



universität
wien

DISSERTATION / DOCTORAL THESIS

Titel der Dissertation /Title of the Doctoral Thesis

On spin systems, height functions and random walks

verfasst von / submitted by

Diederik Gerrit Pieter van Engelenburg, MSc BSc

angestrebter akademischer Grad / in partial fulfilment of the requirements for the degree of
Doktor der Naturwissenschaften (Dr.rer.nat)

Wien, 2023 / Vienna 2023

Studienkennzahl lt. Studienblatt /
degree programme code as it appears on the
student record sheet:

A 796 605 405

Dissertationsgebiet lt. Studienblatt /
field of study as it appears on the student
record sheet:

Mathematik

Betreut von / Supervisor:

Univ.-Prof. Nathanaël Berestycki, PhD

Abstract

The goal of statistical mechanics is to understand the macroscopic behavior of large bodies of interacting particles, starting from microscopic descriptions. As the energy of the system changes these models often undergo a phase transition where the qualitative macroscopic behavior suddenly changes. Both mathematicians and physicists are interested in understanding the behavior at and near the critical point. In this work, we will focus on two aspects of the theory: understanding the graphs on which the models live on the one hand, and continuous spin models on the other hand. It turns out that the geometry of the space and some properties of the models are connected; an instance of *universality*.

Mathematically perhaps the most tractable model is the Uniform Spanning Tree (UST), which is intrinsically related to the loop erased random walk. In the first part of the thesis, we will show a connection between the geometry of certain (random) graphs and this UST. In particular, we prove that for recurrent, reversible graphs, the following conditions are equivalent: (a) existence and uniqueness of the potential kernel, (b) existence and uniqueness of the harmonic measure from infinity, (c) a new anchored Harnack inequality, and (d) one-endedness of the uniform spanning tree. These results are obtained from combinatorial properties of the graph, and most hold in wide generality.

It was conjectured by Aldous and Lyons that, for unimodular random graphs, the wired uniform spanning tree is always one-ended, unless the graph is trivial (is itself two-ended). This was proved in the transient case by Hutchcroft. However, the trivial graphs are always recurrent and the techniques for the transient case do not generalize to the recurrent case. Using the connection to potential kernels and the harmonic measure from infinity, we continue the thesis by showing that the conjecture of Aldous and Lyons holds.

In the second part of the thesis, we focus on continuous abelian spin models and their dual height functions. We will summarize in a relatively elementary and general formalism a result of Sheffield which states that all ergodic Gibbs measures for height functions are extremal. We will then introduce a special loop representation for the XY model, which is reminiscent of the random current representation of the Ising model.

The loop representation connects the planar height function model and the dual spin model together. Using recent results on the so-called delocalization of integer-valued height functions on trivalent planar lattices and the loop representation, we give a new proof of the famous Berezinskii-Kosterlitz-Thouless transition in the XY model.

Finally, we revisit the classical Fourier duality between integer-valued height functions and their dual abelian spin systems. We introduce some new methods to derive general results, including: a universal upper bound on the variance of the height function in terms of the Green's function (a GFF bound), monotonicity of this variance with respect to a natural temperature parameter, delocalization for planar graphs, and the occurrence of a BKT transition in planar spin models.

Zusammenfassung

Das Ziel der statistischen Mechanik ist es, ausgehend von mikroskopischen Beschreibungen das makroskopische Verhalten großer, aus wechselwirkenden Teilchen bestehenden Körper zu verstehen. Wenn sich die Energie des Systems ändert, durchlaufen diese Modelle oft einen Phasenübergang, bei dem sich das qualitative makroskopische Verhalten plötzlich ändert. Sowohl Mathematiker als auch Physiker sind daran interessiert, das Verhalten am und in der Nähe des kritischen Punktes zu verstehen. In dieser Arbeit werden wir uns auf zwei Aspekte der Theorie konzentrieren: einerseits auf das Verständnis der Graphen, auf denen die Modelle leben, und andererseits auf kontinuierliche Spinmodelle. Es stellt sich heraus, dass die Geometrie des Raums und einige Eigenschaften der Modelle miteinander verbunden sind; ein Beispiel für *Universalität*.

Das mathematisch vielleicht am besten nachvollziehbare Modell ist der Uniform Spanning Tree (UST), der eng mit dem Gaußschen freien Feld verbunden ist. Im ersten Teil der Arbeit zeigen wir eine Verbindung zwischen der Geometrie bestimmter (zufälliger) Graphen und diesem UST. Insbesondere beweisen wir, dass für rekurrente, reversible Graphen die folgenden Bedingungen äquivalent sind: (a) Existenz und Einzigartigkeit des Potentialkerns, (b) Existenz und Einzigartigkeit des harmonischen Maßes aus dem Unendlichen, (c) eine neue verankerte Harnack-Ungleichung, und (d) Einseitigkeit des einheitlichen Spannbaums. Diese Ergebnisse ergeben sich aus kombinatorischen Eigenschaften des Graphen, und die meisten gelten in großer Allgemeinheit.

Aldous und Lyons stellten die Vermutung auf, dass für unimodulare Zufallsgraphen der verdrahtete einheitliche Spannbaum immer einseitig ist, es sei denn, der Graph ist trivial (d. h. er hat zwei Enden). Dies wurde für den transienten Fall von Hutchcroft bewiesen. Die trivialen Graphen sind jedoch immer rekurrent und die Techniken für den transienten Fall lassen sich nicht auf den rekurrenten Fall verallgemeinern. Mit Hilfe der Verbindung zu potentiellen Kernen und dem harmonischen Maß aus dem Unendlichen setzen wir die Arbeit fort, indem wir zeigen, dass die Vermutung von Aldous und Lyons zutrifft.

Im zweiten Teil der Arbeit konzentrieren wir uns auf kontinuierliche abelsche Spinmodelle und ihre dualen Höhenfunktionen. Wir werden in einem relativ elementaren und allgemeinen Formalismus ein Ergebnis von Sheffield zusammenfassen, das besagt,

dass alle ergodischen Gibbsmaße für Höhenfunktionen extremal sind. Wir werden dann eine spezielle Schleifendarstellung für das XY-Modell einführen, die an die Zufallsstromdarstellung des Ising-Modells erinnert. Die Schleifendarstellung verbindet das planare Höhenfunktionsmodell und das duale Spinmodell miteinander. Unter Verwendung neuerer Ergebnisse über die so genannte Delokalisierung von ganzzahligen Höhenfunktionen auf dreiwertigen planaren Gittern und der Schleifendarstellung geben wir einen neuen Beweis für den berühmten Berezinskii-Kosterlitz-Thouless-Übergang im XY-Modell.

Schließlich greifen wir die klassische Fourier-Dualität zwischen ganzzahligen Höhenfunktionen und ihren dualen abelschen Spinsystemen wieder auf. Wir führen einige neue Methoden ein, um neue Ergebnisse in hoher Allgemeinheit abzuleiten, darunter: eine universelle obere Schranke für die Varianz der Höhenfunktion in Bezug auf die Greensche Funktion (eine GFF-Schranke), Monotonie dieser Varianz in Bezug auf einen natürlichen Temperaturparameter, Delokalisierung für planare Graphen, und das Auftreten des BKT-Übergangs in planaren Spinmodellen.

Natuur is voor tevredenen of legen.
En dan: wat is natuur nog in dit land?
Een stukje bos, ter grootte van een krant,
Een heuvel met wat villaatjes ertegen.
Geef mij de grauwe, stedelijke wegen,
De' in kaden vastgeklonken waterkant,
De wolken, nooit zo schoon dan als ze, omrand
Door zolderramen, langs de lucht bewegen.
Alles is veel voor wie niet veel verwacht.
Het leven houdt zijn wonderen verborgen
Tot het ze, opeens, toont in hun hogen staat.
Dit heb ik bij mijzelfen overdacht,
Verregend, op een miezerigen morgen,
Domweg gelukkig, in de Dapperstraat.

J.C. Bloem

Acknowledgments

First and foremost, I would like to thank my supervisor Nathanaël Berestycki for giving me the opportunity to come to Vienna and continue my explorations in the area of statistical mechanics, for the help with my projects, and for connecting me to my co-authors and many other people that I have since met. Going to Vienna would not have been possible without the work of Rob during my masters, and the kind help of Wendelin to suggest Vienna in the first place. Enjoying, or finishing perhaps, would not have been possible without the guidance of my unofficial co-supervisor Marcin: I am grateful for exploring together the to us unknown field of spin systems, for introducing me to the world of height functions, and for all your patience. I thank my office mates and fellow students Antoine and Lucas, the climbing crew, my colleagues.

Staying in the eagerly self-proclaimed cultural capital of the (sort of) German-speaking world, would not have been so pleasant without meeting the people I met. Of course, that includes you, Fabian, my lovely goofball, the reason why Vienna has become such a special place. Sascha, thanks for being always so positive, and for organizing so many get-togethers. Zaida, voor je eindeloze energie, lekker klagen met z'n twee, en voor het bijhouden van mijn Nederlands. Mam en pap, Cas en Mirte, voor jullie onuitputtelijke liefde.

Contents

Abstract	i
Zusammenfassung	iii
Acknowledgments	vii
1 Introduction	1
1.1 Geometry of uniform spanning trees	3
1.1.1 Potential theory and random walks	5
1.1.2 Vertex-transitive graphs	7
1.1.3 Coarse geometry	8
1.1.4 Ends in the spanning forest: vertex-transitive graphs	10
1.1.5 Unimodular random rooted graphs	12
1.1.6 Further relations between uniform spanning trees and geometry .	14
1.2 Phase transitions in spin systems	15
1.2.1 Topological phase transition	17
1.3 Fourier–Pontryagin transform and duality	20
1.3.1 BKT-transition through duality	22
1.3.2 The BKT transition revisited	26
1.4 Articles and preprints appearing in this thesis	27
I Geometry and the Uniform Spanning Tree	29
2 Harnack inequality and one-endedness of UST on reversible random graphs	31
2.1 Introduction	31
2.1.1 Some applications	34
2.2 Background and notation	37
2.3 Equivalence between (a) and (b)	39
2.3.1 Base case of equivalence	39

2.3.2	Triangle inequality for the potential kernel	42
2.3.3	Gluing and harmonic measure	46
2.4	Proof of Theorem 2.3	52
2.5	The sublevel set of the potential kernel	54
2.6	Two Harnack inequalities	59
2.6.1	The standing assumptions	59
2.6.2	Elliptic Harnack Inequality	59
2.6.3	(a) implies (c): anchored Harnack inequality	62
2.6.4	(c) implies (a)	66
2.7	Random Walk conditioned to not hit the root	67
2.7.1	Definition and first estimates	68
2.7.2	Hitting probabilities for conditioned walk	72
2.7.3	Harnack inequality for conditioned walk	74
2.7.4	Recurrence of sets	75
2.7.5	Infinite intersection of two conditioned walks	77
2.8	(a) implies (d): One-endedness of the uniform spanning tree	78
2.9	Harmonic measure from infinity on mated-CRT maps	84
2.9.1	Preliminaries and known results.	86
2.9.2	Proof of Theorem 2.49.	87
3	Resolving the conjecture of Aldous and Lyons	89
3.1	Introduction	89
3.2	Boundary theory of recurrent graphs	92
3.2.1	Harmonic measures from infinity	92
3.2.2	Potential kernels and Doob transforms	94
3.2.3	An affine isomorphism	98
3.2.4	The Liouville property for extremal Doob transforms	101
3.2.5	Potential kernels and the uniform spanning tree	102
3.3	Proof of the main theorem	105
3.3.1	Reversible and unimodular graphs	105
3.3.2	The effective resistance is linear on the spine	106
3.3.3	Completing the proof	110
II	Continuous spin models and dual height functions	113
4	Review of height function models	115
4.1	Integer-valued height functions	115
4.2	Ergodicity and extremality	116

4.2.1	The main result: ergodic measures are extremal	117
4.2.2	In case of the Gaussian free field	117
4.2.3	A percolation of sign clusters	118
4.2.4	Switching the clusters	119
4.2.5	Finalizing the proof of Theorem 4.1	121
4.3	Delocalization of height functions	121
4.3.1	Extremal measures	122
4.3.2	The integer-valued Gaussian free field	123
4.3.3	Extension to general height functions	124
5	The BKT-transition in the XY model	129
5.1	Introduction and main result	129
5.2	The dual height function	132
5.2.1	Gibbs measures and delocalization	133
5.2.2	Absolute-value-FKG and dichotomy	134
5.3	Loop representation of currents and path reversal	136
5.3.1	Coupling with the height function	138
5.3.2	The two point-function and path reversal	138
5.4	Delocalization implies no exponential decay	141
5.5	Existence of phase transition in the XY model	146
5.6	No exponential decay implies a power-law lower bound	148
6	Correlation Inequalities	151
6.1	The two point-function and path switching	154
6.2	Application to some inequalities	155
6.3	Limitations of the colored loop representation	158
7	Duality between height functions and spin models	159
7.1	Introduction	159
7.2	General duality	162
7.2.1	Discrete calculus	162
7.2.2	Spin models and random height functions	163
7.3	Upper bound on the variance of the height function	166
7.4	Delocalization implies the BKT phase transition	169
7.4.1	Background	169
7.4.2	Notions of delocalization	170
7.4.3	Setup	171
7.4.4	The implication	172
7.4.5	The case of the XY-model	173

CONTENTS

7.4.6	An equivalence	175
7.5	Monotonicity of variance of the height function	175
7.5.1	Setup	175
7.5.2	Increasing variance	176
7.5.3	Delocalization of roughly planar height function models	177
7.6	GFF covariance for a projection of the spin model	182
7.7	A central limit theorem	183
A	Uniqueness of the potential kernel implies one-endedness of the UST, without finite expected degree	185
B	Duality between height functions and spin models	191
B.1	Proof of the Ginibre inequality	195
B.2	Reflection positivity.	196
B.3	Positive definite functions	197

Introduction

A major problem in physics is to describe the large-scale behavior of a system of particles, starting from a purely local description. Suppose we want to describe the transition from a liquid to a gas, which occurs suddenly when the temperature passes through a critical point. There are two immediate problems. First, there are many particles, so an enormous amount of observables would be needed (for each particle, position and momentum in Newtonian mechanics). Second, there is apparent chaos: even tiny differences between the observations now can lead to huge differences in the future.

The field of thermodynamics was developed for this very purpose: to understand the macroscopic behavior of models, based on physically measurable observables. Instead of tracking each microscopic state and trying to describe the whole system, it is postulated that only a much smaller number of parameters is needed to describe a macroscopic state. We will only be concerned with models that are in *equilibrium*, meaning that the macroscopic observables (of interest) do not change over time.

A special branch of thermodynamics to do this is statistical mechanics, which came to life here in Vienna due to Boltzmann [39], and was later formalized by Gibbs [79]. The idea is to consider a probability distribution over all possible (microscopic) states in which a given system can be. For example, in a pure gas, there are essentially three quantities that are determined at macroscopic level: the volume of the vessel, the number of particles and the (internal) energy. Given these quantities, the goal would be to find a suitable probability distribution over all microscopic states, so that the frequencies in real observations correspond to the distribution. Thus, the chosen probabilities describe the microscopic behavior and provide the desired connection to the large-scale behavior. Moreover, such probabilities can only depend on the macroscopic observables.

In general, we will only work with models where the particles have a fixed position on a crystal in space, but their value can change. A typical example is the ferromagnet: this is a crystal in which every atom has a magnetic charge. The Ising model, a famous simplified model of magnetism introduced by Lenz [98, 124], assigns a positive (+1) or negative (−1) charge to each “atom” of the crystal. Often, it is assumed that only

neighboring atoms interact. In particular, we assign to each microstate σ (which assigns a value of $+1, -1$ to each atom of the crystal) an energy cost $H(\sigma)$, depending only on the underlying lattice. The fact that only neighbors interact is formalized by the assumption that the potential H has the form

$$H(\sigma) = \sum_{x \text{ neighbor } y} F(\sigma_x, \sigma_y).$$

Each microstate σ is drawn with probability proportional to the *Gibbs weight* $e^{-\frac{1}{T}H(\sigma)}$, where T is a fixed temperature.¹

Such a probability measure thus depends on the underlying crystal (which we will often take to be a lattice), the total number of microstates and the temperature. Thus, we have found a suitable framework for studying transitions, at least of magnets. Clearly, the probability weights change continuously with time, but on a large scale a phase transition must occur (if our model describes magnets): at high temperature, the weights for each microstate are essentially equal, so there is no magnetism. However, when the temperature is set to 0, only those microstates that minimize the energy cost H will have any weight, in which case either each atom has positive or each atom has negative charge.

In this thesis, we will be interested not only in the question: how do small changes in the Gibbs weights affect the macroscopic observables? But also in the question: how does the geometry of the underlying graph (crystal) relate to the geometry of the models on top of it?

In the first part of this thesis, we will present a model in which this questions can be answered at least to some extent. Here, we will be interested in the so-called “uniform spanning tree”, which is defined in a way where the potential is highly non-local². In this case, the temperature is completely irrelevant, and there is no phase transition. Rather, the model turns out to be “critical by definition”, at least in some sense. The relationships between random walks, potential theory and spanning trees briefly introduced below, allow to find deep relations between the geometry of the graph and the model.

In the second part, we focus on a completely different model, namely a spin model where each site of a lattice (regular graph, crystal) has with an angle in $[0, 2\pi]$ (think of a compass at every site), with a potential H that penalizes for large angle differences between neighboring sites. Here, the temperature *is* relevant. We will again see

¹In statistical mechanics, this is called the canonical ensemble: in this case the energy of the system is not actually constant, but rather the system interacts with a “heat bath” with which it can exchange energy so that the temperature remains constant.

²Alternatively, it is the exception in which we study the so-called “microcanonical ensemble” (which is defined as the uniform measure over *energy*-minimizing microscopic states, in the setting where the energy is constant).

interesting relationships between the underlying geometry of the lattice, and of certain macroscopic observables. It will come as no surprise to the experienced reader that the two models considered here are related, at least to some extent.

1.1. Geometry of uniform spanning trees

Given a finite graph $G = (V, E)$, a spanning tree is defined as a connected subgraph T of G with the same vertex set, but without any cycles. The uniform spanning tree is then readily defined as the uniform measure on spanning trees. A first natural question, from a combinatorial perspective, is: how many spanning trees does a given graph have? In other words: what is the partition function of the model? Kirchhoff observed already in the nineteenth century that this question is related to a seemingly totally different problem: the eigenvalues of the Laplacian.

Theorem 1.1 (Matrix–tree theorem [105]). *The number of spanning trees of G equals $\frac{1}{n} \prod_{i=2}^n \lambda_i$, where the product is taken over the non-zero eigenvalues of the graph Laplacian Δ .*

This work is the starting point of the fascinating theory connecting spanning trees and potential theory on finite graphs. In this work, we will not focus on these questions, but instead move on to infinite graphs.

For infinite graphs, the notion of uniform spanning tree is not as easily defined. To define the “uniform” spanning *forest* on an infinite graph G , Pemantle [143] proposed to exhaust G by finite subgraphs and take weak limits for appropriately chosen boundary conditions. For two natural choices of such boundary conditions, the *free* and *wired* boundary conditions, he proved that these limits are well-defined and do not depend on the choice of exhaustion. They *do* depend on the choice of boundary condition. The free and wired boundary conditions turn out to be the (only) *extremal* ones.

It is worth noting that the limiting measure will always be supported on spanning subgraphs of G which contain no cycles. However, connectivity is not a local condition and hence it is not at all clear if the limit is supported on trees. For this reason, we generally call the limits free and wired *uniform spanning forest* (FUSF, WUSF). Given this observation, two questions instantly present themselves:

- (1) When do the free and wired uniform spanning forest agree?
- (2) When is the (wired) spanning forest connected (and thus a tree)?

In the setting of the hypercubic lattice \mathbb{Z}^d , Pemantle [143] provided an answer to both questions: he showed that the free and wired limits agree and that the spanning forest is connected if and only if $d \leq 4$. By now, the answers to both questions are understood in quite wide generality as we will explain below.




Figure 1.1: Left: depiction of the topology of some manifold. Right: a finite piece of the manifold is cut out, leaving 4 infinite, connected components.

There is also a slightly different perspective on the uniform spanning forest. It can be seen as a critical model from statistical mechanics, with at least two justifications for this viewpoint. The first is that certain observables of the spanning forests behave much like critical models and the second is through the relation with the (discrete) Gaussian free field.

From this perspective, there is a third natural question related to the spanning forests. Given a vertex x , there is always one path to infinity from x in the tree of the forest containing x , which never comes back to x , but does the removal of x split the tree into multiple infinite components? See also Figure 1.1.

This question is the analogue of the existence of an infinite cluster for percolation at the critical point and as such is interesting to understand. Of course, the question is also natural if the goal is to understand the topology of the spanning forests.

In general, we will say that an infinite graph $G = (V, E)$ has at least k ends, if there exists finite set of vertices $B \subset V$, such that after the removal of B , there are k disjoint infinite components in G . In this light, we say that G is k -ended if it is at least k ended, but not $k + 1$ ended, see again Figure 1.1. Using this terminology, the question above can be rephrased to: how many ends does a component of the uniform spanning tree have?

It turns out that answering this problem is harder than understanding the connectivity of the forest. From a statistical mechanics perspective, this is also what we can expect: there are few cases where absence of percolation at the critical point is known. In case of the wired spanning forest, the direct connection to random walks on the one hand, and to potential theory and the Gaussian free field on the other, provide a relatively deep understanding of the model.

Before we turn to the partial answers of the question about ends, let us briefly explain these connections.

1.1.1 Potential theory and random walks

It is well known that harmonic functions and random walks are related, in particular to potential functions. Here, by “potential” we mean a function which is harmonic on all of the graph, outside of two points where it has value 0 and 1, corresponding to the classical interpretation of an electric current flowing from 0 to 1. The latter connection to currents on graphs and networks was developed first by Kirchoff [105].

To highlight a few details: a *unit flow* from a vertex x to a vertex y on a graph is an anti-symmetric function on the directed edges $\theta : \vec{E} \rightarrow \mathbb{R}$, which satisfies that the in- and outgoing flow at any vertex not x or y must be zero (hence it defines a kind of “transport map”). At x and y the flow will have a source and sink respectively; net outgoing at x the value is 1, net outgoing at y is -1 . As such, we can define the *effective resistance* between x and y in terms of a variational formula:

$$\mathcal{R}_{\text{eff}}(x \leftrightarrow y) := \inf\{\mathcal{E}(\theta) : \theta \text{ a unit flow from } x \text{ to } y\},$$

where \mathcal{E} denotes the $l^2(\vec{E})$ -energy. The flow minimizing the right-hand side is obtained by taking the gradient of the potential, scaled by the effective resistance. This variational notion is known as *Thompson’s principle*. The fact that the potential is the explicit minimizer, implies that the effective resistance has a purely probabilistic interpretation too:

$$\mathcal{R}_{\text{eff}}(x \leftrightarrow y) = \frac{\mathbf{G}_y(x, x)}{\deg(x)},$$

where $\mathbf{G}_y(\cdot, \cdot)$ denotes the *Green function* of the random walk killed at y and $\deg(x)$ is the degree of x in the graph.

As with spanning trees, we can define the effective resistance for infinite graphs using finite exhaustions, but again, care must be taken in the choice of boundary conditions. As before, there are two natural boundaries, the free and wired, giving rise to different effective resistances in general. This time, the relation with potentials and harmonic functions helps to answer the question: when are the resistances the same?

The question can be answered using relatively basic Hodge decompositions in the l^2 -case, as explained in [32] and [131, Chapter 9]. A 1-form on a graph is an anti-symmetric function on the oriented edges. It is called *exact* if it is the gradient of a function on the vertices, and *co-closed* if it is divergence free. The Hodge decomposition tells that the $l^2(\vec{E})$ space of 1-forms has the following orthogonal decomposition:

$$l^2(\vec{E}) = \{\text{exact}\} \oplus \{\text{co-closed}\} \oplus \{dh : h \text{ harmonic and } \mathcal{E}(dh) < \infty\},$$

where dh denotes the gradient of h .

For an oriented edge xy , write $\chi^{xy} := \mathbb{1}_{xy} - \mathbb{1}_{yx}$ for the most basic 1-form (which is also a flow from x to y). On the one hand, it is known that the *wired* effective resistance between the endpoints of two vertices equals the projection of χ onto the space of co-closed forms. On the other hand, the free effective resistance equals the projection of χ on the complement of the exact forms. From this, it follows that the free and wired effective resistance are different precisely when there are non-constant harmonic functions with finite Dirichlet energy.

At this point, it is worth mentioning another well known relation between the uniform spanning tree and potential theory due to Kirchoff [105]. He showed that the probability a given edge e is in the uniform spanning tree, is equal to the effective resistance between the endpoints of the edge. We deduce that if the free and wired effective resistance agree on an infinite graph, then the probability that a given edge is in the spanning forest does not depend on the boundary chosen. Using the extension of Kirchoff's formula to multiple edges by Burton and Pemantle [44], we can answer the question about equality of the free and wired spanning forests:

Theorem 1.2 ([32]). *The free and wired uniform spanning tree are equal if and only if there are no non-constant harmonic functions with finite Dirichlet energy.*

A consequence of this result is that the free and wired spanning forests are the same for all amenable graphs, and hence for all groups of polynomial growth.

Since potential theory is related to random walks on the one hand, and to the uniform spanning tree on the other, it seems reasonable to ask if there is a more direct relation between random walks and spanning trees. It was noted by Aldous and Broder separately [10, 41] that there is an exact way to sample the spanning tree using random walks. An even more insightful sampling algorithm was later discovered by Wilson [171].

Let us describe the latter algorithm for finite graphs $G = (V, E)$. Fix some enumeration of the vertex set $(v_0, \dots, v_{|V|-1})$ and define inductively a sequence of subgraphs $(E_i)_{i \geq 0}$ as follows. Set $E_0 = \{v_0\}$. Given E_{i-1} , run a simple random walk started from v_i , stop it when it hits E_{i-1} (which could be instantaneous), and erase chronologically the loops on the random walk path from v_i to E_{i-1} . Define E_i to be the union of E_{i-1} and this loop-erased path. Call $T = E_n$.

Theorem 1.3 (Wilson [171]). *The tree T is a uniform spanning tree of G .*

It is easy to see that the algorithm generates a spanning tree of G , but that it is distributed as the uniform spanning tree is quite miraculous. In fact, it is a priori even far from obvious that this definition does not depend on the chosen enumeration of V .

Wilson's algorithm extends to recurrent graphs without problem, but was extended to transient graphs in [32]. In this case, the algorithm works more or less the same, the only essential difference is that we take E_0 the chronological loop erasure of a random

walk started from v_0 . This is possible because the random walk is transient. Moreover, there may be positive probability that started from v_i , the random walk never hits E_{i-1} . In this case, the loop erasure of the whole path from v_i towards infinity is added to E_i and the final forest automatically has two disjoint components.

Of course, since uniform spanning forests on transient graphs may depend on the boundary conditions, there is no hope that Wilson's algorithm can be used to generate both the free and the wired forests. However, it is not too hard to see that the algorithm described above corresponds to the wired uniform spanning forest. This is essentially due to the following fact: if we use x as a root for Wilson's algorithm and start a random walk from y far away from x , then the random walk will typically hit the boundary vertex before touching x and the branch from y to x typically goes through the boundary vertex.

Using this extension of Wilson's algorithm, Benjamini, Lyons, Peres and Schramm [32] managed to resolve the question concerning connectivity:

Theorem 1.4 ([32], [130]). *The wired uniform spanning forest of a graph is connected if and only if the traces of two independent random walks have infinitely many intersections with probability one.*

1.1.2 Vertex-transitive graphs

In the remainder of this thesis, vertex-transitive graphs and extensions thereof play a central role. A graph G is said to be *vertex-transitive* if its automorphism group acts transitively on its vertices. In other words, the graph "looks the same" from every vertex. Examples include complete graphs, the tori $(\mathbb{Z}/N\mathbb{Z})^d$ and the hypercubic lattices \mathbb{Z}^d . Another important class of examples are Cayley graphs: for a finitely generated group \mathbb{G} with (symmetric) and finite generating set S , the corresponding Cayley graph has as its vertex set $V = \mathbb{G}$, with edges between $x, y \in V$ if there is some $s \in S$ with $xs = y$.

A finitely generated group is said to have *weakly polynomial growth* if any of its Cayley graphs satisfy that there is some d for which

$$\liminf_{R \rightarrow \infty} \frac{|B_x(R)|}{R^d} = 0$$

where $B_x(R)$ is the R -ball in the Cayley graph. Notice that this fact is independent of the choice of generating set, which justifies the definition in terms of groups.

It turns out that (infinite) groups with weakly polynomial growth can be characterized quite generally. This was part of Gromov's famous program on the classification of finitely generated groups. A powerful result in this context is the following.

Theorem 1.5 (Gromov [84]). *Let G be a finitely generated group of weakly polynomial growth. Then G is virtually nilpotent.*

We will not explain exactly what virtual nilpotence means, but it strongly restricts the coarse geometry of a group. An elegant and somewhat probabilistic proof of Gromov’s theorem is due to Kleiner [106] (see also [156] for a finitary version and [174] for nice lecture notes). A particularly striking consequence is that a group of weak polynomial growth has polynomial growth with an integer exponent:

$$|B_x(R)| \sim R^d$$

for some $d \in \mathbb{N}$. Another consequence is that any group of weakly polynomial growth is either roughly isometric to \mathbb{Z} , or contains \mathbb{Z}^2 as a subgroup.

Of course, the theorem applies only to Cayley graphs, so what about vertex-transitive graphs? There is a nice result by Trofimov [165] which says that every vertex-transitive graph of polynomial growth is “roughly” equal to a Cayley graph. Thus, many results from Cayley graphs extend to the vertex-transitive case. It is worth mentioning that Trofimov’s theorem does not hold for general vertex-transitive graphs: Diestel and Leader [50] provided an example of a vertex-transitive graph and conjectured it was not “roughly” a Cayley graph. Their conjecture was confirmed by [66].

1.1.3 Coarse geometry

Since effective resistance has a geometric interpretation on the one hand, and a probabilistic interpretation in terms of random walks on the other, it is natural to ask how the geometry of the underlying graph relates to the geometry of the uniform spanning forest. The two results above about equality of the free and wired forest, and the connectivity of the latter are examples of such relations.

A general strategy to prove a property or model depends only on the *coarse geometry* of a graph, is to use rough isometries. Without including a precise definition, two metric spaces (X_1, d_1) and (X_2, d_2) are roughly isometric if there is a map between the two spaces which is “coarsely Lipschitz” and “coarsely surjective”, see e.g. [131, Chapter 2]. This can be applied in the context of groups to say that a property of the group defined in terms of its Cayley graph, does not depend on the choice of generating set. In general, rough isometries can provide a strategy for proving that a graph G_1 has some property by finding a better understood graph (or metric space) G_2 roughly isometric to G_1 and by proving G_2 must have the desired property.

Percolation. Let us make a small detour to percolation to give an example of a somewhat related model for which some properties depend only on the coarse geometry of the graph.

Fix a graph $G = (V, E)$ and a parameter $p \in [0, 1]$. Take $x \in V$ fixed. Generate the random subgraph $\omega \subset G$ by including each edge of G in ω with probability p ,

independently for each edge. Write $\theta(p)$ for the probability that the component of x in ω is infinite. It is relatively straightforward to convince oneself that $\theta(p)$ is increasing in p . The critical value p_c is defined as the supremum over all p for which $\theta(p) = 0$. Without going into the beautiful background on percolation, let us mention only two questions concerning the model: when is $0 < p_c < 1$ and what is $\theta(p_c)$?

As mentioned before, the latter question is similar in spirit to the question: how many ends does a component of the WUSF have? Notice also that $\theta(p_c) = 0$ implies that $p_c < 1$, simply because for $p = 1$, the random graph satisfies $\omega = G$ a.s. As we will soon see, for the WUSF, the question on the number of ends can be solved at least in some generality. However, the question of $\theta(p_c) = 0$ is one of the main open problems in probability theory, even in the case where the underlying graph is \mathbb{Z}^3 with nearest neighbor interactions. For the hypercubic lattices \mathbb{Z}^d in the special case of $d = 2$, it was proved by Kesten that $\theta(p_c) = 0$ ([103]), and for $d \geq 11$ it was proved in [67, 89]. However, both these cases rely on methods that cannot work in three dimensions: $d = 2$ relies on planarity, whereas $d \geq 11$ relies on the “lace expansion”.

The question whether or not $p_c < 1$, on the other hand, is by now better understood. Not only is it known that $p_c < 1$ for all Cayley graphs that are not trivial, it is even understood that there is a gap for the possible values of p_c :

Theorem 1.6 ([139]). *There exists some $\epsilon > 0$ such that for all Cayley graphs of superlinear growth, $p_c \leq 1 - \epsilon$.*

The proof of this, and many earlier results relies on a version of Gromov’s theorem 1.5 above. We mention only the strategy to show that $p_c < 1$ for all Cayley graphs of polynomial growth (not roughly \mathbb{Z}). By Gromov’s theorem, any Cayley graph with polynomial growth is either roughly isometric to \mathbb{Z} or has a subgroup isomorphic to \mathbb{Z}^2 . Moreover, it is known that $p_c < 1$ is stable under rough isometries (at least when the graphs involved have bounded degrees) [131]. Thus, to show that $p_c < 1$ for any Cayley graph which has at most polynomial growth, it suffices to prove that $p_c(\mathbb{Z}^2) < 1$. Using planar duality, a counting argument (in the spirit of Peierls’ [141]) gives the desired bound.

This approach at least sheds some hope on answering the question of ends in the wired uniform spanning forest: perhaps it is invariant under rough isometries? Sadly, the answer is no: Hutchcroft [92] recently provided a counter example. However, we will see later that there are other properties which are rough isometry invariant, and which do provide a method to solve the ends question for certain Cayley graphs.

1.1.4 Ends in the spanning forest: vertex-transitive graphs

Let us go back to the question: how many ends does a component of the wired uniform spanning have? It may be tempting to guess that this depends only on the coarse geometry, but as mentioned above, this intuition is false. Nonetheless, to understand what happens for general vertex-transitive graphs, it is worthwhile to first consider the hypercubic lattices $\mathbb{Z}^d, d \geq 1$.

Pemantle [143] showed in this case that the components of the uniform spanning forest are almost surely one-ended if $d \geq 2$. The case $d = 1$ is trivial: the spanning tree equals the whole graph, which is itself two-ended. Thus, the transition between one and two-endedness of the graph happens between dimension 1 and 2, as in the setting of percolation. Heuristically, we could expect that also for general vertex-transitive graphs, the transition happens between dimension 1 and 2. Here, by dimension we mean volume growth exponent.

This heuristic was made rigorous in the fundamental work of Benjamini, Lyons, Peres and Schramm [32] on uniform spanning forests in the context of vertex-transitive graphs. We break the result into two parts: the transient and the recurrent case.

Theorem 1.7 (Benjamini et al. [32]). *Let G be a transient, vertex transitive graph. Each component of the wired uniform spanning forest is one-ended almost surely.*

Note here that the results by Pemantle do not provide an intuition for graphs which do not have polynomial growth, and the result of [32] holds for the *wired* forest. Indeed, it is quite easy to see that it cannot be true for the *free* forest. Consider a 4-regular tree. The free spanning forest must equal the whole tree almost surely, hence is a single connected component which has infinitely many ends.

Let us roughly outline how to obtain Theorem 1.7. A proof due to Hutchcroft [91] (which works in a more general setting) splits into two parts. His biggest contribution is to rule out that the wired spanning forest has more than one two-ended component, which uses a so called “cycle breaking algorithm”. Without providing the details, this is a local update algorithm which allows to say that if there would be two (or more) components with two ends, then there would be one component with three ends, but the latter is impossible for vertex-transitive graphs ([9, Theorems 6.3 and 7.1]). We will outline the second step, which consists in ruling out that there is exactly one component with two ends. A similar argument was already present in [32, Theorem 10.3].

