LECTURE 1

Interacting Particle Systems on $\mathbb{Z}^d$, $d \geq 1$.

Definition, construction, three examples.

§ GOAL

As area of research, IPS started in the 1970s, with pioneers

Spitzer, Dobrushin, Harris, Holley, Stroock, Liggett, Griffeath, Durrett

Over the years, IPS has turned out to be a fertile breeding ground for the development of new ideas and techniques in mathematical statistical physics, including graphical representation, coupling, duality, correlation inequalities.

We start by defining what an IPS is. We focus on spin-flip systems, which constitute a particularly tractable class. Within this class we focus on three examples:

Stochastic Ising Model (SIM)
Voter Model (VM)
Contact Process (CP)

Standard references for IPS on $\mathbb{Z}^d$ are:

T.M. Liggett, *Interacting Particle Systems*,
Grundlehren der mathematische Wissenschaften 276,
Springer, New York, 1985.

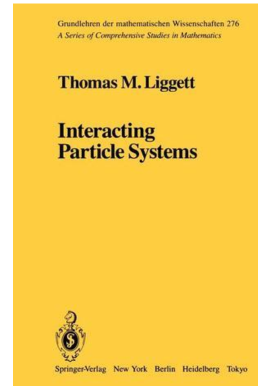

T.M. Liggett, *Stochastic Interacting Systems: Contact, Voter and Exclusion Processes*,
Grundlehren der mathematische Wissenschaften 324,
Springer, Berlin, 1999.

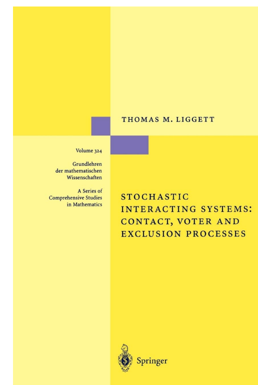

For most of the results to be described below, references can be found in these monographs.

§ DEFINITIONS

An Interacting Particle System (IPS) formally is a Markov process $\xi = (\xi_t)_{t \geq 0}$ on the state space

$$\Omega = \{0,1\}^{\mathbb{Z}^d}, \qquad d \geq 1,$$

where

$$\xi_t = \{\xi_t(x)\colon\ x \in \mathbb{Z}^d\}$$

denotes the configuration at time t, with $\xi_t(x) = 1$ or 0 meaning that there is a 'particle' or a 'hole' at site x at time t, respectively. Alternative interpretations are

$$\begin{aligned} 1 &= \text{spin-up/democrat/infected} \\ 0 &= \text{spin-down/republican/healthy.} \end{aligned}$$

The configuration changes with time, which models how:

- magnetic atoms flip up and down as a result of noise,
- two political parties evolve in an election campaign,
- a virus spreads through a population.

The evolution is specified via a set of local transition rates

$$c(x,\eta), \qquad x \in \mathbb{Z}^d,\ \eta \in \Omega,$$

playing the role of the rate at which the state at site x changes in the configuration η, i.e.,

$$\eta \to \eta^x$$

with η^x the configuration obtained from η by changing the state at site x (either $0 \to 1$ or $1 \to 0$). Since there are only two possible states at each site, the IPS is called a spin-flip system.

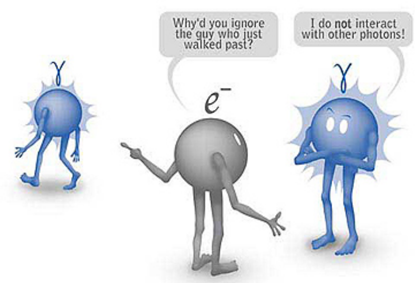

If $c(x,\eta)$ depends on η only via $\eta(x)$, the value of the spin at x, then ξ consists of independent spin-flips. In general, however, the rate to flip the spin at x depends on the spins located in the neighbourhood of x (possibly even on all spins). This dependence models an interaction between the spins at different sites.

In order for ξ to be well-defined, some restrictions must be placed on the local transition rates: $c(x,\eta)$ must depend only weakly on the states at far away sites (formally, $\eta \mapsto c(x,\eta)$ is continuous in the product topology) and must be not too large (formally, bounded away from infinity in some appropriate sense).

