Lecture 1

Introduction and main results

The minicourse to be given over four 50-minute lectures will focus on extremal properties of random walk local time. This turns out to be a particular aspect of the larger area of logarithmically correlated processes that has attracted a lot of attention in recent years. For lack of time, we will focus only on one particular result; namely, the scaling limit of the points avoided by a two-dimensional simple random walk. The main objective of the course is to motivate the students to learn other, and often more difficult, results through self-study of papers and existing review articles.

1.1 Random walk local time.

Throughout we will consider a continuous time Markov chain X on a finite state space that in general will take the form $V \cup \{\varrho\}$, where V is a finite set and ϱ is a distinguished vertex (not belonging to V). The transitions will occur at "constant speed," which means that the chain takes the form $X_t = Z_{N(t)}$, where Z is a discrete-time Markov chain and $\{N(t)\colon t \geqslant 0\}$ is a rate-1 Poisson point process independent of Z. We assume that X (and Z) is irreducible and reversible with respect to measure π and write P^x for the law of the chain started at x, with associated expectation denoted as E^x.

Of our prime interest in these lectures is the *local time* associated with X. This is the two-parameter stochastic process

$$\big\{\ell_t(x)\colon x \in V \cup \{\varrho\},\, t \geqslant 0\big\} \tag{1.1}$$

defined by

$$\ell_t(x) := \frac{1}{\pi(x)} \int_0^t 1_{\{X_s = x\}} \mathrm{d}s \tag{1.2}$$

(Recall that irreducibility forces $\pi(x) > 0$ for all $x \in V \cup \{\varrho\}$.) To explain the normalization, note that X_t will for large t be distributed according to suitably-normalized π which means that the time spent at x grows proportionally to $\pi(x)$.

The main question of interest in these lectures is then:

$$\text{What does } \ell_t \text{ look like at large } t? \tag{1.3}$$

While this is our general objective, we focus on particular questions. For example, we may ask about the size and asymptotic law of

$$\max_{x \in V \cup \{\varrho\}} \ell_t(x) \quad \text{and} \quad \min_{x \in V \cup \{\varrho\}} \ell_t(x) \tag{1.4}$$

Noting that the minimum is zero until the first time all vertices are visited naturally leads us to the notion of the *cover time*,

$$\tau_{\text{cov}} := \inf\Big\{ t \geqslant 0 \colon \min_{x \in V \cup \{\varrho\}} \ell_t(x) > 0 \Big\} \tag{1.5}$$

of which we can then ask how it scales with t and the size of V. Another natural question (which is the one we will focus in these notes) concerns the structure of the set of points not yet visited by X at time t; namely,

$$\mathcal{A}(t) := \big\{ x \in V \cup \{\varrho\} \colon \ell_t(x) = 0 \big\} \tag{1.6}$$

that we will refer to as *avoided points*.

Of course, taking t to be large without changing V will hardly lead to interesting conclusions. We will also treat only one particular class of Markov chains; namely, that arising from the simple symmetric random walk (SRW) on $\mathbb{Z}^d$. So, unless we discuss general aspects of the theory where the setting introduced above is more appropriate, we take V to be a scaled-up and discretized version $D_N \subseteq \mathbb{Z}^d$ of a "nice" continuum domain $D \subseteq \mathbb{R}^d$; i.e., the set (roughly) of the form

$$\big\{ x \in \mathbb{Z}^d \colon x/N \in D \big\} \tag{1.7}$$

or (in some references to prior work) the lattice torus $(\mathbb{Z}/N\mathbb{Z})^d$. The point is now to study the local time for the simple random walk in D_N at times t_N such that $t_N \to \infty$, subject to specific restrictions on growth rate with N as $N \to \infty$.