Proof sketch. The second step is to rule out a single component with two ends, which will be done by contradiction. Recall that if a tree has two ends, there exists a (unique) bi-infinite simple path on the tree. We call such a bi-infinite path the *spine* of the tree, see also Figure 1.2. Since there is only one component with two ends, there is a unique spine




Figure 1.2: Left: the spine of a two-ended tree is depicted in thick red. Right: the first part of Wilson’s algorithm: the loop erased walk from x to infinity (E_0) is in thick black. The random walk started at y (dashed green) hits E_0 not at x . As such, it is impossible that both x and y are in the spine.

in the wired spanning forest. Note that by vertex-transitivity, $\mathbb{P}(x \text{ in the spine}) \geq \delta > 0$ for any x , which shows that

$$\mathbb{P}(x \text{ and } y \text{ in the spine}) \geq \frac{\delta}{2}$$

whenever the distance between x and y is large (this follows by extremality of the spanning forest measure). We will contradict this statement.

Wilson’s algorithm implies that we can do the following: run a random walk started from x (to infinity) and loop erase it, call this E_0 . If x is in the spine, E_0 will be part of the spine. There are two possibilities for y to also be on the spine: either it is already on this path E_0 , or it isn’t. With high probability, E_0 will not contain y when x and y are far apart: indeed, y needs to be touched by a simple random walk started from x , but the graph is transient. If y is not yet in E_0 , then x and y can only be (both) on the spine if a random walk started from y hits x before it either wanders off towards infinity or touches $E_0 \setminus \{x\}$. This event is depicted in Figure 1.2. But again, a random walk started from y has very high probability never to hit x at all if x and y are far apart. This readily implies the contradiction. \square

The argument above clearly does not work in the recurrent case. Here, Gromov’s theorem provides some helpful intuition. Indeed, the latter implies that a Cayley graph which is two-ended must be roughly isometric to \mathbb{Z} , although this fact can be proved more directly as was originally done in [168]. In this case, the uniform spanning tree clearly has two ends. For the converse: is it possible to deduce that G is two-ended if the

uniform spanning tree is? If the answer would be negative, this would imply that even on graphs which look roughly like \mathbb{Z}^2 at large scale, the uniform spanning tree would behave like a “trivial” critical model. Again, Benjamini, Lyons, Peres and Schramm provided the answer for vertex-transitive graphs:

Theorem 1.8 (Benjamini et al. [32]). *The spanning tree of a recurrent, vertex-transitive graph G is one-ended unless G is roughly isometric to \mathbb{Z} .*

The hard part of the proof is to show that the graph is two-ended if the uniform spanning tree is and this is more difficult than in the transient setting.

Theorem 1.8 does provide a link between the geometry of the group and the statistical mechanics model on top of it, but it does not do so using rough isometries. In Chapter 2, we show that if a recurrent, vertex-transitive graph satisfies a type of *rooted* Harnack inequality, then the uniform spanning tree must be one-ended. Vertex-transitive graphs are either roughly isometric to \mathbb{Z} or \mathbb{Z}^2 as mentioned above. The graph \mathbb{Z}^2 satisfies a type of parabolic Harnack inequality and this is stable under rough isometries³, two classical results [83, 118, 152]. This Harnack inequality is not quite rooted, but it is relatively easy to see that any graph roughly isometric to \mathbb{Z}^2 , must also satisfy a rooted version. Hence, the uniform spanning tree on any vertex-transitive graph roughly isometric to \mathbb{Z}^2 is one-ended almost surely.

1.1.5 Unimodular random rooted graphs

A generalization of vertex transitive graphs can be realized by requiring only that the law of the graph is invariant as we move through the graph. To that end, it is convenient to look at rooted graphs: (G, o) where o is a marked vertex, from which we view the graph. A natural topology to work with, is the local topology, also known as Benjamini–Schramm topology, first introduced in [33]. A graph isomorphism $G \mapsto G'$ which maps the root o to o' is an *isomorphism of rooted graphs* (G, o) and (G', o') . The local topology is then induced by the space of rooted isomorphism classes \mathcal{G}_\bullet equipped with the distance

$$d((G, o), (G', o')) := e^{-R},$$

where R is the largest radius so that the graph balls of radius R centered at the respective roots are isomorphic as rooted graphs. Similarly defined is the space $\mathcal{G}_{\bullet\bullet}$ of (isomorphism classes of) doubly rooted graphs .

A probability measure \mathbb{P} on \mathcal{G}_\bullet is said to be unimodular whenever it satisfies the mass transport principle: for every measurable $F : \mathcal{G}_{\bullet\bullet} \rightarrow [0, \infty]$,

$$\mathbb{E} \left[\sum_{x \in V} F(G, o, x) \right] = \mathbb{E} \left[\sum_{x \in V} F(G, x, o) \right].$$

³It is worth mentioning that also the usual Harnack inequality is stable under rough isometries [19], but this result is much harder to prove. The parabolic Harnack inequality always implies the usual one.

For a concise background on this beautiful topic, and many examples, we refer to [9].

Perhaps a canonical example of a unimodular rooted graph, is when the graph itself is finite and fixed, and the root is a uniformly chosen vertex. This basic example captures the intuition behind unimodular graphs rather well: every vertex is equally likely to be the root. With this in mind, unimodularity generalizes vertex transitivity, and is also linked to stationarity under shifts induced by random walk (and thus ergodic theory) using “degree biasing”, see for example [27].

Other important examples which fall in this framework are uniform planar maps, such as UIPT [14], UIPQ [111] (and perhaps tessellations of \mathbb{R}^d) and many perturbations of Cayley graphs, like infinite clusters of certain supercritical statistical mechanics models.

Essentially all questions raised above for vertex-transitive graphs can also be asked in this more general setting. Results relying solely on ergodic arguments extend to the setting of unimodular random graphs almost immediately. For example, the Burton–Keane theorem [43] about the number of infinite components of a percolation model can be extended [9], although some care must be taken in defining amenability. Of course, characterizations such as those by Gromov for Cayley graphs are not available in this setting.

Let us focus our attention for now just on the uniform spanning forests. In their seminal work, Aldous and Lyons conjectured that the behavior known for vertex-transitive graphs should also hold for unimodular random rooted graphs:

Conjecture 1.9 (Aldous–Lyons [9]). *Let (G, o) be a unimodular random rooted graph. The wired uniform spanning tree on (G, o) is one-ended, unless G is two-ended, almost surely.*

The argument by Hutchcroft [91] as sketched above works for transient graph in this setting too (in fact, it was developed for it). The only further requirement needed in his argument, is that the expected degree of the root is finite. Later, Hutchcroft managed to remove this assumption [92]. Of course, the argument does not work for recurrent graphs.

Some interesting recurrent graphs, such as the uniform plane triangulation, are actually planar and hence come with additional structure and toolbox. Planar maps are planar graphs, together with an embedding in \mathbb{S}^2 . In the theory explained above for unimodular random rooted graphs, the isomorphisms between rooted graphs are restricted to homeomorphisms of the rooted planar maps.

The latter, too, was introduced by Benjamini and Schramm [33], who provided a beautiful proof of the fact that “local limits” of planar maps are recurrent under the condition of (uniformly) bounded degrees. Using this setting Angel, Hutchcroft, Nachmias and Ray [13] showed that the uniform spanning tree of any unimodular random

rooted planar map is one-ended, unless in the trivial case where the underlying graph is not (and proved other beautiful insights to unimodular planar maps).

Of course, it is not easy to bootstrap these results beyond planar maps, unless perhaps in some relatively explicit cases (such as taking a planar map, but allowing for non nearest neighbor edges). Moreover, even for planar graphs (not planar *maps*), it is not clear how to apply the framework of [13]. Thus, the conjecture of Aldous and Lyons remained open for recurrent graphs in general. One of the main results of this thesis is the resolution of this conjecture.

1.1.6 Further relations between uniform spanning trees and geometry

Focus for now on fixed recurrent, rooted graphs. Given a random walk on (G, o) , we wonder if it makes sense to talk about a random walk conditioned to never return to its starting point. Since the graph is recurrent, it is not clear if this is always well defined. Define T_z and T_z^+ the first hitting respectively return time of z for the simple random walk. Let $(z_n)_n$ be a sequence of vertices going to infinity. Does

$$a_{z_n}(x) := \frac{\mathbb{P}_x(T_{z_n} < T_o^+)}{\deg(o)\mathbb{P}_o(T_{z_n} < T_o^+)}$$

converge as $n \rightarrow \infty$, independent of the choice of sequence? Does it depend on the choice of the origin?

Let us first mention how such random walks would be related to the uniform spanning tree. For transient graphs, Wilson's algorithm allowed to sample first a path towards infinity. This was instrumental in proving the conjecture of Aldous and Lyons 1.9. For recurrent graphs, it is not clear if there is a simple way to sample such a path towards infinity. However, Wilson's algorithm depends only on the loop erasure of a random walk path, and in particular the loop erasure of the path from o to z_n does not see any of the loops from o to itself. Thus, the loop erasure of this random walk path, is the same as the loop erasure of a random walk path conditioned to first touch z_n . Therefore, if there is an unambiguously defined random walk conditioned to never return to o , then we could potentially use this to sample a path from o to infinity in the spanning tree.

If the underlying graph is a unimodular random rooted graph, this can be used to show that *if* there is an unambiguously defined random walk conditioned to never return to o , then the uniform spanning tree must be one-ended. We do this in Chapter 2. This thus provides another link between potentials on the graph and the uniform spanning tree.

Chapter 2 provides properties of the potentials above and the corresponding random walk conditioned to never return to o . For general recurrent graphs, it turns out that the limit points a of a_{z_n} above, are related to the "harmonic measure from infinity", defined for a finite set B of vertices as follows. Take μ_n the hitting distribution on B of a

random walk started from z_n . Define the “harmonic measure from infinity” as the limit of μ_n , provided that it exists. This limit is well defined (for all finite sets B) precisely when the limit of a_{z_n} exists.

Moreover, a certain type of “rooted” Harnack inequality holds if and only if the limit of a_{z_n} exists and does not depend on the choice z_n . We do not prove that this particular Harnack inequality is rough isometry invariant (although that would be interesting to know). It is known that a stronger, classical Harnack inequality is rough isometry invariant [83, 152]. For vertex-transitive graphs, this can be used to show that the uniform spanning tree is one-ended unless the graph is roughly isometric to \mathbb{Z} .

In Chapter 3 we continue on this track and use it to finally resolve the conjecture by Aldous and Lyons (Conjecture 1.9).

1.2. Phase transitions in spin systems

We switch to a (seemingly) totally different model from statistical mechanics. The previous model described the large scale behavior of uniform spanning trees – perhaps intrinsically mostly a mathematical object. This section presents models with which have a physical motivation to describe ferromagnets. Although there are few immediate correspondences between the two models, some (if not many) of the techniques and objects involved will be the same.

Let $G = (V, E)$ be a finite graph and \mathbb{G} a topological group; most of the time we will take $\mathbb{G} = \mathbb{S}^{n-1}$ for integer n . We will call a function $H : \mathbb{G}^V \rightarrow \mathbb{R}$ a Hamiltonian and define the spin measure μ^{Spin} to be

$$d\mu^{\text{Spin}}(\sigma) = \frac{1}{Z} e^{-\beta H(\sigma)} d\sigma, \tag{1}$$

where Z is the normalizing constant, often also called the partition function, and $\beta > 0$ is the inverse temperature. The Hamiltonian $H(\sigma)$ corresponds to the “energy cost” of the configuration $\sigma \in \mathbb{G}^V$, and thus configurations with low energy cost are favored by μ^{Spin} . If the temperature is high, β is small and the energy difference between configurations plays little role. When the temperature is low, β is large and the energy difference may become influential.

In general we will assume that the Hamiltonian is of the form

$$H(\sigma) := \sum_{xy \in E} \mathcal{U}_{xy}(\sigma_x - \sigma_y),$$

where we use additive notation for the groups. Often, the functions \mathcal{U}_{xy} will be the same on each edge, or differ only by positive scaling. An important example of a Hamiltonian for the group $\mathbb{G} = \mathbb{S}^{n-1} \subset \mathbb{R}^n$ is

$$H(\sigma) := - \sum_{xy \in E} J_{xy} \sigma_x \cdot \sigma_y$$

where $J_{xy} \geq 0$ and $\sigma \cdot \sigma'$ denotes the standard inner product on \mathbb{R}^n . This model is called the classical *spin $O(n)$ model*, with special cases $n = 1$ and $n = 2$ bearing the names *Ising-* and *XY-model* respectively. In this case, the energy H favors spins that are aligned, and is minimized for configurations that are constant.

Note that the group \mathbb{G} acts naturally on the state space: $g \cdot \sigma = g + \sigma$ and that the Hamiltonian is invariant under the action of g . Since we are interested in understanding the occurrence of a phase transition, it is natural to ask: how does the function

$$\beta \mapsto \langle \sigma_x \rangle_\beta^g$$

behave when β changes and $g \in \mathbb{G} \setminus \{0\}$ is some fixed boundary condition? We will call

$$\beta_c := \inf\{\beta \geq 0 : \langle \sigma_x \rangle_\beta^g \neq 0\}$$

the critical temperature and will say that a (non-trivial) *phase transition* occurs if the critical temperature does not equal 0 or ∞ .

In the case when H is ferromagnetic, it is often not hard to show that when β is small, $\langle \sigma_x \rangle_\beta^g = 0$. In particular, in this case the symmetry under the group action is preserved: $\langle g' \sigma_x \rangle_\beta^g = \langle \sigma_x \rangle_\beta^g$ for any $g' \in \mathbb{G}$.

When the group is finite (for example when $\mathbb{G} = \mathbb{Z}_2 \cong \{-1, 1\}$) and $\Gamma = \mathbb{Z}^d$ with nearest neighbor interactions, Peierls' argument [141] guarantees that for β large enough $\langle \sigma_x \rangle_\beta^g \neq 0$, in which case we say that there is spontaneous magnetization or that the symmetry is broken. The intuition behind Peierls' argument is that $g \mapsto \mathcal{U}(g) - \mathcal{U}(0)$ taken over $\mathbb{G} \setminus \{0\}$ has a strict positive minimum because the group is finite. Thus, the energy cost of two neighboring spins that disagree can be made arbitrary high. Therefore, the cost of a large interface of disalignments is high and if the underlying graph is of the form $\Gamma = \mathbb{Z}^d$, $d > 1$, does not occur at each scale.

However, the argument breaks down when the group \mathbb{G} is continuous. This leaves the immediate question: does a phase transition occur at all? It turns out that this depends on the global geometry of the underlying graph Γ . Part of the answer is in the negative, as was rigorously established by Mermin and Wagner in the late sixties:

Theorem 1.10 ([134]). *For the planar square lattice and Hamiltonians invariant under the rotation group $O(n)$ for some $n \geq 2$, there is no symmetry breaking at any temperature.*

In fact, this results holds for *all* recurrent (locally finite) graphs, an elegant proof of which is due to McBryan and Spencer [133]. Their argument also provides power-law upper bounds on the two-point function on the square lattice. The proof is, in both cases, based on the so called *spin wave* theories from physics. Essentially, if we believe that the two dimensional spins behave roughly like

$$e^{i\theta} \approx e^{ih_{\text{GFF}}},$$

for some h_{GFF} a Gaussian free field, then the result easily follows. Establishing an *upper* bound in terms of a spin wave turns out to be relatively elementary.

Given the Mermin–Wagner theorem, two immediate questions arise:

- (1) To what extent does this depend on the recurrence of the graph?
- (2) Is there something else that happens in two dimensions, or is the behavior trivial?

The first question will not be studied in this thesis, but was solved in the special case that the graph $\Gamma = \mathbb{Z}^d$ using reflection positivity by Fröhlich, Simon and Spencer in the eighties:

Theorem 1.11 ([70]). *For $d \geq 3$, for all β large enough, the symmetry in the model is broken: there exists a translation invariant Gibbs measure μ_β for which $\mu_\beta(\sigma_0) \neq 0$.*

For a nice proof of this fact, and introduction to reflection positivity, we refer to the lecture notes by Biskup [38] and to [69]. Reflection positivity has the drawback that it can only be applied to \mathbb{Z}^d (in fact, all the symmetries of the discrete tori $(\mathbb{Z}/L\mathbb{Z})^d$ are needed, and they converge locally to \mathbb{Z}^d). The upshot is that it is quite robust under changing coupling constants and allowing for long(er) range interactions. Altogether, it leaves open the general connection to the geometry of the underlying graph. As for percolation, we think the following is true:

Conjecture 1.12. *The spin $O(n)$ model on any Cayley graph which is not roughly \mathbb{Z} or \mathbb{Z}^2 admits a phase where the symmetry is broken.*

1.2.1 Topological phase transition

For now, we will leave the higher dimensional case and focus only on planar graphs and the question: what happens when the graphs are recurrent? In fact, let us just fix $\Gamma = \mathbb{Z}^2$ with nearest neighbor edges for simplicity. Again, it turns out that the answer to the question depends on the model at hand. Let \mathbb{S}^{n-1} be the spin space for some $n \geq 2$ and take the potential $\mathcal{U}(\sigma_x - \sigma_y) = -\beta \sigma_x \cdot \sigma_y$, corresponding to the classical $O(n)$ spin model described above with inverse temperature $\beta = \frac{1}{T}$.

It turns out that still, the behavior depends on n , at least conjectural. Let us start with the case $n = 2$, in which case the group \mathbb{S} is Abelian. Physicists Berezinskii, Kosterlitz and Thouless [36, 37, 108] predicted that a subtle, topological, type of phase transition occurs. The latter two received a Nobel prize for this discovery.

Let d denote the gradient of a function *in the group* \mathbb{G} , which we will take to be either \mathbb{S} or \mathbb{R} with addition. In other words, we view $d : \mathbb{G}^V \rightarrow \mathbb{G}^E$. Each spin configuration thus defines a map $J : E \rightarrow \mathbb{S}$, which satisfies that there is a $\sigma : V \rightarrow \mathbb{S}$ such that

$d\sigma = J$ in \mathbb{S} . We can next identify \mathbb{S} with $[0, 2\pi)$. However, the induced map $J : E \rightarrow \mathbb{R}$ generally does *not* satisfy that there is a function $f : V \rightarrow \mathbb{R}$ such that $df = J$ in \mathbb{R} . A problem occurs when there are faces on the graph around which the sum of J is larger than 2π . A face where this sum of J is larger than 2π will be called a vortex or anti-vortex depending on the orientation of the face. These discrepancies will be the core of our heuristic explanation of a subtle phase transition which does occur in the XY model.

Let us do a back of the envelope calculation. Suppose that the (anti)-vortices do not play an important role and we could actually find a function $f : V \rightarrow \mathbb{R}$ such that $df = J$ in \mathbb{R} . In this case, the function f must be a Gaussian free field (GFF). In other words,

$$\theta = e^{ih_{\text{GFF}}}$$

with h_{GFF} a real-valued GFF with variance β^{-1} . Such an object is also known as a “spin-wave”.

But what would be the implication of an identity of this form? It follows using standard properties of the GFF that

$$\langle \sigma_x \bar{\sigma}_y \rangle \sim |x - y|^{-\gamma}$$

for some power $\gamma > 0$ which depends on the variance β and (a priori) the lattice.

Obviously, this computation is far from rigorous and the effects of the vortices should be taken into account. In fact, it is *not* true that the vortices do not play an important role in the large scale: renormalization group computations suggest that in the low temperature regime, the spins do behave like $e^{ih_{\text{GFF}}}$, but the effective temperature of the Gaussian field h_{GFF} is believed to be affected by the vortices and may be model dependent [100, 107].

Moreover, in the high temperature regime, the vortices destroy the correlation. The idea is that vortices are energy costly, and therefore, if the temperature is low, the above heuristics works and vortex anti-vortex pairs appear relatively sparsely. The renormalization group argument of [100, 107] suggest that in the “scaling limit”, the spin-wave behavior above is recovered, at least at the level of power-law behavior of the two-point functions. However, if the temperature is high, vortices appear all over the place and induce exponential decay of correlations.

To heuristically verify that vortices are indeed the driving factor of the phase transition, we may wonder what happens if we change the potentials \mathcal{U} in such a way that there are no vortices by construction. Is there is no phase transition at all? For example, if we restrict the potential so that two neighboring spins cannot have an angle difference which is larger than $\frac{\pi}{2}$, then vortices do not appear by definition. Does such a system not have a phase transition? In general, it was conjectured by Patrascioiu and Seiler

[140] that something like this is true. Their conjecture was proved in a very nice way by Aizenman[5], under a stronger restriction on the angle differences than necessary to prevent vortices.

Going back to the XY model, part of the heuristic argument above was made rigorous in a massive breakthrough by Fröhlich and Spencer [71], who provided a power-law *lower* bound for the two-point function. For a recent review paper of their proof, we refer to [104]. Two new approaches to this result will be provided in this thesis.

Theorem 1.13 (Fröhlich and Spencer [71]). *There exists a $\beta_1 < \infty$ such that for all $\beta \geq \beta_1$, there exists a $c(\beta) \in (0, \infty)$ such that*

$$\langle \sigma_x \bar{\sigma}_y \rangle \geq |x - y|^{-c(\beta)}$$

for all $x \neq y$. Moreover, $c(\beta) \rightarrow 0$ as $\beta \rightarrow \infty$.

Most computations to establish the transition (also by Fröhlich and Spencer [71]) rely on an expansion of the spins in terms of a so called spin wave and a Coulomb gas. In order for this to make formal sense, we need to endow the underlying lattice with more topological structure: in general we need to consider a chain of complexes of highest dimension at least 2.

The decomposition amounts to writing the 1-form of spins, viewed as taking value in \mathbb{R} , in their orthogonal decomposition (also known as Hodge decomposition as in this setting, the “harmonic” part vanishes)

$$d\theta := d\varphi + d^*q,$$

where φ is a function on the underlying lattice and q a 2-form. The operators d and d^* refer to the discrete (exterior) derivative and its formal adjoint. The field φ is called the spin-wave, q describes the charges. The latter is often related to a Coulomb gas.

This decomposition is analogous to the Hodge decomposition of white noise: in this case, the two fields φ and q are Gaussian free fields, see for example [51] for the lattice and [16] for the continuum case.

Of course, the fields φ and q are generally not independent, but in the setting of Villain interactions they are. In this case, the spin wave is exactly equal to a Gaussian free field [21, 74, 77]. The remaining task to prove Theorem 1.13 is to show that the interactions of the Coulomb gas are relatively small, and this can be done rigorously using a multi-scale analysis as in [71]. To apply the (now non-rigorous) renormalization group arguments of [100], a more exact understanding of the contributions of the Coulomb gas would be needed.

Let us also remark that the $O(n)$ models with $n \geq 3$ are believed to undergo no phase transition at all in two dimensions, a conjecture attributed to Polyakov [147]. It

must be pointed out that this conjecture has not (even) been settled by physicist, and the only results known to the author are in the “ $n \rightarrow \infty$ ” setting by Kupiainen [113]: for each β , there exists an n large so that the $O(n)$ model has exponential decay at this value of β .

1.3. Fourier–Pontryagin transform and duality

A different way to approach phase transitions in spin models, is to use duality directly. Here, by duality, we mean duality through Fourier transformations of partition and correlation functions. In general, this method turned out to be extremely useful in statistical mechanics, particularly in two dimensions. For the Ising model, at least on the level of partition function, this is known as Kramers–Wannier [110] duality, although nowadays, it is not often explained using Fourier transforms explicitly since their computations lead to stronger results.

In this special case, the primal Ising model maps to dual Ising model on the dual graph, up to a temperature inversion given by

$$\beta \mapsto \log \left(\frac{e^\beta - 1}{e^\beta + 1} \right) =: \beta^*.$$

The point at which the primal and dual Ising model share the same temperature is called the self dual point, and is given by $\beta_{sd} := \log(\sqrt{2} - 1)$. In the case where the underlying lattice is also self dual, the self dual is also the critical point [23, 138], an extremely useful property that lies at the heart of many breakthroughs [46, 55, 90, 159].

A heuristic argument for the latter fact can be seen through the *free energy* of the model, which is defined as

$$f^{\text{Ising}}(\beta) = \lim_{n \rightarrow \infty} \frac{1}{|\Lambda_n|} \log(Z_{\Lambda_n, \beta}^{\text{Ising}}),$$

where Λ_n is the $2n \times 2n$ box on the square lattice and $Z_{\Lambda_n, \beta}^{\text{Ising}}$ the partition function. The free energy is well defined by sub-multiplicative arguments and does not depend on the boundary conditions due to amenability. An important fact about the free energy is that phase transitions of the model are visible as a singularities of the free energy, and the type of singularity determines the type of transition, see [69] for a better explanation.

By identifying the partition function of the primal and dual model (up to an explicit global factor), it follows that

$$f(\beta) = f(\beta^*) + C(\beta),$$

where $C(\beta)$ is explicit and analytic. In conclusion, *if* the Ising model on the square lattice undergoes a single phase transition and *if* this is visible as a singularity in the

free energy, then it must occur at the self dual point. Of course, the assumptions would need mathematical justification.

Beyond the level of partition functions, extensions of Fourier duality to order–disorder variables were studied by Kadanoff and Ceva [101]. These type of observables were used to apply renormalization group arguments in physics, and later rigorously by Smirnov in his breakthrough works, showing the (two-dimensional) Ising model is conformally invariant [46, 90, 159].

More recent related applications are in terms of the Fortuin–Kasteleyn (FK) percolation model [68]. For two parameters $q \in (0, \infty)$ and $p \in [0, 1]$, this is a spin model on the edges of a graph, taking value in the group $\mathbb{Z}_2 \cong \{0, 1\}$, and with Gibbs weights proportional to

$$\phi_{q,p}^{\text{FK}}(\omega) \propto \prod_{e \in E} p^{\omega_e} (1-p)^{1-\omega_e} q^{\#\{\text{clusters in } \omega\}}.$$

If $q = 2$ this is a graphical representation of the Ising model: sampling independently for each *cluster* of ω a uniform element of $\{-1, +1\}$ gives the Ising model [62].

As the emergent symmetry in the definition of $\phi_{p,q}^{\text{FK}}$ may suggest, each ω is dual to $\omega^* := 1 - \omega$ (living on the dual graph) and it turns out that the weight $\phi_{p,q}^{\text{FK}}(\omega)$ equals exactly $\phi_{p',q}^{\text{FK}}(\omega^*)$ with appropriately changed temperature $p' = p'(p, q)$ [56]. Such type of *exact* duality was instrumental already in Smirnov’s work [46, 55, 90, 159]. This duality is much stronger than the Fourier one, as it establishes a bijection between the models, not just equality of partition functions.

These examples are special. On the square lattice, it can be used to show that the phase transition for both models occurs at the point where the model is self dual [23] if $q \geq 1$. The biggest difficulty in establishing this heuristically believable fact is that duality may behave weirdly for non-standard boundary conditions.

However, as we will see next, self duality nor exact combinatorial duality is present for continuous spin models (as far as we know).

The circle group

If an Abelian group \mathbb{G} is finite, its irreducible representations are equal to the characters of the group and the group of characters can be identified with the Pontryagin dual $\widehat{\mathbb{G}} := \text{Hom}(\mathbb{G}, \mathbb{S})$. In the special cases of $\mathbb{G} = \mathbb{Z}/q\mathbb{Z}$, the groups are self-dual. Furthermore, as we have mentioned before, for the Ising and Potts model, the Fourier transform is again an Ising or Potts model with a different temperature.

When the group \mathbb{G} is locally compact and Abelian, the Fourier transform can still be expressed in terms of the Pontryagin dual $\widehat{\mathbb{G}}$. In general, it is no longer true that the groups are self-dual: an Abelian group is compact if and only if its dual is discrete. The case of particular interest to us is when $\mathbb{G} = \mathbb{S}$ and in this case, it is classical that

$\widehat{\mathbb{G}} = \mathbb{Z}$.⁴

The Fourier pairing of two functions $\widehat{f} : \mathbb{Z} \rightarrow \mathbb{C}$ and $f : \mathbb{S} \rightarrow \mathbb{C}$ is then described by

$$\widehat{f}(k) = \int_0^1 f(t) e^{-2\pi i k t} dt \quad f(t) = \sum_{k \in \mathbb{Z}} \widehat{f}(k) e^{i k t},$$

when this makes sense. We will mostly restrict to functions taking value in \mathbb{R} , in which case its Fourier transform is an even function and vice-versa. We say that a function is *positive definite* if its Fourier transform is non negative, so it defines a probability measure on the dual space. In particular, if $f : \mathbb{S} \rightarrow \mathbb{R}$ is itself positive and, moreover, positive definite, its Fourier transform is too.

This is extremely helpful in the context of the $O(2)$ spin models. If $\mathcal{U} : \mathbb{S} \rightarrow \mathbb{R}$ is a symmetric potential and $e^{-\mathcal{U}}$ is positive definite, then the spin model defined by

$$H(\sigma) = \sum_{e \in E} \mathcal{U}(d\sigma_e) \quad \text{and} \quad d\mu(\sigma) \sim e^{-H(\sigma)} d\sigma$$

has a well defined, probabilistic Fourier transform, which we will call the *dual model*.

The latter is described by taking $\mathcal{V} = -\log \widehat{e^{-\mathcal{U}}}$ and defining the Gibbs measure

$$\nu(h) \propto e^{-\sum_{e \in E} \mathcal{V}(dh_e)}$$

with respect to an appropriate counting measure. Some care has to be taken as to where the height function h lives, the details are in Chapter 7. However, when the graph Γ is planar, the dual model will be a proper height function on the dual graph Γ^* .

1.3.1 BKT-transition through duality

Let us go back to the existence of a phase transition. The basic heuristic is that the transition should be visible in the primal and the dual model, and at the same temperature. Again, it may be tempting to look at the free energy, as it is the same for the spin model and its dual height function, however it is a priori not obvious at all *how* the subtle BKT transition would manifest itself at this level. In fact, for the planar XY model it is believed that the free energy $\beta \mapsto f^{\text{Spin}}(\beta)$ is smooth at the critical point, but not analytic (which it should be away from the critical point). This shows, at least on a heuristic level, that the two models must have a single phase transition at the same temperature.

Applying Fourier–Pontryagin duality, we can calculate exactly the characteristic function of the spin model,

$$\langle e^{i(d\sigma, \eta)_2} \rangle_\beta = \frac{Z_\beta^{\text{Height}}(\eta)}{Z_\beta^{\text{Height}}} \quad (2)$$

⁴The reason is as follows: any irreducible (complex) representation of an Abelian group is one-dimensional by Schur’s lemma. Moreover, \mathbb{S} is compact so that any irreducible representation $\rho : \mathbb{S} \rightarrow \text{GL}_{\mathbb{C}}(1)$ takes values in \mathbb{S} again, hence the characters of \mathbb{S} are isomorphic to the automorphisms of \mathbb{S} .

for $\eta : \vec{E} \rightarrow \mathbb{Z}$ a 1-form, where we define the *twisted partition function*

$$Z_\beta^{\text{Height}}(\eta) := \sum_{h:F \rightarrow \mathbb{Z}} \exp(-\mathcal{V}_\beta(dh_e + d\eta_e)),$$

and where \mathcal{V}_β is the dual height function potential as above. In light of (2), it seems plausible that we can indeed transfer information from one model to the other. However, there are two potential issues. First: is it easier to establish a phase transition in the dual model? Second: the ratio of partition functions on the height function side does not easily transfer into an observable in a height function model. Indeed, the twisted partition $Z_\beta^{\text{Height}}(\eta)$ essentially is a sum over 1-forms which are *not* gradients of height functions.

Delocalization

We will assume here that the height function model comes with a natural temperature parameter $\beta > 0$. If the potential \mathcal{V}_β is obtained as the dual potential of the spin $O(2)$ model at inverse temperature β , then it has such a parameterization. Another option is to take the rescaled potential $\mathcal{V}_\beta := \beta\mathcal{V}$.

Write $\nu_\beta^{\text{Height}}$ for the height function measure corresponding to the potential \mathcal{V}_β . In two dimensions, it turns out that it is not always possible to define a height function in a translation invariant way (although translation-invariant measures supported on gradients can always be properly defined). For example, the Gaussian free field on the graph does not allow for a translation invariant, pointwise definition since the two dimensional green function blows up. Of course, this field takes value in \mathbb{R} rather than \mathbb{Z} , so we may wonder if the latter affects such properties. In Chapter 7, we will see that all height function models which are dual to a probabilistic spin model are upper bounded by the Gaussian free field in terms of variances. On the other hand, relatively simple considerations show that for low enough β , we *can* make sense of such measures and, moreover, the variance is finite: $\nu_\beta(h_x^2) < \infty$. In this case, we call the height function *localized*. When there is too much fluctuation and the variance of $(h_o - h_x)$ (which is always well defined) blows up as $x \rightarrow \infty$, we will call the model *delocalized*.

Establishing such a phase transition was considered a major problem, and only a few cases were known up to a few years ago. Fröhlich and Spencer [71] proved a delocalized phase for a few models using delicate multiscale techniques for related Coulomb gasses. Recently, Lammers [114] provided an argument for general potentials based on relatively elementary percolation arguments, under the condition that the potentials are convex functions over the integers. We will present a slight variation of his proof in Chapter 4, to keep the exposition more or less self-contained.

Theorem 1.14 (Lammers [114]). *Let \mathcal{V} be convex and symmetric with $\mathcal{V}(1) \leq \mathcal{V}(0) +$*

$\log(2)$. Suppose Γ is planar, bi-periodic and has degrees bounded above by three. The height function delocalizes.

Thus, the height function undergoes a phase transition on trivalent graphs. In certain cases, a generalization of this theorem to bi-periodic planar graphs was given in [8]. This thesis provides a second approach which works under different conditions, see Chapter 7. Although this extension is interesting in its own right, to prove a BKT phase transition it is not relevant, as is explained at the end of Chapter 5. Still, for integer-valued height functions, it remained open whether or not such a transition actually happens at a single point:

Question 1.15. Is there a $\beta_c \in (0, \infty)$ such that below β_c , the height function is localized, while above β_c , it is delocalized?

In full generality, this question is still open, but in some cases the answer is known. One way to provide a positive answer to the question, is to prove that the map

$$\beta \mapsto \mu_\beta^{\text{Height}}(h_x^2)$$

is monotone in β and we do so in Chapter 7.

At roughly the same time, a different approach to tackle the existence of a critical point was proposed and proved by Lammers [116]. We will postpone the explanation of his approach, as it fits better with the discussion about “loop representations” below.

Loop representation and random currents

The second step to proving the BKT-transition is to transfer delocalization of the height function into properties of the spin model. The apparent problem in equation (2) is also present for the high temperature expansion of van der Waerden [167] for the Ising model: expressing correlation functions gives ratio’s of partition functions which sum over different objects.

There is a slight variation of the high temperature expansion for the Ising model, as proposed in [82], which is now known as the *random current representation*. A *current* for the Ising model is a function $\mathbf{n} : E \rightarrow \{0, 1, \dots\}$ on a graph $\Gamma = (V, E)$. A vertex $x \in V$ is sources of a current whenever $d^*\mathbf{n}_x$ is odd, and denote by $\partial\mathbf{n}$ the set of sources of \mathbf{n} . The random current representation asserts that for sets $A \subset V$

$$\left\langle \prod_{x \in A} \sigma_x \right\rangle_\beta = \frac{\sum_{\mathbf{n}: \partial\mathbf{n}=A} w_\beta(\mathbf{n})}{\sum_{\mathbf{n}: \partial\mathbf{n}=\emptyset} w_\beta(\mathbf{n})}, \quad \text{where} \quad w_\beta(\mathbf{n}) = \prod_{e \in E} \frac{\beta^{\mathbf{n}_e}}{\mathbf{n}_e!}.$$

At first sight, this expression does not help at all: the ratio on the right still involves summations over different objects.

