Dynamical percolation in the plane, and three things I learnt from András

Gábor Pete
http://www.math.bme.hu/∼gabor

Joint works with
Christophe Garban (ENS Lyon)
Alan Hammond (University of Oxford)
Elchanan Mossel (UC Berkeley)
Oded Schramm (Microsoft Research, 1961-2008)

The first thing: expanders

An (n, d, c)-expander is a finite graph $G(V, E)$ on n vertices, degrees bounded by d, and $|\partial S|/|S| \geqslant c > 0$ for all $|S| \leqslant |V|/2$.

Not like a torus, not like a finite tree. Random d-regular graph is good. Non-random examples used to be hard: Margulis '73 from Kazhdan groups.

Formulation using **simple random walk**: take Markov operator $Pf(x) := \mathbf{E}[\, f(X_1) \mid X_0 = x\,]$. Self-adjoint w.r.t. $(f, g)_\pi := \sum_{x \in V} f(x)g(x)\pi(x)$, where π is the stationary measure. Then the condition:

$(P\mathbf{1}_S, \mathbf{1}_S)_\pi \leqslant (1-c)(\mathbf{1}_S, \mathbf{1}_S)_\pi$. More generally, for $\pi(\mathrm{supp} f) \leqslant 1 - \epsilon$,

$$(Pf, f) \leqslant (1 - \delta_1)(f, f) \quad \text{and} \quad (Pf, Pf) \leqslant (1 - \delta_2)(f, f)\,.$$

Equivalently, spectral gap $g := 1 - \lambda_2 > 0$, where

$$\lambda_2 := \sup_{f \perp \mathbf{1}} \frac{(Pf, f)}{(f, f)} = \sup_{f \perp \mathbf{1}} \frac{(Pf, Pf)^{1/2}}{(f, f)^{1/2}}\,.$$

Theorem (Ajtai, Komlós & Szemerédi 1987). Let $(X_i)_{i=0}^{\infty}$ be a stationary reversible chain with P and π and $\lambda_2 < 1$, and let $\pi(A) \leqslant \beta < 1$. Then there exists $\gamma(\lambda_2, \beta) > 0$ with

$$\mathbf{P}\big[\, X_i \in A \text{ for all } i = 0, 1, \ldots, t \,\big] \leqslant C(1-\gamma)^t\,.$$

My proof. Consider the projection $Q : f \mapsto f\mathbf{1}_A$. Then,

$$\begin{aligned}
\mathbf{P}\big[\, X_i \in A \text{ for } i = 0, 1, \ldots, 2t+1 \,\big] &= \big(Q(PQ)^{2t+1}\mathbf{1}, \mathbf{1}\big) \\
&= \big(P(QP)^t Q\mathbf{1}, (QP)^t Q\mathbf{1}\big)\,, \text{ by self-adjointness of } P \text{ and } Q \\
&\leqslant (1-\delta_1)\,\big((QP)^t Q\mathbf{1}, (QP)^t Q\mathbf{1}\big)\,, \text{ by } \pi(\mathrm{supp}(Qg)) \leqslant \beta \\
&\leqslant (1-\delta_1)\,\big(P(QP)^{t-1} Q\mathbf{1}, P(QP)^{t-1} Q\mathbf{1}\big)\,, \text{ by } Q \text{ being a projection} \\
&\leqslant (1-\delta_1)\,(1-\delta_2)\,\big((QP)^{t-1} Q\mathbf{1}, (QP)^{t-1} Q\mathbf{1}\big)\,, \text{ by } \pi(\mathrm{supp}(Qg)) \leqslant \beta \\
&\leqslant (1-\delta_1)\,(1-\delta_2)^t\,\big(Q\mathbf{1}, Q\mathbf{1}\big)\,, \text{ by iterating previous step} \\
&\leqslant (1-\delta_1)\,(1-\delta_2)^t\,\beta\,,
\end{aligned}$$

and done for odd times. For even times, use monotonicity in t. ∎

The second thing: bus paradox

Poisson point process of buses, mean 1 exponential waiting times between.