A technical point for D_N of the form (1.7) is how to interpret the "random walk" at the vertices of D_N that, in $\mathbb{Z}^d$, would have an edge to the complement of D_N. One possibility is to treat this as a *free boundary condition* which means to ignore jumps leading out of D_N. For reasons that will become clear later we use a different "return mechanism," corresponding to the *wired boundary condition*, which is defined as follows: Collapse all the vertices in $\mathbb{Z}^d \setminus D_N$ to one *boundary vertex* ϱ. Then route the edges emanating out of D_N to ϱ. This leads to a domain as in Fig. 1.

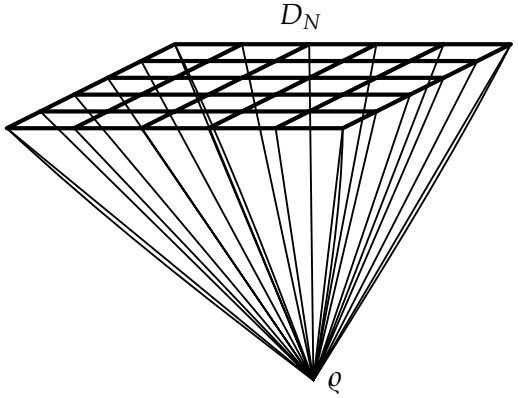

Fig. 1: An illustration of the state space for the random walk. Here D_N is simply an $N \times N$ square while ϱ is a vertex to which all the boundary edges of D_N in $\mathbb{Z}^2$ are re-routed.

Note that the invariant measure $\pi(x)$ equals the degree of the vertex x in the resulting graph which for D_N with wired boundary condition will be equal $2d$ at $x \in D_N$ and equal

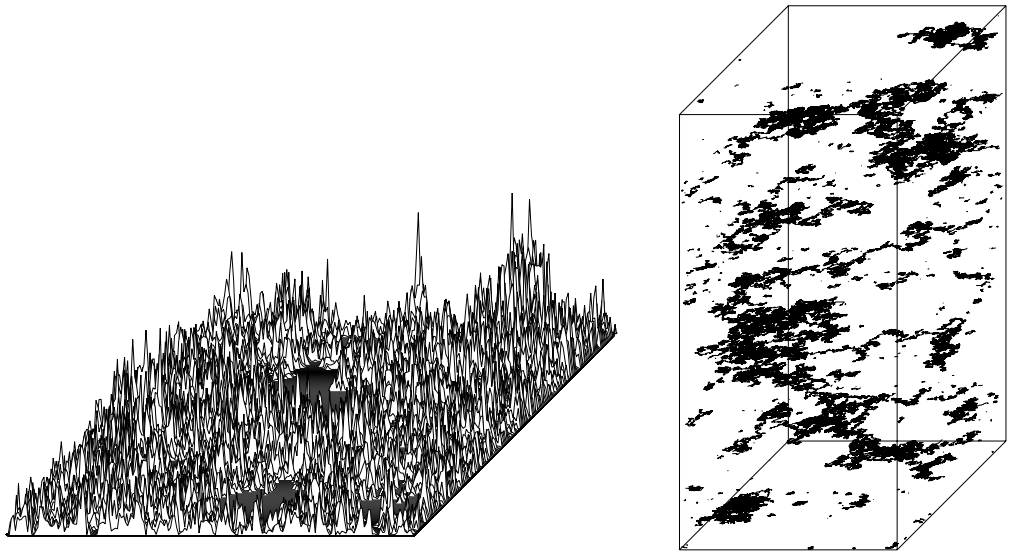

Fig. 2: A sample of the local time (left) and the trajectory of the walk (right) over time (with time axis running upwards) for the random walk on $D_N \cup \{\varrho\}$ as in Fig. 1 with $N = 200$.

to the size of the edge boundary of D_N in $\mathbb{Z}^d$ at ϱ. The chain X is then a constant-speed continuous-time random walk on the resulting finite graph that runs just as the simple random walk on D_N and, after each exit, returns back to D_N through a randomly-chosen boundary edge. (For readers worried that this might lead to the local time building up near the boundary, Fig. 2 and our Theorem 1.7 show that this is not the case.)