However, it turns out that products of correlation functions are handled more easily. This is a consequence of the so-called “switching lemma” [82] (see also [53, 56] for an introduction). In particular, the square of a two-point function can be expressed as the probability of some (percolation) event for random currents; and the problem of equation (2) is solved:

$$\left\langle \prod_{x \in A} \sigma_x \right\rangle_\beta^2 = \frac{\sum_{\substack{\mathbf{n}_1: \partial \mathbf{n}_1 = \emptyset \\ \mathbf{n}_2: \partial \mathbf{n}_2 = \emptyset}} w_\beta(\mathbf{n}_1) w_\beta(\mathbf{n}_2) \mathbf{1}_{\mathcal{F}_A}}{\sum_{\substack{\mathbf{n}_1: \partial \mathbf{n}_1 = \emptyset \\ \mathbf{n}_2: \partial \mathbf{n}_2 = \emptyset}} w_\beta(\mathbf{n}_1) w_\beta(\mathbf{n}_2)}$$

for some appropriate event \mathcal{F}_A depending on $\mathbf{n}_1 + \mathbf{n}_2$. There are also other consequences of the switching lemma; it paves a way to provide unified proofs of many correlation inequalities such as Simon’s inequality [158] and the Mermin–Wagner inequality [134].

This random current representation has been instrumental in the understanding of the Ising model, perhaps mostly in dimensions $d > 2$ [53]. In the eighties, Aizenman and others used the representation to prove triviality and continuity of the Ising model in $d > 4$ respectively $d \geq 4$ [1, 3]. A different approach to triviality was proposed by Fröhlich [72], although it is now understood that the different representations are (roughly) the same. Recently, the current representation was used to provide a new proof of sharpness for the Ising model [60], continuity of the phase transition at the critical point [6], and away from the critical point [150]. It was also instrumental in the breakthrough by Aizenman and Duminil-Copin, showing that the Ising model at the critical dimension $d = 4$ is trivial [2]. This list is, of course, by no means exhaustive.

A minor downside to the representation, perhaps, is that it is rather specific to the Ising model, although it does allow to extend some results to the φ^4 model [1, 2, 72]. There is also a very recent extension to a more direct “tangled” current representation for the φ^4 model [85]. But still, for continuous spin systems, we have not seen how to get around the problem of equation (2). Still, the classical XY model was known to have random walk representations somewhat in the spirit of the random current representation, going back to Symanzik [161] and Fröhlich, Simon and Spencer [70], but the representation did not provide a link between the height function (i.e. dual model) and the spin model.

In Chapter 5, we will provide a version of the random walks / loop representation mentioned, which does provide this link. The idea is to do the same high temperature expansion as for the Ising model (using the Abelian nature of the spins), in order to find a random current representation of the height function model. A current now is a function on the oriented edges $\mathbf{n} : \vec{E} \rightarrow \mathbb{Z}$ which we think of as a collection of arrows and which can be coupled with the height function by taking $\nabla h_{xy} = \mathbf{n}_{xy} - \mathbf{n}_{yx}$. A loop representation ω can be obtained by connecting at every vertex the incoming and outgoing edges of the current.

Doing the combinatorics, it turns out that this loop representation does satisfy a switching lemma, and the two-point function of the XY model can be expressed as an expectation in the loop model. Finally, this allows to remove the excitations in (2) (for well chosen correlation functions) and this can be used to say that if the height function delocalizes, the spin model must have algebraic decay of correlations.

This proves the fact that if the height function undergoes a phase transition, then so does its dual spin model, but it does not prove that the converse is also true, nor that the transition happens at the same point. The reason is that there are certain topological events for the loops, which could, theoretically, imply there are large loops, but the height function is still localized. Recently, Lammers provided an answer to this problem by confirming that such topological events do not occur [116]. In essence, his methods rely on exploring the loops described above “one by one” (and hence really as a random walk), in a way that allows to glue together different “crossing” events. A particular consequence of these works, together with correlation inequalities on the spin side, is that the spin model and height function must undergo a phase transition at exactly the same temperature.

1.3.2 The BKT transition revisited

As was explained after (2), calculating the spin-spin correlations using the duality transform is not straightforward. However, if we change our perspective from the spin model to the height function, we can compute the characteristic functions related to the height function. This results in ratios of (excited) partition functions on the spin side. Because the Pontryagin dual of \mathbb{Z} is isomorphic to \mathbb{S} , it turns out that the excitations can be differentiated away, and an exact duality relation becomes visible on the level of covariances (see Lemma 7.3). In case of the 2-dimensional torus, it states that for any $\omega : \vec{E} \rightarrow \mathbb{R}$ a 1-form,

$$\mathbb{E}[(dh, \omega)] + \mathbb{E}[(\mathcal{U}'(d\theta), \omega)] = C(\omega, \omega)$$

for some explicit constant C depending on \mathcal{U} . Here, we view dh as the gradient of the height function on the dual graph, while we view $d\theta$ as the gradient on the primal graph.

Although stated here for the planar case, versions of this duality formula hold for all graphs. Our perspective follows closely the lines of [131] where similar (and more advanced) techniques are used to study random walks, the uniform spanning tree and harmonic functions. Higher dimensional analogues of the planar duality often impose further topological constraints on the objects studied: for example, the hypercubic lattice can be viewed as a (CW-)complex and comes with the related algebraic tools, see e.g. [21, 74, 77]. But, a graph itself is only a 1-dimensional complex and it is therefore interesting to ask what happens when no further topological properties exist.

The covariance duality immediately yields that the covariance function of gradient of a height function and \mathcal{U}' of the spin model is exactly the same (up to a sign). This highlights a direct link between a possible phase transition occurring in the two models and it seems a natural flow of information can be established. The latter is exactly the content of Chapter 7. One of our conclusions is a form of equivalence of phase transition in the spin model and height function, as in Theorem 7.7. This equivalence is presented in terms of susceptibility of a non-classical correlation function. In case of the XY model, this implies the existence of a BKT phase transition in the sense that the two-point function $\langle \sigma_x \bar{\sigma}_y \rangle_\beta$ decays algebraically above the critical point.

1.4. Articles and preprints appearing in this thesis

- Chapter 2 is based on a preprint [34] with Nathanaël Berestycki (submitted).
- Chapter 3 is based on the preprint [63] with Tom Hutchcroft (submitted).
- Chapter 4 has not appeared anywhere and contains a review of some material for height functions. It is partially based on work with Marcin Lis, which is yet to appear.
- Chapter 5 is based on an article together with Marcin Lis [64],
- Chapter 6 appeared in the original preprint of [64], but before publication we found a slight modification which did not need the theory presented in this chapter. We think it may present an useful tool for future use.
- Chapter 7 is based on a recent preprint with Marcin Lis [65] (submitted).

Part I.

**Geometry and the Uniform Spanning
Tree**

Harnack inequality and one-endedness of UST on reversible random graphs

2.1. Introduction

Let (G, o) be a random unimodular rooted graph, which is almost surely recurrent (with $\mathbb{E}(\deg(o)) < \infty$). The **wired Uniform Spanning Tree** (UST for short) on G is defined to be the unique weak limit of the uniform spanning tree on any finite exhaustion of the graph, with wired boundary conditions. The existence of this limit is well known, see e.g. [131]. (In fact, since the graph is assumed to be recurrent, the wired or free boundary conditions give the same weak limit). The UST is a priori a spanning forest of the graph G , but since G is recurrent this spanning forest consists in fact a.s. of a single spanning tree which we denote by \mathcal{T} (see e.g. [143]). We say that \mathcal{T} is **one-ended** if the removal of any finite set of vertices A does not disconnect \mathcal{T} into at least two infinite connected components. Intuitively, a one-ended tree consists of a unique semi-infinite path (the spine) to which finite bushes are attached.

The question of the one-endedness of the UST (or the components of the UST, when the graph is not assumed to be recurrent) has been the focus of intense research ever since the seminal work of Benjamini, Lyons, Peres and Schramm [32]. Among many other results, these authors proved (in Theorem 10.1) that on every unimodular vertex-transitive graph, and more generally on a network with a transitive unimodular automorphism group, every component is a.s. one-ended unless the graph is itself roughly isometric to \mathbb{Z} (in which case it and the UST are both two-ended). (This was extended by Lyons, Morris and Schramm [130] to graphs that are neither transitive nor unimodular but satisfy a certain isoperimetric condition slightly stronger than uniform transience). More generally, a conjecture attributed to Aldous and Lyons is that every unimodular one-ended graph is such that every component of the UST is a.s. one-ended. This has been proved in the planar case in the remarkable paper of Angel, Hutchcroft, Nachmias and Ray [13] (Theorem 5.16) and in the transient case by results of Hutchcroft [91, 92]. The conjecture therefore remains open in the recurrent case, which is the focus

of this chapter.

Let us motivate further the question of the one-endedness of the UST. It can in some sense be seen as the analogue¹ of the question of percolation at the critical value. To see this, note that when the UST is one-ended, every edge can be oriented towards the unique end, so that following the edges forward from any given vertex w , we have a unique semi-infinite path starting from w obtained by following the edges forward successively. Observe that this forward path necessarily eventually arrives at the spine and moves to infinity along it. Given a vertex v , we may define the past $\mathbf{Past}(v)$ of v to be the set of vertices w for which the forward path from w contains v ; it is natural to view $\mathbf{Past}(v)$ as the analogue of a connected component in percolation. From this point of view, the a.s. one-endedness of the tree is equivalent to the finiteness of the past (i.e., connected component in this analogy) of every vertex, as anticipated. We further note that on a unimodular graph, the expected value of the size of the past is however always infinite, as shown by a simple application of the mass transport principle. This confirms the view that the past displays properties expected from a critical percolation model. In fact, Hutchcroft proved in [93] that the two models have same critical exponents in sufficiently high dimension.

In this paper we give necessary and sufficient conditions for the one-endedness of the UST on a recurrent, unimodular graph. These are, respectively: (a) existence of the potential kernel, (b) existence of the harmonic measure from infinity, and finally (c) an anchored Harnack inequality. We illustrate our results by showing that they give straightforward proofs of the aforementioned result of Benjamini, Lyons, Peres and Schramm [32] in the recurrent case (which is one of the most difficult aspects of the proof of the whole theorem, and is in fact stated as Theorem 10.6). We also apply our results to some unimodular random graphs of interest such as the Uniform Infinite Planar Triangulation (UIPT) and related models of infinite planar maps, for which we deduce the Harnack inequality.

To state these results, we first recall the following definitions. Our results can be stated for reversible environments or **reversible random graphs**, i.e., random rooted graphs such that if X_0 is the root and X_1 the first step of the random walk conditionally given G and X_0 then (G, X_0, X_1) and (G, X_1, X_0) have the same law. As noted by Benjamini and Curien in [27], any unimodular graph (G, o) with $\mathbb{E}(\deg(o)) < \infty$ satisfies this reversibility condition after biasing by the degree of o . Conversely, any reversible random graph gives rise to a unimodular rooted random graph after unbiasing by the degree of the root. This biasing/unbiasing does not affect any of the results below since they are almost sure properties of the graph. Note also that again by results in

¹We thank Tom Hutchcroft for this wonderful analogy.

[27], a *recurrent* rooted random graph whose law is stationary for random walk is in fact necessarily reversible. See also Hutchcroft and Peres [96] for a nice discussion and Aldous and Lyons [9] for a systematic treatment.

For a nonempty set $A \subset v(G)$ we define the **Green function** by setting for $x \in v(G) \setminus A$ and $y \in v(G)$:

$$\mathbf{G}_A(x, y) = \mathbb{E}_x \left[\sum_{n=0}^{T_A-1} \mathbb{1}_{\{X_n=y\}} \right], \quad (1)$$

where T_A denotes the hitting time of A , and $G_A(x, y) = 0$ for $x \in A$. Let

$$g_A(x, y) := \frac{\mathbf{G}_A(x, y)}{\deg(y)}$$

denote the normalised Green function. (Note that due to reversibility, $g_A(x, y) = g_A(y, x)$.)

Let A_n be any (sequence) of finite sets of vertices such that $d(A_n, o) \rightarrow \infty$ as $n \rightarrow \infty$. Here, by $d(A_n, o)$, we just mean the minimal distance of any vertex in A_n to o . It is natural to construct the **potential kernel** of the infinite graph G by an approximation procedure; we set

$$a_{A_n}(x, y) := g_{A_n}(y, y) - g_{A_n}(x, y) \quad (2)$$

In this manner, the potential kernel compares the number of visits to y , starting from x versus y , until hitting the far away set A_n . We are interested in existence and uniqueness of limits for a_{A_n} as $n \rightarrow \infty$. In this case we call the unique limit the potential kernel of the graph G . We will see that the existence and uniqueness of this potential kernel turns out to be equivalent to a number of very different looking properties of the graph. This definition of the potential kernel differs slightly from the one appearing in [118] for \mathbb{Z}^2 , because we work with a more convenient normalization for graphs that are not transitive.

We move on to harmonic measure from infinity. Let A be a fixed finite, nonempty set of vertices. Let $\mu_n(\cdot)$ denote the harmonic measure on A , started from A_n if we wire all the vertices in A_n . The **harmonic measure from infinity**, if it exists, is the limit of μ_n (necessarily a probability measure on A).

Now let us turn to Harnack inequality. We say that (G, o) satisfies an **(anchored) Harnack inequality** (AHI) if there exists an exhaustion $(V_R)_{R \geq 1}$ of the graph (i.e. V_R is a finite subset of vertices and $\cup_{R \geq 1} V_R = v(G)$), and there exists a nonrandom constant $C > 0$, such that the following holds. For every function $h : v(G) \rightarrow \mathbb{R}_+$ which is harmonic except possibly at 0, and such that $h(0) = 0$:

$$\max_{x \in \partial V_R} h(x) \leq C \min_{x \in \partial V_R} h(x). \quad (3)$$

The word *anchored* in this definition refers to the fact that the exhaustion is allowed to depend on the choice of root o , and the functions are not required to be harmonic there. (As we show in Remark 2.10, a consequence of our results is that an anchored Harnack inequality automatically implies the Elliptic Harnack inequality (EHI) on a suitably defined sequence of growing sets.)

We now state the main theorem.

Theorem 2.1. *Suppose (G, o) is a recurrent reversible random graph (or equivalently after unbiaseding by the degree of the root, (G, o) is recurrent unimodular random graph with $\mathbb{E}(\deg(o)) < \infty$). The following properties are equivalent.*

- (a) *Almost surely, the pointwise limit of the truncated potential kernel $a_{A_n}(x, y)$ exists and does not depend on the choice of A_n .*
- (b) *Almost surely, the weak limit of the harmonic measure μ_n from A_n exists and does not depend on A_n .*
- (c) *Almost surely, (G, o) satisfies an anchored Harnack inequality.*
- (d) *The uniform spanning tree \mathcal{T} is a.s. one-ended.*

Furthermore, if any of these conditions hold, a suitable exhaustion for the anchored Harnack inequality is provided by the sublevel sets of the potential kernel, see Sections 2.5 and 2.6.

2.1.1 Some applications

Strengthening of [32]. Before showing some applications of this result, let us point out that Theorem 2.1 complements and strengthens some of the results of Benjamini, Lyons, Peres and Schramm [32]. In that paper, the (easy) implication (d) implies (b) was noted. We therefore in particular obtain a converse in the reversible case. One can furthermore easily see using their results that on any recurrent planar graph with bounded face degrees (e.g., any recurrent triangulation) (d) holds, i.e., the uniform spanning tree is a.s. one-ended: indeed, for such a graph, there is a rough embedding from the planar dual to the primal, which is assumed to be recurrent, and therefore the planar dual must be recurrent too by Theorem 2.17 in [131]. By Theorem 12.4 in [32] this implies that the uniform spanning tree (on the primal) is a.s. one-ended, and so (d) holds. (In fact, Theorem 5.16 in [13] shows that the bounded face degree assumption is not needed).

Applications to planar maps. Therefore, in combination with [32], Theorem 2.1 above applies in particular to unimodular, recurrent triangulations such as the UIPT,

or similar maps such as the UIPQ. This therefore implies that these maps have a well-defined potential kernel, harmonic measure from infinity, and satisfy the anchored Harnack inequality. As shown in Remark 2.10, this also implies the **elliptic Harnack inequality** (for sublevel sets of the potential kernel, see Theorem 2.24 for a precise statement). We point out that the elliptic Harnack inequality should not be expected to hold on usual metric balls, but can only be expected on growing sequences of sets which take into account the “natural conformal embedding” of these maps. This is exactly what the potential kernel and its sublevel sets allows us to do.

More general implications. We already mention that the equivalence between (a) and (b) is valid more generally, for instance for any locally finite, recurrent graph. The implication (a) \implies (c) to the Harnack inequality (c) is then valid under the additional assumption that the potential kernel grows to infinity (something which we can prove assuming unimodularity). We recall that (d) implies (b) is also true for deterministic graphs, as proved in [32].

Remark 2.1. Many of the arguments in this chapter are true for deterministic graphs. The unimodularity (or reversibility) of the graph with respect to random walk is only used in Lemma 2.20, whose main use is to show that the potential kernel, if it exists, diverges to infinity along any sequence going to infinity (see Lemma 2.21). This property is used for instance in both directions of the relations between (c) and (d), since both go via (a). The unimodularity (or stationarity) is also used to prove that the walks conditioned not to return to the origin satisfy the infinite intersection property, a key aspect of the proof one-endedness. Finally this is also proved to show that if there is a bi-infinite path in the UST then it must essentially be almost space-filling, which is the other main argument of the proof of one-endedness.

Deterministic case of the Aldous–Lyons conjecture. As previously mentioned, Theorem 2.1 can be applied to give a direct proof of the one-endedness of the UST for recurrent vertex-transitive graphs not roughly equivalent to \mathbb{Z} , which is essentially Theorem 10.6 in [32].

Corollary 2.2. *Suppose G is a fixed recurrent, vertex-transitive graph. If G is one-ended then the UST is also a.s. one-ended. Otherwise G is roughly isometric to \mathbb{Z} .*

Proof. Note that G must be unimodular, otherwise G is nonamenable and so cannot be recurrent (see [160]). Note also that the volume growth of the graph is at most polynomial (as otherwise the walk cannot be recurrent). By results of Trofimov [165], the graph is therefore roughly isometric to a Cayley graph Γ . Since it is recurrent (as recurrence is preserved under rough isometries, see Theorem 2.17 and Proposition 2.18

of [131]), we deduce by a classical theorem of Varopoulos (see e.g. Theorem 1 and its corollary in [168]) that Γ is a finite extension of \mathbb{Z} or \mathbb{Z}^2 and is therefore (as is relatively easily checked) roughly isometric to either of these lattices. Since either of these lattices enjoy the Parabolic Harnack Inequality (PHI), which is, by a consequence of a result proved by Grigoryan [83] and Saloff-Coste [152] independently, preserved under rough isometries (see also [48]), we see that G itself satisfies PHI and therefore also the Elliptic Harnack Inequality (EHI): for any $R > 1$, if h is harmonic in the metric ball $B(2R)$ of radius $2R$ around the origin, then $\sup_{B(R)} h(x) \leq C \inf_{B(R)} h(x)$. (In fact, by a deep recent result of Barlow and Murugan, EHI is now known directly to be stable under rough isometries [19], but here we can appeal to the much simpler stability of PHI. We recommend the following textbooks for related expository material: [112], [17] and [169].)

Suppose that G is not roughly isometric to \mathbb{Z} , therefore it is roughly isometric to \mathbb{Z}^2 . Let us show that G satisfies the anchored Harnack inequality (3), with the exhaustion sequence simply obtained by considering metric balls $V_R = B(R)$. Let h be nonnegative harmonic on G except at 0. Since G is rough isometric to \mathbb{Z}^2 , we can cover ∂V_R with a fixed number (say K) of balls of radius $R/10$, such that the union of these balls is connected (here we used two-dimensionality). Let $x, y \in \partial V_R$, we can find $x = x_0, \dots, x_K = y$ with $d(x_i, x_{i+1}) \leq R/10$, and $d(x_i, o) > 2R/10$. Exploiting the EHI in each of the K balls $B(x_i, 2R/10)$ inductively (since h is harmonic in each of these balls), we find that $h(x) \leq C^K h(y)$. Since x, y are arbitrary in ∂V_R , this proves the anchored Harnack inequality (3). \square

We also show that the one-endedness of the UST holds for unimodular recurrent random graphs if we in addition assume that they are strictly subdiffusive; that is, we settle the Aldous–Lyons conjecture in that case. (This encompasses many models of random planar maps, but can of course hold on more general graphs, see in particular [121], recalled also in Remark 2.5, for sufficient conditions guaranteeing this).

Theorem 2.3. *Suppose (G, o) is reversible, almost surely recurrent and strictly subdiffusive (i.e., satisfies (SD) below). Then (G, o) satisfies (a)–(d).*

This applies e.g. for high-dimensional incipient infinite percolation cluster, as explained after Remark 2.5. The proof of Theorem 2.3 takes as an input the results of Benjamini, Duminil–Copin, Kozma and Yadin [30] which shows that for strictly subdiffusive unimodular graphs there are no nonconstant harmonic functions of linear growth, and the trivial observation that the effective resistance between points is at most linear in the distance between these points. We believe it should be possible to use the same idea to prove the result assuming only diffusivity: to do this, it would suffice to

prove that the effective resistance grows strictly sublinearly, except on graphs roughly isometric to \mathbb{Z} .

Random walk conditioned to avoid the origin. The existence of the potential kernel allows us to define (by h -transform) a random walk conditioned to never touch a given point (even though this is of course a degenerate conditioning on recurrent graphs). We study some properties of the conditioned walk and show among other things that two independent conditioned walks must intersect infinitely often, a fact which plays an important role in the proof of Theorem 2.1 for the equivalence between (a) and (d). We conclude the chapter with a finer study of this conditioned walk on CRT-mated random planar maps. In this case we are able to show that the hitting probability of a point far away from the origin by the conditioned walk remains bounded away from 1 in the limit as the point diverges to infinity (and is bounded away from 0 for “almost all” such points). See Theorem 2.49 for a precise statement. We also discuss a conjecture (see (49)) which, if true, would show a significant difference of behaviors with respect to the more standard case of \mathbb{Z}^2 (where these hitting probabilities converge to 1/2, as surprisingly shown in [148]).

2.2. Background and notation

Before we begin with the proofs of our theorems, we need to introduce the main notations that we will use throughout this text.

A **graph** G consists of a countable collection of vertices $v(G)$ and edges $e(G) \subset \{\{x, y\} : x, y \in v(G)\}$ and we will always assume that the vertex degrees are finite. We will work with undirected graphs, but will sometimes take the directed edges $\vec{e}(G) = \{(x, y) : \{x, y\} \in e(G)\}$.

The graph G comes with a natural metric $d(x, y)$, which is the **graph distance**, i.e. the minimal length of a path between two vertices x and y . For $n \in \mathbb{N}$, we will denote by

$$B(y, n) = \{x \in v(G) : d(x, y) \leq n\},$$

the **metric ball** of radius n . For a set $A \subset v(G)$, we will write ∂A for its outer boundary in $v(G)$, that is

$$\partial A = \{x \in v(G) \setminus A : \text{there exists a } y \in A \text{ with } x \sim y\}.$$

We will make extensive use of the graph **Laplacian** which we normalise as follows:

$$\Delta f(x) = \sum_{y \sim x} c(x, y)(f(y) - f(x)), \quad (4)$$

for functions $f : v(G) \rightarrow \mathbb{R}$ (here $c(x, y)$ is the conductance of the edge (x, y) , which is typically equal to one in this paper, except in Section 2.7 where we consider random walk conditioned to avoid the origin forever). A function $h : v(G) \rightarrow \mathbb{R}$ is called **harmonic** at x if $(\Delta h)(x) = 0$.

Let $X = (X_n)_{n \geq 0}$ denote the simple random walk on G , with its law written as \mathbb{P} and \mathbb{P}_x to mean $\mathbb{P}(\cdot \mid X_0 = x)$. For a set $A \subset v(G)$, we define the **hitting time** $T_A = \inf\{n \geq 0 : X_n \in A\}$ and $T_x := T_{\{x\}}$ whenever $A = \{x\}$ consists of just one element. We will write T_A^+ for the **first return time** to a set A . Suppose that G is a connected graph. The **effective resistance** is defined through

$$\mathcal{R}_{\text{eff}}(x \leftrightarrow y) := \frac{\mathbf{G}_x(y, y)}{\deg(y)}.$$

Recall the useful identity

$$\mathcal{R}_{\text{eff}}(x \leftrightarrow y) = \frac{1}{\deg(y) \mathbb{P}_y(T_x < T_y^+)} \quad (5)$$

The proof is obvious from the definition of effective resistance when we use the obvious identity

$$\mathbf{G}_x(y, y) = \frac{1}{\mathbb{P}_y(T_x < T_y^+)},$$

which can be seen by considering the number of excursions from y to y , which is a geometric random variable by the Markov property.

For infinite graphs G , we will say that a sequence of subgraphs $(G_n)_{n \geq 1}$ of G is an **exhaustion** of G whenever G_n is finite for each n and $v(G_n) \rightarrow G$ as $n \rightarrow \infty$. Fix some exhaustion $(G_n)_{n \geq 1}$ of an infinite graph G and define the graph G_n^* as G_n , together with the identification of G_n^c , where we have deleted all self-loops created in the process. For two vertices $x, y \in v(G)$ we recall that

$$\mathcal{R}_{\text{eff}}(x \leftrightarrow y) = \lim_{n \rightarrow \infty} \mathcal{R}_{\text{eff}}(x \leftrightarrow y; G_n^*),$$

see for instance [131, Section 9.1]. As is well known, the effective resistance defines a metric (see for instance exercise 2.67 in [131]).

Later, we will often work with the metric $\mathcal{R}_{\text{eff}}(\cdot \leftrightarrow \cdot)$ on $v(G)$, instead of the standard graph distance. We introduce the notation

$$\mathcal{B}_{\text{eff}}(x, R) = \{y \in v(G) : \mathcal{R}_{\text{eff}}(x \leftrightarrow y) \leq R\} \quad (6)$$

for the closed ball with respect to the effective resistance metric. Notice that, in general, this metric space is *not* a length space - making it somewhat inconvenient.

Another result that we will need to use a few times is the ‘last exit decomposition’, or rather two versions thereof which can be proved similarly to [118, Proposition 4.6.4].

Lemma 2.4 (Last Exit Decomposition). *Let G be a graph and $A \subset B \subset v(G)$ finite. Then for all $x \in A$ and $b \in \partial B$ we have*

$$\mathbb{P}_x(X_{T_{B^c}} = b) = \sum_{z \in A} \mathbf{G}_{B^c}(x, z) \mathbb{P}_z(T_{B^c} < T_A^+, X_{T_{B^c}} = b).$$

Moreover, for $x \in B$ we have

$$\mathbb{P}_x(T_A < T_{B^c}) = \sum_{z \in A} \mathbf{G}_{B^c}(x, z) \mathbb{P}_z(T_{B^c} < T_A^+).$$

2.3. Equivalence between (a) and (b)

2.3.1 Base case of equivalence

We will say that a sequence of finite sets of vertices $(A_n)_{n \geq 1}$ ‘goes to infinity’ whenever $d(A_n, o) \rightarrow \infty$ as $n \rightarrow \infty$. Here, by $d(A_n, o)$, we just mean the minimal distance of any vertex in A_n to o . Recall the definition of a_{A_n} , which also satisfies

$$a_{A_n}(x, y) = g_{A_n}(y, y) - g_{A_n}(x, y) = \frac{1}{\deg(y)} \frac{\mathbb{P}_x(T_{A_n} < T_y)}{\mathbb{P}_y(T_{A_n} < T_y^+)}. \quad (7)$$

Clearly, both the numerator and the denominator tend to 0 as n tends to infinity by recurrence of the underlying graph G . When a sequence of subsets A_n has been chosen we will write a_n instead of a_{A_n} with a small abuse of notations.

The goal of this section is to prove the equivalence between (a) and (b) in Theorem 2.1 (in the base case where the set A consists of two points; this will be extended to arbitrary finite sets in Section 2.3.3). First, we show that subsequential limits of a_n always exist.

Lemma 2.5. *Let $(A_n)_{n \geq 1}$ be some sequence of finite sets of vertices going to infinity. There exists a subsequence $(n_k)_{k \geq 1}$ going to infinity such that for all $x, y \in v(G)$ the limit*

$$a(x, y) := \lim_{k \rightarrow \infty} a_{n_k}(x, y)$$

exists in $[0, \infty)$. Moreover, $a(x, y) > 0$ precisely when the removal of y from G does not disconnect x from A_{n_k} for all k large enough.

Proof. Fix $y \in v(G)$ and suppose first that for all $u \sim y$ we have $\mathbb{P}_u(T_{A_n} < T_y) > 0$ for all n large enough (i.e., y does not disconnect a portion of the graph from infinity).

Let $x \in v(G)$ and fix n so large that A_n does not contain y, x or any of the neighbors of y . For each $u \sim y$, we can force the random walk started from x to go through u before touching A_n or y to get

$$\mathbb{P}_x(T_{A_n} < T_y) \geq \mathbb{P}_x(T_u < T_{A_n} \wedge T_y) \mathbb{P}_u(T_{A_n} < T_y). \quad (8)$$

Upon taking $u \sim y$ such that it maximizes $\mathbb{P}_u(T_{A_n} < T_y)$ and by recurrence of G we get the existence of $c(x, y) > 0$ for which

$$a_n(x, y) = \frac{\mathbb{P}_x(T_{A_n} < T_y)}{\sum_{u \sim y} \mathbb{P}_u(T_{A_n} < T_y)} \geq \frac{\mathbb{P}_x(T_u < T_{A_n} \wedge T_y)}{\deg(y)} \geq c(x, y) > 0.$$

The same reasoning as in (8) but in the other direction gives

$$\mathbb{P}_x(T_{A_n} < T_y) \leq \frac{\mathbb{P}_u(T_{A_n} < T_y)}{\mathbb{P}_u(T_x < T_{A_n} \wedge T_y)}.$$

Hence, using again recurrence of G we get that there is some $C(x, y) < \infty$ such that (upon taking the right u)

$$a_n(x, y) \leq \frac{\mathbb{P}_u(T_{A_n} < T_y)}{\mathbb{P}_u(T_x < T_{A_n} \wedge T_y) \sum_{u \sim y} \mathbb{P}_u(T_{A_n} < T_y)} \leq C(x, y) < \infty.$$

We deduce that for fixed x, y , subsequential limits of $a_n(x, y)$ exist and the existence of subsequential limits for all x, y simultaneously follows from diagonal extraction.

The existence of subsequential limits in the general case is the same as we can always lower bound $a_n(x, y)$ by 0 and the upper bound does not change.

Now, if $x \in v(G)$ is such that the removal of y disconnects x from A_{n_k} , then $a_{n_k}(x, y) = 0$. Suppose thus that x is such that the removal of y does not disconnect x from A_{n_k} for all k large enough. In this case, we can restrict ourselves to just the component of G with y removed, in which both A_{n_k} and x are as the hitting probabilities are the same in this case. Hence, we are back in the situation above and $a_{n_k}(x, y) \geq c(x, y) > 0$. \square

We next present a result, which shows that any subsequential limit appearing in Lemma 2.5 must satisfy a certain number of properties.

Proposition 2.6. *Let $a(x, y)$ be any subsequential limit as in Lemma 2.5. Then $a : v(G) \rightarrow \mathbb{R}_+$ satisfies*

(i) for each $y \in v(G)$

$$\Delta a(\cdot, y) = \delta_y(\cdot) \quad \text{and} \quad a(y, y) = 0,$$

where we recall that Δ is defined in (4) and is normalised so that $\Delta f(x) = \sum_y (f(y) - f(x))$.

(ii) for all $x, y \in v(G)$ we have

$$a(x, y) = \lim_{k \rightarrow \infty} \mathbb{P}_{A_{n_k}}(T_x < T_y) \mathcal{R}_{\text{eff}}(x \leftrightarrow y),$$

where \mathbb{P}_A refers to the law of a random walk starting from A , when all of the vertices in A have been wired together.

The equivalence between (a) and (b) of Theorem 2.1 (in the base case where the finite set B on which we need to define harmonic measure consists of two points) is then obvious, and we collect it here:

Corollary 2.7. *Let G be a recurrent graph. Then*

$$\text{hm}_{x,y}(x) := \lim_{n \rightarrow \infty} \mathbb{P}_{A_n}(T_x < T_y)$$

exists for all $x, y \in v(G)$ and is independent of the sequence $(A_n)_n$ if and only if the potential kernel is uniquely defined. Furthermore, in this case,

$$a(x, y) = \text{hm}_{x,y}(x) \mathcal{R}_{\text{eff}}(x \leftrightarrow y).$$

Proof of Proposition 2.6. The proof of item (i) is rather elementary. Fix $y \in v(G)$ and $n \geq 1$. Since $x \mapsto \mathbb{P}_x(T_{A_n} < T_y)$ is a harmonic function outside of y and A_n by the simple Markov property, we get that $x \mapsto a_n(x, y)$ is harmonic outside y and A_n , see (7). It follows that $x \mapsto a(x, y)$ is harmonic at least away from y . Furthermore, note that $a_n(y, y) = 0$ by definition and

$$\sum_{u \sim y} a_n(u, y) = \frac{\sum_{u \sim y} \mathbb{P}_u(T_{A_n} < T_y)}{\sum_{u \sim y} \mathbb{P}_u(T_{A_n} < T_y)} = 1$$

so $\Delta a_n(\cdot, y)|_{=y} = 1$. This finishes the proof of (i).

For part (ii), we notice first that by properties of the electrical resistance,

$$\sum_{u \sim y} \mathbb{P}_u(T_{A_n} < T_y) = \deg(y) \mathbb{P}_y(T_{A_n} < T_y^+) = \frac{1}{\mathcal{R}_{\text{eff}}(y \leftrightarrow A_n)},$$

which allows us to write

$$a_n(x, y) = \mathcal{R}_{\text{eff}}(y \leftrightarrow A_n) \mathbb{P}_x(T_{A_n} < T_y). \quad (9)$$

Identify the vertices in A_n and delete possible self-loops created in the process. The resulting graph G'_n is then still recurrent. Let $\mathbf{G}_y(\cdot, \cdot)$ denote the Green function on this graph when the walk is killed at y . We can also express the effective resistance in terms of the normalised Green function: that is,

$$\mathcal{R}_{\text{eff}}(y \leftrightarrow A_n) = \frac{\mathbf{G}_y(A_n, A_n)}{\deg(A_n)}$$

Using the Markov property and since G'_n is reversible,

$$\begin{aligned} a_n(x, y) &= \mathbb{P}_x(T_{A_n} < T_y) \frac{\mathbf{G}_y(A_n, A_n)}{\deg(A_n)} = \frac{\mathbf{G}_y(x, A_n)}{\deg(A_n)} \\ &= \frac{\mathbf{G}_y(A_n, x)}{\deg(x)} = \mathbb{P}_{A_n}(T_x < T_y) \mathcal{R}_{\text{eff}}(x \leftrightarrow y; G'_n) \end{aligned} \quad (10)$$

by using the same argument in the other direction, and where the effective resistance in the last line is calculated in G'_n .

Since the graph G is recurrent, it follows that $\mathcal{R}_{\text{eff}}(x \leftrightarrow y; G'_n)$ converges to $\mathcal{R}_{\text{eff}}(x \leftrightarrow y; G)$ as $n \rightarrow \infty$ (as the free and wired effective resistances agree). We deduce that

$$a(x, y) = \lim_{k \rightarrow \infty} \mathbb{P}_{A_{n_k}}(T_x < T_y) \mathcal{R}_{\text{eff}}(x \leftrightarrow y),$$

which finishes part (ii). \square

Remark 2.2. We wish to point out that, in general, the potential kernels are *not* symmetric (even if they are uniquely defined).

2.3.2 Triangle inequality for the potential kernel

Before we start of the proof of the remaining implications, we need some preliminary estimates on the potential kernel, showing that it satisfies a form of triangle inequality. This plays a crucial role throughout the rest of this paper. We also need a decomposition of the potential kernel in order to prove that for reversible graphs, the potential kernel (if it is well defined) satisfies the growth condition.