Liggett 1985

SHIFT-INVARIANT ATTRACTIVE SYSTEMS

- Shift-invariant

Typically it is assumed that

$$c(x, \eta) = c(x + y, \tau_y \eta) \qquad \forall\, y \in \mathbb{Z}^d$$

with τ_y the shift of space over y, i.e.,

$$(\tau_y \eta)(x) = \eta(x - y), \qquad x \in \mathbb{Z}^d.$$

This property says that the flip rate at x only depends on the configuration η seen relative to x, which is natural when the interaction between spins is homogeneous in space.

- Attractive

Another useful assumption is that the interaction favors spins that are alike, i.e.,

$$\eta \preceq \eta' \to \begin{cases} c(x,\eta) \leq c(x,\eta') & \text{if } \eta(x) = \eta'(x) = 0, \\ c(x,\eta) \geq c(x,\eta') & \text{if } \eta(x) = \eta'(x) = 1, \end{cases}$$

where $\preceq$ denotes the partial order in Ω. This property says that the spin at x flips up faster in η' than in η when η' is everywhere larger than η, and flips down slower.

In other words, the dynamics preserves $\preceq$. Spin-flip systems with this property are called attractive.

§ THREE EXAMPLES Liggett 1985

1. SIM: Stochastic Ising model

   This model is defined on $\Omega = \{-1, 1\}^{\mathbb{Z}^d}$ with rates

   $$c(x, \eta) = \exp[-\beta\eta(x) \sum_{y \sim x} \eta(y)], \qquad \beta \geq 0,$$

   which means that spins prefer to align with the majority of the neighbouring spins.

2. VM: Voter model

   This model is defined on $\Omega = \{0, 1\}^{\mathbb{Z}^d}$ with rates

   $$c(x, \eta) = \frac{1}{2d} \sum_{y \sim x} 1_{\{\eta(y) \neq \eta(x)\}},$$

   which means that sites choose a random neighbour at rate 1 and adopt the opinion of that neighbour.

3. CP: Contact process

This model is defined on $\Omega = \{0,1\}^{\mathbb{Z}^d}$ with rates

$$c(x,\eta) = \begin{cases} \lambda \sum_{y \sim x} \eta(y), & \text{if } \eta(x) = 0, \\ 1, & \text{if } \eta(x) = 1, \end{cases} \qquad \lambda \geq 0,$$

which means that infected sites become healthy at rate 1 and healthy sites become infected at rate λ times the number of infected neighbours.

EXERCISE:

Check that these three examples indeed are shift-invariant and attractive.

In the sequel we will discuss each model in some detail. We will see that the properties

shift-invariant
attractive

allow for a number of interesting conclusions concerning their equilibrium, as well as their convergence to equilibrium.

EXERCISE: (= digression)

Look up the following notions:

(1) Stochastic ordering of two IPSs.

(2) Ordered coupling of two IPSs.

(3) Strassen theorem about stochastic ordering being equivalent to ordered coupling.

§ CONVERGENCE TO EQUILIBRIUM

Write [0] and [1] to denote the configurations $\eta \equiv 0$ and $\eta \equiv 1$, respectively. These are the smallest, respectively, the largest configurations in the partial order, and hence

$$[0] \preceq \eta \preceq [1], \qquad \forall\, \eta \in \Omega.$$

Since the dynamics preserves the partial order (see below), we obtain information about what happens when the system starts from any $\eta \in \Omega$ by comparing with what happens when it starts from [0] or [1].

An IPS can be described by semigroups of transition kernels

$$(P_t)_{t\geq 0}.$$

Formally, P_t is an operator acting on $C_b(\Omega)$, the space of bounded continuous functions on Ω, as

$$(P_t f)(\eta) = \mathbb{E}_\eta[f(\xi_t)], \qquad \eta \in \Omega,\ f \in C_b(\Omega).$$

If this definition holds on a dense subset of $C_b(\Omega)$, then it uniquely determines P_t. Formally, we can write $P_t = \mathrm{e}^{tL}$ with L the generator of the IPS:

$$(Lf)(\eta) = \sum_{x\in\mathbb{Z}^d} c(x,\eta)[f(\eta^x) - f(\eta)].$$

EXERCISE:

Check that P_0 is the identity and that $P_{s+t} = P_t \circ P_s$ for all $s,t \geq 0$ (where $\circ$ denotes composition).