1.2 The case for $d = 2$.

In these lectures we will focus on the above setting in spatial dimension $d = 2$. To motivate this, let us recount some of the basic facts about random walk on $\mathbb{Z}^d$. The key difference arises already in the celebrated Pólya theorem that says

$$\text{SRW on } \mathbb{Z}^d \text{ is } \begin{cases} \text{transient}, & \text{in } d \geqslant 3 \\ \text{recurrent}, & \text{in } d = 1, 2 \end{cases} \tag{1.8}$$

The transience can be thought of as a "short memory" (or decaying-autocorrelation) property that very often makes a number of arguments easier to handle.

The discrepancy between the recurrent and transient regime typically manifests itself in the analytic form of the conclusions. To give an example, note that for the cover time of the lattice torus $(\mathbb{Z}/N\mathbb{Z})^d$ the following holds:

$$\tau_{\text{cov}} \asymp \begin{cases} N^d \log N, & \text{in } d \geqslant 3 \\ N^2 (\log N)^2, & \text{in } d = 2 \\ N^2, & \text{in } d = 1 \end{cases} \tag{1.9}$$

where the first two lines are true as sharp asymptotics (with a known constant of proportionality) because the cover time concentrates strongly around its expected value. This

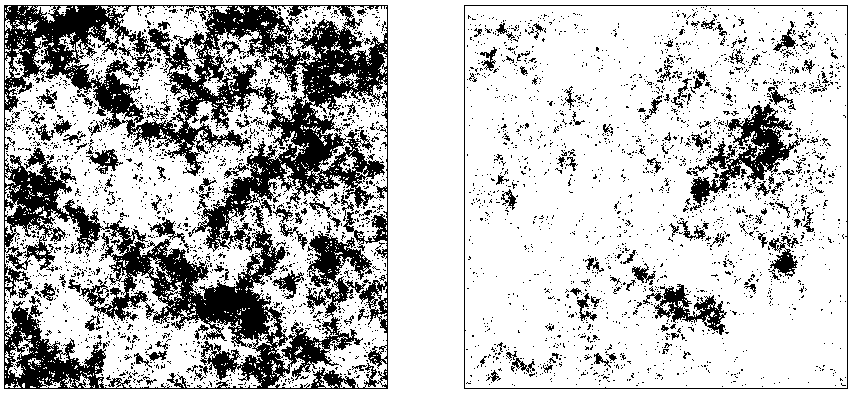

Fig. 3: Samples of the set of avoided points for the random walk on D_N run for time proportional to $\theta = 0.1$ (left) and $\theta = 0.3$ (right) fraction of the expected cover time. (The same run of the random walk is used for both figures.)

concentration fails in $d = 1$ as τ_{cov}/N^2 tends in law (as $N \to \infty$) to a non-degenerate random variable.

The formula $N^d \log N$ in $d \geqslant 3$ is easy to understand: One needs N^d time to visit most of the vertices but then a coupon-collector reasoning need to be employed to sweep out the outliers. A similar reasoning can be used in $d = 2$; the extra $\log N$ appears because once a vertex is hit, it will be visited order $\log N$ times before it is left for good. (Much more is known in fact; thanks to a result of D. Belius [13], we know a full limit law of suitably centered τ_{cov} in all $d \geqslant 3$. In $d = 2$, the corresponding asymptotic is the subject of active research by several groups.)

Moving to the set of avoided points, the natural time scales to consider are those proportional to the cover time. So we will specialize (1.6) to the form

$$\mathcal{A}_N(\theta) := \{x \in D_N \colon \ell_{\theta E^o \tau_{\text{cov}}}(x) = 0\} \tag{1.10}$$

We now ask about the asymptotic properties of this set, specifically, the number of points and the way they are distributed in D_N, in the limit as $N \to \infty$.