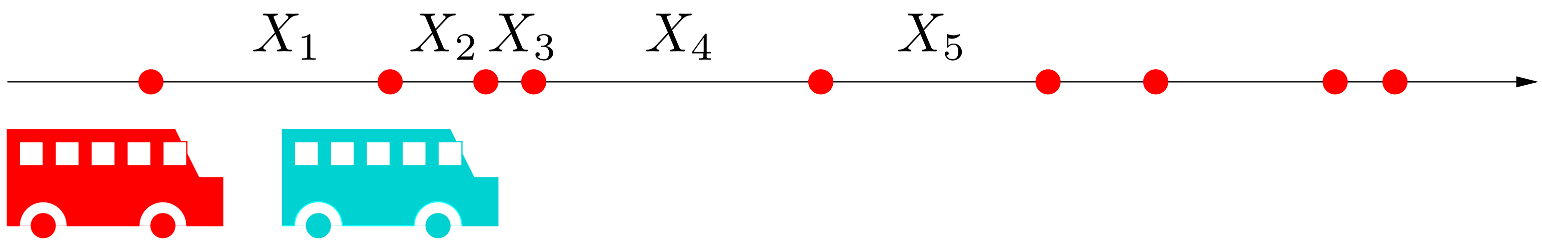

The second thing: bus paradox

Poisson point process of buses, mean 1 exponential waiting times between.

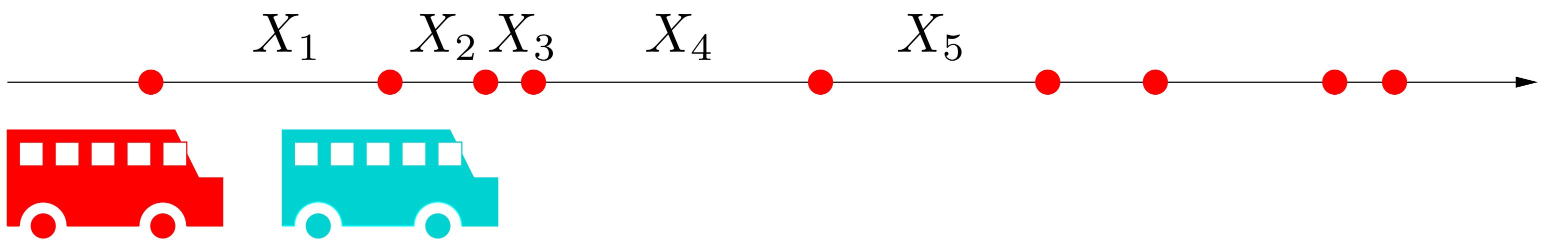

When we arrive at the bus stop, although middle of a waiting time, still mean 1 exponential time to go.

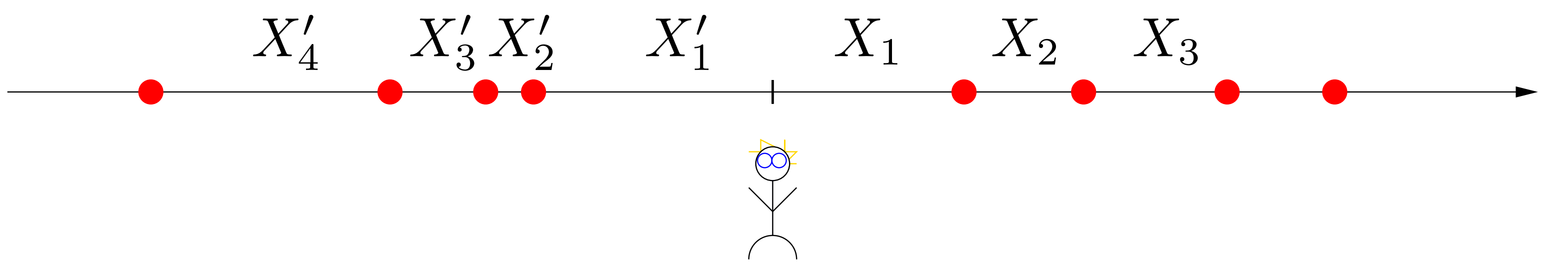

Our waiting period is doubled! *The bus we are waiting for is always late.*

Size biasing: longer waiting periods are more likely to contain a given point.

The third thing: Holley's proof of FKG inequality

Bond percolation on a graph $G(V, E)$: each edge in E is chosen open with probability p, closed with $1 - p$, independently. Gives random $\omega \subset E$.

An event $\mathcal{A}$ is increasing if $\omega \in \mathcal{A}$, $\omega \subset \omega' \implies \omega' \in \mathcal{A}$.

For instance, $G = K_n$ complete graph, ω is Erdős-Rényi $G(n, p)$.
$\mathcal{A} = \{\omega$ has Hamilton cycle$\}$ and $\mathcal{B} = \{\omega$ is non-planar$\}$ are increasing.