Alternatively, the semigroup can be viewed as acting on the space of probability measures μ on Ω via the duality relation

$$\int_\Omega f \,\mathrm{d}(\mu P_t) = \int_\Omega (P_t f)\,\mathrm{d}\mu, \qquad f \in C_b(\Omega).$$

LEMMA 1.1

Let $(P_t)_{t\geq 0}$ *denote the semigroup of transition kernels that is associated with* ξ*. Write* $\delta_\eta P_t$ *to denote the law of* ξ_t *conditional on* $\xi_0 = \eta$ *(which is a probability distribution on* Ω*). Then*

$t \mapsto \delta_{[0]} P_t$ is stochastically increasing,
$t \mapsto \delta_{[1]} P_t$ is stochastically decreasing.

PROOF

For $t, h \geq 0$,

$$\delta_{[0]} P_{t+h} = (\delta_{[0]} P_h) P_t \succeq \delta_{[0]} P_t,$$
$$\delta_{[1]} P_{t+h} = (\delta_{[1]} P_h) P_t \preceq \delta_{[1]} P_t,$$

where we use that

$$\delta_{[0]} P_h \succeq \delta_{[0]}, \qquad \delta_{[1]} P_h \preceq \delta_{[1]},$$

for any $h \geq 0$, and we use the Strassen theorem and the coupling representation that goes with the partial order. □

COROLLARY 1.2

Both

$$\underline{\nu} = \lim_{t\to\infty} \delta_{[0]} P_t = \text{lower stationary law},$$
$$\overline{\nu} = \lim_{t\to\infty} \delta_{[1]} P_t = \text{upper stationary law},$$

exist as probability distributions on Ω *and are equilibria for the dynamics. Any other equilibrium* π *satisfies* $\underline{\nu} \preceq \pi \preceq \overline{\nu}$.

PROOF

This is immediate from Lemma 1.1 and the sandwich

$$\delta_{[0]} P_t \preceq \delta_\eta P_t \preceq \delta_{[1]} P_t, \qquad \eta \in \Omega,\ t \geq 0.$$

$\square$

The class of all equilibria for the dynamics is a convex set in the space of signed bounded measures on Ω. An element of this set is called extremal when it is not a proper linear combination of any two distinct elements in the set, i.e., is not of the form

$$p\nu_1 + (1-p)\nu_2, \qquad p \in (0,1),\ \nu_1 \neq \nu_2.$$

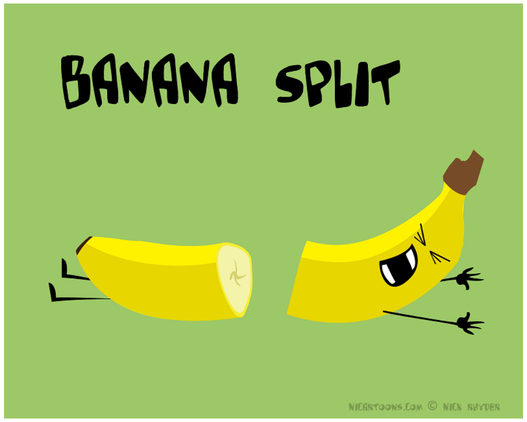

LEMMA 1.3

Both $\underline{\nu}$ and $\overline{\nu}$ are extremal.

PROOF

We give the proof for $\overline{\nu}$ only. Suppose that

$$\overline{\nu} = p\nu_1 + (1-p)\nu_2, \qquad \nu_1 \neq \nu_2,\ p \in (0,1).$$

Since ν_1 and ν_2 are equilibria, by Corollary 1.2 we have

$$\int_\Omega f \mathrm{d}\nu_1 \leq \int_\Omega f \mathrm{d}\overline{\nu}, \qquad \int_\Omega f \mathrm{d}\nu_2 \leq \int_\Omega f \mathrm{d}\overline{\nu},$$

for any f non-decreasing. Since

$$\int_\Omega f \mathrm{d}\overline{\nu} = p\int_\Omega f \mathrm{d}\nu_1 + (1-p)\int_\Omega f \mathrm{d}\nu_2$$

and $p \in (0,1)$, both inequalities must be equalities. Integrals of non-decreasing functions determine the measure w.r.t. which is being integrated, and so it follows that $\nu_1 = \overline{\nu} = \nu_2$. $\square$

EXERCISE

Prove that integrals of non-decreasing functions determine the measure.

COROLLARY 1.4

The following three properties are equivalent (for shift-invariant and attractive spin-flip systems):

1. ξ *is ergodic (i.e.,* $\delta_\eta P_t$ *has the same limiting distribution as* $t \to \infty$ *for all* η*).*
2. *There is a unique stationary distribution,*
3. $\underline{\nu} = \overline{\nu}$.