We start with a simple and well known application of the optional stopping theorem:

Lemma 2.8. *Let A be some finite set and suppose that $x, y \in A$. Then*

$$\frac{\mathbf{G}_{A^c}(x, y)}{\deg(y)} = \mathbb{E}_x[a(X_{T_{A^c}}, y)] - a(x, y).$$

Proof. This is Proposition 4.6.2 in [118], but we include for completeness since its proof is simple. Let $x, y \in A$ and notice that

$$M_n := a(X_n, y) - \sum_{j=0}^{n-1} \frac{\delta_y(X_j)}{\deg(y)}$$

is a martingale. Applying the optional stopping theorem at $T_{A^c} \wedge n$, we obtain

$$a(x, y) = \mathbb{E}_x[M_0] = \mathbb{E}_x[a(X_{n \wedge T_{A^c}}, y)] - \frac{1}{\deg(y)} \mathbb{E}_x \left[\sum_{j=0}^{(n \wedge T_{A^c})-1} \delta_y(X_j) \right],$$

Taking $n \rightarrow \infty$, since A is finite, we deduce from dominated (resp. monotone) convergence that

$$\mathbb{E}_x[a(X_{n \wedge T_{A^c}}, y)] \rightarrow \mathbb{E}_x[a(X_{T_{A^c}}, y)], \quad \frac{1}{\deg(y)} \mathbb{E}_x \left[\sum_{j=0}^{(n \wedge T_{A^c})-1} \delta_y(X_j) \right] \rightarrow \frac{\mathbf{G}_{A^c}(x, y)}{\deg(y)},$$

showing the result. \square

Proposition 2.9. *Let $x, y, z \in v(G)$ be three vertices. We have the identity*

$$\frac{\mathbf{G}_z(x, y)}{\deg(y)} = a(x, z) - a(x, y) + a(z, y).$$

Proof. Fix $x, y, z \in v(G)$ and let $(A_n)_{n \geq 1}$ be some sequence of finite sets of vertices going to infinity². Glue together A_n on the one hand, and the vertices of $B(o, m)^c$ on the other hand. Delete all self-loops created in the process and write ∂_m for the vertex corresponding to $B(o, m)^c$. Let \tilde{X}_k be the simple random walk on the graph obtained from gluing A_n and ∂_m . We define for $w, w' \in B(o, m) \cup \{\partial_m\}$ the function

$$a_{m,n}(w, w') := \mathcal{R}_{\text{eff}}(\{\partial_m, w'\} \leftrightarrow A_n) \mathbb{P}_w(T_{A_n} < T_{w'} \wedge T_{\partial_m}).$$

By recurrence and (9), we have that $a_{m,n}(w, w') \rightarrow a_n(w, w')$ as $m \rightarrow \infty$, for all w, w' .

Fix n so large that x, y and z are not in A_n . Let m be so large that x, y, z and A_n are in $B(o, m)$. Define $E_{n,m} = \{A_n, z, \partial_m\}$. Then, as in Lemma 2.8,

$$a_{m,n}(x, y) = \mathbb{E}_x[a_{m,n}(\tilde{X}_{T_{E_{n,m}}}, y)] - \frac{\mathbf{G}_{E_{n,m}}(x, y)}{\deg(y)} \quad (11)$$

On the other hand, by definition of $E_{m,n}$ we have

$$\begin{aligned} \mathbb{E}_x[a_{m,n}(X_{T_{E_{n,m}}}, y)] &= \mathbb{P}_x(T_z < T_{A_n} \wedge T_{\partial_m}) a_{m,n}(z, y) \\ &\quad + \mathbb{P}_x(T_{A_n} < T_z \wedge T_{\partial_m}) a_{m,n}(A_n, y) \\ &\quad + \mathbb{P}_x(T_{\partial_m} < T_{A_n} \wedge T_z) a_{m,n}(\partial_m, y), \end{aligned}$$

where a priori the hitting probabilities are calculated on the graph where A_n and ∂_m are glued. However, as we are only interested in the first hitting time of either of these sets, it does not matter and we can calculate the probabilities also for the random walk on the graph G . Notice that, by definition, $a_{m,n}(\partial_m, y) = 0$. Plugging this back into (11) we obtain

$$\begin{aligned} a_{m,n}(x, y) &= \mathbb{P}_x(T_z < T_{A_n} \wedge T_{\partial_m}) a_{m,n}(z, y) \\ &\quad + \mathbb{P}_x(T_{A_n} < T_z \wedge T_{\partial_m}) a_{m,n}(A_n, y) - \frac{\mathbf{G}_{E_{n,m}}(x, y)}{\deg(y)}. \end{aligned}$$

We have already observed that $a_{m,n}(w, y) \rightarrow a_n(w, y)$ for each w as $m \rightarrow \infty$. Then, by recurrence of G and monotone convergence, we get

$$a_n(x, y) = \mathbb{P}_x(T_z < T_{A_n}) a_n(z, y) + \mathbb{P}_x(T_{A_n} < T_z) a_n(A_n, y) - \frac{\mathbf{G}_{\{A_n, z\}}(x, y)}{\deg(y)}. \quad (12)$$

²Although the proof here relies on the assumption that the limit of a_{A_n} does not depend on the choice A_n , we note for future reference that it also applied if we replace a by any subsequential limit of a_{A_n} . See also Remark 2.4.

Next, we wish to take $n \rightarrow \infty$. The left-hand side converges to $a(x, y)$ as $n \rightarrow \infty$, by definition of the potential kernel. The first term on the right-hand side converges to $a(z, y)$ by the same argument and recurrence of the graph G . Using once more monotone convergence, we find

$$\frac{\mathbf{G}_{\{A_n, z\}}(x, y)}{\deg(y)} \rightarrow \frac{\mathbf{G}_z(x, y)}{\deg(y)} \quad (13)$$

as n goes to infinity. We are left to deal with the term $\mathbb{P}_x(T_{A_n} < T_z)a_n(A_n, y)$, which we claim converges to $a(x, z)$.

From the definition of a_n , together with the representation in (9), we find

$$a_n(A_n, y) = \mathcal{R}_{\text{eff}}(y \leftrightarrow A_n) \mathbb{P}_{A_n}(T_{A_n} < T_y) = \mathcal{R}_{\text{eff}}(y \leftrightarrow A_n).$$

Thus, using again the same representation of $a_n(x, z)$, we see that

$$\begin{aligned} a_n(A_n, y) \mathbb{P}_x(T_{A_n} < T_z) &= \mathcal{R}_{\text{eff}}(y \leftrightarrow A_n) \mathbb{P}_x(T_{A_n} < T_z) \\ &= a_n(x, z) \frac{\mathcal{R}_{\text{eff}}(y \leftrightarrow A_n)}{\mathcal{R}_{\text{eff}}(z \leftrightarrow A_n)}. \end{aligned}$$

Using the triangle inequality for the effective resistance, we notice that

$$\frac{\mathcal{R}_{\text{eff}}(y \leftrightarrow A_n)}{\mathcal{R}_{\text{eff}}(z \leftrightarrow y) + \mathcal{R}_{\text{eff}}(y \leftrightarrow A_n)} \leq \frac{\mathcal{R}_{\text{eff}}(y \leftrightarrow A_n)}{\mathcal{R}_{\text{eff}}(z \leftrightarrow A_n)} \leq \frac{\mathcal{R}_{\text{eff}}(y \leftrightarrow z) + \mathcal{R}_{\text{eff}}(z \leftrightarrow A_n)}{\mathcal{R}_{\text{eff}}(z \leftrightarrow A_n)}.$$

By recurrence of G , the left and right hand side converge to 1 as $n \rightarrow \infty$. In particular, we deduce that

$$a_n(A_n, y) \mathbb{P}_x(T_{A_n} < T_z) \rightarrow a(x, z)$$

as $n \rightarrow \infty$. Plugging this, together with (13) back into (12) we conclude:

$$a(x, y) = a(x, z) + a(z, y) - \frac{\mathbf{G}_z(x, y)}{\deg(y)}$$

as desired. □

Remark 2.3. Proposition 2.9 is an extensions of results known for the lattice \mathbb{Z}^2 , see Proposition 4.6.3 in [118] and the discussion thereafter. As far as we know, these proofs are based on precise asymptotic behavior of the potential kernel, a tool we do not seem to have.

Remark 2.4. The statement of Proposition 2.9 is also valid for an arbitrary subsequential limit $a(\cdot, \cdot)$ of $a_n(\cdot, \cdot)$, even when a proper limit is not known to exist. In particular, it shows that given such a subsequential limit $a(\cdot, y)$ there is a unique way to coherently define $a(\cdot, z)$. For this reason, if $\lim_{n \rightarrow \infty} a_n(x, y)$ is shown to exist for a fixed y and *all* $x \in v(G)$, it follows that this limit exists for *all* $x, y \in v(G)$ simultaneously. This will be used in Theorem 2.12.

Corollary 2.10. *For each $x, z \in v(G)$ and all $\epsilon > 0$ there exists an $N = N(\epsilon, x, z)$ such that for all y with $d(x, y) \geq N$ we have*

$$|a(x, y) - a(z, y)| \leq \epsilon$$

and in particular $\lim_{n \rightarrow \infty} a(x, y_n) - a(z, y_n) = 0$ for any sequence $(y_n)_{n \geq 1}$ going to infinity.

Notice that Corollary 2.10 does *not* say that $a(y_n, x) - a(y_n, z) \rightarrow 0$ as $n \rightarrow \infty$ in general! Indeed, a similar argument shows that $a(y_n, x) - a(y_n, z) \rightarrow a(z, x) - a(x, z)$, which is nonzero in general.

Proof. Fix $x, z \in v(G)$ and suppose by contradiction that there is some $\epsilon > 0$, such that for infinitely many $n \geq 1$ (but in fact we can with a small abuse of notation assume for all $n \geq 1$ after taking a subsequence), there is some y_n with $d(x, y_n) \geq n$ for which

$$|a(z, y_n) - a(x, y_n)| > \epsilon.$$

By Proposition 2.9 and $\deg(\cdot)$ -reversibility of the Simple Random Walk we have

$$a(x, y_n) - a(z, y_n) = a(x, z) - \frac{\mathbf{G}_z(x, y_n)}{\deg(y_n)} = a(x, z) - \frac{\mathbf{G}_z(y_n, x)}{\deg(x)}.$$

Take $A_n = \{y_n\}$ and recall (see e.g. (10)) that

$$a_n(x, z) = \frac{\mathbf{G}_z(y_n, x)}{\deg(x)}.$$

Therefore

$$a(x, y_n) - a(z, y_n) = a(x, z) - a_n(x, z).$$

Since this converges to zero as $n \rightarrow \infty$, we get the desired contradiction. \square

We immediately deduce that the harmonic measures from infinity of $\{x, y\}$ and $\{z, y\}$ are very similar if y is far away from x and z .

Corollary 2.11. *Fix $x, z \in v(G)$. For every $\epsilon > 0$, there exists an $N = N(x, z, \epsilon)$ such that for all y with $d(x, y) \geq N$ we have*

$$|\mathrm{hm}_{x,y}(x) - \mathrm{hm}_{z,y}(z)| \leq \frac{\epsilon + \mathcal{R}_{\mathrm{eff}}(z \leftrightarrow x)}{\mathcal{R}_{\mathrm{eff}}(x \leftrightarrow y)}.$$

Proof. Fix $x, z \in v(G)$ and $\epsilon > 0$. Let N_0 be so large that Corollary 2.10 holds, i.e. so that for every y with $d(x, y) \geq N_0$,

$$|a(x, y) - a(x, z)| \leq \epsilon.$$

Recall from Corollary 2.7 the expression

$$a(x, y) = \text{hm}_{x,y}(x) \mathcal{R}_{\text{eff}}(x \leftrightarrow y).$$

so that

$$\text{hm}_{x,y}(x) - \text{hm}_{z,y}(z) = \frac{a(x, y) - a(z, y)}{\mathcal{R}_{\text{eff}}(x \leftrightarrow y)} + \frac{\text{hm}_{z,y}(z)(\mathcal{R}_{\text{eff}}(x \leftrightarrow y) - \mathcal{R}_{\text{eff}}(z \leftrightarrow y))}{\mathcal{R}_{\text{eff}}(x \leftrightarrow y)}.$$

Last, using the triangle inequality for the effective resistance twice (and symmetry $\mathcal{R}_{\text{eff}}(x \leftrightarrow y) = \mathcal{R}_{\text{eff}}(y \leftrightarrow x)$), we find

$$|\mathcal{R}_{\text{eff}}(x \leftrightarrow y) - \mathcal{R}_{\text{eff}}(z \leftrightarrow y)| \leq \mathcal{R}_{\text{eff}}(x \leftrightarrow z).$$

Plugging this all together and defining $N \geq N_0$ so large that $\mathcal{R}_{\text{eff}}(x \leftrightarrow y) \geq \frac{1}{\epsilon} \mathcal{R}_{\text{eff}}(x \leftrightarrow z)$ for all y with $d(x, y) \geq N$, gives that

$$|\text{hm}_{x,y}(x) - \text{hm}_{z,y}(z)| \leq \epsilon + \epsilon,$$

which is the desired result. \square

2.3.3 Gluing and harmonic measure

We suppose throughout this section that the potential kernel is well defined in the sense that the subsequential limits appearing in Lemma 2.5 are all equal. By Corollary 2.7, this implies that the harmonic measure from infinity is well defined for two-point sets.

Let $B \subset v(G)$ be a set. Glue together all vertices in B and delete all self-loops that were created in the process. We denote the graph induced by the gluing G_B . Note that G_B need not be a simple graph, even when G was.

We will prove in this section that, if the potential kernel is well defined on G , it is also well defined on G_B , whenever B is a finite set. Furthermore, we will prove an explicit expression of the potential kernel on the graph G_B in the case where B is a finite set. These results are an extension of results on the lattice \mathbb{Z}^2 , see for instance [118, Chapter 6], but we will use different arguments, following from the expression for the potential kernel in terms of harmonic measure from infinity as in Corollary 2.7.

Theorem 2.12 (Gluing Theorem). *Suppose $a(x, y) = \lim_{n \rightarrow \infty} a_n(x, y)$ exists for all $x, y \in v(G)$ and does not depend on the choice of the sequence of sets A_n going to infinity. Let $B \subset v(G)$ be a finite set, whose removal does not disconnect G , and suppose $x \in B$. Then*

$$q_B(w) := \lim_{n \rightarrow \infty} \mathcal{R}_{\text{eff}}(B \leftrightarrow A_n) \mathbb{P}_w(T_{A_n} < T_B) \quad (14)$$

exists and is given by

$$q_B(w) = a(w, x) - \mathbb{E}_w[a(X_{T_B}, x)]; \quad w \in v(G_B) \setminus \{B\}; q_B(B) = 0. \quad (15)$$

Extending q_B to $v(G)$ in the natural way (i.e., using (15) with $w \in v(G)$), we have

$$(\Delta q_B)(w) = \text{hm}_B(w) := \lim_{z \rightarrow \infty} \mathbb{P}_z(X_{T_B} = w); \quad w \in B \quad (16)$$

where the Laplacian Δ is calculated on G via (4).

Note in particular, that in the expression (15) for q_B , any choice of $x \in B$ gives the same value and so is irrelevant. We will prove this theorem in the two subsequent subsections, proving first (14) and (15) in Section 2.3.3, and then (16) in Section 2.3.3.

Before we give the proof, we first state some corollaries. The first one is that the harmonic measure from infinity is well defined for the arbitrary finite set B (subject to the assumption that the removal of B does not disconnect G).

Corollary 2.13. *Fix a finite set $B \subset v(G)$ as in Theorem 2.12. Let A_n be a set of vertices tending to infinity. Then for any $x \in B$,*

$$\text{hm}_B(x) = \lim_{n \rightarrow \infty} \mathbb{P}_{A_n}(X_{T_B} = x) \quad (17)$$

exists and is positive for all $x \in B$ such that the removal of $B \setminus \{x\}$ does not disconnect x from infinity.

Proof. Fix $w \notin B$, then arguing as in (9) and (10) we get

$$\mathbb{P}_{A_n}(T_w < T_B) = \frac{\mathcal{R}_{\text{eff}}(A_n \leftrightarrow B) \mathbb{P}_w(T_{A_n} < T_B)}{\mathcal{R}_{\text{eff}}(w \leftrightarrow B; G_{A_n})} \rightarrow \frac{q_B(w)}{\mathcal{R}_{\text{eff}}(w \leftrightarrow B)}$$

as $n \rightarrow \infty$. This limit is by definition the desired value of $\text{hm}_{B \cup \{w\}}(w)$. Note furthermore that $q_B(w)$ is strictly positive by Lemma 2.5.

Applying the same reasoning but with B changed into $B' = B \setminus \{x\}$ (with $x \in B$) and $w = x$, shows that the limit in (17) exists. Furthermore, if the removal of B' does not disconnect x from ∞ , we see that $q_{B'}(w) > 0$ again, and so $\text{hm}_B(x) > 0$. \square

Next, we show that the potential kernel can only be well defined if the graph G is one-ended.

Corollary 2.14. *If the potential kernel is well defined, G is one-ended.*

Proof. Intuitively, on multiple-ended graphs there isn't a single harmonic measure from infinity since there are several ways of converging to infinity. Suppose G has more than one end. Let x_1, x_2, \dots, x_M be some finite number of vertices, such that removing them from $v(G)$ and looking at the induced graph, we have (at least) two infinite components. Write $B_n = B(o, n)$ and choose n large enough that $x_1, \dots, x_M \in B_n$. Consider the graph G_{B_n} resulting from gluing B_n together as in the theorem. Clearly, the removal of B_n creates at least two infinite components. Pick a vertex z of B_n^c and suppose it is in

one infinite component. Let $(\{w_i\})_{i \geq 1}$ be any sequence of vertices going to infinity in an infinite component that does not contain z . Then $\mathbb{P}_{w_i}(T_z < T_{B_n}) = 0$ (for each i), yet this converges by Corollary 2.13 to $\text{hm}_{B_n \cup z}(z) > 0$ since the removal of B_n does not disconnect z from infinity. This is the desired contradiction. \square

Theorem 2.12 a priori only shows that the potential kernel with ‘pole’ B is well defined when B does not disconnect G . We can, however, extend it to arbitrary finite sets B and to an arbitrary second variable y .

Corollary 2.15. *Let $B \subset v(G)$ be any finite set. The potential kernel $a_B : v(G_B)^2 \rightarrow \mathbb{R}_+$ is well defined in the sense that the limit*

$$a^{G_B}(w, y) = \lim_{n \rightarrow \infty} \mathbb{P}_w(T_{A_n} < T_y; G_B) \mathcal{R}_{\text{eff}}(w \leftrightarrow y; G_B),$$

exist for all $w, y \in v(G_B)$ and does not depend on the choice of sequence of sets A_n . Here, the probability and effective resistance are calculated on the graph G_B .

Proof. We start with taking \bar{B} as the hull (in the sense of complex analysis, meaning we “fill it in” with respect to the point at infinity) of B , defined by adding to B all the points in $v(G_B)$ that belong to finite connected components of $v(G_B) \setminus B$. Since G is one-ended by Corollary 2.14, \bar{B} does not disconnect G . By Theorem 2.12, we have that for any sequence of sets $(A_n)_{n \geq 1}$ going to infinity, the limit

$$a^{G_{\bar{B}}}(w, \bar{B}) := q_{\bar{B}}(w) = \lim_{n \rightarrow \infty} \mathbb{P}_w(T_{A_n} < T_{\bar{B}}) \mathcal{R}_{\text{eff}}(\bar{B} \leftrightarrow A_n)$$

exists for each $w \in v(G_{\bar{B}})$ and does not depend on the choice of sequence of vertices A_n going to infinity. Moreover, this limit also trivially exists (and is zero) if w is in one of the finite components of $v(G) \setminus B$.

Hence we deduce that actually for all $w \in v(G_B)$ we have that the limit

$$a^{G_B}(w, B) = \lim_{n \rightarrow \infty} \mathbb{P}_w(T_{A_n} < T_B) \mathcal{R}_{\text{eff}}(A_n \leftrightarrow B)$$

exists and does not depend on the choice of sequence of sets going to infinity A_n . Now, by Proposition 2.9 (see Remark 2.4) we get that for all $w, y \in v(G_B)$ the limit

$$a^{G_B}(w, y) = \lim_{n \rightarrow \infty} \mathbb{P}_w(T_{A_n} < T_y; G_B) \mathcal{R}_{\text{eff}}(A_n \leftrightarrow y; G_B)$$

exists and does not depend on the choice of the sequence A_n . This is the desired result. \square

Proof of (14) and (15)

Proof. Fix $(A_n)_{n \geq 1}$ a sequence of finite sets of vertices going to infinity. For a finite set $B \subset v(G)$ and $x \in B$, we will define the function $q_B : v(G_B) \rightarrow \mathbb{R}_+$ through

$$q_B(w) = a(w, x) - \mathbb{E}_w[a(X_{T_B}, x)],$$

and $q_B(B) = 0$, whenever the potential kernel on G is well defined. We will prove (14) using induction on the number of vertices m in B . To be more precise, we will show that for any recurrent graph G for which the potential kernel is well defined (in other words, $\lim_{n \rightarrow \infty} a_{A_n}(x, y) = a(x, y)$ and does not depend on the sequence A_n) for any set $B \subset v(G)$ with $|B| = m$ and $v(G) \setminus B$ connected, we have that (14) holds. The base case $m = 1$ holds trivially.

Let $m \in \mathbb{N}$ and suppose that for any recurrent graph G on which the potential kernel is well defined and for any subset $B \subset v(G)$ with $|B| = m$ and $v(G) \setminus B$ connected we have that (14) and (15) are satisfied for each $x \in B$.

In this case,

$$q_B(w) = \lim_{n \rightarrow \infty} \mathbb{P}_w(T_{A_n} < T_B) \mathcal{R}_{\text{eff}}(B \leftrightarrow A_n)$$

by assumption exists and does not depend on the sequence $(A_n)_{n \geq 1}$, so we also have that $a^{G_B}(\cdot, B) = q_B(\cdot)$ by (9). Remark 2.4 then shows us that $a^{G_B}(\cdot, y)$ is well defined for any $y \in v(G_B)$ and hence we know that the potential kernel is well defined on G_B too.

Induction. Let G be a recurrent graph for which the potential kernel is well defined and let $B \subset v(G)$ be a finite set such that $|B| = m + 1$ and $v(G) \setminus B$ is connected. Fix $x \in B$. We split into two cases, depending on x :

- (i) the removal of x from G disconnects all components of $B \setminus \{x\}$ from infinity in G
or
- (ii) it does not.

We begin with the easy case. Suppose we are in situation (i). We have that for all $w \notin B$ (for n large enough)

$$\mathbb{P}_{A_n}(T_w < T_x) = \mathbb{P}_{A_n}(T_w < T_B) \quad \text{and} \quad \mathcal{R}_{\text{eff}}(w \leftrightarrow B) = \mathcal{R}_{\text{eff}}(w \leftrightarrow x).$$

The limit on the left-hand side exists as the potential kernel is well defined, see Corollary 2.7, and hence $\lim_{n \rightarrow \infty} \mathbb{P}_{A_n}(T_w < T_B) \mathcal{R}_{\text{eff}}(w \leftrightarrow B)$ exists and equals $a(w, x)$. Moreover, we also have

$$q_B(w) = a(w, x) - \mathbb{E}_w[a(X_{T_B}, x)] = a(w, x) - a(x, x) = a(w, x)$$

which proves the result for this choice of x .

We move on to the more interesting case (ii). Since we are not in case (i), we can find a set $B' \subset B$ with $|B'| = m$ and $v(G) \setminus B'$ connected (indeed, since we are not in case (i), there is at least a path going from some vertex in B to infinity, without touching x , and removing from B the last vertex in B visited by this path provides such a set B'). Take y to be the vertex such that $\{y\} = B \setminus B'$.

Since $|B'| = m$, we have by the induction hypothesis that the potential kernel $a^{G_{B'}}(\cdot, \cdot)$ is well defined. Pick $w \in v(G)$ such that $w \notin B$, which we can view also as a vertex in G_B and $G_{B'}$. Fix n so large that both B and w are not in A_n . Using (9) we have that

$$a_{A_n}^{G_{B'}}(w, B') = \mathcal{R}_{\text{eff}}(B' \leftrightarrow A_n) \mathbb{P}_w(T_{A_n} < T_{B'}).$$

We focus on the probability appearing on the right-hand side. By the law of total probability and the strong Markov property of the simple random walk, we have

$$\begin{aligned} \mathbb{P}_w(T_{A_n} < T_{B'}) &= \mathbb{P}_w(T_y < T_{A_n} < T_{B'}) + \mathbb{P}_w(T_{A_n} < T_y \wedge T_{B'}) \\ &= \mathbb{P}_w(T_y < T_{A_n} \wedge T_{B'}) \mathbb{P}_y(T_{A_n} < T_{B'}) + \mathbb{P}_w(T_{A_n} < T_{B'} \wedge T_y). \end{aligned}$$

Since G (and hence $G_{B'}$) is recurrent, we have that $\mathcal{R}_{\text{eff}}(B' \leftrightarrow A_n) \sim \mathcal{R}_{\text{eff}}(x \leftrightarrow A_n) \sim \mathcal{R}_{\text{eff}}(B \leftrightarrow A_n)$ where $a_n \sim b_n$ means $a_n/b_n \rightarrow 1$ as $n \rightarrow \infty$. Taking $n \rightarrow \infty$ in the above identity after multiplying by $\mathcal{R}_{\text{eff}}(B' \leftrightarrow A_n)$ and using once more recurrence, we deduce that

$$a^{G_{B'}}(w, B') = \mathbb{P}_w(T_y < T_{B'}) a^{G_{B'}}(y, B') + \lim_{n \rightarrow \infty} \mathbb{P}_w(T_{A_n} < T_B) \mathcal{R}_{\text{eff}}(A_n \leftrightarrow B),$$

because the potential kernel on $G_{B'}$ is well defined by assumption. This implies in particular that

$$\lim_{n \rightarrow \infty} \mathbb{P}_w(T_{A_n} < T_B) \mathcal{R}_{\text{eff}}(B \leftrightarrow A_n) = a^{G_{B'}}(w, B') - \mathbb{P}_w(T_y < T_{B'}) a^{G_{B'}}(y, B')$$

exists and does not depend on the sequence A_n and, thus, we deduce that $a^{G_B}(w, B)$ is well defined and satisfies

$$a^{G_B}(w, B) = a^{G_{B'}}(w, B') - \mathbb{P}_w(X_{T_B} = y) a^{G_{B'}}(y, B'). \quad (18)$$

We are left to prove that $q_B(w) = a^{G_B}(w, B)$. By the induction hypothesis (because $x \in B'$) we know that

$$a^{G_{B'}}(w, B') = q_{B'}(w) = a(w, x) - \mathbb{E}_w[a(X_{T_{B'}}, x)].$$

Using this in (18) we get

$$\begin{aligned} a^{G_B}(w, B) &= a(w, x) - \mathbb{E}_w[a(X_{T_{B'}}, x)] - \mathbb{P}_w(X_{T_B} = y) \left(a(y, x) - \mathbb{E}_y[a(X_{T_{B'}}, x)] \right) \\ &= a(w, x) - \mathbb{P}_w(X_{T_B} = y) a(y, x) - \sum_{z \in B'} \mathbb{P}_w(X_{T_{B'}} = z) a(z, x) \\ &\quad + \sum_{z \in B'} \mathbb{P}_w(X_{T_B} = y) \mathbb{P}_y(X_{T_{B'}} = z) a(z, x) \\ &= a(w, x) - \sum_{z \in B} \mathbb{P}_w(X_{T_B} = z) a(z, x), \end{aligned}$$

where in the last line we used for $z \in B'$ the equality

$$\mathbb{P}_w(X_{T_B} = z) = \mathbb{P}_w(X_{T_{B'}} = z) - \mathbb{P}_w(X_{T_B} = y)\mathbb{P}_y(X_{T_{B'}} = z),$$

which holds due to the strong Markov property for the random walk. But of course, this is the same as

$$a^{G_B}(w, B) = a(w, x) - \mathbb{E}_w[a(X_{T_B}, x)],$$

so indeed we have that $a^{G_B}(w, B) = q_B(w)$, which finishes the induction argument. \square

Proof of (16)

Let $B \subset v(G)$ be a finite set, such that its removal does not disconnect G . So far, we have shown that the potential kernel is well defined on the graph G_B and hence that the harmonic measure from infinity is well defined, see Corollary 2.13. In this section, we will prove (16); the third statement of Theorem 2.12. First, let us introduce some notation that will only be used here. If G is a graph and $B \subset v(G)$ a (finite) set, then we will write Δ for the Laplacian on G and Δ^{G_B} for the Laplacian on G_B .

Proof of (16). Let G be a recurrent graph on which the potential kernel is well defined, and suppose that $B \subset v(G)$ is a finite set such that $v(G) \setminus B$ is connected. Fix $x \in B$. We split into two cases:

- (i) the removal of x disconnects $B \setminus \{x\}$ from infinity in G or
- (ii) it does not.

In the first case, we have that $\text{hm}_B(x) = 1$ and also that $q_B(w) = a(w, x)$ for all $w \in B$ (indeed, for $w \notin B$ this follows immediately from (15) and for $w \in B \setminus \{x\}$ we have that $q_B(w) = 0 = a(w, x)$ in this case). Hence, we deduce

$$\delta_x(\cdot) = \Delta(a(\cdot, x)) = \Delta(q_B(\cdot)),$$

which shows the result in case (i).

In case (ii), take $B' = B \setminus \{x\}$. We will show that

$$\Delta_w^{G_{B'}}(a^{G_{B'}}(w, B') - \mathbb{P}_w(T_x < T_{B'})a^{G_{B'}}(x, B')) \Big|_{w=x} = \text{hm}_B(x), \quad (19)$$

where $\Delta_u^{G_{B'}}$ is the Laplacian acting on the function with variable u . Let us first explain how this shows the final result. As in (18) and (15) we know that (when q_B is viewed as a function on $v(G_{B'})$)

$$q_B(w) = a^{G_{B'}}(w, B') - \mathbb{P}_w(T_x < T_{B'})a^{G_{B'}}(x, B').$$

Moreover, when $w \in B$, we have

$$q_B(w) = a(w, x) - \mathbb{E}_w[a(X_{T_B}, x)] = a(w, x) - a(w, x) = 0.$$

Hence, actually,

$$(\Delta^{G_{B'}} q_B)(x) = \sum_{\substack{w \sim x \\ w \in v(G_{B'})}} q_B(w) = \sum_{\substack{w \sim x \\ w \in v(G)}} q_B(w) = (\Delta q_B)(x),$$

so that (19) implies the final result.

To prove (19), recall from (5) that³

$$\sum_{\substack{u \sim x \\ u \in v(G_{B'})}} \mathbb{P}_u(T_x < T_{B'}) = \frac{1}{\mathcal{R}_{\text{eff}}(x \leftrightarrow B')},$$

and that $\Delta^{G_{B'}}(a^{G_{B'}}(\cdot, B')) = \delta_{B'}(\cdot)$ by Proposition 2.6. Using these two facts, we get

$$\begin{aligned} \Delta_w^{G_{B'}}(a^{G_{B'}}(w, x_2) - \mathbb{P}_w(T_x < T_{B'})a^{G_{B'}}(x, B')) \Big|_{w=x_1} \\ &= -a^{G_{B'}}(x, B') \sum_{\substack{u \sim x \\ u \in v(G_{B'})}} (\mathbb{P}_u(T_x < T_{B'}) - 1) \\ &= a^{G_{B'}}(x, B') \sum_{\substack{u \sim x \\ u \in v(G_{B'})}} \mathbb{P}_u(T_{B'} < T_x) \\ &= \frac{a^{G_{B'}}(x, B')}{\mathcal{R}_{\text{eff}}(x \leftrightarrow B')} = \text{hm}_{B',x}(x). \end{aligned}$$

The last equality follows from Corollary 2.7, which allows us to write

$$a^{G_{B'}}(x, B') = \text{hm}_{x,B'}(x) \mathcal{R}_{\text{eff}}(x \leftrightarrow B').$$

This shows (16) and therefore concludes the proof of Theorem 2.12. In turn this finishes the proof that (a) is equivalent to (b) in Theorem 2.1 (see e.g. Corollary 2.13). \square

2.4. Proof of Theorem 2.3

Before proceeding with the remaining equivalences we give a proof that (a) holds under the assumption of Theorem 2.3. Recall that a random graph (G, o) is **strictly subdiffusive** whenever there exists a $\beta > 2$ such that

$$\mathbf{E}[d(o, X_n)^\beta] \leq Cn. \quad (\text{SD})$$

We collect the following theorem of [30]. The main theorem from that paper shows that, assuming subdiffusivity, strictly sublinear harmonic functions must be constant.

³Of course, to be precise we would need to calculate the probabilities and effective resistances on the graph $G_{B'}$, but since this makes no difference in the current setting, we skip the extra notation.

In fact, as already mentioned in that paper (see Example 2.10), the arguments in that paper also show that assuming *strict* subdiffusivity, even harmonic functions of at most linear growth must be constant. It is this extension which we use here, and which we quote below.

Theorem 2.16 (Theorem 3 in [30]). *Let (G, o, X_1) be a strictly subdiffusive (SD), recurrent, stationary environment. A.s., every harmonic function on G that is of at most linear growth is constant.*

We now give the proof of Theorem 2.3 using this result.

Proof of Theorem 2.3 assuming Theorem 2.1. Let (G, o) be a unimodular graph that is almost surely strictly subdiffusive (SD) and recurrent, satisfying $\mathbb{E}[\deg(o)] < \infty$. Then degree biasing (G, o) gives a reversible environment and hence, almost surely, all harmonic functions on (G, o, X_1) that are at most linear are constant due to Theorem 2.16. After degree unbiasing, the same statement is true for (G, o) .

We will prove that this implies that statement (a) of Theorem 2.1 holds, which (by assumption) implies (a)–(d) must be satisfied.

Let $a_1, a_2 : v(G)^2 \rightarrow \mathbb{R}_+$ be two potential kernels arising as subsequential limits in the sense of Lemma 2.5. Fix $y \in v(G)$. By Proposition 2.6 we have that $a_i(\cdot, y)$ is of the form

$$a_i(x, y) = \mathcal{R}_{\text{eff}}(x \leftrightarrow y)H_i(x),$$

with $0 \leq H_i(x) \leq 1$ for each x and $i = 1, 2$. Define next the map $h : v(G) \rightarrow \mathbb{R}$ through

$$h(x) = a_1(x, y) - a_2(x, y).$$

Clearly, h is harmonic everywhere outside y by choice of the a_i 's and linearity of the Laplacian. Since $\Delta a_1(\cdot, y) = \Delta a_2(\cdot, y)$ by Proposition 2.6, we also get that $\Delta h(y) = 0$ and we deduce that h is harmonic everywhere.

Next, we notice

$$|h(x)| \leq |H_1(x) - H_2(x)|\mathcal{R}_{\text{eff}}(y \leftrightarrow x) \leq 2d(y, x),$$

implying that h is (at most) linear. Thus h must be constant. Since $h(y) = a_1(y, y) - a_2(y, y) = 0$, it follows that $h(x) = 0$ and hence we finally obtain $a_1(x, y) = a_2(x, y)$ for all $x \in v(G)$. Since $y \in v(G)$ was arbitrary, we deduce the desired result. \square

Remark 2.5. Strict subdiffusivity on the UIPQ was obtained by Benajmini and Curien in the beautiful paper [28]. A result of Lee [121, Theorem 1.10] gives a more general condition which guarantees strict subdiffusivity (essentially, the graph needs to be planar with at least cubic volume growth).

As an example of application of Theorem 2.3 consider the Incipient Infinite percolation Cluster (IIC) of \mathbb{Z}^d for sufficiently large d . By a combination of Theorem 1.2 in [109] and Theorem 1.1 in [122], one can check that the strict subdiffusivity (SD) is satisfied in all sufficiently high dimensions. The recurrence is easier to check. (Note that a weaker form of subdiffusivity can be deduced by combining [109] with [18]). In fact, it was already checked earlier that in high dimensions the backbone of the IIC is one-ended ([166]), implying also the UST is one-ended in this case.