Harris-FKG: increasing events positively correlated, $\mathbf{P}[\mathcal{A} \cap \mathcal{B}] \geqslant \mathbf{P}[\mathcal{A}]\, \mathbf{P}[\mathcal{B}]$.
Same for decreasing; increasing and decreasing are negatively correlated.

Dynamical coupling proof. Let $\omega_0 = \emptyset$ and $\hat{\omega}_0 = \binom{V}{2}$. Note that $\hat{\omega}_0 \in \mathcal{A}$.
Each edge has Poisson clock; when rings, open with p, close with $1 - p$, in both configurations, except that in $\hat{\omega}_t$ keep it open if closing killed $\mathcal{A}$.

Get $\omega_t \subset \hat{\omega}_t$ for all $t \geqslant 0$. If $\omega_t \in \mathcal{B}$, then also $\hat{\omega}_t \in \mathcal{B}$.

But $\omega_t \xrightarrow{d} \mathbf{P}[\,\cdot\,]$ and $\hat{\omega}_t \xrightarrow{d} \mathbf{P}[\,\cdot \mid \mathcal{A}]$, hence $\mathbf{P}[\mathcal{B}] \leqslant \mathbf{P}[\mathcal{B} \mid \mathcal{A}]$, and done. ∎

Bernoulli(p) bond and site percolation on planar lattices

Each site (or bond) is chosen open with probability p, closed with $1 - p$, independently. Consider open connected clusters. $\theta(p) := \mathbf{P}_p[0 \longleftrightarrow \infty]$.

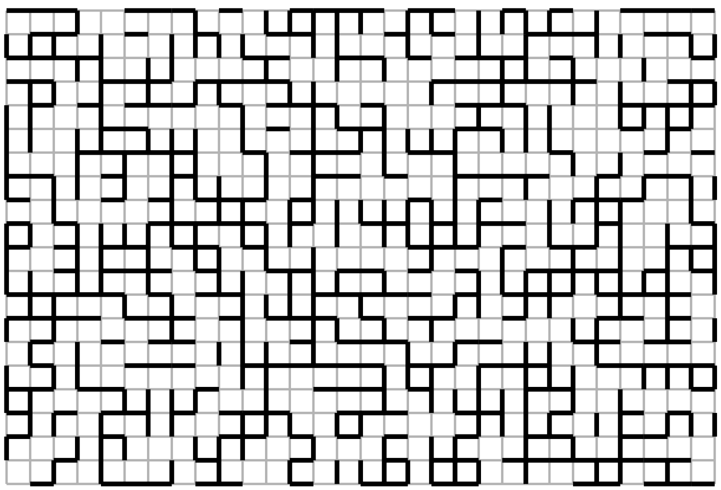

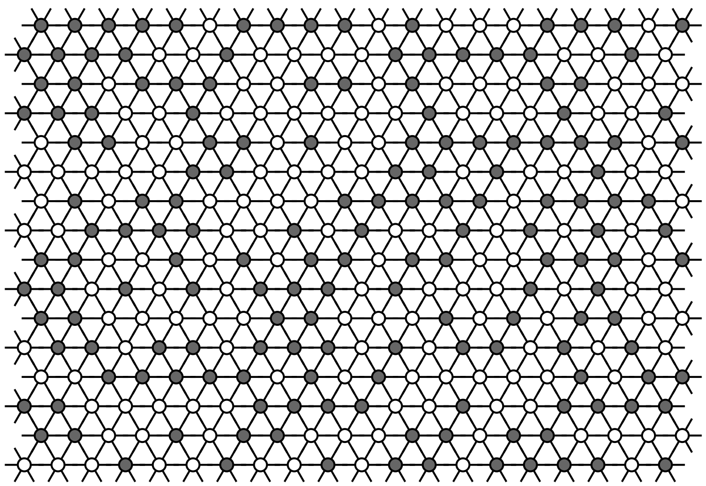

Bernoulli(p) bond and site percolation on planar lattices

Each site (or bond) is chosen open with probability p, closed with $1-p$, independently. Consider open connected clusters. $\theta(p) := \mathbf{P}_p[0 \longleftrightarrow \infty]$.

