd edn. Springer, Berlin.

[5] LIGGETT, T. M. (1985). *Interacting Particle Systems.* Springer, New York.

[6] PITTS, S. J., YOUNG, T. AND ANDERSEN, H. C. (2000). Facilitated spin models, mode coupling theory, and ergodic–nonergodic transitions. *J. Chem. Phys.* **113,** 8671–8679.

[7] RICHARDSON, D. (1973). Random growth in tessellation. *Proc. Camb. Phil. Soc.* **74,** 515–528.