We point out that the result should apply in dimension two (even for non-nearest neighbor walk), or for the IIC of spread-out percolation, although we do not know if strict subdiffusivity has been checked in that case.

2.5. The sublevel set of the potential kernel

Let (G, o) be some recurrent, rooted, graph for which the potential kernel is well defined in the sense that $a_n(x, y)$ obtains a limit and this does not depend on the choice of the sequence $(A_n)_{n \geq 1}$ of finite sets of vertices going to infinity.

Fix $z \in v(G)$ and $R \in \mathbb{R}_+$. Recall the notation in (6) for the ball with respect to the effective resistance metric:

$$\mathcal{B}_{\text{eff}}(z, R) = \{x \in v(G) : \mathcal{R}_{\text{eff}}(z \leftrightarrow x) \leq R\}$$

We also introduce the notation for the sublevel set of $a(\cdot, z)$ through

$$\Lambda_a(z, R) = \{x \in v(G) : a(x, z) \leq R\}.$$

In case $z = o$, we will drop the notation for z and write $\mathcal{B}_{\text{eff}}(R)$, $\Lambda_a(R)$ for $\mathcal{B}_{\text{eff}}(o, R)$, $\Lambda_a(o, R)$ respectively. Although $a(\cdot, \cdot)$ fails to be a distance as it lacks to be symmetric, it *is* what we call a quasi-distance as it does satisfy the triangle inequality due to Proposition 2.9. On 2-connected graphs (where the removal of any single vertex does not disconnect the graph), we have that $a(x, y) = 0$ precisely when $x = y$. In particular, this is true for triangulations.

Let us first explain *why* we care about the sublevel sets of the potential kernel and why we will prefer it over the effective resistance balls. We will call a set $A \subset v(G)$ **simply connected** whenever it is connected (that is, for any two vertices x, y in A , there exists a path connecting x and y , using only vertices inside A) and when removing A from the graph does not disconnect a part of the graph from infinity. We make the following observation, which holds because $x \mapsto a(x, o)$ is harmonic outside of o .

Observation. The set $\Lambda_a(R)$ is simply connected.


This is not true, in general, for $\mathcal{B}_{\text{eff}}(R)$. Introduce the **hull** $\overline{\mathcal{B}_{\text{eff}}(z, R)}$ of $\mathcal{B}_{\text{eff}}(z, R)$ as the set $\mathcal{B}_{\text{eff}}(z, R)$ together with the *finite* components of $v(G) \setminus \mathcal{B}_{\text{eff}}(z, R)$. Even though

$\overline{\mathcal{B}_{\text{eff}}(z, R)}$ does not have any more “holes”, we notice that still, it is not evident (or true in general) that $\overline{\mathcal{B}_{\text{eff}}(z, R)}$ is *connected*. See Figure 2.1a for an example.

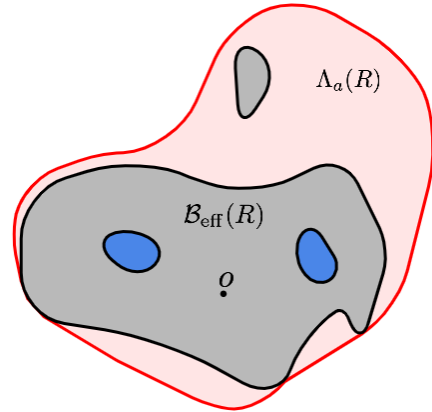
We do notice that $\overline{\mathcal{B}_{\text{eff}}(z, R)} \subset \Lambda_a(z, R)$ as

$$a(x, z) = \text{hm}_{x,z}(x) \mathcal{R}_{\text{eff}}(x \leftrightarrow z) \leq \mathcal{R}_{\text{eff}}(x \leftrightarrow z),$$

by Corollary 2.7. See also Figure 2.1b for a schematic picture.



(a) Example of a graph where $\mathcal{B}_{\text{eff}}(z, R)$ is not connected for each R : the effective resistance between x and y equals $1/2$, whereas the resistance between x and v_i equals $5/8$.



(b) A schematic drawing. In dark gray, we see the set $\mathcal{B}_{\text{eff}}(R)$. The blue parts are $\overline{\mathcal{B}_{\text{eff}}(R)} \setminus \mathcal{B}_{\text{eff}}(R)$. The red area (and everything inside) is then the sublevel set $\Lambda_a(R)$.

We thus get that the sets $\Lambda_a(R)$ are more regular than the sets $\mathcal{B}_{\text{eff}}(R)$ and if G is planar, they correspond to Euclidean simply connected sets.

In this section, we are interested in some properties of $\Lambda_a(R)$, that we will need to prove our Harnack inequalities. We now state the main result, which shows that $\lim_{z \rightarrow \infty} a(z, x) = \infty$, under the additional assumption that the underlying rooted graph is random and (stationary) reversible.

Proposition 2.17. *Suppose (G, o) is a reversible random graph, that is a.s. recurrent and for which the potential kernel is a.s. well defined. Almost surely, the sets $\Lambda_a(z, R)$ are finite for each $R \geq 1$ and all $z \in v(G)$, and hence $(\Lambda_a(z, R))_{R \geq 1}$ defines an exhaustion of G .*

Although we expect this proposition to hold for all graphs where the potential kernel is well defined, we do not manage to prove the general case. In addition, the proof actually yields something slightly stronger which may not necessarily hold in full generality.

Note also that for all $R \geq 0$ we have $v(G) \setminus \Lambda_a(R)$ is non-empty because $x \mapsto a(x, o)$ is unbounded (to see this, assume it is bounded and use recurrence and the optional stopping theorem to deduce that $a(x, o)$ would be identically zero, which is not possible since the Laplacian is nonzero at o). We introduce the following definition, that we will use throughout the remaining document.

Definition 2.18. Let $\delta \in [0, 1]$ and $x \in v(G)$.

- We call x (δ, o) -good if $\text{hm}_{o,x}(x) \geq \delta$. We will omit the notation for the root if it is clear from the context.
- We call the rooted graph (G, o) δ -good if for all $\epsilon > 0$, there exist infinitely many $(\delta - \epsilon, o)$ -good vertices.
- We call the rooted graph (G, o) uniformly δ -good if all vertices are (δ, o) -good.

Note that if the graph (G, o) is uniformly δ -good for some $\delta > 0$, then actually $\Lambda_a(\delta R) \subset \mathcal{B}_{\text{eff}}(R)$, so that the sets $\Lambda_a(\delta R)$ are finite for each R . It turns out that the graph (G, o) being δ -good is also enough, which is the content of Lemma 2.21 below.

Although the definition of δ -goodness is given in terms of rooted graphs (G, o) , the next (deterministic) lemma shows that the definition is actually invariant under the choice of the root, and hence we can omit the root and say “ G is δ -good” instead.

Lemma 2.19. *Suppose $\delta > 0$ is such that (G, o) is δ -good, then also (G, z) is δ -good for each $z \in v(G)$.*

Proof. Fix $z \in v(G)$ and let $\delta > 0$ be such that (G, o) is δ -good. Fix $0 < \epsilon < \delta$ and denote by $G_{\alpha,o}$ the set of (α, o) -good vertices. Take $\epsilon_1, \epsilon_2 > 0$ such that $\epsilon_1 + \epsilon_2 = \epsilon$. Then $G_{\delta-\epsilon_1,o}$ has infinitely many points by assumption.

By Corollary 2.11, we can take $R_0 := R_0(z, o, \epsilon_2)$ so large that for all $x \notin B(o, R_0)$ we have

$$|\text{hm}_{x,z}(x) - \text{hm}_{x,o}(x)| < \epsilon_2.$$

This implies that any vertex $x \in G_{\delta-\epsilon_1,o} \cap B(o, R_0)^c$ must in fact be $(\delta - \epsilon, z)$ -good since $\epsilon = \epsilon_1 + \epsilon_2$. This shows the desired result as ϵ was arbitrary. \square

The next lemma shows the somewhat interesting result that reversible environments are always δ -good, with δ arbitrary close to $\frac{1}{2}$.

Lemma 2.20. *Suppose that (G, o, X_1) is a recurrent reversible random rooted graph (that is a.s. infinite) on which the potential kernel is a.s. well defined. Then a.s. (G, o) is $\frac{1}{2}$ -good.*

Proof. In this proof we will write \mathbf{P}, \mathbf{E} to denote probability respectively expectation with respect to the law of the random rooted graph (G, o) . In compliance with the rest of the document, we will write \mathbb{P}, \mathbb{E} to denote the probability respectively expectation w.r.t. the law of the simple random walk, conditional on (G, o) .

By Lemma 2.19, we note that (G, o) being δ -good is independent of the root and hence for each $\delta > 0$, the event

$$\mathcal{A}_\delta = \{(G, o) \text{ is } \delta\text{-good}\}$$

is invariant under re-rooting, that is

$$(G, o) \in \mathcal{A}_\delta \iff (G, x) \in \mathcal{A}_\delta \text{ for all } x \in v(G).$$

A natural approach to go forward would be to use that any unimodular law is a mixture of ergodic laws [9, Theorem 4.7]. We will not use this, as there is an even simpler argument in this case.

We will use the invariance under re-rooting to prove that \mathcal{A}_δ has probability one. Suppose, to the contrary, that the event \mathcal{A}_δ does not occur with probability one, so that $\mathbf{P}(\mathcal{A}_\delta) \in [0, 1)$. Then we can condition the law \mathbf{P} on \mathcal{A}_δ^c to obtain again a reversible law $\mathbf{P}(\cdot | \mathcal{A}_\delta^c)$ (it is here that we use the invariance under re-rooting of \mathcal{A}_δ , see for example [49, Exercise 15] or [9]), under which \mathcal{A}_δ has probability zero. However, we will show that $\mathbf{P}(\mathcal{A}_\delta) > 0$ always holds when $\delta < \frac{1}{2}$, independent of what the exact underlying reversible law \mathbf{P} is - as long as the potential kernel is a.s. well defined and the graph is a.s. recurrent. Now, this implies that we actually need to have $\mathbf{P}(\mathcal{A}_\delta) = 1$, which is the desired result.

Fix $\delta < \frac{1}{2}$. We thus still need to prove that $\mathbf{P}(\mathcal{A}_\delta) > 0$, which we do by contradiction. Assume henceforth that $\mathbf{P}(\mathcal{A}_\delta) = 0$. By reversibility, we get for each $n \in \mathbb{N}$ the equality

$$\mathbf{E}[\text{hm}_{o, X_n}(X_n)] = \mathbf{E}[\text{hm}_{X_n, o}(o)] = \frac{1}{2},$$

due to the fact that (G, o, X_n) has the same law as (G, X_n, o) , which is reversibility (here, the expectation is both with respect to the environment and the walk).

As $\mathbf{P}(\mathcal{A}_\delta) = 0$, we can assume that a.s. there exists a (random) $N = N(G, o) \in \mathbb{N}$, such that for all $x \notin B(o, N)$ we have

$$\text{hm}_{x, o}(x) \leq \delta$$

Also, note that the environment is a.s. null-recurrent (as is the case for any connected, infinite recurrent graph, which follows e.g. by uniqueness of invariant of measures for recurrent graphs, Theorem 1.7.6 in [137], in conjunction with Theorem 1.7.7 of [137]). Hence we have that, (G, o) -a.s.

$$\mathbb{P}(X_n \text{ in } B(o, N)) \rightarrow 0,$$

whenever $n \rightarrow \infty$. Moreover, notice that for each n we have

$$\mathbb{E}[\text{hm}_{o, X_n}(X_n)] \leq \mathbb{P}(X_n \text{ not in } B(o, N))\delta + \mathbb{P}(X_n \text{ in } B(o, N)).$$

Since $\text{hm}_{o, X_n}(X_n) \in [0, 1]$, we can apply Fatou's lemma (applied to *just* the expectation with respect to the law of (G, o) , so that we can use the just found inequality) from which we deduce that

$$\frac{1}{2} = \limsup_{n \rightarrow \infty} \mathbf{E}[\text{hm}_{o, X_n}(X_n)] \leq \delta,$$

which is a contradiction as $\delta < \frac{1}{2}$. \square

We next show that for any δ -good (rooted) graph, the set $\Lambda_a(R)$ is finite for each $R \geq 1$. Combined with Lemma 2.20, this implies Proposition 2.17 in case of reversible environments. However, Lemma 2.20 shows more than just this fact. Indeed, $\Lambda_a(R)$ being finite need not imply that (G, o) is δ -good for some $\delta > 0$.

Lemma 2.21. *If (G, o) is δ -good for some $\delta > 0$, then $\Lambda_a(o, R)$ is finite for each $R \geq 1$.*

Proof. Let $\delta > 0$ and suppose that G is δ -good. We will show that for each $R \geq 1$, there exists an $M \geq 1$ such that for all $x \notin B(o, M)$ we have

$$a(x, o) \geq \frac{\delta^2 R}{8}.$$

This implies the final result.

By assumption on δ -goodness, for each $R \geq 1$ there exists a vertex $x_R \notin \mathcal{B}_{\text{eff}}(o, R)$ such that

$$\text{hm}_{x_R, o}(x_R) \geq \frac{\delta}{2}$$

This implies by Corollary 2.7 that $a(x_R, o) \geq \frac{\delta}{2} \mathcal{R}_{\text{eff}}(o \leftrightarrow x_R) \geq \frac{\delta R}{2}$.

Fix $R \geq 1$ and define the set $B_R = \{o, x_R\}$. By Theorem 2.12, we get for all x the decomposition

$$a(x, o) = q_{B_R}(x) + \mathbb{E}_x[a(X_{T_{B_R}}, o)],$$

where $q_{B_R}(\cdot)$ is the potential kernel on the graph G_{B_R} , which we recall is the graph G , with B_R glued together. Since potential kernels are non-negative, we can focus our attention to the right-most term.

Take $M = M(o, x_R, \delta)$ so large that for all $x \notin B(o, M)$

$$|\text{hm}_{B_R}(x_R) - \mathbb{P}_x(X_{T_{B_R}} = x_R)| \leq \frac{\delta}{4},$$

which is possible as the potential kernel is well defined, see Proposition 2.6 and Corollary 2.7. We deduce that for all $x \notin B(o, M)$

$$a(x, o) \geq \mathbb{E}_x[a(X_{T_{B_R}}, o)] \geq \frac{\delta^2 R}{8},$$

as desired. \square

2.6. Two Harnack inequalities

We are now ready to prove the equivalence between (c) and (a). The first part of this section deals with a classical Harnack inequality, whereas the second part of this section provides a variation thereof, where the functions might have a single pole. The first Harnack inequality (Theorem 2.24 below) does not involve Theorem 2.1.

Recall that $\Lambda_a(z, R)$ is the sublevel set $\{x \in v(G) : a(x, z) \leq R\}$ (for R not necessarily integer valued) and that $a(z, x)$ defines a quasi distance on G . Also recall the notation $\mathcal{B}_{\text{eff}}(z, R) = \{x : \mathcal{R}_{\text{eff}}(z \leftrightarrow x) \leq R\}$, for the (closed) ball with respect to the effective resistance distance.

2.6.1 The standing assumptions

Throughout this section we will work with deterministic graphs G , which satisfy a certain number of assumptions.

Definition 2.22 (Standing assumptions). We will say that G satisfies the **standing assumptions** whenever it is infinite, recurrent, the potential kernel is well defined and the level sets $(\Lambda_a(z, R))_{R \geq 1}$ are finite for some (hence all by Proposition 2.9) $z \in v(G)$.

We will *not* use that (G, o) is random reversible in this section, other than to verify that it satisfies the standing assumptions 2.22. The remainder of this section works for all (deterministic) graphs that satisfy the standing assumptions.

Lemma 2.23. *Let (G, o) be a random unimodular graph with $\mathbf{E}[\deg(o)] < \infty$, for which a.s. the potential kernel is uniquely defined. Then (G, o) a.s. satisfies the standing assumptions.*

Proof. Proposition 2.17 implies that any unimodular random graph with $\mathbf{E}[\deg(o)] < \infty$ that is a.s. recurrent and for which the potential kernel is a.s. well defined, the level sets $\Lambda_a(z, R)$ are finite for all R and $z \in v(G)$. \square

Remark 2.6. Note for instance that this implies that the UIPT therefore satisfies the standing assumptions.

2.6.2 Elliptic Harnack Inequality

We first show that under the standing assumptions (Definition 2.22), a version of the elliptic Harnack inequality holds, where the constants are uniform over all graphs that satisfy the standing assumptions. Recall the definition of the ‘‘hull’’ $\overline{\mathcal{B}_{\text{eff}}(z, R)}$ introduced in Section 2.5.

Theorem 2.24 (Harnack Inequality). *There exist $M, C > 1$ such that the following holds. Let G be a graph satisfying the Standing Assumptions 2.22. For all $z \in v(G)$, all $R \geq 1$ and all $h : \Lambda_a(z, MR) \cup \partial\Lambda_a(z, MR) \rightarrow \mathbb{R}_+$ that are harmonic on $\Lambda_a(z, MR)$ we have*

$$\max_{x \in \mathcal{B}_{\text{eff}}(z, R)} h(x) \leq C \min_{x \in \mathcal{B}_{\text{eff}}(z, R)} h(x) \quad (\text{H})$$

Remark 2.7. In case the rooted graph (G, o) is in addition uniformly δ -good for some δ (that is, $\text{hm}_{x,o}(x) \geq \delta$ for each x , see Definition 2.18), then we have that

$$\Lambda_a(\delta R) \subset \mathcal{B}_{\text{eff}}(R) \subset \Lambda_a(R),$$

and hence the Harnack inequality above becomes a standard “elliptic Harnack inequality” for the graph equipped with the effective resistance distance. (As will be discussed below, we conjecture that many infinite models of random planar maps, including the UIPT, satisfy the property of being δ -good for some nonrandom $\delta > 0$.)

The harmonic exit measure.

In the proof, we fix the root $o \in v(G)$, but it plays no special role. Define for $k \in \mathbb{N}$, $x \in \Lambda_a(k)$ and $b \in \partial\Lambda_a(k)$ the “harmonic exit measure”

$$\mu_k(x, b) = \mathbb{P}_x(X_{T_k} = b),$$

where T_k is the first hitting time of $\partial\Lambda_a(k)$. We will write

$$\mathbf{G}_k(x, y) := \mathbf{G}_{\Lambda_a(k)^c}(x, y) \quad (20)$$

where we recall the definition of the Green function in (1). The following proposition shows that changing the starting points $x, y \in \mathcal{B}_{\text{eff}}(R)$, does not significantly change the exit measure $\mu_k(\cdot, b)$. The Harnack inequality will follow easily from this proposition (in fact, it is equivalent).

Proposition 2.25. *There exist constants $\tilde{C}, M > 1$ such that for all G satisfying the Standing Assumptions 2.22, all $R \geq 1$ and all $x, y \in \overline{\partial\mathcal{B}_{\text{eff}}(R)}$ we have*

$$\frac{1}{\tilde{C}} \mu_{MR}(y, b) \leq \mu_{MR}(x, b) \leq \tilde{C} \mu_{MR}(y, b)$$

for each $b \in \partial\Lambda_a(MR)$.

We first prove the following lemma, giving an estimate on the number of times the simple random walk started from x visits y , before exiting the set $\Lambda_a(MR)$.

Lemma 2.26. *For all $M_0 > 1$ and all $M \geq M_0 + 3$ there exists $C = C(M, M_0) > 1$ such that for all G satisfying the Standing Assumptions 2.22 and for all $R \geq 1$ we have*

$$\frac{R}{C} \leq \frac{\mathbf{G}_{MR}(x, y)}{\deg(y)} \leq CR$$

for all $x \in \partial\Lambda_a(M_0R)$ and $y \in \overline{\partial\mathcal{B}_{\text{eff}}(R)}$.

Proof. Fix $M_0 > 1$ and let $M \geq M_0 + 3$. Let G be any graph satisfying the standing assumptions 2.22. Let $R \geq 1$, take $x \in \partial\Lambda_a(M_0R)$ and $y \in \overline{\partial\mathcal{B}_{\text{eff}}(R)}$. Notice that, by Lemma 2.8, we can write

$$\frac{\mathbf{G}_{MR}(x, y)}{\deg(y)} = \mathbb{E}_x[a(X_{T_{MR}}, y)] - a(x, y). \quad (21)$$

Let $z \in \Lambda_a(MR)$. Recalling that $a(\cdot, \cdot)$ is a quasi metric that satisfies the triangle inequality due to Proposition 2.9, we have, by assumption on x and y and the expression for the potential kernel in terms of harmonic measure and effective resistance (Corollary 2.7), that

$$a(z, y) \leq a(z, o) + a(o, y) \leq MR + \mathcal{R}_{\text{eff}}(o \leftrightarrow y) = (M + 1)R. \quad (22)$$

Going back to (21) and upper-bounding $-a(x, y) \leq 0$, we find the desired upper bound:

$$\frac{\mathbf{G}_{MR}(x, y)}{\deg(y)} \leq (M + 1)R.$$

For the lower bound, fix again $z \in \partial\Lambda_a(MR)$. From Theorem 2.12 (and the fact that $\mathcal{B}_{\text{eff}}(R) \subset \Lambda_a(R)$) we obtain the equality

$$a(z, y) - \mathbb{E}_z[a(X_{T_R}, y)] = a(z, o) - \mathbb{E}_z[a(X_{T_R}, o)].$$

It follows that

$$a(z, y) \geq a(z, o) - \mathbb{E}_z[a(X_{T_R}, o)] = (M - 1)R. \quad (23)$$

On the other hand, invoking the triangle inequality (as in (22)), we have

$$a(x, y) \leq a(x, o) + a(o, y) \leq (M_0 + 1)R.$$

The lower-bound now follows from (21) and (23) as

$$\frac{\mathbf{G}_{MR}(x, y)}{\deg(y)} \geq (M - 1)R - (M_0 + 1)R = (M - M_0 - 2)R.$$

Since $M \geq M_0 + 3$, we can take $C = C(M, M_0)$ depending *only* on M, M_0 such that we get the result. \square

Proof of Proposition 2.25. Take $M_0 > 1$, $M = M(M_0)$ and $C > 1$ as in Lemma 2.26. Let G be a graph satisfying the standing assumptions 2.22. Fix $R \geq 1$ and let $x, y \in \overline{\partial\mathcal{B}_{\text{eff}}(R)}$. For $b \in \Lambda_a(MR)$ we use the last-exit decomposition (Lemma 2.4) to see

$$\mu_{MR}(x, b) = \sum_{z \in \partial\Lambda_a(M_0R)} \frac{\mathbf{G}_{MR}(x, z)}{\deg(z)} \deg(z) \mathbb{P}_z(X_{T_{MR}} = b; T_{MR} < T_{M_0R}^+).$$

By Lemma 2.26, we have for each $z \in \partial\Lambda_a(M_0R)$

$$\frac{\mathbf{G}_{MR}(z, x)}{\deg(x)} \leq CR \leq C^2 \frac{\mathbf{G}_{MR}(z, y)}{\deg(y)}.$$

We thus get, defining $\tilde{C} = C^2$, and using $\deg(\cdot)$ -reversibility of the simple random walk that

$$\begin{aligned} \mu_{MR}(x, b) &\leq \tilde{C} \sum_{z \in \partial\Lambda_a(M_0R)} \frac{\mathbf{G}_{MR}(y, z)}{\deg(z)} \deg(z) \mathbb{P}_z(X_{T_{MR}} = b; T_{MR} < T_{M_0R}^+) \\ &= \tilde{C} \mu_{MR}(y, b), \end{aligned}$$

showing the final result. \square

Proof of Theorem 2.24. The proof of Theorem 2.24 is easy now. Indeed, let $C, M > 1$ large enough, as in Proposition 2.25 and take any graph G satisfying the standing assumptions and $R \geq 1$. Take $h : \Lambda_a(MR) \cup \partial\Lambda_a(MR) \rightarrow \mathbb{R}_+$ a function harmonic on $\Lambda_a(MR)$. Using the maximum principle for harmonic functions, we deduce that it is enough to prove

$$\max_{x \in \overline{\partial\mathcal{B}_{\text{eff}}(R)}} h(x) \leq C \min_{x \in \overline{\partial\mathcal{B}_{\text{eff}}(R)}} h(x).$$

Take $x, y \in \overline{\partial\mathcal{B}_{\text{eff}}(R)}$. By optional stopping and Proposition 2.25 we have

$$\begin{aligned} h(x) &= \mathbb{E}_x[h(X_{T_{MR}})] = \sum_{b \in \partial\Lambda_a(MR)} h(b) \mu_{MR}(x, b) \\ &\leq \tilde{C} \sum_{b \in \partial\Lambda_a(MR)} h(b) \mu_{MR}(y, b) = \tilde{C} h(y), \end{aligned}$$

showing the result. \square

2.6.3 (a) implies (c): anchored Harnack inequality

Sometimes, one wants to apply a version of the Harnack inequality to functions that are harmonic on a big ball, but not in some vertex inside this ball (the pole). Clearly, we can only hope to compare the value of harmonic function in points that are “far away” from the pole, say on the boundary of a ball centered at the pole.

This “anchored” inequality does not always follow from the Harnack inequality as stated in Theorem 2.24. As an example, think of the graph \mathbb{Z} with nearest neighbor connections. Pick any two positive real numbers α, β satisfying $\alpha + \beta = 1$. Then the function h that maps x to $\alpha(-x)$ when x is negative and to βx when x is positive, is harmonic everywhere outside of 0, with $\Delta h(0) = 1$. This implies that no form of “anchored Harnack inequality” can hold.

We next present a reformulation of (a) implies (c) in Theorem 2.1. We will use it to prove results for the “conditioned random walk” as introduced in Section 2.7.

Theorem 2.27 (Anchored Harnack Inequality). *There exists a $C < \infty$ such that the following holds. Let G be a graph satisfying the Standing Assumptions 2.22. For $z \in v(G)$, $R \geq 1$ and all $h : v(G) \rightarrow \mathbb{R}_+$ that are harmonic outside of z and satisfy $h(z) = 0$, we have*

$$\max_{x \in \partial \Lambda_a(z, R)} h(x) \leq C \min_{x \in \partial \Lambda_a(z, R)} h(x). \quad (\text{aH})$$

Remark 2.8. Actually, we will prove that for each $z \in v(G)$ and $R \geq 1$, there exists $\Psi_z(R) \geq R$ such that for all harmonic functions $h : \Lambda_a(z, \Psi_z(R)) \cup \partial \Lambda_a(z, \Psi_z(R)) \rightarrow \mathbb{R}_+$ that are harmonic on $\Lambda_a(z, \Psi_z(R)) \setminus \{z\}$ and $h(z) = 0$, we have

$$\max_{x \in \partial \Lambda_a(z, R)} h(x) \leq C \min_{x \in \partial \Lambda_a(z, R)} h(x).$$

As before, if the graph is uniformly δ -good for some $\delta > 0$, we can actually take $\Psi_z(R) = MR$ for some $M = M(\delta)$ depending *only* on δ .

Proof of Theorem 2.27

The proof will be somewhat similar to the proof of Theorem 2.24. Again, we will prove it for a given root vertex o to simplify our writing, but it will not matter which vertex we choose. For $k \in \mathbb{N}$, we will write again $T_k = T_{\Lambda_a(k)^c}$ for the first time the random walk exists the sublevel-set $\Lambda_a(k)$. Fix $k \in \mathbb{N}$ and $x \in \Lambda_a(k)$. Define, given a graph G satisfying the standing assumptions 2.22 and a root vertex o the exit measure

$$\nu_k(x, b) = \mathbb{P}_x(X_{T_k} = b, T_k < T_o),$$

for $b \in \partial \Lambda_a(k)$. We begin by showing that, taking x, y in $\Lambda_a(R)$, the exit measures $\nu_k(x, \cdot)$ and $\nu_k(y, \cdot)$ are similar up to division by $a(x, o), a(y, o)$ respectively, when k is large enough. Although it might seem at first slightly counterintuitive that that we need to divide by $a(x, o)$, this actually means that the *conditional* exit measures $\mathbb{P}_w(X_{T_k} = b \mid T_k < T_o)$ for $w = x, y$ are comparable.

Proposition 2.28. *There exists a $C < \infty$ such that for all graphs G satisfying the Standing Assumptions 2.22 with root o , for each $R \geq 1$, there exists a constant $\Psi(R) \geq R$ such that for all $x, y \in \Lambda_a(R) \setminus \Lambda_a(1)$ and all $b \in \partial\Lambda_a(\Psi(R))$ we have*

$$\frac{\nu_{\Psi(R)}(x, b)}{a(x, o)} \leq C \frac{\nu_{\Psi(R)}(y, b)}{a(y, o)}.$$

In order to prove this proposition, we will first prove a few preliminary lemmas. We assume here that the underlying graphs satisfy the standing assumptions 2.22. The next result offers bounds on the probability that the random walk goes “far away” before hitting o in terms of the potential kernel.

Lemma 2.29. *For each $z \in v(G) \setminus \Lambda_a(1)$ and all $M > a(z, o)$, we have*

$$\frac{a(z, o)}{M+1} \leq \mathbb{P}_z(T_M < T_o) \leq \frac{a(z, o)}{M}.$$

Proof. This is a straightforward consequence of the optional stopping theorem. Indeed, since $(a(X_{n \wedge T_M \wedge T_o}, o))_{n \geq 0}$ is an a.s. bounded martingale,

$$a(z, o) = \mathbb{E}_z[a(T_M \wedge T_o, o)]$$

and because $M \leq a(w, o) \leq M+1$ for each $w \in \partial\Lambda_a(M)$ and $a(o, o) = 0$, we find

$$\frac{a(z, o)}{M+1} \leq \mathbb{P}_z(T_M < T_o) \leq \frac{a(z, o)}{M},$$

which are the desired bounds. \square

Lemma 2.30. *For each $R \geq 1$, there exist $M > M_0 > R$ such that for all $x \in \Lambda_a(R)$ and $z \in \Lambda_a(M_0)$,*

$$\frac{1}{10} \leq \frac{\mathbf{G}_{B_M}(z, x)}{\deg(x)a(x, o)} \leq 2,$$

where $B_M = \{o\} \cup \Lambda_a(M)^c$.

Proof. Fix $R \geq 1$ and $x, y \in \Lambda_a(R) \setminus \Lambda_a(0)$. Take $M_0 = M_0(R)$ at least so large that for all $w \notin \Lambda_a(M_0)$ and all $z \in \Lambda_a(R) \setminus \Lambda_a(0)$ we have

$$\frac{1}{2} \leq \frac{\mathbb{P}_w(T_z < T_o)}{\text{hm}_{z,o}(z)} \leq 2. \quad (24)$$

This is possible because $\mathbb{P}_w(T_z < T_o)$ converges to $\text{hm}_{z,o}(z)$ for all z , $\Lambda_a(R)$ is finite, and uniformity in w outside $\Lambda_a(M_0)$ follows just as in Corollary 2.10 (otherwise, we can construct a sequence w_M of vertices going to infinity such that $\mathbb{P}_{w_M}(T_z < T_o)$ does not converge to $\text{hm}_{z,o}(z)$). Fix next $M = 5M_0$ and $B_M = \{o\} \cup \Lambda_a(M)^c$.

Take $z \in \Lambda_a(M_0)$. By choice of M and Lemma 2.29, we have

$$\mathbb{P}_z(T_M < T_o) \leq \frac{M_0}{M} \leq \frac{1}{5}. \quad (25)$$

Using the strong Markov property of the walk we get

$$\mathbf{G}_{B_M}(z, x) = \mathbf{G}_o(z, x) - \mathbb{P}_z(T_M < T_o) \sum_{b \in \partial\Lambda_a(M)} \mathbb{P}_z(X_{T_M} = b \mid T_M < T_o) \mathbf{G}_o(b, x). \quad (26)$$

The definition of the Green function and Corollary 2.7 allow us to write

$$\frac{\mathbf{G}_o(z, x)}{\deg(x)} = \mathbb{P}_z(T_x < T_o) \mathcal{R}_{\text{eff}}(x \leftrightarrow o) \quad \text{and} \quad a(x, o) = \text{hm}_{x,o}(x) \mathcal{R}_{\text{eff}}(x \leftrightarrow o),$$

which implies that

$$\frac{\mathbf{G}_o(z, x)}{\deg(x)a(x, o)} = \frac{\mathbb{P}_z(T_x < T_o)}{\text{hm}_{x,o}(x)} \quad \text{and} \quad \frac{\mathbf{G}_o(b, x)}{\deg(x)a(x, o)} = \frac{\mathbb{P}_b(T_x < T_o)}{\text{hm}_{x,o}(x)},$$

for each $b \in \Lambda_a(M)$. Thus (26) is equivalent to

$$\begin{aligned} \frac{\mathbf{G}_{B_M}(z, x)}{\deg(x)a(x, o)} &= \frac{\mathbb{P}_z(T_x < T_o)}{\text{hm}_{x,o}(x)} \\ &\quad - \mathbb{P}_z(T_M < T_o) \sum_{b \in \partial\Lambda_a(M)} \mathbb{P}_z(X_{T_M} = b \mid T_M < T_o) \frac{\mathbb{P}_b(T_x < T_o)}{\text{hm}_{o,x}(x)}. \end{aligned} \quad (27)$$

Hence, by (25) and using (24) twice with $w = z$ and $w = b$ respectively in (27) we get

$$\frac{1}{2} - \frac{2}{5} \leq \frac{G_{o, \partial M}(z, x)}{\deg(x)a(x, o)} \leq 2,$$

which is the desired result. \square

Proof of Proposition 2.28. Just as in the proof of Proposition 2.25, we use the last-exit decomposition to see

$$\frac{\nu_M(x, b)}{a(x, o)} = \sum_{z \in \partial\Lambda_a(k)} \frac{\mathbf{G}_{o, \partial M}(x, z)}{a(x, o) \deg(z)} \deg(z) \mathbb{P}_z(X_{T_M} = b; T_M < T_k^+).$$

This implies that

$$\frac{\nu_M(x, b)}{a(x, o)} \leq 20 \frac{\nu_M(y, b)}{a(y, o)}.$$

We are left to define $\Psi(R) = M$ and $C = 20$ (which thus does not depend on the graph) to obtain the desired result. \square

Finishing the proof of Theorem 2.27 is now straightforward. Indeed, fix $C > 1$ as in Proposition 2.28. Given a (rooted) graph G satisfying the standing assumptions 2.22, take Ψ also as in Proposition 2.28. Let $R \geq 1$ and $h : \Lambda_a(\Psi(R)) \rightarrow \mathbb{R}_+$ harmonic outside o ; with $h(o) = 0$. Fix $x, y \in \Lambda_a(R)$. By optional stopping, which holds as $\Lambda_a(\Psi(R))$ is finite,

$$\begin{aligned} h(x) &= \int_{\partial\Lambda_a(\Psi(R))} h(b) \nu_{\Psi(R)}(x, b) \\ &\leq C \frac{a(x, o)}{a(y, o)} \int_{\partial\Lambda_a(\Psi(R))} h(b) \nu_{\Psi(R)}(y, b) = C \frac{a(x, o)}{a(y, o)} h(y). \end{aligned}$$

This shows the desired result when $x, y \in \partial\Lambda_a(R)$. \square

2.6.4 (c) implies (a)

Let $(V_R)_R$ be any sequence of connected subsets of $v(G)$ satisfying $o \in V_R \subset V_{R+1}$, $|V_R| < \infty$ for all R and $\cup_{R \geq 1} V_R = v(G)$.

Proposition 2.31. *Suppose that the (rooted) graph (G, o) satisfies the anchored Harnack inequality with respect to the sequence $(V_R)_{R \geq 1}$ and some (non-random) constant C : for all $h : v(G) \rightarrow \mathbb{R}_+$ harmonic outside possibly o and such that $h(o) = 0$,*

$$\max_{x \in \partial V_R} h(x) \leq C \min_{x \in \partial V_R} h(x).$$

In this case, the potential kernel $a(x, o)$ is well defined.