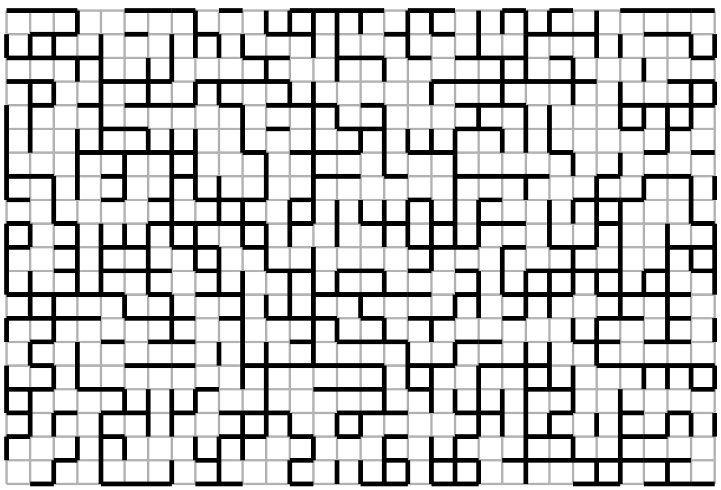

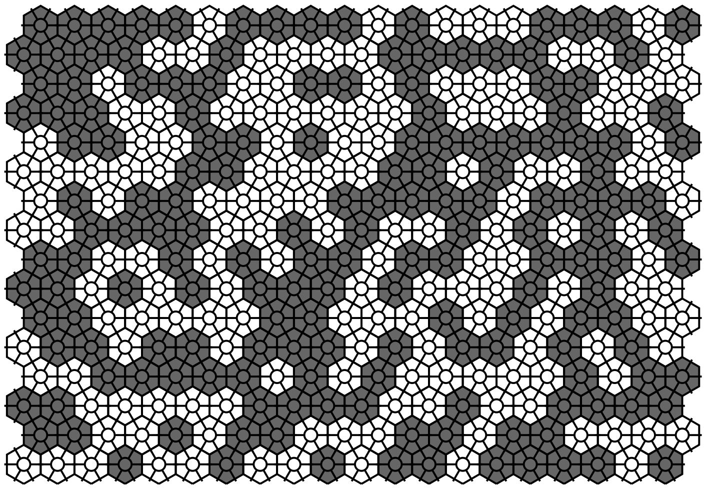

Bernoulli(p) bond and site percolation on planar lattices

Each site (or bond) is chosen open with probability p, closed with $1-p$, independently. Consider open connected clusters. $\theta(p) := \mathbf{P}_p[0 \longleftrightarrow \infty]$.

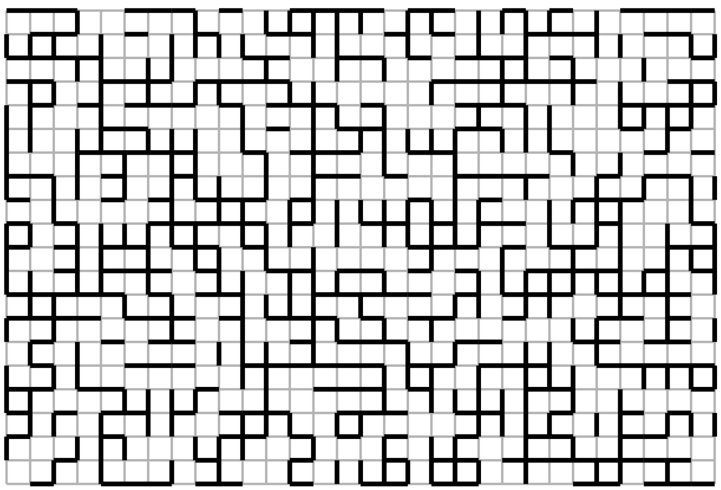

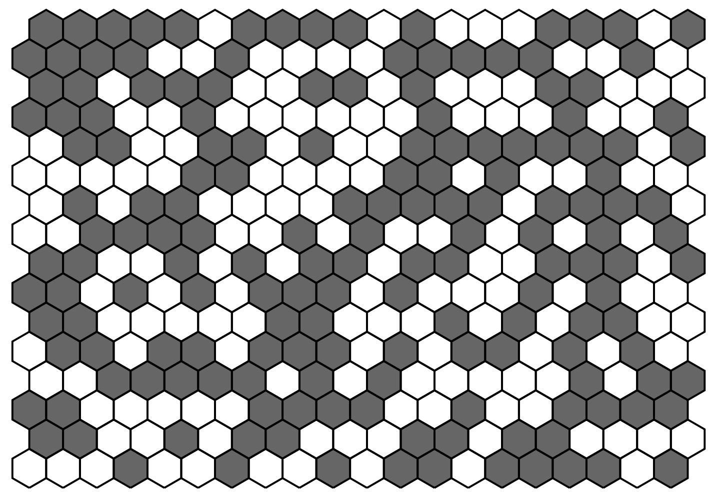

Theorem (Harris 1960 and Kesten 1980).