Also in this problem the recurrence/transience dichotomy manifests itself strongly in the conclusions. Indeed, in $d \geqslant 3$ the set $\mathcal{A}_N(\theta)$ partitions into a collection of independent small "islands" which, thanks to a result of J. Miller and P. Sousi [52] from 2017 can even be nailed to the form

$$\mathcal{A}_N(\theta) \overset{\text{a.s.}}{\approx} \text{Bernoulli}(N^{-\theta d}) \tag{1.11}$$

for a non-trivial interval of $\theta \in [0, 1]$. Here the squiggly equality represents a coupling in total variational distance to the set where the Bernoulli process equals 1.

In contrast to this, in $d = 2$, the set $\mathcal{A}_N(\theta)$ scales to a *random fractal*, as shown in Fig. 3. The point of these lectures is to make sense of a limit of these pictures as $N \to \infty$.

1.3 Link to Gaussian Free Field.

Our method to control the local time will rely on a close connection between the local time of a Markov chain and a Gaussian process called *Gaussian Free Field*. We will now introduce this concept in the general setting of Markov chains on $V \cup \{\varrho\}$ introduced above. The connection itself will be discussed in Lecture 2.

We start with some standard definitions. For any $x \in V \cup \{\varrho\}$ we introduce the *first hitting time* of x by X as

$$H_x := \inf\{t \geqslant 0 \colon X_t = x\} \tag{1.12}$$

Notice that $H_x = 0$ P^x-a.s. We then use this to define the *Green function*

$$G^V \colon (V \cup \{\varrho\}) \times (V \cup \{\varrho\}) \to [0, \infty) \tag{1.13}$$

by the formula

$$G^V(x,y) := E^x\big(\ell_{H_\varrho}(y)\big) = \frac{1}{\pi(y)} \int_0^\infty P^x\big(X_t = y,\, H_\varrho > t\big)\mathrm{d}t \tag{1.14}$$

where the second expression is based on writing $\ell_{H_\varrho}(y) = \frac{1}{\pi(y)} \int_0^{H_\varrho} 1_{\{X_t = t\}}\mathrm{d}t$ and applying Tonelli's theorem. We now pose our first exercise:

Exercise 1.1 *Show that, if viewed as a matrix, G^V is symmetric and positive semidefinite.*

The reason why we highlight these properties is that they make G^V a covariance. This is enough to make sense of:

Definition 1.2 *The Discrete Gaussian Free Field (DGFF) on V is a Gaussian process*

$$\{h_x^V \colon x \in V \cup \{\varrho\}\} \tag{1.15}$$

such that

$$\forall x, y \in V \cup \{\varrho\} \colon \quad \mathbb{E}h_x^V = 0 \quad \wedge \quad \mathbb{E}(h_x^V h_y^V) = G^V(x,y) \tag{1.16}$$

(We will use $\mathbb{P}$ and $\mathbb{E}$ for probability and expectation associated with the DGFF.)

Note that the definitions ensure that $G^V(\varrho, y) = 0 = G^V(x, \varrho)$ for any x and y. This along with $\mathbb{E}h_\varrho = 0$ forces

$$h_\varrho^V = 0 \quad \mathbb{P}\text{-a.s.} \tag{1.17}$$

which also explains the special role the "boundary vertex" ϱ plays in the whole setup. In particular, our h^V corresponds to the case of Dirichlet boundary conditions.

The reason for calling this the "Discrete" GFF is to make a distinction between the corresponding concept in the continuum, called the "Continuum" GFF with the shorthand CGFF. While the latter is not a prime target of interest in our notes, we will make some references to it when we discuss the scaling limits in Lectures 3 and 4.