We take some inspiration from [153], although the strategy goes back in fact to a paper of Ancona [11]. Pick some sequence $e = (e_R)_{R \geq 1}$ on $v(G)$ satisfying $e_R \in \partial V_R$.

Lemma 2.32. *Suppose G satisfies the anchored Harnack inequality. Let $R \geq 1$ and suppose that h, g are two positive, harmonic functions on $\Lambda_a(\Psi(R)) \setminus \{o\}$ vanishing at o . We have*

$$\max_{x \in V_R \setminus \{o\}} \frac{h(x)}{g(x)} \leq C^2 \frac{h(e_R)}{g(e_R)}.$$

Proof. Fix $R \geq 1$ and let h, g be as above. Write $T_R = T_{\partial V_R}$. By optional stopping, $h(o) = 0$ and the Harnack inequality, we get

$$h(x) = \mathbb{P}_x(T_R < T_o) \mathbb{E}_x[h(X_{T_R}) \mid T_R < T_o] \leq Ch(e_R) \mathbb{P}_x(T_R < T_o)$$

for all $x \in V_R \setminus \{o\}$. Similarly, we obtain

$$g(x) \geq \frac{1}{C} g(e_R) \mathbb{P}_x(T_R < T_o)$$

for $x \in V_R \setminus \{o\}$. Combining this, we find

$$\frac{1}{C} \frac{h(x)}{h(e_R)} \leq \mathbb{P}_x(T_R < T_o) \leq C \frac{g(x)}{g(e_R)},$$

showing the final result. □

Proof of Proposition 2.31. We follow closely Section 3.2 in [153]. We will show that whenever $h_1, h_2 : v(G) \rightarrow [0, \infty)$ are harmonic functions on $v(G) \setminus \{o\}$, vanishing at o , such that $h_1(e_1) = h_2(e_1)$, we have $h_1 = h_2$. The result then follows as we can pick $h_1(\cdot)$ and $h_2(\cdot)$ to be two subsequential limits of $a_{A_n}(\cdot, o)$ (for possibly different sequences (A_n) going to infinity), and rescaling so that they are equal at e_1 .

Consider $h_1, h_2 : v(G) \rightarrow [0, \infty)$ harmonic functions on $v(G) \setminus \{o\}$, vanishing at o . Assume without loss of generality that $h_1(e_1) = h_2(e_1) = 1$. By Lemma 2.32 we get that there is some appropriate (large) M which does not depend on h_1, h_2 , for which

$$\frac{1}{M} \frac{h_1(x)}{h_1(e_R)} \leq \frac{h_2(x)}{h_2(e_R)} \leq M \frac{h_1(x)}{h_1(e_R)}, \quad (28)$$

for all $x \in V_R$ and $R \geq 1$. It follows that (setting $x = e_1$)

$$\frac{1}{M}h_1(e_R) \leq h_2(e_R) \leq Mh_1(e_R).$$

Using this in (28) and letting $R \rightarrow \infty$, we obtain

$$\frac{1}{M^2} \leq \frac{h_2(x)}{h_1(x)} \leq M^2, \tag{29}$$

for all $x \in v(G) \setminus \{o\}$. Define recursively, for $i \geq 3$,

$$h_i(x) = h_{i-1}(x) + \frac{1}{M^2 - 1}(h_{i-1}(x) - h_1(x)). \tag{30}$$

It is straightforward to check that h_i is non-negative (as follows from an iterated version of (29)) and harmonic outside o . Since M did not depend on h_1, h_2 , and because $h_i(e_1) = 1$ also, we obtain that

$$\frac{1}{M^2} \leq \frac{h_i(x)}{h_1(x)} \leq M^2. \tag{31}$$

On the other hand, it is straightforward to check that the recursion (30) can be solved explicitly to get:

$$h_i(x) = \left(\frac{M^2}{M^2 - 1} \right)^{i-2} (h_2(x) - h_1(x)) + h_1(x).$$

Unless $h_1(x) = h_2(x)$, this grows exponentially, which is incompatible with (31). Therefore $h_1(x) = h_2(x)$. \square

Remark 2.9. The proof above makes it clear that if the potential kernel is uniquely defined (i.e. if (a) holds), then any function $h : v(G) \rightarrow \mathbb{R}_+$ satisfying $\Delta h(x) = 0$ for all $x \in v(G) \setminus \{o\}$ and for which $h(o) = 0$, is of the form $\alpha a(x, o)$ for some $\alpha \geq 0$.

Remark 2.10. If G is reversible, and satisfies the anchored Harnack inequality, then it satisfies (a) as a consequence of the above. It therefore satisfies the standing assumptions: in particular, by Theorem 2.24 holds so it also satisfies the Elliptic Harnack Inequality (EHI). We have therefore proved that anchored Harnack inequality (AHI) \implies (EHI) at least for reversible random graphs, which is not a priori obvious.

2.7. Random Walk conditioned to not hit the root

Let (G, o) be a rooted graph. We will assume **throughout this section** that it satisfies the *standing assumptions* of Definition 2.22, i.e., it is recurrent, the potential kernel is well defined and the potential kernel tends to infinity. In this section, we will define what we call the conditioned random walk (CRW), which is the simple random walk on

G , conditioned to never hit the root o (or any other vertex). Of course, a priori this does not make sense as the event that the simple random walk X will never hit o has probability zero. However, we can take the Doob $a(\cdot, o)$ -transform and use this to define the CRW. We make this precise below.

We apply some of the results derived earlier to answer some basic questions about CRW. For example: is there a connection between the harmonic measure from infinity and the hitting probability of points (and sets)? What is the probability that the CRW will ever hit a given vertex? Do the traces of two independent random walks intersection infinitely often? Does the random walk satisfy a Harnack inequality? Does it satisfy the Liouville property? The answers will turn out to be yes for all of the above, and the majority of this section is devoted to proving such statements. These properties play a crucial role in our proof of one-endedness in the next section.

In a series of papers studying the conditioned random walk ([47, 75, 149], see also the lecture notes by Popov [148]), the following remarkable observation about the CRW $(\hat{X}_t, t \geq 0)$ on \mathbb{Z}^2 was made. Let

$$\hat{q}(y) = \mathbb{P}(\hat{X}_t = y \text{ for some } t \geq 0) = \mathbb{P}(\hat{T}_y < \infty),$$

then $\lim_{y \rightarrow \infty} \hat{q}(y) = 1/2$, even though asymptotically the conditioned walk \hat{X} looks very similar to the unconditioned walk.

One may wonder if such a fact holds in the generality of stationary random graphs for which the potential kernel is well defined. This question was in fact an inspiration for the rest of the paper. Unfortunately, we are not able to answer this question in generality, but believe it should not be true in general. In fact, on random planar maps in the universality class of Liouville quantum gravity with parameter $\gamma \in (0, 2)$ (which includes the CRT-mated maps discussed below), we expect

$$0 < \liminf_{y \rightarrow \infty} \hat{q}(y) < 1/2 < \limsup_{y \rightarrow \infty} \hat{q}(y) < 1, \tag{32}$$

with every possible value in the interval between $\liminf_{y \rightarrow \infty} \hat{q}(y)$ and $\limsup_{y \rightarrow \infty} \hat{q}(y)$ a possible subsequential limit. See also Conjecture 2.50. We will prove the upper-bound of (32) and a form of the lower bound on CRT-mated maps in Theorem 2.49. The fact that every possible value between $\liminf_{y \rightarrow \infty} \hat{q}(y)$ and $\limsup_{y \rightarrow \infty} \hat{q}(y)$ will have a subsequential limit converging to it, holds in general and will be proved in Proposition 2.38.

2.7.1 Definition and first estimates

Instead of the graph distance or effective resistance distance, we will work with the quasi distance $a(x, y)$. Recall the definition $\Lambda_a(y, R) := \{x \in v(G) : a(x, y) \leq R\}$ and $\Lambda_a(R) = \Lambda_a(o, R)$. We will fix $y = o$, but we note that in the random setting, it is of no

importance that we perform our actions on the root (in that setting, everything here is conditional on some realization (G, o)).

We can thus define the **conditioned random walk** (CRW), denoted by \widehat{X} , as the so called Doob h -transform of the simple random walk, with $h(x) = a(x, o)$. To avoid unnecessarily loaded notations, we will in fact denote $a(x) = a(x, o)$ in the rest of this section.

To be precise, let $p(x, y)$ denote the transition kernel of the simple random walk on G . Then the transition kernel of the CRW is defined as

$$\widehat{p}(x, y) = \begin{cases} \frac{a(y)}{a(x)}p(x, y), & x \neq 0 \\ 0, & \text{else} \end{cases}.$$

It is a standard exercise to show that \widehat{p} indeed defines a transition kernel. To include the root o as a possible starting point for the CRW, we will let \widehat{X}_1 have the law $\mathbb{P}_o(\widehat{X}_1 = x) = a(x)$, and then take the law of the CRW afterwards. In this case, we can think of the CRW as the walk conditioned to never return to o .

We now collect some preliminary results, starting with transience, and showing that the walk conditioned to hit a far away region before returning to the origin converges to the conditioned walk, as expected.

We will write \widehat{T}_A for the first hitting time of a set $A \subset v(G)$ by the conditioned random walk, and \widehat{T}_x when $A = \{x\}$. We will also denote $\widehat{T}_R = \widehat{T}_{v(G) \setminus \Lambda_a(R)}$. We recall that $a(\cdot, \cdot)$ satisfies a triangle inequality (see Proposition 2.9) and hence we have the growth condition

$$a(x) \leq a(y) + 1 \tag{33}$$

for two neighboring sites x, y since $a(x, y) \leq 1$ in this case.

Proposition 2.33. *Let $x \in v(G) \setminus \{o\}$ and \widehat{X} the CRW starting from x . Then*

(i) *The walk \widehat{X} is transient.*

(ii) *The process $n \mapsto 1/a(\widehat{X}_{n \wedge \widehat{T}_N})$ is a martingale, where $N = \{y : y \sim o\}$*

Proof. The proof of (ii) is straightforward since $1/a(\widehat{X}_{n \wedge \widehat{T}_N})$ is the Radon–Nikodym derivative of the usual simple random walk with respect to the conditioned walk. (i) then follows from the fact that $a(y) \rightarrow \infty$ along at least a sequence of vertices. Indeed, fix $2 < r < R$ large and $y \in v(G) \setminus \Lambda_a(r + 1)$. By optional stopping (since $1/a(y)$ is bounded)

$$\frac{1}{a(y)} = \mathbb{E}_y \left[\frac{1}{a(\widehat{X}_{\widehat{T}_R \wedge \widehat{T}_r})} \right] \geq \frac{1}{r+1} \mathbb{P}_y(\widehat{T}_r < \widehat{T}_R) + \frac{1}{R+1} \mathbb{P}_y(\widehat{T}_R \leq \widehat{T}_r).$$

Rearranging gives

$$\mathbb{P}_y(\widehat{T}_r < \widehat{T}_R) \leq \frac{\frac{1}{a(y)} - \frac{1}{R+1}}{\frac{1}{r+1} - \frac{1}{R+1}}. \quad (34)$$

Taking $R \rightarrow \infty$, we see that $\mathbb{P}_y(\widehat{T}_r < \infty) \leq (r+1)/(a(y)) < 1$, showing that the chain is transient. \square

We now check (as claimed earlier) that the conditioned walk \widehat{X} can be viewed as a limit of simple random walk conditioned on an appropriate event of positive (but vanishingly small) probability.

Lemma 2.34. *Uniformly over all choices of $m \geq 1$ and paths $\varphi = (\varphi_0, \dots, \varphi_m) \subset \Lambda_a(R)$, as $R \rightarrow \infty$,*

$$\begin{aligned} \mathbb{P}_x((X_0, \dots, X_m) = (\varphi_0, \dots, \varphi_m) \mid T_R < T_o^+) \\ = \mathbb{P}_x((\widehat{X}_0, \dots, \widehat{X}_m) = (\varphi_0, \dots, \varphi_m))(1 + o(1)). \end{aligned}$$

Proof. The proof is similar to [148, Lemma 4.4]. Assume here that $x \neq o$ for simplicity. The proof for $x = o$ follows after splitting into first taking one step and, comparing this, and then do the remainder. Let us first assume that the end point φ_m of φ lies in $\partial\Lambda_a(R)$. Then

$$\mathbb{P}_x((\widehat{X}_0, \dots, \widehat{X}_m) = \varphi) = \frac{a(\varphi_m)}{a(\varphi_0)} \mathbb{P}_x((X_0, \dots, X_m) = \varphi).$$

Since $\varphi_m \in \partial\Lambda_a(R)$, we know that $a(\varphi_m) \in (R, R+1]$ due to (33). By optional stopping, we see

$$a(x) = \mathbb{P}_x(T_R < T_o) \mathbb{E}_x[a(X_{T_R}) \mid T_R < T_o],$$

and also $a(X_{T_R}) \in (R, R+1]$. We thus find that

$$\mathbb{P}_x(T_R < T_o) = \frac{a(x)}{R} (1 + o_R(1)). \quad (35)$$

Combining this, we get

$$\mathbb{P}_x((X_0, \dots, X_m) = \varphi \mid T_R < T_o) = \frac{\mathbb{P}_x((X_0, \dots, X_m) = \varphi)}{a(x)} R(1 + o(1)).$$

Now let φ be an arbitrary path in $\Lambda_a(R)$ starting from x , then by the Markov property,

$$\begin{aligned} \mathbb{P}_x((X_0, \dots, X_m) = \varphi \mid T_R < T_o) &= \mathbb{P}_x((X_0, \dots, X_m) = \varphi) \mathbb{P}_{\varphi_m}(T_R < T_o) / \mathbb{P}_x(T_R < T_o) \\ &= \mathbb{P}_x((X_0, \dots, X_m) = \varphi) a(\varphi_m) / a(x) (1 + o(1)) \\ &= \mathbb{P}_x((\widehat{X}_0, \dots, \widehat{X}_m) = \varphi) (1 + o(1)), \end{aligned}$$

as desired. \square

The Green Function

We can find an explicit expression for the Green function associated to \widehat{X} . To that end, we define for $x, y \in v(G) \setminus \{o\}$

$$\widehat{\mathbf{G}}(x, y) = \mathbb{E}_x \left[\sum_{n=0}^{\infty} \mathbb{1}_{\widehat{X}_n=y} \right],$$

which is well defined as \widehat{X} is transient (also, the well-definition would follow from the proof below, which provides yet another way to see that the CRW is transient).

Proposition 2.35. *Let $x, y \in v(G) \setminus \{o\}$. Then*

$$\frac{\widehat{\mathbf{G}}(x, y)}{\deg(y)} = \frac{a(y, o)}{a(x, o)} \frac{\mathbf{G}_o(x, y)}{\deg(y)} = \frac{a(y, o)}{a(x, o)} (a(x, o) - a(x, y) + a(o, y)).$$

Proof. Fix $x, y \in v(G) \setminus \{o\}$. For definiteness we take the exhaustion $\Lambda_a(R)$ of G here, but we need not to, any exhaustion would work. Define for $R \geq 1$ the truncated Green function:

$$\widehat{\mathbf{G}}_R(x, y) := \mathbb{E}_x \left[\sum_{n=0}^{\widehat{T}_R-1} \mathbb{1}_{\widehat{X}_n=y} \right].$$

We denote $A_R = (\Lambda_a(R))^c \cup \{o\}$ and will show that

$$\widehat{\mathbf{G}}_R(x, y) = \frac{a(y)}{a(x)} \mathbf{G}_{A_R}(x, y), \quad (36)$$

from which the result follows when R goes to infinity. Fix $R \geq 1$ and notice the following standard equality, which follows from the Markov property of the CRW:

$$\widehat{\mathbf{G}}_R(x, y) = \frac{\mathbb{P}_x(\widehat{T}_y < \widehat{T}_R)}{\mathbb{P}_y(\widehat{T}_y^+ < \widehat{T}_R)}$$

We first deal with the numerator. From the definition of the CRW we get

$$\mathbb{P}_x(\widehat{T}_y < \widehat{T}_R) = \frac{a(y)}{a(x)} \mathbb{P}_x(T_y < T_R \wedge T_o). \quad (37)$$

Indeed, just sum over all paths φ taking x to y , and which stay inside $\Lambda_a(R) \setminus \{o\}$. Then each path has as endpoint y , and the probability that the simple random walk will take any of these paths is nothing but $\mathbb{P}_x(T_y < T_R \wedge T_o)$.

We can deal with the denominator in a similar fashion, only this time we note that the beginning and end point are the same. Hence, the $a(y)$ -terms cancel and we get

$$\widehat{\mathbf{G}}_R(x, y) = \frac{a(y)}{a(x)} \frac{\mathbb{P}_x(T_y < T_o \wedge T_R)}{\mathbb{P}_y(T_y^+ < T_o \wedge T_R)} = \frac{a(y)}{a(x)} \mathbf{G}_{A_R}(x, y).$$

This shows the first equality appearing in Proposition 2.35 upon taking $R \rightarrow \infty$. The second statement follows from Proposition 2.9. \square

2.7.2 Hitting probabilities for conditioned walk

Suppose \widehat{X} and \widehat{Y} are two independent CRW's. We will begin by describing hitting probabilities of points and sets and use this to prove that the traces of \widehat{X} and \widehat{Y} intersect infinitely often a.s.

We begin giving a description of the hitting probability of a vertex y by the CRW started from x . Although it is a rather straightforward consequence of the expression for the Green function of the CRW, it is still remarkably clean.

Lemma 2.36. *Let $x, y \in v(G) \setminus \{o\}$, then*

$$\mathbb{P}_x(\widehat{T}_y < \infty) = \frac{\text{hm}_{y,o}(y)\mathbb{P}_y(T_x < T_o)}{\text{hm}_{x,o}(x)}.$$

Proof. Note that for $x \neq y$ we have

$$\widehat{\mathbf{G}}(x, y) = \mathbb{P}_x(\widehat{T}_y < \infty)\widehat{\mathbf{G}}(y, y),$$

so that by Proposition 2.35 and Corollary 2.7 we find

$$\begin{aligned} \mathbb{P}_x(\widehat{T}_y < \infty) &= \frac{\widehat{\mathbf{G}}(x, y)}{\widehat{\mathbf{G}}(y, y)} = \frac{a(y, o)}{a(x, o)} \frac{\frac{\mathbf{G}_o(x, y)}{\deg(y)}}{a(y, o) + a(o, y)} \\ &= \frac{\text{hm}_{y,o}(y)\mathcal{R}_{\text{eff}}(o \leftrightarrow y) \frac{\mathbf{G}_o(x, y)}{\deg(y)}}{\text{hm}_{x,o}(x)\mathcal{R}_{\text{eff}}(o \leftrightarrow x)\mathcal{R}_{\text{eff}}(o \leftrightarrow y)} \\ &= \frac{\text{hm}_{y,o}(y)\mathbb{P}_y(T_x < T_o)}{\text{hm}_{x,o}(x)}, \end{aligned}$$

as desired. □

Since the potential kernel is assumed to be well defined, we also have that $\mathbb{P}_y(T_x < T_o) \rightarrow \text{hm}_{o,x}(x)$ as $y \rightarrow \infty$ due to Corollary 2.7, and hence we deduce immediately the next result.

Corollary 2.37. *Write $\widehat{q}(y) = \mathbb{P}_o(\widehat{T}_y < \infty)$. We have that*

$$\liminf_{y \rightarrow \infty} \widehat{q}(y) = \liminf_{y \rightarrow \infty} \text{hm}_{o,y}(y)$$

and the same with 'limsup' instead of 'liminf'.

In particular, it is true that on *transitive* graphs that are recurrent and for which the potential kernel is well defined, by symmetry one always has $\widehat{q}(y) \rightarrow \frac{1}{2}$. This gives another proof to a result of [148] on the square lattice once it has been established that the potential kernel is uniquely defined. There are multiple ways to show the latter, including using the tools from this paper, e.g., by proving an anchored Harnack inequality

as in Corollary 2.2, or by showing that sublinear harmonic functions are constant, and showing that the effective resistance grows sublinearly (in fact logarithmically).

We can now prove that the subsequential limits of the hitting probabilities $\widehat{q}(y)$ define an interval, as promised before. Note that this proposition is fairly general: it does not require the underlying graphs to be unimodular, only for the graph to satisfy the standing assumption (Definition 2.22, i.e. recurrence, existence of potential kernel and convergence to infinity of the potential kernel).

Proposition 2.38. *Let $o \in V$ be fixed and $\widehat{q}(y) = \mathbb{P}_o(\widehat{T}_y < \infty)$.*

For each $q \in [\liminf_{y \rightarrow \infty} \widehat{q}(y), \limsup_{y \rightarrow \infty} \widehat{q}(y)]$, there exists a sequence of vertices $(y_n)_{n \geq 1}$ going to infinity such that

$$\lim_{n \rightarrow \infty} \widehat{q}(y_n) = q.$$

Proof. Assume that there exist $q_1 < q_2$ such that there are sequences $(y_n^1)_{n \geq 1}$ and $(y_n^2)_{n \geq 1}$ going to infinity for which $\lim_{n \rightarrow \infty} \widehat{q}(y_n^i) = q_i$, but there does not exist a sequence y_n going to infinity for which $q_1 < \lim_{n \rightarrow \infty} \widehat{q}(y_n) < q_2$. We will derive a contradiction. We do so via the following claim.

Claim. *For each $\epsilon > 0$, there exists an $N = N(G, o, \epsilon)$ such that for each neighboring vertices $x, y \notin B(o, N)$, we have*

$$|\widehat{q}(x) - \widehat{q}(y)| < \epsilon.$$

To see this claim is true, we use Lemma 2.36 and Corollary 2.11 to get the existence of N_1 such that

$$|\widehat{q}(z) - \text{hm}_{o,z}(z)| < \frac{\epsilon}{4} \tag{38}$$

for all $z \notin B(o, N_1)$. Next, pick N_2 such that all $z \notin B(o, N_2)$ have $\mathcal{R}_{\text{eff}}(o \leftrightarrow z) > \frac{4}{\epsilon}$. Let $x, y \notin B(o, N_1 \vee N_2)$ be neighbors. Due to (33) we have $a(x) - a(y) \leq 1$ and by the triangle inequality for effective resistance also $\mathcal{R}_{\text{eff}}(o \leftrightarrow y) \leq \mathcal{R}_{\text{eff}}(o \leftrightarrow x) + 1$. Hence, using the expression $a(x) = \text{hm}_{x,o}(x) \mathcal{R}_{\text{eff}}(o \leftrightarrow x)$ of Corollary 2.7, we deduce that

$$\text{hm}_{x,o}(x) \frac{\mathcal{R}_{\text{eff}}(o \leftrightarrow y) - 1}{\mathcal{R}_{\text{eff}}(o \leftrightarrow y)} - \text{hm}_{y,o}(y) \leq \frac{a(x) - a(y)}{\mathcal{R}_{\text{eff}}(y \leftrightarrow o)} \leq \frac{1}{\mathcal{R}_{\text{eff}}(o \leftrightarrow y)},$$

which implies by choice of N_2 that in fact

$$\text{hm}_{x,o}(x) - \text{hm}_{o,y}(y) < \frac{\epsilon}{2}. \tag{39}$$

Thus, taking together equations (38) and (39) we obtain

$$\widehat{q}(x) - \widehat{q}(y) \leq \epsilon.$$

Since x, y are arbitrary neighbors, this implies the claim when taking $N = N_1 \vee N_2$.

By Corollary 2.14, we know that the graph G is one-ended as the potential kernel is assumed to be well defined. Take $\epsilon > 0$ so small that $q_2 > q_1 + 3\epsilon$. By assumption on q_1, q_2 , we thus have that for each n large enough, there exist two neighboring vertices $x, y \notin B(o, n)$ satisfying

$$\widehat{q}(y) > q_2 - \epsilon > q_1 + 2\epsilon > \widehat{q}(x) + \epsilon,$$

so that $\widehat{q}(y) > \widehat{q}(x) + \epsilon$, a contradiction. \square

2.7.3 Harnack inequality for conditioned walk

Notice that the conditioned random walk viewed as a Doob h -transform may be viewed as a random walk on the original graph G but with new conductances by

$$\widehat{c}(x, y) = a(x)a(y)$$

for each edge $\{x, y\} \in e(G)$. Indeed the symmetry of this function is obvious, as is non-negativity, and since a is harmonic for the original graph Laplacian Δ ,

$$\pi(x) := \sum_{y \sim x} \widehat{c}(x, y) = \sum_{y \sim x} a(y)a(x) = \deg(x)a(x)^2,$$

we get that the random walk associated with these conductances coincides indeed with our Doob h -transform description of the conditioned walk.

We can thus consider the network (G, \widehat{c}) , which is transient by Proposition 2.33. It will be useful to consider the graph Laplacian $\widehat{\Delta}$, associated with these conductances, defined by setting

$$(\widehat{\Delta}h)(x) = \sum_{y \sim x} \widehat{c}(x, y)(h(y) - h(x)).$$

for a function h defined on the vertices of G , although h does not need to be defined at o . We will say that a function $h : v(G) \setminus \{o\} \rightarrow \mathbb{R}$ is harmonic (w.r.t. the network (G, \widehat{c})) whenever $\widehat{\Delta}h \equiv 0$. This is of course equivalent to

$$h(x) = \mathbb{E}_x[h(\widehat{X}_1)]$$

for each $x \in v(G) \setminus \{o\}$.

It might be of little surprise that the anchored Harnack inequality (Theorem 2.27) implies (in fact, it is equivalent but this will not be needed) to an elliptic Harnack inequality on the graph G with conductance function \widehat{c} , at least when viewed from the root (i.e., for exhaustion sequences centered on the root o).

Proposition 2.39. *There exists a $C > 1$ such that the following holds. Suppose the graph G satisfies the standing assumptions. Let $\widehat{h} : v(G) \setminus \{o\} \rightarrow \mathbb{R}_+$ be harmonic with respect to (G, \widehat{c}) . Then for each $R \geq 1$,*

$$\max_{x \in \partial\Lambda_a(R)} \widehat{h}(x) \leq C \min_{x \in \partial\Lambda_a(R)} \widehat{h}(x).$$

Alternatively, the max and the min could be taken over $\Lambda_a(R)$ instead of $\partial\Lambda_a(R)$.

Proof. Since the graph follows the standing assumptions it satisfies the anchored Harnack inequality of Theorem 2.27. Furthermore, $\hat{h}(x)$ is $\hat{\Delta}$ -harmonic if and only if

$$h(x) = \begin{cases} a(x)\hat{h}(x) & \text{if } x \neq o \\ 0 & \text{if } x = o \end{cases}$$

is harmonic for Δ away from o . Thus we can apply Theorem 2.27 to it at $z = o$. Since also $|a(x) - R| \leq 1$ for $x \in \partial\Lambda_a(R)$, this anchored Harnack inequality implies the anchored Harnack inequality for \hat{h} immediately. To obtain the corresponding inequality where the extrema are taken on $\Lambda_a(R)$, we use the maximum principle (see Section 2.1 in [131]) with respect to the \hat{c} conductances; note that these extrema may not be attained at o . \square

As a corollary we obtain the Liouville property for $\hat{X}: (G, \hat{c})$ does not carry any non-constant, bounded harmonic functions. This implies in turn that the invariant σ -algebra \mathcal{I} of the CRW is trivial.

Corollary 2.40. *The network (G, \hat{c}) satisfies the Liouville property, that is: any function $h : v(G) \setminus \{o\} \rightarrow \mathbb{R}$ that is harmonic and bounded must be constant.*

Proof. Let h be a bounded, harmonic function with respect to (G, \hat{c}) . Define the function

$$\hat{h} = h - \inf_{x \in v(G)} h(x),$$

which is non-negative and harmonic. Moreover, for each $\epsilon > 0$, there exists an x_ϵ such that $\hat{h}(x_\epsilon) \leq \epsilon$. Take R_ϵ so large that $x_\epsilon \in \Lambda_a(R_\epsilon)$. By the Harnack inequality (Proposition 2.39) we deduce that for all $x \in \Lambda_a(R_\epsilon)$,

$$0 \leq \hat{h}(x) \leq C\hat{h}(x_\epsilon) \leq C\epsilon.$$

Since ϵ is arbitrary, and C does not depend on R_ϵ nor ϵ , this shows the desired result. \square

2.7.4 Recurrence of sets

We will say that a set $A \subset v(G)$ is recurrent for the chain \hat{X} whenever there exist $x \in v(G)$ such that

$$\mathbb{P}_x(\hat{X}_n \in A \text{ i.o.}) = 1,$$

where i.o. is short-hand for ‘infinitely often’. Since (G, \hat{c}) satisfies the Liouville property, such probabilities are 0 or 1, hence the definition of A being recurrent is independent of the choice of x . If a set is not recurrent, it is called transient. Since \hat{X} is transient, any

finite set A is transient too. Notice, by the way, that the definition above is equivalent to saying that A is recurrent whenever $\mathbb{P}_x(\widehat{T}_A < \infty) = 1$ for all $x \in v(G)$.

We capture next some results, relating recurrence and transience of sets to the harmonic measure from infinity. Recall Definition 2.18 of δ -good points: x is δ -good whenever $\text{hm}_{x,o}(x) \geq \delta$.

Lemma 2.41. *If A has infinitely many δ -good points for some $\delta > 0$, then A is recurrent for \widehat{X} .*

Proof. This follows from a Borel-Cantelli argument. Indeed, fix $x \in v(G)$. Let δ be as in the assumption. Take $(g_i)_{i=1}^\infty$ a sequence of δ -good points in A , with $a(g_i) > i$ (which we can clearly find as $\Lambda_a(i)$ is finite whereas A has infinitely many good points).

We will define two sequences $(R_i)_{i \geq 1}$ and $(M_i)_{i \geq 1}$. Set $M_0 = 0$ and $R_0 = 0$. Suppose we have defined R_i, M_{i-1} already. Set $a_i = a(g_{R_i})$ and $\Lambda_i = \Lambda_a(a_i)$, and note that by definition $g_{R_i} \in \Lambda_i$. Take M_i so large (and greater than R_i) that

$$\mathbb{P}_z(\widehat{X} \text{ ever hits } \Lambda_i) \leq \frac{\delta}{4}, \text{ uniformly over } z \in \Lambda_a(M_i)^c \quad (40)$$

This is possible since Λ_i is finite and \widehat{X} is transient by Proposition 2.33 and more precisely the hitting probabilities of a finite set converge to zero (see (34)). Next, let R_{i+1} be so large (and greater than M_i) that

$$\frac{\mathbb{P}_y(T_x < T_o)}{\text{hm}_{o,x}(x)} \geq 1/2, \quad (41)$$

for $y = g_{R_{i+1}}$ and all $x \in \Lambda_a(M_i)$. This is possible because $\Lambda_a(M_i)$ is finite and hitting probabilities converge to harmonic measure from infinity, by Corollary 2.7. We can also require without loss of generality that $g_{R_{i+1}} \in \Lambda_a(M_i)^c$.

Suppose that $x \in \Lambda_a(M_{i-1})$ is arbitrary. We first claim that from x it is reasonably likely that the conditioned walk \widehat{X} will hit $y = g_{R_i}$. Indeed, note that by Lemma 2.36, and since y is δ -good and (41) holds,

$$\mathbb{P}_x(\widehat{T}_y < \infty) = \text{hm}_{y,o}(y) \frac{\mathbb{P}_y(T_x < T_o)}{\text{hm}_{x,o}(x)} \geq \delta/2$$

On the other hand, conditionally on hitting $y = g_{R_i}$, the conditioned walk \widehat{X} is very likely to do so before exiting $\Lambda_a(M_i)$ (let us call τ_i this time). Indeed, by the strong Markov property at τ_i and (40),

$$\mathbb{P}_x(\widehat{T}_y > \tau_i, \widehat{T}_y < \infty) \leq \sup_{z \in \Lambda_a(M_i)^c} \mathbb{P}_z(\widehat{X} \text{ ever hits } \Lambda_i) \leq \delta/4.$$

Therefore,

$$\mathbb{P}_x(\widehat{T}_y < \tau_i) \geq \frac{\delta}{2} - \frac{\delta}{4} = \frac{\delta}{4}.$$

Let E_i be the above event, i.e., $E_i = \{\widehat{T}_{g_{R_i}} < \tau_i\}$. Since $x \in \Lambda_a(M_{i-1})$ in the above lower bound is arbitrary, it follows from the strong Markov property at time τ_{i-1} that $\mathbb{P}(E_i | \mathcal{F}_{\tau_{i-1}}) \geq \delta/4$, where $(\mathcal{F}_n)_{n \geq 0}$ is the filtration of the conditional walk. By Borel–Cantelli we conclude immediately that E_i occurs infinitely often a.s. (for the conditioned walk), which concludes the proof. \square

2.7.5 Infinite intersection of two conditioned walks

We finish this section by showing that two independent conditioned random walks have traces that intersect infinitely often (for simplicity here the CRW’s are conditioned to not hit the same root o). We manage to prove this under two (different) additional assumptions. We start by adding the assumption that (G, o) is random and reversible.

Proposition 2.42. *Suppose that (G, o) is a reversible random graph, such that a.s. it is recurrent and a.s. the potential kernel is well defined. Let \widehat{X}, \widehat{Y} be two independent CRW’s started from $x, y \in v(G)$ respectively, avoiding o . Then a.s.*

$$\mathbb{P}(|\{\widehat{X}_n : n \in \mathbb{N}\} \cap \{\widehat{Y}_n : n \in \mathbb{N}\}| = \infty) = 1.$$

Proof. Suppose that (G, o) has infinitely many $\frac{1}{3}$ -good vertices, and call the set of such vertices $A := A(G, o)$. Since there are various sources of randomness here, it is useful to recall that \mathbb{P} the underlying probability measure \mathbb{P} is always conditional on the rooted graph (G, o) . Then by Lemma 2.41, we know that

$$\mathbb{P}(|\{\widehat{X}_n : n \in \mathbb{N}\} \cap A| = \infty) = 1.$$

Now, consider the set $B = \{\widehat{X}_n : n \in \mathbb{N}\} \cap A$. By definition, every point in B is $\frac{1}{3}$ -good. Since \widehat{Y} is independent of \widehat{X} (when conditioned on (G, o)), we can use Lemma 2.41 again to see that on an event of \mathbb{P} -probability 1,

$$\mathbb{P}(|\{\widehat{Y}_n : n \in \mathbb{N}\} \cap B| = \infty | \widehat{X}) = 1$$

Taking expectation w.r.t. \widehat{X} we deduce that the traces of \widehat{X} and \widehat{Y} intersect infinitely often \mathbb{P} -almost surely, conditioned on (G, o) having infinitely many $\frac{1}{3}$ -good vertices. However, Lemma 2.20 implies that, under our assumptions on (G, o) , this happens with \mathbb{P} -probability one, showing the desired result. \square

A consequence of the infinite intersection property is that the (random) network (G, \widehat{c}) is a.s. Liouville. Therefore we get a new proof of the already obtained (in Corollary 2.40) Liouville property for the conditioned walk, but this time without using the Harnack inequality. On the other hand, [29] proved that for planar graphs, the Liouville property is in fact equivalent to the infinite intersection property and this results extends without any additional arguments to the case of planar networks.

By Proposition 2.39 and Corollary 2.40 we thus also obtain as a corollary of [29] the infinite intersection property for planar networks such that the potential kernel tends to infinity.

Proposition 2.43. *Suppose G is a (not necessarily random reversible) planar graph satisfying the standing assumptions. Let \widehat{X} and \widehat{Y} be two independent CRW's avoiding o , started from $x, y \in v(G)$ respectively. Then*

$$\mathbb{P}(|\{\widehat{X}_n : n \in \mathbb{N}\} \cap \{\widehat{Y}_n : n \in \mathbb{N}\}| = \infty) = 1.$$

Remark 2.11. It will be useful for us to recall that the infinite intersection property implies that one walk intersects the loop-erasure of the other:

$$\mathbb{P}(|\{\text{LE}(\widehat{X})_n : n \in \mathbb{N}\} \cap \{\widehat{Y}_n : n \in \mathbb{N}\}| = \infty) = 1,$$

where $\text{LE}(\widehat{X})$ is the Loop Erasure of \widehat{X} and \widehat{X}, \widehat{Y} are two independent CRW's that don't hit the root o , started from x, y respectively. See [132] for this result.