$p_c(\mathbb{Z}^2, \text{bond}) = p_c(\Delta, \text{site}) = 1/2$, and $\theta(1/2) = 0$.
For $p > 1/2$, there is a.s. a unique infinite cluster.

Dynamical percolation

Each variable is resampled according to an independent Poisson(1) clock, keeping $\mathbf{P}_{p_c}$ stationary. This is a Markov process $\{\omega(t) : t \in (-\infty, \infty)\}$.

An exceptional time for percolation is a (random) t at which there is an infinite cluster in $\omega(t)$; that is, where this almost sure static property fails.

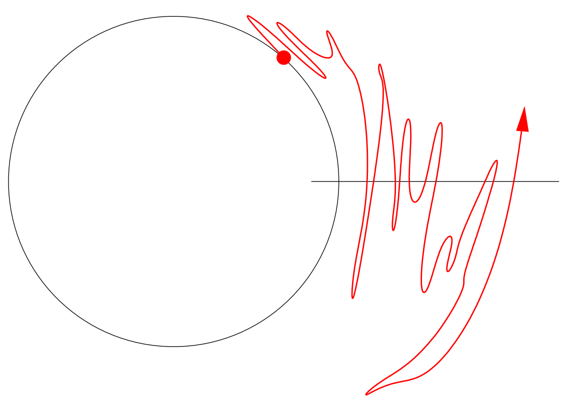

Toy example: Brownian motion on circle hitting 0 vs. a uniform random point.

Exceptional times form a random Cantor set, of Lebesgue measure zero (because of Fubini) and Hausdorff-dimension $1/2$.

If the static event is (A) not extremely unlikely, and (B) has fast decorrelation in time (very sensitive to noise), then we might see exceptional times.

Dynamical percolation results

Theorem (Häggström, Peres & Steif 1997).

- No exceptional times when $p \neq p_c$.
- No exceptional times at p_c for bond percolation on T_k and $\mathbb{Z}^d$, $d \geqslant 19$.

The latter is essentially due to Hara-Slade '90, proving that high d behaves like a tree: the off-critical exponent in $\theta(p_c + \epsilon) = \epsilon^{\beta+o(1)}$ is $\beta = 1$:

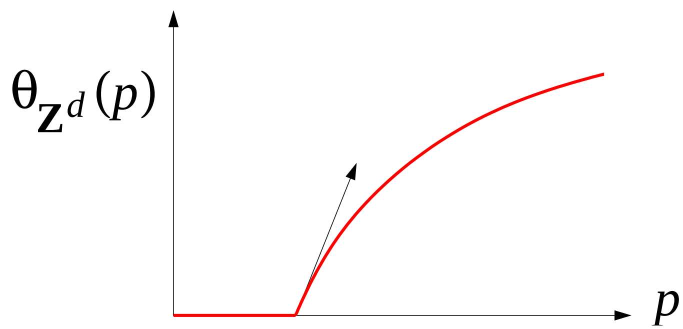

The transition into the supercritical world is too smooth, not enough driving force into exceptional behavior.

Theorem (Schramm & Steif 2005).

- There are exceptional times (a.s.) for critical site percolation on the triangular lattice.
- They have Hausdorff dimension in $[1/6, 31/36]$.

Just like before, upper bound 31/36 comes from $\beta = 5/36$ (conformal invariance and SLE$_6$ critical exponents, Smirnov, Schramm, Lawler, Werner 2000-01 $+$ Kesten's scaling relation 1987)

Theorem (Garban, P & Schramm 2008).

- There are exceptional times also on $\mathbb{Z}^2$.
- On the triangular grid they have Hausdorff dimension $31/36$.
- On the triangular grid, there are exceptional times with an infinite white **and** an infinite black cluster simultaneously. $(1/9 \leqslant \text{dim} \leqslant 2/3)$

Lower bounds on the exceptional set

(A) Likeliness of static event: one-arm probability $\alpha_1(R) := \mathbf{P}[\,0 \longleftrightarrow R\,]$ ($= R^{-5/48+o(1)}$ on Δ, by Lawler-Schramm-Werner 2001)

(B) Decorrelation of crossing events:

A site (or bond) is pivotal in ω, if flipping it changes the existence of a left-right crossing. Equivalent to an alternating four-arm event.