When we specialize ourselves to the random walk on $V := D_N$ and take the resulting DGFF h^{D_N} at face value, we are naturally led to ask a number of questions about its extremal properties similar as those for the local time asked above. For instance:

$$\text{What is the growth rate/scaling limit of } \max_{x \in D_N} h_x^{D_N}? \tag{1.18}$$

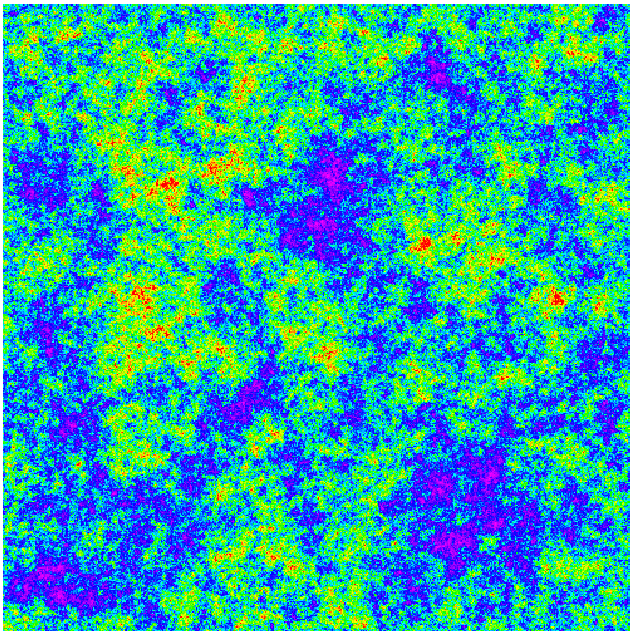

Fig. 4: A sample of DGFF on 500×500-square color coded so that the red regions are those with large positive values and purple regions are those with large negative values. The values in-between are coded according to usual ordering of colors by the wave-length.

(Since the field is symmetric, the minimum scales as the negative of the maximum.) Another question to ask is:

$$\text{What is the cardinality/scaling limit of } \left\{ x \in D_N : h_x^{D_N} \geqslant \lambda \max_{x \in D_N} h_x^{D_N} \right\}? \tag{1.19}$$

Here, for $\lambda \in (0,1)$, we call the points in the set the *λ-thick points* of h^{D_N}.

Similarly as for the local time, the DGFF samples in spatial dimensions $d \geqslant 3$ are not nearly so interesting as in spatial dimension two. Indeed, in $d = 2$ the field itself is a random fractal which we demonstrate in Fig. 4.

Looking at the figure, the reader will surely notice the fractal curves separating the mostly-green and mostly-blue regions; these are known to be described by the SLE_4-curves thanks to a celebrated work by O. Schramm and S. Sheffield [55]. Our interest in the present text are the yellow-to-red regions, where the field is unusually large.

A key problem in making any of the above mathematically reasonable is the fact that the DGFF in $d = 2$ becomes increasingly singular as the side of the underlying domain increases. This has to do with the fact that the field is *logarithmically correlated*; we will elaborate on what this means in Lecture 2.

1.4 Main result on DGFF.

We are now ready to make precise statements of the main results to be discussed in detail throughout the rest of the course. We start with those for the DGFF. First we identify the continuum regions D to which our results apply:

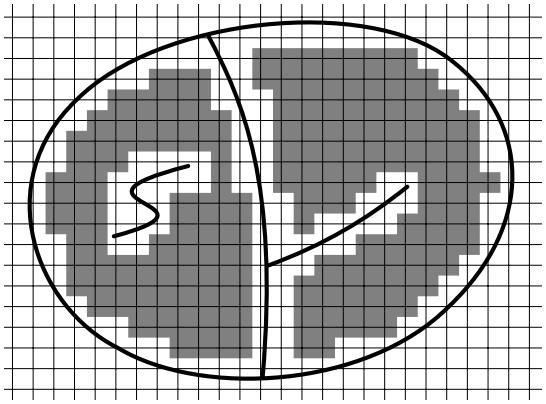

Fig. 5: An illustration of an admissible approximation D_N (marked by the lattice side in the dark region) of an admissible domain $D \subseteq \mathbb{R}^2$ (bounded by the thick lines).