2.8. (a) implies (d): One-endedness of the uniform spanning tree

In this section we show that the uniform spanning tree is one ended, provided that the underlying graph is unimodular. In particular, we prove that (a) implies (d) in Theorem 2.1.

Theorem 2.44. *Suppose that (G, o) is a reversible, recurrent graph for which the potential kernel is a.s. well defined and such that $a(x) \rightarrow \infty$ along any sequence $x \rightarrow \infty$. Then the uniform spanning tree is one-ended almost surely.*

Before proving this theorem, we start with a few preparatory lemmas. We will write \mathcal{T} to denote the uniform spanning tree and begin by recalling the following “path reversal” for the simple random walk, a standard result. In what follows, fix the vertex $o \in v(G)$, but it plays no particular role other than to simplify the notation.

Lemma 2.45 (Path reversal). *Let $o, u \in v(G)$. For any subset of paths \mathcal{P}*

$$\mathbb{P}_u((X_n : n \leq T_o) \in \mathcal{P} \mid T_o < T_u^+) = \mathbb{P}_o((X_n : n \leq T_u) \in \mathcal{P}' \mid T_u < T_o^+),$$

where a path $\varphi \in \mathcal{P}'$ if and only if the reversal of the path is in \mathcal{P} .

See Exercise (2.1d) in [131]. The next result says that the random walk started from o and stopped when hitting u , conditioned to hit u before returning to o looks locally like a conditioned random walk when u is far away. This is an extension of Lemma 2.34 and its proof is similar.

2.8. (A) IMPLIES (D): ONE-ENDEDNESS OF THE UNIFORM SPANNING TREE

Lemma 2.46. *For each $M \in \mathbb{N}$ and $\epsilon > 0$, there exists an L such that for all $u \notin \Lambda_a(L)$ and uniformly over all paths φ going from o to $\partial\Lambda_a(M)$,*

$$\mathbb{P}_o((X_0, \dots, X_{T_M}) = \varphi \mid T_u < T_o^+) = \mathbb{P}_o((\widehat{X}_0, \dots, \widehat{X}_{T_M}) = \varphi) \pm \epsilon.$$

Proof. Fix $M \in \mathbb{N}$ and $\epsilon > 0$. Let φ be some path o to $\Lambda_a(M)$ not returning to o . Denote by $\varphi_{end} \in \partial\Lambda_a(M)$ the endpoint of such a path. By the Markov property for the simple walk

$$\mathbb{P}_o((X_o, \dots, X_{T_M}) = \varphi, T_u < T_o^+) = \mathbb{P}_o((X_o, \dots, X_{T_M}) = \varphi) \mathbb{P}_{\varphi_{end}}(T_u < T_o).$$

Now, take L so large that uniformly over $x \in \Lambda_a(M)$ with $x \neq o$,

$$\frac{\mathbb{P}_x(T_u < T_o)}{\deg(o) \mathbb{P}_o(T_u < T_o^+)} = a(x) \pm \epsilon$$

By definition, we have that

$$\mathbb{P}_o(X_1 = \varphi_1) = \frac{1}{\deg(o)},$$

yet $\mathbb{P}_o(\widehat{X}_1 = \varphi_1) = a(\varphi_1)$. Therefore, and by definition of the h -transform,

$$\mathbb{P}_o((X_o, \dots, X_{T_M}) = \varphi, T_u < T_o^+) = \mathbb{P}_o((\widehat{X}_o, \dots, \widehat{X}_{T_M}) = \varphi) \frac{1}{\deg(o) a(\varphi_{end})} \mathbb{P}_{\varphi_{end}}(T_u < T_o),$$

so that after dividing both sides through $\mathbb{P}_o(T_u < T_o^+)$, we have

$$\mathbb{P}_o((X_o, \dots, X_{T_M}) = \varphi \mid T_u < T_o^+) = \mathbb{P}_o((\widehat{X}_o, \dots, \widehat{X}_{T_M}) = \varphi) \pm \epsilon$$

as desired. \square

We will say that the graph satisfies an infinite intersection property for the CRW whenever

$$\mathbb{P}(|\{\widehat{X}_n : n \in \mathbb{N}\} \cap \{\text{LE}(\widehat{Y})_n : n \in \mathbb{N}\}| = \infty) = 1 \quad (\text{cIP})$$

where \widehat{X} and \widehat{Y} are independent.

Next, under the assumption (cIP) it holds that as $u \rightarrow \infty$, a simple random walk started at u is very unlikely to hit $\text{LE}(\widehat{Y})$ in o . This is the key property which gives one-endedness of the UST.

Lemma 2.47. *Suppose (cIP) holds, then*

$$\limsup_{M \rightarrow \infty} \sup_{u \notin \Lambda_a(M)} \mathbb{P}_u(X_{T_{\text{LE}(\widehat{Y})}} = o) = 0,$$

where X is a simple random walk started at u and \widehat{Y} an independent conditioned walk started at o .

Proof. Let A be any simple path from o to infinity in G . Then

$$\mathbb{P}_u(X_{T_A} = o) \leq \mathbb{P}_u(\{X_n : n \leq T_o\} \cap A = \{o\} \mid T_o < T_u^+). \quad (42)$$

To see this, it is useful to recall that the successive excursions (or loops) from u to u forms a sequence (Z_1, Z_2, \dots) of i.i.d. paths (with a.s. finite length). Let N be the index of the first excursion which touches o . Then the law of Z_N , up to its hitting time of o , is that of $\mathbb{P}_u(\cdot \mid T_o < T_u^+)$. Furthermore, on the event $\{X_{T_A} = o\}$ it is necessarily the case that:

- Z_1, \dots, Z_{N-1} avoid A .
- Z_N touches A for the first time in o .

When we ignore the first point above, we therefore obtain the upper-bound (42).

By Lemma 2.45, the right hand side is equal to

$$\mathbb{P}_o(\{X_n : n \leq T_u\} \cap A = \{o\} \mid T_u < T_o^+).$$

Therefore, it suffices to show that this converges uniformly to zero over $u \in \Lambda_a(M)^c$, as $M \rightarrow \infty$.

Let \widehat{X} be a CRW, started at o . Fix $\epsilon > 0$ and let M be some integer to be fixed later. Take $L = L(M, \epsilon)$ large enough so that

$$\mathbb{P}_o(\{X_n : n \leq T_M\} \cap A = \{o\} \mid T_u < T_o^+) \leq \mathbb{P}_o(\{\widehat{X}_n : n \leq T_M\} \cap A = \{o\}) + \frac{\epsilon}{2}, \quad (43)$$

for all $u \notin \Lambda_a(L)$, which is possible by Lemma 2.46 (note that L depends only on ϵ and M , in particular does not depend on the choice of A). Next, take M so large that

$$\mathbb{P}(\{\widehat{X}_n : n \leq T_M\} \cap \{\text{LE}(\widehat{Y})_n : n \in \mathbb{N}\} \neq \{o\}) \geq 1 - \frac{\epsilon}{2}, \quad (44)$$

where \widehat{Y} is an independent CRW. This is possible by the intersection property (cIP) as the expression in (44) is increasing in M . Hence for $u \notin \Lambda_a(L)$, combining (43) and (44), conditioning on $\text{LE}(\widehat{Y})$,

$$\mathbb{P}_u(X_{T_{\text{LE}(\widehat{Y})}} = o) \leq \epsilon$$

As ϵ was arbitrary, this shows the result. \square

Wilson's algorithm rooted at infinity

Recall Wilson's algorithm for recurrent graphs: let $I = (v_0, v_1, \dots)$ be any enumeration of the vertices $v(G)$. Fix $E_0 = \{v_0\}$ and define inductively E_{i+1} given E_i , to be E_i together with the loop erasure of an (independent) simple random walk started at v_{i+1}

2.8. (A) IMPLIES (D): ONE-ENDEDNESS OF THE UNIFORM SPANNING TREE

and stopped when hitting E_i . Set $E = E(I) = \cup_{i \geq 0} E_i$. Then Wilson's algorithm tells us that the spanning tree E is in fact a uniform spanning tree (i.e., its law is the weak limit of uniform spanning trees on exhaustions) and in particular, its law does not depend on I , see Wilson [171] for finite graphs and e.g. [131] for infinite recurrent graphs.

Since the conditioned random walk is well defined, we can also start differently: namely take again some enumeration $I = (v_0, \dots)$ of $v(G)$. Define $F_0 = \text{LE}(\widehat{X})$, started at v_0 say and let F_{i+1} be F_i together with the loop erasure of a simple random walk started at v_{i+1} and stopped when hitting F_i . Define $F = F(I) = \cup_{i \geq 0} F_i$. It is not hard to see that again, F is a spanning tree of G (the idea is that the loops formed by the walk coming back to the origin are erased anyway, so one might as well consider the conditioned walk). This is called "Wilson's algorithm rooted at infinity". A similar idea was first introduced for transient graphs in [32] and later defined for \mathbb{Z}^2 .

Lemma 2.48 (Wilson's algorithm rooted at infinity). *The spanning tree F is a uniform spanning tree.*

Proof. Begin with o and let $(z_n)_{n \geq 0}$ be some sequence of vertices going to infinity in G . Apply Wilson's algorithm with the orderings $I_n := (o, z_n, v_2, \dots) \equiv v(G)$, then the law of the first branch E_1 equals $\text{LE}(X^{z_n \rightarrow o})$ by construction, where $X^{z_n \rightarrow o}$ is (the trace of) a random walk started at z_n and stopped when hitting o . This law converges to $\text{LE}(\widehat{X})$ as $i \rightarrow \infty$ due to first the path-reversal (Lemma 2.45) and then Lemma 2.46. Since Wilson's algorithm is independent of the ordering of $v(G)$, the result follows. \square

Orienting the UST. When the UST is one-ended, it is always possible to unambiguously assign a consistent orientation to the edges (from each vertex there is a unique forward edge) such that the edges are oriented towards the unique end of the tree. Although we do not of course know *a priori* that the UST is one-ended, it will be important for us to show that the tree inherits such a consistent orientation from Wilson's algorithm rooted at infinity. Furthermore, we need to show this orientation does not depend on the ordering used in the algorithm. To see this, consider an exhaustion G_n of the graph. Perform Wilson's algorithm (with initial boundary given by the boundary of G_n) and some given sequence of vertices. When adding the branch containing the vertex x to the tree by performing a loop-erased walk starting from x , orient these edges uniquely from x to the boundary.

We point out that it is not entirely clear *a priori* that this orientation converges, or that the limit of the orientation does not depend on the exhaustion (indeed on \mathbb{Z} the oriented tree converges but the orientation depends on the exhaustion, though the UST itself doesn't), nor is it immediately clear that the law of the oriented tree doesn't depend on the sequence of vertices. But this follows readily from the fact that the loops at x from a random walk starting from x are all erased, so that the branch containing

x is obtained by loop-erasing a random walk conditioned to hit the boundary before returning to x , a process which has a limit as $n \rightarrow \infty$, is transient, and does not depend on the exhaustion used when we assume that the potential kernel is well defined. Thus the law of this oriented tree, call it $\vec{\mathcal{T}}_n$, has a limit $\vec{\mathcal{T}}$ as $n \rightarrow \infty$. Obviously, $\vec{\mathcal{T}}$ can be described directly in the infinite graph by adding to the construction of Lemma 2.48 the orientation we get from Wilson's algorithm. When seen like this, it might not be immediately clear that the law of $\vec{\mathcal{T}}$ doesn't depend on the ordering of vertices for Wilson's algorithm. To see this, observe that the orientation of $\vec{\mathcal{T}}_n$ is identical to the one where all edges of \mathcal{T}_n are oriented towards ∂G_n , and the law of \mathcal{T}_n itself does not depend on the ordering, as discussed before. Hence $\vec{\mathcal{T}}_n$ does not depend on the ordering of vertices, and taking limits, neither does $\vec{\mathcal{T}}$.

Note that if x, y are two vertices on a bi-infinite path of $\vec{\mathcal{T}}$, then it makes sense to ask if y is in the past of x or vice-versa: exactly one of these alternatives must hold.

We are now ready to start with the proof of Theorem 2.44.

Proof of Theorem 2.44. Notice that if G is a graph satisfying the standing assumptions (Definition 2.22) and is moreover planar or random and unimodular then (almost surely), G satisfies the intersection property for CRW (cIP) due to Propositions 2.43 and 2.42 respectively.

Suppose (G, o) is reversible, and satisfies the standing assumptions a.s. For a vertex x of G , consider the event $\mathcal{A}_2(x)$ that there are two *disjoint and simple* paths from x to infinity in the UST \mathcal{T} , in other words there is a bi-infinite path going through x . Note that it is sufficient to prove

$$\mathbb{P}(\mathcal{A}_2(x)) = 0$$

for each $x \in v(G)$ a.s., where we remind the reader that here \mathbb{P} is conditional given the graph (i.e., it is an average over the spanning tree \mathcal{T}). Indeed, for the tree \mathcal{T} to be more than one-ended, there must at least be some simple path in \mathcal{T} which goes to infinity in both directions. By biasing and unbiasing by the degree of the root to get a unimodular graph, it is sufficient to prove that $\mathbb{P}(\mathcal{A}_2(o)) = 0$ a.s. Therefore it is sufficient to prove $\mathbf{P}(\mathcal{A}_2(o)) = 0$, where we remind the reader that \mathbf{P} is averaged also over the graph. We first outline the rough idea before giving the details. Suppose for contradiction that $\mathbf{P}(\mathcal{A}_2(o)) \geq \epsilon > 0$. If this is the case then it is possible for both $\mathcal{A}_2(o)$ and $\mathcal{A}_2(x)$ to hold simultaneously, for many other vertices – including vertices far away from o . However, \mathcal{T} is connected (since G is recurrent) and by Theorem 6.2 and Proposition 7.1 in [9], \mathcal{T} is at most two-ended. Therefore the bi-infinite paths going through x and o must coincide: essentially, the bi-infinite path containing o must be almost space-filling.

Suppose x is in the past of o (which we can assume without loss of generality by reversibility). Using Wilson's algorithm rooted at infinity to sample first the path from

2.8. (A) IMPLIES (D): ONE-ENDEDNESS OF THE UNIFORM SPANNING TREE

o and then that from x , the event $\mathcal{A}_2(o) \cap \mathcal{A}_2(x)$ requires a very unlikely behavior: namely, a random walk starting from x must hit the loop-erasure of the conditioned walk starting from o exactly at o . This is precisely what Lemma 2.47 shows is unlikely, because of the infinite intersection properties.

Let us now give the details. Given G , we sample k independent random walks (X^1, \dots, X^k) from o , independently of \mathcal{T} , where $k = k(\epsilon)$ will be chosen below. Observe that by stationarity of (G, o) , we have for every $n \geq 0$,

$$\mathbf{P}(\mathcal{A}_2(X_n^i)) = \mathbf{P}(\mathcal{A}_2(o)) \geq \epsilon.$$

First we show that we can choose k such that for every n , there is i and j such that $\mathcal{A}_2(X_n^i) \cap \mathcal{A}_2(X_n^j)$ holds with \mathbf{P} -probability at least $\epsilon/2$. Indeed fix $n \geq 0$ arbitrarily for now, write $E_i = \mathcal{A}_2(X_n^i)$. Then by the Bonferroni inequalities,

$$\mathbf{P}\left(\bigcup_{i=1}^k E_i\right) \geq \sum_{i=1}^k \mathbf{P}(E_i) - \sum_{1 \leq i \neq j \leq k} \mathbf{P}(E_i \cap E_j)$$

so that

$$\sum_{1 \leq i \neq j \leq k} \mathbf{P}(E_i \cap E_j) \geq k\epsilon - \mathbf{P}\left(\bigcup_{i=1}^k E_i\right) \geq k\epsilon - 1.$$

Choose $k = \lceil 2/\epsilon \rceil$, then we deduce that for some $1 \leq i < j \leq k$,

$$\mathbf{P}(E_i \cap E_j) \geq \binom{k}{2}^{-1}.$$

By stationarity (rerooting at the endpoint of the i th walk), and the Markov property of the walk, this implies

$$\mathbf{P}(\mathcal{A}_2(o) \cap \mathcal{A}_2(X_{2n})) \geq \binom{k}{2}^{-1}. \quad (45)$$

When $\mathcal{A}_2(o) \cap \mathcal{A}_2(X_{2n})$ occurs, both o and X_{2n} are on some bi-infinite path, the two paths must coincide. By symmetry (i.e., reversibility) and invariance of the oriented tree $\vec{\mathcal{T}}$ with respect to the ordering of vertices,

$$\mathbf{P}(\mathcal{A}_2(o) \cap \mathcal{A}_2(X_{2n}); X_{2n} \in \mathbf{Past}(o)) \geq \delta := (1/2) \binom{k}{2}^{-1}. \quad (46)$$

Let \widehat{Y} denote a conditioned walk starting from o and let $\text{LE}(\widehat{Y})$ denote its loop-erasure, and let Z be a random walk starting from a different vertex x . Now, pick M large enough that for any $x \in \Lambda_a(M)^c$

$$\mathbb{P}_x(Z_{T_{\text{LE}(\widehat{Y})}} = o) \leq \delta/3, \quad (47)$$

which we may by Lemma 2.47. Even though M is random (depending only on the graph), observe that as $n \rightarrow \infty$,

$$\mathbb{P}(X_{2n} \in \Lambda_a(M)) \rightarrow 1$$

since G is a.s. null recurrent (as is any recurrent infinite graph). Therefore by dominated convergence,

$$\mathbf{P}(X_{2n} \in \Lambda_a(M)) \rightarrow 1.$$

It follows using (46) that we may choose n large enough that

$$\mathbf{P}(\mathcal{A}_2(o) \cap \mathcal{A}_2(X_{2n}) \cap \{X_{2n} \in \mathbf{Past}(o)\} \cap \{X_{2n} \notin \Lambda_a(M)\}) \geq 2\delta/3. \quad (48)$$

To conclude, we pick n as above, and use Wilson's algorithm rooted at infinity (Lemma 2.48) by first sampling the path from o (which is nothing else by $\text{LE}(\widehat{Y})$) and then sampling the path in $\vec{\mathcal{T}}$ from $x = X_{2n}$, by loop-erasing a random walk Z from this point, stopped at the time T where it hits $\text{LE}(\widehat{Y})$. As mentioned above, When $\mathcal{A}_2(x)$ and $\mathcal{A}_2(o)$ occur and x is in the past of o , since \mathcal{T} is at most two-ended (by [9]), it must be that $Z_T = o$. (If we do not specify that $x \in \mathbf{Past}(o)$ there might otherwise also be the possibility that x itself was directly on the loop-erasure of the conditioned walk). Hence, using (48) and (47),

$$\begin{aligned} 2\delta/3 &\leq \mathbf{P}(\mathcal{A}_2(o) \cap \mathcal{A}_2(X_{2n}) \cap \{X_{2n} \in \mathbf{Past}(o)\} \cap \{X_{2n} \notin \Lambda_a(M)\}) \\ &\leq \mathbf{E}(1_{\{Z_T=o\}} 1_{\{X_{2n} \notin \Lambda_a(M)\}}) \\ &\leq \mathbf{E}(\mathbb{P}_{X_{2n}}(Z_T = o) 1_{X_{2n} \notin \Lambda_a(M)}) \leq \delta/3, \end{aligned}$$

after conditioning on X_{2n} . This is a contradiction, and concludes the proof of Theorem 2.44 (and hence also that (a) implies (d) in Theorem 2.1). \square

Furthermore, (d) is already known by [32, Theorem 14.2] to imply (b), which we have already shown is equivalent to (a). This finishes the proof of Theorem 2.1.

2.9. Harmonic measure from infinity on mated-CRT maps

Let \mathbf{P} denote the law of the whole plane mated-CRT map $G = G^1$ with parameter $\gamma \in (0, 2)$ and with root o . We will not give a precise definition of these maps here and instead refer the reader for instance to [88] or [35]. Since $\mathbf{E}[\deg(o)] < \infty$, the potential kernel is well defined (either because it is planar, or because it is strictly subdiffusive). We now discuss a more quantitative statement concerning the harmonic measure from infinity which underlines substantial differences with the usual square lattice.

We will write $B_{\text{euc}}(x, n)$ for the ball of vertices $z \in v(G)$ such that the Euclidean distance between z and x (w.r.t. the natural embedding) is at most n .

Theorem 2.49. *There exists a $\delta = \delta(\gamma) > 0$ such that the following holds. Almost surely, there exists an $N \geq 1$ such that for all $x \notin B_{\text{euc}}(N)$ we have that*

$$\text{hm}_{o,x}(x) \leq 1 - \delta.$$

In particular,

$$\mathbf{P} \left(\frac{1}{2} \leq \limsup_{y \rightarrow \infty} \widehat{q}(y) \leq 1 - \delta \right) = 1.$$

In fact, we expect the following stronger result to hold:

Conjecture 2.50. *For some (nonrandom) $a, b > 0$, almost surely*

$$a = \liminf_{y \rightarrow \infty} \widehat{q}(y) \leq \limsup_{y \rightarrow \infty} \widehat{q}(y) = 1 - b. \quad (49)$$

In fact, sharp values for a, b can be conjectured by considering the minimal and maximal exponents for the LQG volume of a Euclidean ball of radius ε in a γ -quantum cone, which all decay polynomially as $\varepsilon \rightarrow 0$ (see Lemma A.1 in [35]). We also conjecture that this holds for other random planar maps in the universality class of Liouville quantum gravity with parameter $\gamma \in (0, 2)$, such as the UIPT.

Based on this we conjecture that $\max(a, b) < 1/2$. This would show a stark contrast with the square lattice \mathbb{Z}^2 where we recall that $a = b = 1/2$ (see e.g. [148]). The upper bound in (49) is of course stated in Theorem 2.49 so that the lower bound in (49) is what we are asking about. While we are not able to prove this, we may use the unimodularity of the law \mathbf{P} is unimodular, to prove a slightly weaker lower bound:

Corollary 2.51. *Let $\delta > 0$ as in the previous theorem. Then, almost surely, the asymptotic fraction of δ -good points equals one or in other words, a.s.,*

$$\liminf_{n \rightarrow \infty} \frac{1}{|B(n)|} |\{x \in B(n) : \text{hm}_{o,x}(x) < \delta\}| = 0.$$

Proof. Let $\tilde{\mathbf{P}}$ denote the law \mathbf{P} after degree biasing. We write (\tilde{G}, \tilde{o}) for the random graph with law $\tilde{\mathbf{P}}$.

On the one hand, by reversibility of $\tilde{\mathbf{P}}$, we know that

$$\tilde{\mathbf{P}}(\text{hm}_{\tilde{o}, X_n}(X_n) > 1 - \delta) = \tilde{\mathbf{P}}(\text{hm}_{\tilde{o}, X_n}(\tilde{o}) > 1 - \delta) = \tilde{\mathbf{P}}(\text{hm}_{\tilde{o}, X_n}(X_n) < \delta).$$

On the other hand, by Theorem 2.49 and the reversed Fatou's lemma, we have

$$\limsup_{n \rightarrow \infty} \tilde{\mathbf{P}}(\text{hm}_{\tilde{o}, X_n}(X_n) > 1 - \delta) = 0,$$

thus

$$\lim_{n \rightarrow \infty} \tilde{\mathbf{P}}(\text{hm}_{\tilde{o}, X_n}(X_n) < \delta) = 0.$$

The result now follows by contradiction: indeed, suppose that with positive probability, there is a positive asymptotic fraction of vertices $x \in B(n)$ which have $\text{hm}_{\tilde{o}, x}(x) < \delta$, then the random walk will spend a positive fraction of time in these points, giving a contradiction. \square

2.9.1 Preliminaries and known results.

We collect some known results about mated-CRT maps which are needed for the proof of Theorem 2.49.

Lemma 2.52. *There exist $C = C(\gamma) < \infty$ and $\alpha = \alpha(\gamma) > 0$, such that for all $n \in \mathbb{N}$,*

$$\mathbf{P} \left(\frac{1}{C} \log(n) \leq \mathcal{R}_{\text{eff}}(o \leftrightarrow \partial B_{\text{euc}}(o, n)) \leq C \log(n) \right) \geq 1 - \frac{1}{\log(n)^\alpha}.$$

Proof. This is Proposition 3.1 in [87]. □

Lemma 2.53. *There exists a $C = C(\gamma) < \infty$ and $\alpha = \alpha(\gamma) > 0$ such that with \mathbf{P} -probability at least $1 - n^{-\alpha}$, for all $x \in B_{\text{euc}}(3n)$ and all $s \in [1/3, 1]$*

$$\max_{z \in \partial B_{\text{euc}}(x, sn)} h(z) \leq C \min_{x \in B_{\text{euc}}(x, sn)} h(z)$$

whenever $h : B_{\text{euc}}(x, 3n) \cup \partial B_{\text{euc}}(x, 3n) \rightarrow \mathbb{R}_+$ is harmonic outside of possibly x and $\partial B_{\text{euc}}(x, 3n)$.

Proof. This is the content of Proposition 3.8 [35]. □

Lemma 2.54. *There exist $C = C(\gamma) < \infty$ and $\alpha = \alpha(\gamma)$ such that with \mathbf{P} -probability at least $1 - n^{-\alpha}$, for all $x \in B_{\text{euc}}(3n)$,*

$$\mathcal{R}_{\text{eff}}(x \leftrightarrow \partial B_{\text{euc}}(x, n)) \geq \frac{1}{C} \log(n).$$

Proof. This follows from Lemma 4.2 in [35]. □

Proposition 2.55. *Let (G, o) have the law of the mated-CRT map with parameter γ . There exist constants $C = C(\gamma)$ and $\alpha = \alpha(\gamma) > 0$ such that*

$$\mathbf{P} \left(\frac{1}{C} \log(n) \leq a(x, o) \leq C \log(n) \text{ for all } x \in B_{\text{euc}}(o, 2n) \setminus B_{\text{euc}}(o, n) \right) \geq 1 - \frac{1}{\log(n)^\alpha}.$$

Proof of Proposition 2.55. By Lemma 2.8 we know that for each $n \in \mathbb{N}$

$$\mathcal{R}_{\text{eff}}(o \leftrightarrow \partial B_{\text{euc}}(o, 2n)) = \mathbb{E}_o[a(X_{T_{2n}}, o)]$$

(where we recall that \mathbb{E}_o is the expectation solely on the random walk). Now, fix n and let \mathcal{E}_n be the intersection of both events in Lemmas 2.52 and 2.53, which are properties of the graph only. Note that \mathcal{E}_n holds with high probability over the mated-CRT maps (possibly by suitably changing the values of the constants).

Then, as $x \mapsto a(x, o)$ is harmonic outside o , conditional on \mathcal{E}_{2n} , we know that whenever $x \in B_{\text{euc}}(o, 2n) \setminus B_{\text{euc}}(o, n)$,

$$\frac{1}{C} a(X_{T_{2n}}, o) \leq a(x, o) \leq C a(X_{T_{2n}}, o),$$

so that taking (random walk) expectations,

$$\frac{1}{C^2} \log(n) \leq a(x, o) \leq C^2 \log(n).$$

This is the desired result. \square

2.9.2 Proof of Theorem 2.49.

Take throughout the proof the constants C, α such that Lemmas 2.54, 2.53 and 2.52 and Proposition 2.55 hold simultaneously with the same constants.

Proof. The second statement follows immediately from the first statement, from the identity

$$\limsup_{y \rightarrow \infty} \widehat{q}(y) = \limsup_{y \rightarrow \infty} \text{hm}_{y,o}(y),$$

in Corollary 2.37, and from the fact that for each $\epsilon > 0$, there are infinitely many $(\frac{1}{2} - \epsilon)$ -good vertices by Lemma 2.20. We are thus left to prove the first statement.

To that end, fix N_0 so large that for all $n \geq N_0$,

$$\frac{n^{2/\alpha}}{(n-1)^{2/\alpha}} \leq 3.$$

Define next for $m \geq 1$ the event

$$E_m \text{ the event that } a(x, o) \leq C \log(m) \text{ for all } x \in B_{\text{euc}}(m). \quad (50)$$

By Proposition 2.55, we know that $\mathbf{P}(E_m^c) \leq \log(m)^{-\alpha}$ and therefore,

$$\sum_{n=1}^{\infty} \mathbf{P}(E_{e^{n^{2/\alpha}}}^c) < \infty.$$

By Borel-Cantelli, this implies that there is some (random) $N_1 = N_1(G, o) < \infty$ such that $E_{e^{n^{2/\alpha}}}$ occurs for all $n \geq N_1$. Suppose without loss of generality that $N_1 \geq N_0$ almost surely. In this case, it follows that

$$a(x, o) \leq C \log |x| \quad \text{for all } x \notin B_{\text{euc}}(o, N_1). \quad (51)$$

Next, define the events

H_m the event that for all $x \in B_{\text{euc}}(3m) \setminus B_{\text{euc}}(m)$ and for all

$h : v(G) \rightarrow \mathbb{R}_+$ harmonic outside of x ,

$$\max_{z \in \partial B_{\text{euc}}(x, |x|)} h(z) \leq C \min_{z \in \partial B_{\text{euc}}(x, |x|)} h(z)$$

and

R_m the event that for all $x \in B_{\text{euc}}(3m)$, $\mathcal{R}_{\text{eff}}(x \leftrightarrow \partial B_{\text{euc}}(x, m)) \geq \frac{1}{C} \log(m)$.

By Lemmas 2.53 and 2.54 respectively, it holds that $\mathbf{P}(H_m^c) \leq m^{-\alpha}$ and $\mathbf{P}(R_m) \leq m^{-\alpha}$. Therefore, using again a Borel-Cantelli argument, there exists some (random) $N_2 \geq N_1 \geq N_0$ such that almost surely, for all $n \geq N_2$ the events $H_{n^{2/\alpha}}$ and $R_{n^{2/\alpha}}$ occur. In particular, we know that almost surely,

$$\mathcal{R}_{\text{eff}}(x \leftrightarrow \partial B_{\text{euc}}(x, |x|)) \geq \frac{1}{C} \log |x| \quad \text{for all } x \notin B_{\text{euc}}(o, N_2) \quad (52)$$

and almost surely

$$\begin{aligned} \text{For all } x \notin B_{\text{euc}}(o, N_2), \text{ for all } h : v(G) \rightarrow \mathbb{R}_+ \text{ harmonic outside of } x \\ \max_{z \in \partial B_{\text{euc}}(x, |x|)} h(z) \leq C \min_{z \in \partial B_{\text{euc}}(x, |x|)} h(z). \end{aligned} \quad (53)$$

Take $x \notin B_{\text{euc}}(o, N_2)$. Assume without loss of generality that $\text{hm}_{o,x}(x) \leq \frac{1}{2}$, as otherwise we are done. Then

$$\mathcal{R}_{\text{eff}}(o \leftrightarrow x) \leq 2 \text{hm}_{o,x}(x) \mathcal{R}_{\text{eff}}(o \leftrightarrow x) = 2a(x, o) \leq 2C \log |x|, \quad (54)$$

where we used Corollary 2.7 in the equality and (51) in the last inequality.

Furthermore, as $z \mapsto a(z, x)$ is harmonic outside of x , applying (53) first and then (52) gives

$$a(o, x) \geq \frac{1}{C} \mathbb{E}_x[a(X_{T_{B_{\text{euc}}(x, |x|)}}, x)] = \frac{1}{C} \mathcal{R}_{\text{eff}}(x \leftrightarrow \partial B_{\text{euc}}(x, |x|)) \geq \frac{1}{C^2} \log |x|.$$

Combining the last equation with (54), we find

$$\text{hm}_{o,x}(o) = \frac{a(o, x)}{\mathcal{R}_{\text{eff}}(o \leftrightarrow x)} \geq \frac{1}{2C^3},$$

which shows the final result. \square

Resolving the conjecture of Aldous and Lyons

3.1. Introduction

The **uniform spanning tree** of a finite connected graph G is defined by picking uniformly at random a connected subgraph of G containing all vertices but no cycles. To go from finite to infinite graphs, it is possible to exhaust G by finite subgraphs and take weak limits with appropriate boundary conditions. For two natural such choices of boundary conditions, known as **free** and **wired** boundary conditions, Pemantle [143] proved that these infinite-volume limits are always well-defined independently of the choice of exhaustion, and that the choice of boundary conditions also does not affect the limit obtained when $G = \mathbb{Z}^d$. Since connectivity of a subgraph is not a closed condition, these weak limits might be supported on configurations that are *forests* rather than trees, and indeed Pemantle proved for \mathbb{Z}^d that the limit is connected if and only if $d \leq 4$. For a general infinite, connected, locally finite graph the infinite-volume limit of the UST with free boundary conditions is called the **free uniform spanning forest** (FUSF) and the infinite volume limit with wired boundary conditions is called the **wired uniform spanning forest** (WUSF); when the two limits are the same we refer to them simply as the uniform spanning forest (USF). In their highly influential work [32], Benjamini, Lyons, Peres and Schramm resolved the connectivity question for the WUSF in large generality: the wired uniform spanning tree is a single tree if and only if two random walks intersect infinitely often. The connectivity of the FUSF appears to be a much more subtle question and, outside of the case that the two forests are the same, is understood only in a few examples [13, 95, 144, 162]. For *recurrent* graphs, which are the main topic of this chapter, the infinite-volume limit of the UST is always defined independently of boundary conditions and a.s. connected [32, Proposition 5.6], so that we can unambiguously refer to the uniform spanning tree (UST) of an infinite, connected, locally finite, recurrent graph G .

After connectivity, the next most basic topological property of the USF is the number of **ends** its components have. Here, we say that a graph has at least m ends whenever

there exists some finite set of vertices W such that $G \setminus W$ has at least m infinite connected components. The graph is said to be m -ended if it has at least m but not $m + 1$ ends. Understanding the number of ends of the USF turns out to be rather more difficult than connectivity, with a significant literature now devoted to the problem. For *Cayley graphs*, it follows from abstract principles [13, Section 3.4] that every component has 1, 2, or infinitely many ends almost surely, and for *amenable* Cayley graphs such as \mathbb{Z}^d (for which the WUSF and FUSF always coincide) it follows by a Burton-Keane [43] type argument that every component has either one or two ends almost surely; see [131, Chapter 10] for detailed background. For the *wired* uniform spanning forest on transitive graphs, a complete solution to the problem was given by Benjamini, Lyons, Peres, and Schramm [32] and Lyons, Morris, and Schramm [130], who proved that every component of the WUSF of an infinite transitive graph is one-ended almost surely unless the graph in question is rough-isometric to \mathbb{Z} . Before going forward, let us emphasize that the recurrent case of this result [32, Theorem 10.6] is established using a completely different argument to the transient case, with the tools available for handling the two cases being largely disjoint.

Beyond the transitive setting, various works have established mild conditions under which every component of the WUSF is one-ended almost surely, applying in particular to planar graphs with bounded face degrees [95] and graphs satisfying isoperimetric conditions only very slightly stronger than transience [92, 130]. These proofs are quantitative, and recent works studying critical exponents for the USF of \mathbb{Z}^d with $d \geq 3$ [12, 93, 97] and Galton-Watson trees [99] can be thought of as a direct continuation of the same line of research.

In parallel to this deterministic theory, Aldous and Lyons [9] observed that the methods of [32] also apply to prove that the WUSF has one-ended components on any transient *unimodular random rooted graph* of bounded degree, and the second author of the present chapter later gave new proofs of this result with different methods that removed the bounded degree assumption [91, 93]. It is also proven in [94, 164] that every component of the *free* uniform spanning forest of a unimodular random rooted graph is infinitely ended a.s. whenever the free and wired forests are different. Here, unimodular random rooted graphs comprise a very large class of random graph models including Benjamini-Schramm limits of finite graphs [33], Cayley graphs, and (suitable versions of) Galton-Watson trees, as well as e.g. percolation clusters on such graphs; See Section 3.3.1 for definitions and e.g. [9, 49] for detailed background.