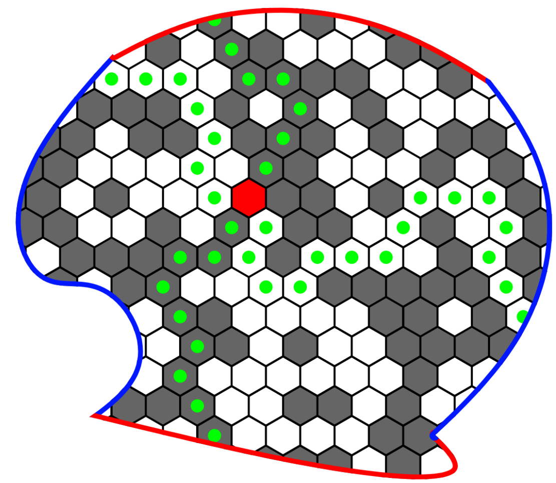

In a piecewise smooth conformal rectangle (a quad) $\mathcal{Q}$ and lattice of mesh $1/n$,

$$\mathbf{E}|\mathrm{Piv}_{\mathcal{Q},n}| \asymp_{\mathcal{Q}} n^2\,\alpha_4(n)$$

($= n^{3/4+o(1)}$ on Δ, by Smirnov-Werner 2001)

Dynamics certainly needs long enough time to flip pivotals in order to start making changes. But it is not obvious that hitting many pivotals is enough for full decorrelation. This needs discrete Fourier analysis, as started by Benjamini-Kalai-Schramm 1998.

Two natural questions on the exceptional set

How do exceptional infinite clusters look like? The first one? A typical one?

There is an "infinite critical cluster" in the static world, Kesten's Incipient Infinite Cluster measure (1986): for $H \subset \Delta$ and ω^H configuration in H, the limit $\mathsf{IIC}(\omega^H) = \lim_{R\to\infty} \mathbf{P}[\,\omega^H \,|\, 0 \leftrightarrow R\,]$ exists. Many other natural definitions give the same measure (Járai 2003).

Two natural questions on the exceptional set

How do exceptional infinite clusters look like? The first one? A typical one?

There is an "infinite critical cluster" in the static world, Kesten's Incipient Infinite Cluster measure (1986): for $H \subset \Delta$ and ω^H configuration in H, the limit $\mathsf{IIC}(\omega^H) = \lim_{R\to\infty} \mathbf{P}[\,\omega^H \,|\, 0 \leftrightarrow R\,]$ exists. Manyl other natural definitions give the same measure (Járai 2003).

What is the hitting time tail $\mathbf{P}\big[\,\mathscr{E} \cap [0,t] = \emptyset\,\big]$?

To answer the first question, we needed to answer the second one:

Theorem (Hammond, Mossel & P. 2011). The hitting time tail is exponentially small.

Theorem (Hammond, P & Schramm 2012). Configuration at a "typical" exceptional time has the law of IIC, but the First Exceptional Time Infinite Cluster (FETIC) is thinner.

Proof of exponential tail for First Exc Time

Dynamical percolation in B_R is just continuous time random walk on the hypercube $\{0,1\}^{B_R}$, with rate 1 clocks on the edges. On $\{0,1\}^n$, discrete time random walk has spectral gap $1/n$, but in continuous time, the gap is uniformly positive, so could try to use the Ajtai-Komlós-Szemerédi lemma.

However, $\mathbf{P}[\,0 \longleftrightarrow R\,]$ is tiny, so we don't want to hit that set. But $\mathbf{P}[\,\mathscr{E}_R \cap [0,1] \neq \emptyset\,]$ is uniformly positive!

So, first idea: Markov chain $\{\omega[2t, 2t+1] : t = 0, 1, 2\}$ on a huge state space. This again has a uniform spectral gap. However, it's not reversible!

Instead, use dynamic $\to$ static projection, $g(\omega) := \mathbf{E}[\, f(\omega[0,1]) \mid \omega_0 = \omega\,]$ inside a suitable proof of AKSz '87.

Exponential lower bound: $\mathbf{P}\big[\,\mathscr{E} \cap [0,t] = \emptyset\,\big] \geqslant \mathbf{P}[\,\mathscr{E} \cap [0,1] = \emptyset\,]^t$, by dynamical FKG inequality. Esoteric proof by dynamics on dynamical percolation configurations, a second time coordinate... But proved earlier by Liggett using infinitesimal generators.