Definition 1.3 (Admissible domain) *A set $D \subseteq \mathbb{R}^2$ is an admissible domain if it is bounded, open and ∂D has a finite number of connected components each of which has a positive Euclidean diameter.*

Since a bounded simply connected open subset of $\mathbb{R}^2$ has a connected boundary, it follows that any such set is an admissible domain by the above definition. However, we also allow for non-trivial arcs in the interior. While connectedness is not required, the fact that our processes will trivially factor over connected components means that we only need to work with connected admissible D.

Next we will specify more precisely the way we allow ourselves to discretize D. While (1.7) seems to be a canonical choice, the problem is that this choice may result in a discrete set that "looks" quite different than D itself; particularly, from the perspective of harmonic analysis. Writing d_∞ for the infinity distance on $\mathbb{Z}^2$, we instead use:

Definition 1.4 (Admissible approximations) *A sequence $\{D_N\}_{N\geqslant 1}$ of non-empty subsets of $\mathbb{Z}^2$ is an admissible approximation of an admissible domain $D \subseteq \mathbb{R}^d$ if*

$$\forall N \geqslant 1\colon \quad D_N \subseteq \big\{x \in \mathbb{Z}^d\colon d_\infty(x/N, \mathbb{R}^2 \smallsetminus D) > 1/N\big\} \tag{1.20}$$

and, for all $\delta > 0$ there exists $N_0 \geqslant 1$ such that

$$\forall N \geqslant N_0\colon \quad D_N \supseteq \big\{x \in \mathbb{Z}^d\colon d_\infty(x/N, \mathbb{R}^2 \smallsetminus D) > \delta\big\} \tag{1.21}$$

To illustrate this on an example, Fig. 5 shows an admissible lattice approximation of an admissible domain. In Lecture 2 we will give the reasons why definitions need to be set up this way.

In order to set the scales, next we note that for x deep inside D_N we will have

$$G^{D_N}(x, x) = g \log N + O(1) \tag{1.22}$$

where the constant of proportionality equals

$$g := \frac{1}{2\pi} \tag{1.23}$$

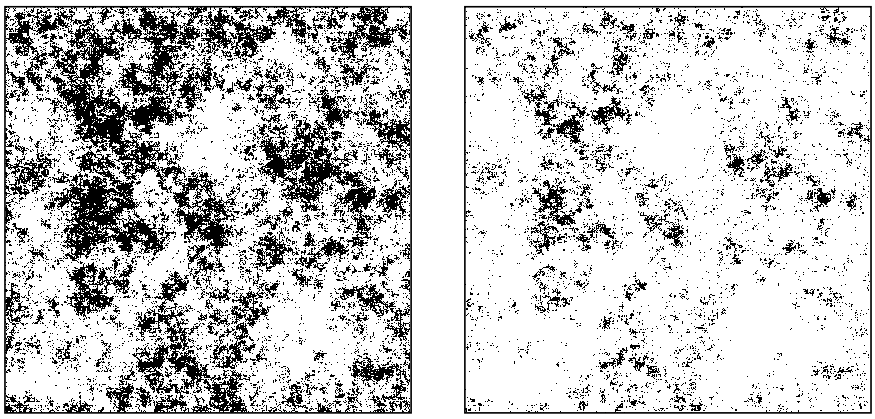

Fig. 6: The level sets at $\lambda := 0.1$ (left) and $\lambda = 0.3$ (right) multiple of the expected maximum of DGFF on a square box of side-length 500..

For the maximum of the DGFF we then get the asymptotic

$$\max_{x\in D_N} h_x^{D_N} = 2\sqrt{g}\,\log N + O(\log\log N) \tag{1.24}$$

Interestingly, the same asymptotic applies even for i.i.d. normals with variance $g\log N$; however, while the leading order is the same in the two cases, the constant multiplying $\log\log N$ is already different.