PROOF

The claim is obvious in view of the sandwich of the configurations between [0] and [1]. □

REMARK

If $\underline{\nu} \neq \overline{\nu}$, then there is no guarantee that $\lim_{t\to\infty} \mu P_t = \nu$ exists for arbitrary μ. In fact, stronger assumptions than attractiveness are needed to make that happen. We do know that any convergent subsequence has a limit ν such that $\underline{\nu} \preceq \nu \preceq \overline{\nu}$.

§ Example 1: Stochastic Ising Model

For $\beta = 0$, $c(x,\eta) = 1$ for all x and η, in which case the dynamics consists of independent spin-flips, up and down at rate 1. In that case $\overline{\nu} = \underline{\nu} = (\frac{1}{2}\delta_{-1} + \frac{1}{2}\delta_{+1})^{\otimes \mathbb{Z}^d}$.

For $\beta > 0$ the dynamics has a tendency to align spins. For small β this tendency is weak, for large β it is strong. It turns out that in $d \geq 2$ there is a critical value $\beta_d \in (0,\infty)$ such that

$$\begin{aligned} \beta \leq \beta_d &: \quad \underline{\nu} = \overline{\nu}, \\ \beta > \beta_d &: \quad \underline{\nu} \neq \overline{\nu}. \end{aligned}$$

The proof uses the so-called Peierls argument. In the first case there is a unique ergodic equilibrium, which depends on β and is denoted by ν_β.

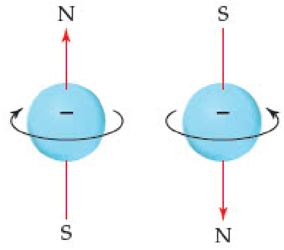

In the second case there are two extremal equilibria, both of which depend on β and are denoted by

$$\begin{aligned} \nu_\beta^+ &= \text{plus state} \quad \text{with} \quad m_\beta^+ = \textstyle\int_\Omega \eta(0)\nu_\beta^+(d\eta) > 0, \\ \nu_\beta^- &= \text{minus-state} \quad \text{with} \quad m_\beta^- = \textstyle\int_\Omega \eta(0)\nu_\beta^-(d\eta) < 0, \end{aligned}$$

which are called the magnetised states. Note that ν_β^+ and ν_β^- are images of each other under the swapping of $+1$'s and -1's and so $m_\beta^+ = -m_\beta^- = m_\beta$.

It can be shown that in $d = 2$ all equilibria are a convex combination of ν_β^+ and ν_β^-, while in $d \geq 3$ other equilibria are possible as well (e.g. not shift-invariant) when β is large enough. It turns out that $\beta_1 = \infty$, i.e., in $d = 1$ the SIM is ergodic for all $\beta > 0$. It is known that $\beta_2 = \frac{1}{2}\log(1 + \sqrt{2})$.

Example 2: Voter Model

Note that [0] and [1] are both traps for the dynamics (if all sites have the same opinion, then no change of opinion occurs), and so

$$\underline{\nu} = \delta_{[0]}, \qquad \overline{\nu} = \delta_{[1]}.$$

It turns out that in $d = 1, 2$ these are the only extremal equilibria, while in $d \geq 3$ there is a 1-parameter family of extremal equilibria

$$(\nu_\rho)_{\rho \in [0,1]}$$

with ρ the density of 1's, i.e., $\nu_\rho(\eta(0) = 1) = \rho$. This fact is remarkable because the VM has no parameter. For $\rho = 0$ and $\rho = 1$ these equilibria coincide with $\delta_{[0]}$ and $\delta_{[1]}$, respectively.

REMARK

The dichotomy $d = 1, 2$ versus $d \geq 3$ is directly related to simple random walk being recurrent in $d = 1, 2$ and transient in $d \geq 3$.

This property has to do with the fact that the VM is dual to a system of coalescing random walks.

EXERCISE

Give the graphical representation of the VM and indicate how duality is obtained via time reversal.

§ Example 3: Contact Process

Note that [0] is a trap for the dynamics (if all sites are healthy, then no infection will ever occur), and so

$$\underline{\nu} = \delta_{[0]}.$$

For small λ infection is transmitted slowly, for large λ rapidly. It turns out that in $d \geq 1$ there is a critical value $\lambda_d \in (0, \infty)$ such that

$$\begin{aligned} \lambda \leq \lambda_d: &\quad \overline{\nu} = \delta_{[0]} &= \text{extinction, no epidemic},\\ \lambda > \lambda_d: &\quad \overline{\nu} \neq \delta_{[0]} &= \text{survival, epidemic}. \end{aligned}$$

LEMMA 1.5 Liggett 1985, Durrett 1988

(i) $d\lambda_d \leq \lambda_1$.