Local time measure for exceptional times

Want measure on exceptional set $\mathscr{E}$. Simplest possible idea:

$$\overline{M}_r(\omega_s) := \frac{\mathbf{1}\{0 \leftrightarrow r\}}{\mathbf{P}[\,0 \leftrightarrow r\,]}, \quad \overline{\mu}_r[a,b] := \int_a^b \overline{M}_r(\omega_s)\,ds, \quad \overline{\mu}[a,b] := \lim_{r\to\infty} \overline{\mu}_r[a,b].$$

Although $\overline{M}_r(\omega)$ is a martingale, $\overline{\mu}_r[a,b]$ is NOT, and we couldn't prove convergence. A bit more complicated idea:

$$M_H(\omega) := \lim_{R\to\infty} \frac{\mathbf{P}[\,0 \leftrightarrow R \,|\, \omega^H\,]}{\mathbf{P}[\,0 \leftrightarrow R\,]} = \lim_{R\to\infty} \frac{\mathbf{P}[\,\omega^H \,|\, 0 \leftrightarrow R\,]}{\mathbf{P}[\,\omega^H\,]} = \frac{\mathsf{IIC}(\omega^H)}{\mathbf{P}[\,\omega^H\,]}.$$

$$M_r(\omega_s) := M_{B_r}(\omega_s), \quad \mu_r[a,b] := \int_a^b M_r(\omega_s)\,ds, \quad \mu[a,b] := \lim_{r\to\infty} \mu_r[a,b].$$

Now $\mu_r[a,b]$ is a MG. The limit $\mu[a,b]$ exists a.s. and in L^2, supported inside $\mathscr{E}$.

Theorem (Hammond, P & Schramm 2012). At a μ-typical time (quenched or annealed), the configuration has the distribution of IIC.

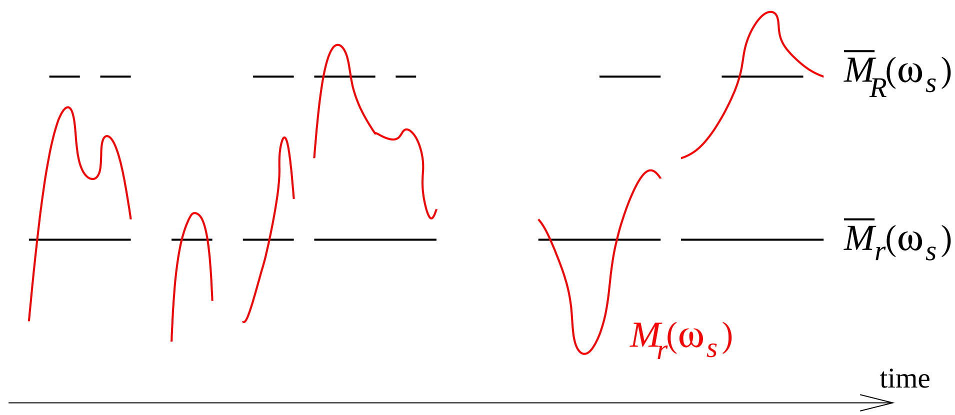

The two sequences are closely related; if L^2-limit $\overline{\mu}$ exists, it is μ.

Question: Does the L^2-limit $\overline{\mu}$ exist?

Question: Is $\operatorname{supp}(\mu) = \mathscr{E}$? Is μ the $31/36$-dimensional Minkowski content of $\mathscr{E}$?

FETIC versus IIC

Mutual singularity should hold, but we only show that there is some ω^{B_r} such that $\lim_{R\to\infty} \mathsf{FETIC}_R(\omega^{B_r}) \neq \lim_{R\to\infty} \mathsf{IIC}_R(\omega^{B_r})$.

The configuration at a typical switch time for $\{0 \longleftrightarrow R\}$ is size-biased by the number of pivotals. Because of the many pivotals far from the origin, inside B_r this bias becomes negligible as $R \to \infty$, so we still have IIC.

The configuration at FET_R is further size-biased by the length of the non-connection interval ending at the switch time.

FETIC versus IIC

Mutual singularity should hold, but we only show that there is some ω^{B_r} such that $\lim_{R\to\infty} \mathsf{FETIC}_R(\omega^{B_r}) \neq \lim_{R\to\infty} \mathsf{IIC}_R(\omega^{B_r})$.

The configuration at a typical switch time for $\{0 \longleftrightarrow R\}$ is size-biased by the number of pivotals. Because of the many pivotals far from the origin, inside B_r this bias becomes negligible as $R \to \infty$, so we still have IIC.

The configuration at FET_R is further size-biased by the length of the non-connection interval ending at the switch time.

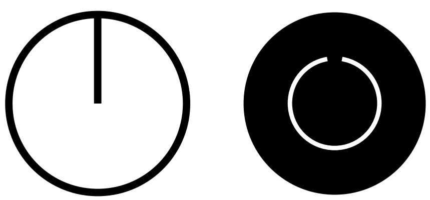

Reversing time, expected reconnection time from THIN_R is much larger than from FULL_R. Their probabilities under IIC_R are the same, hence under FETIC_R are different. But, these configurations are not visible in the limit.