We now define the set of λ-thick points again as

$$\mathcal{T}_N(\lambda) := \bigl\{x\in D_N\colon h_x^{D_N} \geqslant 2\sqrt{\lambda}g\,\log N\bigr\},\quad \lambda\in(0,1). \tag{1.25}$$

As noted earlier, this set is expected to look like a random fractal which readily confirmed by simulation, see Fig. 6.

A question to address next is in what sense we can take a scaling limit of the pictures in Fig. 6. The natural choice of Hausdorff distance is out because, after scaling by $1/N$ and taking $N\to\infty$, these sets are everywhere dense and so have vanishing asymptotic distance to domain D itself. We will therefore use a different approach: We associate with each set a point process that records both the position and the value of the field and then take the limit of this process itself.

For instance, for the set (1.25) the point process would take the form

$$\sum_{x\in D_N} \delta_{x/N}\otimes\delta_{h_x^{D_N}-2\sqrt{g}\,\lambda\log N} \tag{1.26}$$

where the tensor product of delta-measures is just a convenient way to write a delta-measure at the corresponding two-coordinate quantity. The point is that, while having total mass of D_N, this measure does give us access to the cardinality of $\mathcal{T}_N(\lambda)$ by integrating it against the function $f(x,h) := 1_{[0,\infty)}(h)$. Of course, (1.26) by itself would not allow for a reasonable limit as $N\to\infty$ as one still needs to normalize the measure in such a way that a limit in law is possible. This is what we do in:

Theorem 1.5 (B.-Louidor [20]) *Let $\{D_N\}_{N\geqslant 1}$ be admissible approximations of an admissible domain $D \subseteq \mathbb{R}^2$. There exists a family of a.s-finite random Borel measures $\{Z^D_\lambda \colon \lambda \in (0,1)\}$ on D such that for any positive sequence $\{a_N\}_{n\geqslant 1}$ with*

$$\lambda := \lim_{N\to\infty} \frac{a_N}{2\sqrt{g}\log N} \in (0,1) \tag{1.27}$$

and with

$$K_N := \frac{N^2}{\sqrt{\log N}} \mathrm{e}^{-\frac{a_N^2}{2g\log N}} \tag{1.28}$$

we have

$$\frac{1}{K_N} \sum_{x\in D_N} \delta_{x/N} \otimes \delta_{h^{D_N}_x - a_N} \xrightarrow[N\to\infty]{\text{law}} Z^D_\lambda(\mathrm{d}x) \otimes \mathrm{e}^{-\alpha\lambda h}\mathrm{d}h \tag{1.29}$$

where $\alpha := 2/\sqrt{g}$ and the weak convergence is relative to vague topology on $\overline{D}\times\mathbb{R}$. Moreover, a.e. sample of Z^D_λ charges every non-empty open set and is diffuse (i.e., has no atoms).

A few remarks are on order: First, the statement (1.29) means that integrating the measure on the left with respect to any number of continuous functions $\overline{D}\times\mathbb{R}\to\mathbb{R}$ with compact support results in a family of random variables that converge jointly in law to the corresponding integrals with respect to the measure on the right-hand side. As it turns out, due to linearity of the integral and the Cramèr-Wold device, it suffices to check convergence for integrals of just one function at the time.

Next observe that (1.27–1.28) give

$$K_N = N^{2(1-\lambda^2)+o(1)}, \quad N\to\infty. \tag{1.30}$$

Although the total mass of the measure in (1.27) diverges proportionally to $N^{2\lambda^2+o(1)}$, the mode of convergence ensures that the mass it puts on any compact subset of $\overline{D}\times\mathbb{R}$ is tight and admits a distributional limit.