(ii) $2d\lambda_d \geq 1$.

(iii) $\lambda_1 < \infty$.

Note that (i–iii) combine to yield that $0 < \lambda_d < \infty$ for all $d \geq 1$, so that the phase transition occurs at a non-trivial value of the infection rate parameter.

EXERCISE

Give the proof of (i–ii) with the help of coupling.

REMARK

Sharp estimates are available for λ_1, but these require heavy machinery. Numerically, $\lambda_1 \approx 1.6494$. A series expansion of λ_d in powers of $1/2d$ is known up to several orders, but again the proof is very technical.

§ THE COX-GREVEN FINITE SYSTEMS SCHEME

As a prelude to Lectures 2-4, in which we take a closer look at SIM, VM, CP on finite random graphs, we describe what is known about these processes on a large finite torus in $\mathbb{Z}^d$,

$$\Lambda_N = [0, N)^d \cap \mathbb{Z}^d, \qquad N \in \mathbb{N},$$

endowed with periodic boundary conditions.

The behaviour on Λ_N is different from that on $\mathbb{Z}^d$. In particular, there is an N-dependent characteristic time scale α_N on which the process notices that Λ_N differs from $\mathbb{Z}^d$, resulting in different behaviour for short, moderate and long times.

A systematic study was initiated in

Cox, Greven 1990, Cox, Greven, Shiga 1995+1998

WARNING: The text on pages 28–36 is technical.

§ SIM ON THE TORUS

Since $|\Lambda_N| < \infty$, we have

$$\underline{\nu}^N = \overline{\nu}^N = \nu_\beta^N \text{ with } \int_\Omega \eta(0)\nu_\beta^N(\mathrm{d}\eta) = 0 \qquad \forall\, \beta \in (0,\infty),$$

i.e., on any finite lattice eventually the average magnetisation vanishes. An interesting question is: How long does it take the SIM to loose its magnetisation and what does it do along the way?

Let

$$\mathcal{M}_t^N = \frac{1}{|\Lambda_N|} \sum_{x \in \Lambda_N} \xi_t^N(x)$$

denote the magnetisation at time t. Suppose that the law of ξ_0^N is the restriction to Λ_N of the equilibrium measure ν_β^- on $\mathbb{Z}^d$, which has magnetisation m_β^-.

THEOREM 1.6 Cox, Greven 1990

Bovier, Eckhoff, Gayrard, Klein 2002

(a) *For* $\beta < \beta_d$ *and any* $T_N \to \infty$*,*

$$\lim_{N\to\infty} \mathcal{L}\left[\mathcal{M}^N_{T_N}\right] = \delta_0.$$

(b) *For* $\beta > \beta_d$*,*

$$\lim_{N\to\infty} \mathcal{L}\left[\mathcal{M}^N_{s\alpha_N}\right] = m_\beta^{Z_s}, \qquad Z_0 = -,$$

where $(Z_s)_{s\geq 0}$ *is the Markov chain on* $\{-,+\}$ *jumping at rate* 1*, and* α_N *is the average crossover time between the magnetisations associated with* ν_β^- *and* ν_β^+ *on* $\mathbb{Z}^d$ *restricted to* Λ_N*.*

For $\beta > \beta_d$ it can further be shown that $(\xi^N_{s\alpha_N})_{s\geq 0}$ converges in distribution to $\nu_\beta^{Z_s}$ as $N \to \infty$.

The computation of α_N is hard and belongs to the area of metastability. It is expected that

$$\alpha_N = \exp\left[\kappa_d(\beta) N^{d-1}(1 + o(1))\right]$$

with $\kappa_d(\beta)$ the free energy of the so-called Wulff droplet of volume $\frac{1}{2}$ in $\mathbb{R}^d$ representing the barrier between ν_β^-, ν_β^+.

The proof remains a challenge.

Schonmann, Shlosman 1998

§ VM ON THE TORUS

Since $|\Lambda_N| < \infty$, we have

$$\underline{\nu}^N = [0]_N, \qquad \overline{\nu}^N = [1]_N,$$

i.e., on any finite lattice eventually consensus is reached. An interesting question is: How long does it take the VM to reach consensus and what does it do along the way?