So, need thinning only in bounded neighbourhood of origin. But then: typical reconnection time is tiny for large R, so will difference in reconnection time be detectable?

FETIC versus IIC

Mutual singularity should hold, but we only show that there is some ω^{B_r} such that $\lim_{R\to\infty}$ FETIC$_R(\omega^{B_r}) \neq \lim_{R\to\infty}$ IIC$_R(\omega^{B_r})$.

The configuration at a typical switch time for $\{0 \longleftrightarrow R\}$ is size-biased by the number of pivotals. Because of the many pivotals far from the origin, inside B_r this bias becomes negligible as $R \to \infty$, so we still have IIC.

The configuration at FET$_R$ is further size-biased by the length of the non-connection interval ending at the switch time.

For any $\omega = \omega^{B_R}$ satisfying $\{0 \longleftrightarrow R\}$, get THIN$_r(\omega)$ by thinning inside B_r.

Want to show that the reconnection time $T = T_{r,R}$ started from THIN$_r(\omega^{B_R})$ is larger in expectation than $N = N_{r,R}$, the one started from the normal ω^{B_R}, uniformly as $R \to \infty$. (While both are very small.)

Because of the thinning, there is some $\epsilon(r) \to 0$ and $g(r) \to \infty$ with

$$\mathbf{P}\big[\, T > g(r) \;\big|\; T > \epsilon(r) \,\big] \;>\; c_1 \,. \tag{1}$$

Also, from stochastic domination, $\mathbf{P}[\, T > \epsilon(r) \,] \;\geqslant\; \mathbf{P}[\, N > \epsilon(r) \,]\,.$ (2)

(1) would be hard, so our thinning is different, hence (2) doesn't quite hold.

Write $X^\epsilon = X\,\mathbf{1}_{\{X > \epsilon(r)\}}$. Note that size-biased $\widehat{N}$ times Unif$[0,1]$ is FET.

A size-biasing lemma: $\mathbf{P}\big[\, \widehat{N} > \epsilon(r) \,\big] = \frac{\mathbf{E}[\, N^\epsilon \,]}{\mathbf{E}[\, N \,]} > c_2$ and $\mathbf{E}\big[\, \widehat{N} \,\big] < C_1$ imply

$$\mathbf{E}\big[\, N \;\big|\; N > \epsilon(r) \,\big] \;<\; C_2 \,. \tag{3}$$

From these three,

$$\mathbf{E}[\, T^\epsilon \,] \;\geqslant\; c_1\, g(r)\, \mathbf{P}[\, N > \epsilon(r) \,] \;\geqslant\; c_1\, g(r)\, \frac{\mathbf{E}[\, N^\epsilon \,]}{C_2} \,,$$

hence

$$\mathbf{E}T \;\geqslant\; \mathbf{E}[\, T^\epsilon \,] \;\gg_r\; \mathbf{E}[\, N^\epsilon \,] \;\geqslant\; c_2\, \mathbf{E}N \,.$$

What about the tail of left-right connection?

GPS '08: $\mathbf{E}\big[\, f_{\mathcal{Q},n}(\omega_0)\, f_{\mathcal{Q},n}(\omega_{tn^{-3/4+o(1)}})\,\big] - \mathbf{E}[\, f_{\mathcal{Q},n}\,]^2 \asymp_{\mathcal{Q}} t^{-2/3}$ as $t \to \infty$, uniformly in n, hence natural to rescale time like this.

In fact, there exists a scaling limit of dynamical percolation [Garban, P. & Schramm 2013], so one can either talk about the rescaled finite chains, "uniformly in n", or about the scaling limit process.

Another theorem of HMP gives $\mathbf{P}\big[\, f_{\mathcal{Q}}(\omega_s) = 1 \ \forall\, s \in [0,t]\,\big] \leqslant t^{-2/3+o(1)}$.

In fact, by cutting $\mathcal{Q}$ vertically into L slabs: $\leqslant t^{-2L/3+o(1)}$ for any L, superpolynomial decay.

Exponential lower bound is easy from dynamical FKG inequality.

Conjecture. $\mathbf{P}\big[\, f_{\mathcal{Q}}(\omega_s) = 1$ for all $s \in [0,t]\,\big] = \exp(-t^{2/3+o(1)})$.

Supported by a *very* non-rigorous renormalization argument. Especially the upper bound is questionable.