Third, as a corollary of Theorem 1.5 we get a limit law for the total size of the level set of λ-thick points:

Corollary 1.6 *For the setting of Theorem 1.5,*

$$\frac{1}{K_N}\#\bigl\{x\in D_N\colon h^{D_N}_x \geqslant a_N\bigr\} \xrightarrow[N\to\infty]{\text{law}} (\alpha\lambda)^{-1} Z^D_\lambda(D) \tag{1.31}$$

This generalizes a result of Daviaud [30] from 2006 who obtained the leading order growth of the level set without identifying the subleading terms and/or a limit law. The above results will be discussed and, for $\lambda < 1/\sqrt{2}$, proved in Lecture 3.

1.5 Main result on the local time.

Moving to our results on the local time, consider the random walk on $D_N\cup\{\varrho\}$ as described above. We will for simplicity focus only on the avoided points at times proportional to the cover time of $D_N\cup\{\varrho\}$. As it turns out, the easiest format to state this is under a different time parametrization. For each $t\geqslant 0$, let

$$\tau_\varrho(t) := \inf\bigl\{s\geqslant 0\colon \ell_s(\varrho)\geqslant t\bigr\} \tag{1.32}$$

Then set

$$L_t(x) := \ell_{\tau_\varrho(t)}(x), \quad x \in D_N \cup \{\varrho\} \tag{1.33}$$

Since $L_t(\varrho) = t$ a.s., this is the parametrization of the local time by the local time at ϱ.

As it turns out $E^\varrho(L_t(x)) = t$ for each $x \in D_N$, so (1.9) tells us that the cover time with happen on time scales such that $t \asymp (\log N)^2$. This motivates the parametrization (1.34) in our second main result to be discussed in these lectures:

Theorem 1.7 (Abe-B. [3]) *Suppose $\{t_N\}_{N\geqslant 1}$ is a positive sequence such that*

$$\theta := \lim_{N\to\infty} \frac{t_N}{2g(\log N)^2} > 0 \tag{1.34}$$

Then setting

$$\widehat{K}_N := N^2 \mathrm{e}^{-\frac{t_N}{g\log N}} \tag{1.35}$$

for $\theta \in (0,1)$ we have

$$\frac{1}{\widehat{K}_N} \sum_{x\in D_N} 1_{\{L_{t_N}(x)=0\}} \delta_{x/N} \xrightarrow[N\to\infty]{\text{law}} Z^D_{\sqrt{\theta}} \tag{1.36}$$

where the random measures $\{Z^D_\lambda \colon \lambda \in (0,1)\}$ are as in Theorem 1.5. The limit (and, for large enough N, the sequence of measures on the left) vanishes when $\theta > 1$.

Observe that from (1.34–1.35) we get

$$\widehat{K}_N = N^{2(1-\theta)+o(1)}, \quad N \to \infty \tag{1.37}$$

This decays to zero when $\theta > 1$ and so $\theta = 1$ corresponds to the scaling of the cover time. (Note that we make no claims in this case as that requires going beyond the leading order asymptotic; see Appendix of these notes for more discussion.)

The punchline of Theorem 1.7 is that, in the parametrization by the local time at ϱ, the set of avoided points at θ-multiple of the cover time is asymptotically distributed *exactly* as the $\sqrt{\theta}$-thick points of the DGFF. We note that, even though we are looking at the points where the local time vanishes, the time-parametrization matters. Indeed, a follow-up joint work with Y. Abe and S. Lee [4] shows that the limit measure is different in the natural time parametrization. See Appendix for a precise statement.

As for the DGFF, we get information about the cardinality of the avoided set:

Corollary 1.8 *For the setting of Theorem 1.7,*

$$\frac{1}{\widehat{K}_N} \#\bigl\{x \in D_N \colon L_{t_N}(x) = 0\bigr\} \xrightarrow[N\to\infty]{\text{law}} Z^D_{\sqrt{\theta}}(D) \tag{1.38}$$

Theorem 1.7 will be extracted from Theorem 1.5 and the connection of the local time and the DGFF discussed in Lecture 2. The actual proof will be given in Lecture 4.