Let

$$\mathcal{O}_t^N = \frac{1}{|\Lambda_N|} \sum_{x \in \Lambda_N} \xi_t^N(x)$$

denote the fraction of 1-opinions at time t. Suppose that the law of ξ_0^N is the restriction to Λ_N of a shift-invariant and ergodic probability measure on $\mathbb{Z}^d$ with mean $\theta \in [0, 1]$.

THEOREM 1.7 Cox, Greven 1990

(a) *For* $d = 1, 2$ *and any* $T_N \to \infty$,

$$\lim_{N\to\infty} \mathcal{L}\Big[\mathcal{O}^N_{T_N}\Big] = (1-\theta)\delta_0 + \theta\delta_1.$$

(b) *For* $d \geq 3$,

$$\lim_{N\to\infty} \mathcal{L}\Big[\mathcal{O}^N_{s\alpha_N}\Big] = Z_s, \qquad Z_0 = \theta,$$

where $\alpha_N = |\Lambda_N|$ *and* $(Z_s)_{s\geq 0}$ *is the Fisher-Wright diffusion on* $[0,1]$ *with diffusion constant* $1/G_d$, *the inverse of the average number of visits to* 0 *of simple random walk on* $\mathbb{Z}^d$.

EXERCISE

Give a heuristic explanation of (a).

§ CP ON THE TORUS

Since $|\Lambda_N| < \infty$, we have

$$\underline{\nu}^N = \overline{\nu}^N = [0]_N \qquad \forall\, \lambda \in (0,\infty),$$

i.e., on a finite lattice every infection eventually becomes extinct, irrespective of the infection rate.

An interesting question is the following: Starting from $[1]_N$, how long does it take the CP to reach $[0]_N$? In particular, we want to know the extinction time

$$\tau_{[0]_N} = \inf\{t \geq 0 \colon\, \xi_t^N = [0]_N\}.$$

We expect this time to grow slowly with N when $\lambda < \lambda_d$ and rapidly with N when $\lambda > \lambda_d$, where λ_d is the critical infection threshold for the infinite lattice $\mathbb{Z}^d$.

Let

$$\mathcal{I}_t^N = \frac{1}{|\Lambda_N|} \sum_{x \in \Lambda_N} \xi_t^N(x)$$

denote the fraction of infected vertices at time t. Suppose that $\xi_0^N = [1]_N$.

THEOREM 1.8 Cox, Greven 1990

(a) *For* $\lambda < \lambda_d$ *and any* $T_N \to \infty$,

$$\lim_{N\to\infty} \mathcal{L}\left[\mathcal{I}_{T_N}^N\right] = \delta_0.$$

(b) *For* $\lambda > \lambda_d$,

$$\lim_{N\to\infty} \mathcal{L}\left[\mathcal{I}_{s\alpha_N}^N\right] = Z_s, \qquad Z_0 = 1,$$

where $\alpha_N = \mathbb{E}_{[1]_N}(\tau_{[0]_N})$ *and* $(Z_s)_{s\geq 0}$ *is the Markov chain on* $\{0,1\}$ *that jumps from* 1 *to* 0 *at rate* 1 *and is absorbed in* 0.

THEOREM 1.9

Durrett, Liu 1988, Durrett, Schonmann 1988, Mountford 1993+1999

There exist $C_-(\lambda), C_+(\lambda) \in (0, \infty)$ *such that*

$$\lambda < \lambda_d: \qquad \lim_{N \to \infty} \frac{\alpha_N}{\log |\Lambda_N|} = C_-(\lambda),$$

$$\lambda > \lambda_d: \qquad \lim_{N \to \infty} \frac{\log \alpha_N}{|\Lambda_N|} = C_+(\lambda).$$

In the subcritical phase the extinction time grows logarithmically fast with the volume of Λ_N, while in the supercritical phase it grows exponentially fast. This is a rather dramatic dichotomy.

EXERCISE

Explain heuristically where the dichotomy comes from, i.e., give a physical rather than a mathematical reason for the difference between the two phases.

EXERCISE

Why is it plausible that the distribution of the extinction time is exponential on the scale of its mean?

Rough polynomial bounds on α_N are available in $d = 1$ at $\lambda = \lambda_1$.

Duminil-Copin, Tassion, Teixeira 2017