The aforementioned works [9, 91, 92, 94, 164] completely resolved the problem of the number of ends of the WUSF and FUSF for *transient* unimodular random rooted graphs, but the recurrent case remained open. Besides the fact that the transient methods do not apply, a further complication of the recurrent case is that it is possible for the UST

to be either one-ended or two-ended according to the geometry of the graph: indeed, the UST of \mathbb{Z}^2 is one-ended while the UST of \mathbb{Z} is two-ended.

Aldous and Lyons conjectured [9, p. 1485] that the dependence of the number of ends of the UST on the geometry of the graph is as simple as possible: The UST of a recurrent unimodular random rooted graph is one-ended if and only if the graph is. The fact that two-ended unimodular random rooted graphs have two-ended USTs is trivial; the content of the conjecture is that one-ended unimodular random rooted graphs have one-ended USTs. Previously, the conjecture was resolved under the assumption of planarity in [13], while in [32, 34] it was proved (without using the planarity assumption) that the UST of a recurrent unimodular random rooted graph is one-ended precisely when the “harmonic measure from infinity” is uniquely defined. In this chapter we resolve the conjecture.

Theorem 3.1. *Let (G, o) be a recurrent unimodular random rooted graph and let T be the uniform spanning tree of G . Then T has the same number of ends as G a.s.*

To see that the theorem is not true without unimodularity, consider taking the line graph \mathbb{Z} and adding a path of length 2^n connecting $-n$ to n for each n , making the graph one-ended. Kirchoff’s effective resistance formula implies that the probability that the additional path connecting $-n$ to n is included in the UST is at most $n/(2^n + n)$, and a simple Borel-Cantelli argument implies that the UST is two-ended almost surely. Similar examples show that Theorem 3.1 does not apply to unimodular random rooted *networks*, since we can use edges of very low conductance to make the network one-ended while having very little effect on the geometry of the UST.

About the proof. We stress again that the tools used in the transient case do not apply at all to the recurrent case, and we are forced to use completely different methods that are specific to the recurrent case. We build on [34] which proved that the “harmonic measures from infinity” are uniquely defined if and only if the uniform spanning tree is one-ended; A self-contained treatment of (a slight generalization of) the results of [34] that we will need is given in Appendix A. The set of harmonic measures from infinity can be thought of as a “boundary at infinity” for the graph, analogously to the way the Martin boundary is used in transient graphs. It is implicit in [34] that these measures correspond to the ways in which a random walk “conditioned to never return to the root” can escape to infinity. We develop these ideas further in Section 3.2, in which we make this connection precise. We then apply these ideas inside an ergodic-theoretic framework to prove that *if* the UST has two ends, then the effective resistance must grow linearly along the unique bi-infinite path in the tree, which implies in particular that graph distances must also grow linearly. To conclude, we argue that this can only happen when the graph has linear volume growth, which is known to be equivalent to two-endedness for unimodular random rooted graphs [31, 40].

3.2. Boundary theory of recurrent graphs

In this section we develop the theory of harmonic measures from infinity on recurrent graphs, their associated potential kernels and Doob transforms, and how this relates to the spanning tree. Much of the theory we develop here is a direct analogue for recurrent graphs of the theory of Martin boundaries of transient graphs [61, 172]. Some results on recurrent boundary theory can be found in [102]. This theory is interesting in its own right, and we were surprised to find how little attention has been paid to these notions outside of some key motivating examples such as \mathbb{Z}^2 [75, 149].

All of the results in this section will concern deterministic infinite, connected, recurrent, locally finite graphs and can be extended to general locally finite networks; applications of the theory to unimodular random rooted graphs will be given in Section 3.3.

3.2.1 Harmonic measures from infinity

Let $G = (V, E)$ be an infinite, connected, locally finite, recurrent graph. For each $v \in V$ we write \mathbf{P}_v for the law of the simple random walk on G started at v , and for each set $A \subseteq V$ write T_A and T_A^+ for the first visit time of the random walk to A and first positive visit time of the random walk to A respectively. Given a probability measure μ on V , we also write \mathbf{P}_μ for the law of the random walk started at a μ -distributed vertex.

A **harmonic measure from infinity** $h = (h_B : B \subset V \text{ finite})$ on G is a collection of probability measures on V indexed by the finite subsets B of V with the following properties:

1. h_B is supported on ∂B for each $B \subset V$, where ∂B is the set of elements of B that are adjacent to an element of $V \setminus B$.
2. For each pair of finite sets $B \subseteq B'$, h_B and $h_{B'}$ satisfy the consistency condition

$$h_B(u) = \sum_{v \in B'} h_{B'}(v) \mathbf{P}_v(X_{T_B} = u) \tag{1}$$

for every $u \in B$.

We denote the space of harmonic measures from infinity by \mathcal{H} , which (identifying the measures h_B with their probability mass functions) is a compact convex subset of the space of functions $\{\text{finite subsets of } V\} \rightarrow \mathbb{R}^V$ when equipped with the product topology. As mentioned above, the space \mathcal{H} plays a role for recurrent graphs analogous to that played by the Martin boundary for transient graphs; the analogy will become clearer once we introduce potential kernels in the next subsection. We say that the harmonic measure from infinity is **uniquely defined** when \mathcal{H} is a singleton.

If μ_n is a sequence of probability measures on V converging vaguely to the zero measure in the sense that $\mu_n(v) \rightarrow 0$ as $n \rightarrow \infty$ for each fixed $v \in V$ then any subsequential limit of the collections $(\mathbf{P}_{\mu_n}(X_{T_B} = \cdot) : B \subset V \text{ finite})$ belongs to \mathcal{H} , with these collections themselves satisfying every property of a harmonic measure from infinity other than the condition that h_B is supported on ∂B for every finite B . (Indeed, the consistency condition (1) follows from the strong Markov property of the random walk.) In fact every harmonic measure from infinity can be written as such a limit.

Lemma 3.2. *If $h \in \mathcal{H}$ is a harmonic measure from infinity then there exists a sequence of finitely supported probability measures $(\mu_n)_{n \geq 1}$ on V such that $\mu_n(v) \rightarrow 0$ for every $v \in V$ and*

$$h_B(\cdot) = \lim_{n \rightarrow \infty} \mathbf{P}_{\mu_n}(X_{T_B} = \cdot) \quad \text{for every } B \subset V \text{ finite.} \quad (2)$$

Proof. Fix $h \in \mathcal{H}$. Let $V_1 \subset V_2 \subset V_3 \cdots$ be an increasing sequence of finite subsets of V with $\bigcup_i V_i = V$, and for each $n \geq 1$ let $\mu_n = h_{V_n}$. It follows from the consistency condition (1) that

$$h_B(\cdot) = \mathbf{P}_{\mu_n}(X_{T_B} = \cdot) \quad \text{for every } B \subset V_n,$$

and the claim follows since every finite set is eventually contained in V_n . □

Since \mathcal{H} is a weakly compact subspace of the set of functions from finite subsets of V to \mathbb{R}^V , which is a locally convex topological vector space, it is a Choquet-simplex: Every element can be written as a convex combination of the extremal points. In particular, if \mathcal{H} has more than one point then it must have more than one extremal point. This will be useful to us because extremal points of \mathcal{H} are always limits of harmonic measures from sequences of single vertices. Indeed, identifying each vertex $v \in V$ with the collection of harmonic measures $(\mathbf{P}_v(X_{T_B} = \cdot) : B \subset V \text{ finite})$ allows us to think of $V \cup \mathcal{H}$ as a compact Polish space containing V (in which V might not be dense), and we say that a sequence of *vertices* $(v_n)_{n \geq 0}$ converges to a point $h \in \mathcal{H}$ if $h_B(\cdot) = \lim_{n \rightarrow \infty} \mathbf{P}_{v_n}(X_{T_B} = \cdot)$ for every $B \subset V$ finite.

Lemma 3.3. *If $h \in \mathcal{H}$ is extremal, there exists a sequence of vertices $(v_n)_{n \geq 0}$ such that v_n converges to h as $n \rightarrow \infty$.*

Proof. Let \mathcal{I} be the set of functions $h : \{B \subset V \text{ finite}\} \rightarrow \mathbb{R}^V$ of the form

$$h_B(\cdot) = \mathbf{P}_\mu(X_{T_B} = \cdot) \quad \text{for every } B \subset V \text{ finite}$$

for some finitely supported measure μ on V . Lemma 3.2 implies that $\bar{\mathcal{I}} = \mathcal{I} \cup \mathcal{H}$ is a compact convex subset of the space of all functions $\{B \subset V \text{ finite}\} \rightarrow \mathbb{R}^V$ equipped with the product topology, which is a locally convex topological vector space. By the

Krein-Milman theorem, a subset W of $\mathcal{I} \cup \mathcal{H}$ has closure containing the set of extremal points of $\mathcal{I} \cup \mathcal{H}$ if and only if $\mathcal{I} \cup \mathcal{H}$ is contained in the closed convex hull of W . Thus, if we define \mathcal{I}_{ext} to be the set of functions $h : \{B \subset V \text{ finite}\} \rightarrow \mathbb{R}^V$ of the form

$$h_B(\cdot) = \mathbf{P}_z(X_{T_B} = \cdot) \quad \text{for every } B \subset V \text{ finite}$$

for some $z \in V$ then \mathcal{I} is clearly contained in the convex hull of \mathcal{I}_{ext} , so that $\mathcal{I} \cup \mathcal{H}$ is contained in the closed convex hull of \mathcal{I}_{ext} and, by the Krein-Milman theorem, the set of extremal points of $\mathcal{I} \cup \mathcal{H}$ is contained in the closure of \mathcal{I}_{ext} .

Now, observe that for any non-trivial convex combination of an element of \mathcal{I} and an element of \mathcal{H} , there must exist a finite set of vertices B and a point z in the interior of B (i.e., in B and not adjacent to any element of $V \setminus B$) such that $h_B(z) \neq 0$; indeed, if the element of \mathcal{I} corresponds to some finitely supported measure μ , then any B containing the support of μ in its interior and any z in the support of μ will do. Since no element of \mathcal{H} can have this property, it follows that non-trivial convex combinations of elements of \mathcal{I} and \mathcal{H} cannot belong to \mathcal{H} and hence that extremal points of \mathcal{H} are also extremal in $\mathcal{I} \cup \mathcal{H}$. It follows that the set of extremal points of \mathcal{H} is contained in the closure of \mathcal{I}_{ext} , which is equivalent to the claim. \square

Remark 3.1. The converse to this lemma is *not* true: A limit of a sequence of Dirac measures need not be extremal. For example, if we construct a graph from \mathbb{Z} by attaching a very long path between $-n$ and n for each $n \geq 1$ and take z_n to be a point in the middle of this path for each n , the sequence $(z_n)_{n \geq 1}$ will converge to a non-extremal element of \mathcal{H} that is the convex combination of the limits of $(n)_{n \geq 1}$ and $(-n)_{n \geq 1}$.

3.2.2 Potential kernels and Doob transforms

The arguments in [34] heavily rely on a correspondence between the harmonic measure from infinity and its **potential kernel**. One important feature of the potential kernel is that, given a vertex $o \in V$ and a point $h \in \mathcal{H}$, it provides a sensible way to “condition the random walk to converge to h before returning to o ”. We begin by discussing how conditioning the random walk to hit a particular vertex before returning to o can be described in terms of Doob transforms before developing the analogous limit theory.

Doob transforms and non-singular conditioning. Suppose that we are given two distinct vertices o and z in an infinite, connected, locally finite recurrent graph G . Letting $\mathbf{G}_z(x, y)$ be the expected number of times a random walk started at x visits y before hitting z , we can compute that the function

$$a(x) = \frac{\mathbf{G}_z(o, o)}{\deg(o)} - \frac{\mathbf{G}_z(x, o)}{\deg(o)} = \frac{\mathbf{P}_x(T_o > T_z)\mathbf{G}_z(o, o)}{\deg(o)}$$

is harmonic at every vertex other than o and z , and has

$$\Delta a(o) = 0 - \deg(o)\mathbf{E}_o[a(X_1)] = -\mathbf{P}_o(T_o^+ > T_z)\mathbf{G}_z(o, o) = -1,$$

where Δ denotes the graph Laplacian

$$\Delta f(x) = \deg(x)f(x) - \sum_{y \sim x} f(y) = \deg(x)\mathbf{E}_x[f(X_0) - f(X_1)]$$

(terms in this sum are counted with appropriate multiplicity if there is more than one edge between x and y). Moreover, the quantity $a(x)$ is strictly positive at every vertex x that is neither equal to o nor disconnected from z by o in the sense that every path from x to z must pass through o . Observe that the trivial identity

$$\begin{aligned} \mathbf{P}_o((X_0, \dots, X_n) = (x_0, \dots, x_n)) &= \prod_{i=1}^n p(x_{i-1}, x_i) \\ &= \frac{1}{a(x_n)} a(x_1) p(o, x_1) \prod_{i=2}^n \frac{a(x_i)}{a(x_{i-1})} p(x_{i-1}, x_i) \end{aligned} \quad (3)$$

holds for every sequence of vertices x_0, \dots, x_n with $x_0 = o$ and $a(x_i) > 0$ for every $i > 0$. Since $a(z) = G_z(o, o) = \mathbf{P}_o(T_z < T_o^+)^{-1}$ it follows that

$$\mathbf{P}_o((X_0, \dots, X_n) = (x_0, \dots, x_n) \mid T_z < T_o^+) = a(x_1) p(o, x_1) \prod_{i=2}^n \frac{a(x_i)}{a(x_{i-1})} p(x_{i-1}, x_i) \quad (4)$$

for every sequence of vertices x_0, \dots, x_n with $x_0 = o$, $x_n = z$, and $x_i \notin \{o, z\}$ for every $0 < i < n$ (which implies that $a(x_i) > 0$ for every $1 \leq i \leq n$). Now, the fact that a is harmonic off of $\{o, z\}$ and has $\Delta a(o) = -1$ implies that we can define a stochastic matrix with state space $\{x \in V : x = o \text{ or } a(x) > 0\}$ by

$$\widehat{p}^a(x, y) = \begin{cases} \frac{a(y)}{a(x)} p(x, y) & x \notin \{o, z\} \\ a(y) & x = o \\ \mathbf{1}(y = z) & x = z, \end{cases}$$

and if we define the **Doob transformed walk** \widehat{X}^a to be the Markov chain with this transition matrix started from o then it follows from (4) that $(\widehat{X}_n^a)_{n=0}^{T_z}$ has law equal to the conditional law of the simple random walk $(X_n)_{n=0}^{T_z}$ started at o and conditioned to hit z before returning to o . Moreover, letting $\widehat{\mathbf{P}}_o^a$ denote the law of \widehat{X}^a , it follows from the definition of \widehat{X}^a that

$$\begin{aligned} \widehat{\mathbf{P}}_o^a \left((\widehat{X}_0^a, \dots, \widehat{X}_n^a) = (x_0, \dots, x_n) \right) &= \prod_{i=1}^n \widehat{p}^a(x_{i-1}, x_i) = a(x_1) \prod_{i=2}^n \frac{a(x_i)}{a(x_{i-1})} p(x_{i-1}, x_i) \\ &= a(x_n) \prod_{i=2}^n p(x_{i-1}, x_i) = \deg(o) a(x_n) \mathbf{P}_o((X_0, \dots, X_n) = (x_0, \dots, x_n)) \end{aligned} \quad (5)$$

for every sequence x_0, \dots, x_n with $x_0 = o$ and $x_i \notin \{o, z\}$ for every $0 < i < n$.

Defining the potential kernel. Let $\mathcal{R}_{\text{eff}}(x \leftrightarrow y) := \deg(x)\mathbb{P}_x(T_y < T_x^+)$ be the effective resistance between two vertices x, y . We now define the **potential kernel** a^h associated to a point $h \in \mathcal{H}$ via the formula

$$a^h(x, y) = h_{x,y}(x)\mathcal{R}_{\text{eff}}(x \leftrightarrow y) \quad (6)$$

where we write $h_{x,y} = h_{\{x,y\}}$, so that $a^h(x, x) = 0$ for each $x \in V$. The fact that this is a sensible definition owes largely to the following lemma.

Lemma 3.4. *For each $h \in \mathcal{H}$, the potential kernel $a^h(x, y) = h_{x,y}(x)\mathcal{R}_{\text{eff}}(x \leftrightarrow y)$ satisfies*

$$\Delta a^h(\cdot, y) = -\mathbf{1}(\cdot = y), \quad (7)$$

so that the potential kernel $a^h(\cdot, y)$ is harmonic away from y and subharmonic at y .

Proof. Since the map $h \mapsto a^h$ is affine and the equality (7) is linear, it suffices to prove the lemma in the case that h is extremal. By Lemma 3.3, there exists a sequence of vertices $(v_n)_{n \geq 1}$ such that v_n converges to h . For each $n \geq 1$ we define

$$a^n(x, y) = \frac{\mathbf{G}_{v_n}(y, y)}{\deg(y)} - \frac{\mathbf{G}_{v_n}(x, y)}{\deg(y)}$$

and claim that

$$a^h(x, y) = \lim_{n \rightarrow \infty} a^n(x, y) \quad (8)$$

for every $x, y \in V$. (Note that this limit formula is often taken as the *definition* of the potential kernel.) We will prove (8) with the aid of three standard identities for the Greens function:

1. By the strong Markov property, $\mathbf{G}_z(x, y)$ is equal to $\mathbf{P}_x(T_y < T_z)\mathbf{G}_z(y, y)$ for every three distinct vertices x, y , and z .
2. By the strong Markov property, $\mathbf{G}_x(y, y)$ is equal to $\mathbf{P}_x(T_y < T_x^+)^{-1}$ for every pair of distinct vertices x and y . It follows in particular that $\deg(y)^{-1}\mathbf{G}_x(y, y) = \mathcal{R}_{\text{eff}}(x \leftrightarrow y)$ and, since the effective resistance is symmetric in x and y , that $\deg(y)^{-1}\mathbf{G}_x(y, y) = \deg(x)^{-1}\mathbf{G}_y(x, x)$.
3. By time-reversal, $\deg(x)\mathbf{G}_z(x, y)$ is equal to $\deg(y)\mathbf{G}_z(y, x)$ for every three distinct vertices x, y , and z .

Applying these three identities in order yields that

$$\begin{aligned} a^n(x, y) &= \frac{\mathbf{G}_{v_n}(y, y)}{\deg(y)}\mathbf{P}_x(T_y > T_{v_n}) \\ &= \frac{\mathbf{G}_y(v_n, v_n)}{\deg(v_n)}\mathbf{P}_x(T_y > T_{v_n}) = \frac{\mathbf{G}_y(x, v_n)}{\deg(v_n)} = \frac{\mathbf{G}_y(v_n, x)}{\deg(x)} \end{aligned}$$

whenever x, y , and v_n are distinct. Applying the first and second identities a second time then yields that

$$a^n(x, y) = \mathbf{P}_{v_n}(T_x < T_y) \frac{\mathbf{G}_y(x, x)}{\deg(x)} = \mathbf{P}_{v_n}(T_x < T_y) \mathcal{R}_{\text{eff}}(x \leftrightarrow y) \quad (9)$$

whenever x, y , and v_n are distinct. This is easily seen to imply the claimed limit formula (8). \square

In light of this lemma, we define \mathcal{P}_o to be the space of non-negative functions $a : V \rightarrow [0, \infty)$ with $a(o) = 0$ and $\Delta a(x) = -\mathbf{1}(x = o)$, so that $a^h(\cdot, o)$ belongs to \mathcal{P}_o for each $o \in V$ and $h \in \mathcal{H}$ by Lemma 3.4. We will later show that the map $h \mapsto a^h$ is an affine isomorphism between the two convex spaces \mathcal{H} and \mathcal{P}_o . We first describe how elements of \mathcal{P}_o can be used to define Doob transformed walks.

Doob transforms and singular conditioning. We now define the Doob transform associated to an element of the space \mathcal{P}_o . Given $a \in \mathcal{P}_o$, we define \widehat{X}^a to be the Doob a -transform of the simple random walk X on G , so that \widehat{X}^a has state space $\{x \in V : x = o \text{ or } a(x) > 0\}$ and transition probabilities given by

$$\widehat{p}^a(x, y) := \begin{cases} \frac{a(y)}{a(x)} p(x, y) & \text{if } x \neq o \\ a(y) & \text{if } x = o, y \sim o \end{cases}$$

where p is the transition kernel for the simple random walk. Similarly, given $h \in \mathcal{H}$, we write $\widehat{X}^h = \widehat{X}^{a^h(\cdot, o)}$ where a^h is the potential kernel associated to h . Informally, we think of \widehat{X}^h as the walk that is “conditioned to go to h before returning to o ”. (In particular, when the harmonic measure from infinity is unique and \mathcal{H} and \mathcal{P}_o are singleton sets, we think of the associated Doob transform as the random walk conditioned to never return to o .) We write $\widehat{\mathbf{P}}_o^a$ or $\widehat{\mathbf{P}}_o^h$ for the law of \widehat{X}^a or \widehat{X}^h as appropriate.

As before, it follows from this definition that if $a \in \mathcal{P}_o$ and we write $X[0, m]$ for the initial segment consisting of the first m steps of the random walk X then

$$\begin{aligned} \widehat{\mathbf{P}}_o^a(\widehat{X}^a[0, m] = \gamma) &= \prod_{i=1}^m \widehat{p}^a(\gamma_{i-1}, \gamma_i) = a(\gamma_1) \prod_{i=2}^m \frac{a(\gamma_i)}{a(\gamma_{i-1})} p(\gamma_{i-1}, \gamma_i) \\ &= a(\gamma_m) \prod_{i=2}^m p(\gamma_{i-1}, \gamma_i) = \deg(o) a(\gamma_m) \mathbf{P}_o(X[0, m] = \gamma) \end{aligned} \quad (10)$$

for every finite path $\gamma = (\gamma_0, \dots, \gamma_m)$ with $\gamma_0 = o$ and $\gamma_i \neq o$ for every $i > 0$. Summing over all paths γ that begin at o , end at some point $x \neq o$, and do not visit o or x at any intermediate point yields in particular that if $h \in \mathcal{H}$ then

$$\widehat{\mathbf{P}}_o^h(\widehat{X} \text{ hits } x) = \deg(o) a^h(x, o) \mathbf{P}_o(T_x < T_o^+) = h_{o,x}(x), \quad (11)$$

where the last equality follows from (6) and the definition of the effective resistance.

Lemma 3.5. *Let $G = (V, E)$ be a recurrent graph and let $a \in \mathcal{P}_o$. Then the associated Doob-transformed walk \widehat{X}^a is transient.*

Proof. One can easily verify from the definitions that the sequence of reciprocals $(a(\widehat{X}_n^a)^{-1})_{n \geq 1}$ is a non-negative martingale with respect to its natural filtration, and hence converges almost surely to some limiting random variable, which it suffices to prove is zero almost surely. It follows from the identity (10) that

$$\begin{aligned} \widehat{\mathbf{P}}_o^a(a(\widehat{X}_n^a) \leq M) \\ = \sum_v \mathbf{1}(a(v) \leq M) \deg(o)a(v) \mathbf{P}_o(X_n = v, T_o^+ > n) \leq M \deg(o) \mathbf{P}_o(T_o^+ > n), \end{aligned}$$

for every $n, M \geq 1$. Since G is recurrent, the right hand side tends to zero as $n \rightarrow \infty$ for each fixed M . It follows that $\limsup_{n \rightarrow \infty} a(\widehat{X}_n^a) = \infty$ almost surely, and hence that $\lim_{n \rightarrow \infty} a(\widehat{X}_n^a) = \infty$ almost surely since the limit is well-defined almost surely. This implies that \widehat{X}^a is transient. \square

3.2.3 An affine isomorphism

Let $G = (V, E)$ be recurrent, fix $o \in V$, and let \mathcal{P}_o denote the set of positive functions $a : V \rightarrow [0, \infty)$ with $a(o) = 0$ that satisfy $\Delta a(\cdot) = -\mathbf{1}(\cdot = o)$. As we have seen, for each $h \in \mathcal{H}$ the potential kernel $a^h(\cdot, o)$ defines an element of \mathcal{P}_o . Moreover, the map sending $h \mapsto a^h(\cdot, o)$ is affine in the sense that if $h = \theta h_1 + (1 - \theta)h_2$ then $a^h(\cdot, o) = \theta a^{h_1}(\cdot, o) + (1 - \theta)a^{h_2}(\cdot, o)$. We wish to show that this map defines an affine isomorphism between \mathcal{H} and \mathcal{P}_o in the sense that it is bijective (in which case its inverse is automatically affine). We begin by constructing the inverse map from \mathcal{P}_o to \mathcal{H} .

Lemma 3.6. *Let $G = (V, E)$ be a infinite, connected, locally finite recurrent graph and let $o \in V$. For each $a \in \mathcal{P}_o$ there exists a unique $h \in \mathcal{H}$ satisfying*

$$h_B(u) = \widehat{\mathbf{P}}_o^a(\widehat{X}^a \text{ visits } B \text{ for the last time at } u)$$

for every finite set B containing o . Moreover, this h satisfies $a^h(x, o) = a(x)$ for every $x \in V$.

Proof of Lemma 3.6. Fix $a \in \mathcal{P}_o$. We define a the family of probability measures $h = (h_B : B \subset V \text{ finite})$ by

$$h_B(u) = \widehat{\mathbf{P}}_o^a(\widehat{X}^a \text{ visits } B \text{ for the last time at } u)$$

for every $u \in B$ if $o \in B$ and

$$h_B(u) = \widehat{\mathbf{P}}_o^a(\widehat{X}^a \text{ visits } B \cup \{o\} \text{ for the last time at } u) \\ + \widehat{\mathbf{P}}_o^a(\widehat{X}^a \text{ visits } B \cup \{o\} \text{ for the last time at } o) \mathbf{P}_o(X_{T_B} = u)$$

for every $u \in B$ if $o \notin B$, so that if $o \notin B$ then

$$h_B(u) = \sum_{v \in B \cup \{o\}} h_{B \cup \{o\}}(v) \mathbf{P}_v(X_{T_B} = u)$$

for every $u \in V$. We claim that this defines an element of \mathcal{H} . It is clear that h_B is a probability measure that is supported on ∂B for each finite set $B \subset V$; we need to verify that it satisfies the consistency property (1). Once it is verified that $h \in \mathcal{H}$, the fact that $a = a^h(\cdot, o)$ follows easily from the definition of a^h together with the identity (10), which together yield that

$$a^h(v, o) = h_{v,o}(v) \mathcal{R}_{\text{eff}}(v \leftrightarrow o) = \frac{\widehat{\mathbf{P}}_o^a(\widehat{X}^a \text{ visits } \{o, v\} \text{ for the last time at } v)}{\deg(o) \mathbf{P}_o(T_v < T_o^+)} \\ = \frac{\widehat{\mathbf{P}}_o^a(\widehat{X}^a \text{ hits } v)}{\deg(o) \mathbf{P}_o(T_v < T_o^+)} = \frac{\deg(o) a(v) \mathbf{P}_o(T_v < T_o^+)}{\deg(o) \mathbf{P}_o(T_v < T_o^+)} = a(v)$$

for each $v \in V$.

We now prove that h satisfies the consistency property (1). We will prove the required identity in the case $o \in B$, the remaining case $o \notin B$ following from this case and the definition. Let $B \subseteq B'$ be finite sets with $o \in B$ and let $(V_n)_{n \geq 1}$ be an exhaustion of V by finite sets such that $B' \subseteq V_n$ for every $n \geq 1$. Writing $V_n^c = V \setminus V_n$ for each $n \geq 1$ and τ_n for the first time the walk visits V_n^c , we have that

$$h_B(u) = \lim_{n \rightarrow \infty} \widehat{\mathbf{P}}_o^a(\widehat{X}[0, \tau_n] \text{ last visits } B \text{ at } u) \\ = \lim_{n \rightarrow \infty} \sum_{b \in V_n^c} \widehat{\mathbf{P}}_o^a(\widehat{X}[0, \tau_n] \text{ last visits } B \text{ at } u, \widehat{X}_{\tau_n} = b)$$

and hence by (10) and time-reversal that

$$h_B(u) = \lim_{n \rightarrow \infty} \sum_{b \in V_n^c} \deg(o) a(b) \mathbf{P}_o(X[0, \tau_n] \text{ last visits } B \text{ at } u, X_{\tau_n} = b) \\ = \lim_{n \rightarrow \infty} \sum_{b \in V_n^c} \deg(b) a(b) \mathbf{P}_b(X_{T_B} = u, T_o < T_{V_n^c}^+). \quad (12)$$

It follows from this together with the strong Markov property that

$$h_B(u) = \lim_{n \rightarrow \infty} \sum_{v \in B'} \sum_{b \in V_n^c} \deg(b) a(b) \mathbf{P}_b(X_{T_{B'}} = v, X_{T_B} = u, T_o < T_{V_n^c}^+) \\ = \lim_{n \rightarrow \infty} \sum_{v \in B'} \sum_{b \in V_n^c} \deg(b) a(b) \mathbf{P}_b(X_{T_{B'}} = v, T_{B'} < T_{V_n^c}^+) \mathbf{P}_v(X_{T_B} = u, T_o < T_{V_n^c}^+).$$

Now, we have by the strong Markov property that for each $b \in V_n^c$ and $v \in B'$

$$\mathbf{P}_b(X_{T_{B'}} = v, T_o < T_{V_n^c}^+) = \mathbf{P}_b(X_{T_{B'}} = v, T_{B'} < T_{V_n^c}^+) \mathbf{P}_v(T_{V_n^c} > T_o).$$

and by recurrence that $\lim_{n \rightarrow \infty} \mathbf{P}_v(T_o < T_{V_n^c}^+) = 1$, so that

$$h_B(u) = \lim_{n \rightarrow \infty} \sum_{v \in B'} \sum_{b \in V_n^c} \deg(b) a(b) \mathbf{P}_b(X_{T_{B'}} = v, T_o < T_{V_n^c}^+) \mathbf{P}_v(X_{T_B} = u).$$

The claimed identity (1) follows from this together with the identity (12) applied to the larger set B' . \square

Theorem 3.7. *Let G be an infinite, recurrent, locally finite graph, and let $o \in V$. The map $h \mapsto a^h(\cdot, o)$ is an affine isomorphism $\mathcal{H} \rightarrow \mathcal{P}_o$. In particular, this map identifies extremal elements of \mathcal{H} with extremal elements of \mathcal{P}_o .*

Proof. It remains only to prove that $h \mapsto a^h$ is injective. To prove this it suffices by definition of a^h to prove that h_B is determined by $(h_{x,o}(x) : x \in \partial B)$ for each finite set $B \subset V$ containing the vertex o . Fix one such set B . We have by definition of \mathcal{H} that

$$h_{x,o}(x) = \sum_{y \in \partial B} h_B(y) \mathbf{P}_y(T_x < T_o) = \sum_{y \in \partial B} A(x, y) h_B(y)$$

for each $x \in \partial B$ where $A(x, y) := \mathbf{P}_y(T_x < T_o)$ for each $x, y \in \partial B$, so that it suffices to prove that the matrix A (which is indexed by ∂B) is invertible. Define a matrix Q indexed by ∂B by

$$Q(x, y) = \mathbf{P}_y(T_{\partial B}^+ < T_o, X_{T_{\partial B}^+} = x).$$

Then we have by the strong Markov property that

$$\begin{aligned} A(x, y) - \mathbf{1}(x = y) \mathbf{P}_x(T_x^+ \geq T_o) &= \mathbf{P}_y(T_x^+ < T_o) \\ &= \sum_{z \in \partial B} \mathbf{P}_z(T_x < T_o) Q(z, y) = \mathbf{P}_x(T_x^+ \geq T_o) Q(x, y) + \sum_{z \in \partial B} \mathbf{P}_z(T_x^+ < T_o) Q(z, y) \end{aligned}$$

and hence inductively that

$$\mathbf{P}_y(T_x^+ < T_o) = \mathbf{P}_x(T_x^+ \geq T_o) \sum_{i=1}^n Q^n(x, y) + \sum_{z \in \partial B} \mathbf{P}_z(T_x^+ < T_o) Q^n(z, y)$$

for every $n \geq 1$. Since Q is irreducible and substochastic, we can take the limit as $n \rightarrow \infty$ to obtain that

$$A(x, y) = \mathbf{1}(x = y) \mathbf{P}_x(T_x^+ \geq T_o) + \mathbf{P}_y(T_x^+ < T_o) = \mathbf{P}_x(T_x^+ \geq T_o) \sum_{i=0}^{\infty} Q^i(x, y)$$

for every $x, y \in \partial B$. It follows by a standard argument that the matrix A is invertible with inverse $A^{-1} = \mathbf{P}_x(T_x^+ \geq T_o)^{-1} (1 - Q)$ as required. \square

3.2.4 The Liouville property for extremal Doob transforms

In this section we prove a kind of tail-triviality property of the Doob-transformed walk corresponding to an extremal point $h \in \mathcal{H}$. Letting $G = (V, E)$ be a graph, we recall that an event $A \subseteq V^{\mathbb{N}}$ is said to be *invariant* if $(x_0, x_1, \dots) \in A$ implies that $(x_1, x_2, \dots) \in A$ for every $(x_0, x_1, \dots) \in V^{\mathbb{N}}$.

Theorem 3.8. *Let $G = (V, E)$ be an infinite, connected, recurrent, locally finite graph and let $o \in V$. If $h \in \mathcal{H}$ is extremal then the Doob transformed random walk \hat{X}^h does not have any non-trivial invariant events: If $A \subseteq V^{\mathbb{N}}$ is an invariant event then $\hat{\mathbf{P}}_o^h(A) \in \{0, 1\}$.*

Proof. It suffices to prove the corresponding statement for \hat{X}^a when a is an extremal element of \mathcal{P}_o . Suppose not, so that A is a non-trivial invariant event. We have by Levy's 0-1 law that

$$\mathbf{P}_o(\hat{X}^a \in A \mid \hat{X}_1^a, \dots, \hat{X}_n^a) \rightarrow \mathbf{1}(\hat{X}^a \in A) \text{ almost surely as } n \rightarrow \infty. \quad (13)$$

Moreover, we also have by invariance that

$$\hat{\mathbf{P}}_x^a(\hat{X}^a \in A) = \sum_{y \in V} \frac{a(y)}{a(x)} p(x, y) \hat{\mathbf{P}}_y^a(\hat{X}^a \in A)$$

and that

$$\hat{\mathbf{P}}_o^a(\hat{X}^a \in A) = \sum_{y \in V} a(y) \hat{\mathbf{P}}_y^a(\hat{X}^a \in A).$$

Since similar inequalities hold when we replace A by A^c it follows that we can write a as a non-trivial convex combination of two elements of \mathcal{P}_o

$$a(x) = \hat{\mathbf{P}}_o^a(\hat{X}^a \in A) \cdot \frac{a(x) \hat{\mathbf{P}}_x^a(\hat{X}^a \in A)}{\hat{\mathbf{P}}_o^a(\hat{X}^a \in A)} + \hat{\mathbf{P}}_o^a(\hat{X}^a \notin A) \cdot \frac{a(x) \hat{\mathbf{P}}_x^a(\hat{X}^a \notin A)}{\hat{\mathbf{P}}_o^a(\hat{X}^a \notin A)},$$

these two factors being different by (13), contradicting extremality of a . \square

Remark 3.2. Underlying this proposition is the fact that once we fix $a \in \mathcal{P}_o$, we can identify \mathcal{P}_o with the Martin boundary of the conditioned walk \hat{X}^a . Theorem 3.8 is the recurrent version of the fact that Doob transforming by an extremal element of the Martin boundary yields a process with trivial invariant sigma-algebra.

For our purposes, the most important output of the Liouville property is the following proposition, which lets us easily tell apart the trajectories of two different Doob transformed walks \hat{X}^h and $\hat{X}^{h'}$ by looking at any infinite subset of their traces (and, in particular, from their loop-erasures).

Proposition 3.9. *Let h, h' be distinct extremal elements of \mathcal{H} and let \widehat{X}^h be the Doob-transformed simple random walk corresponding to h . Then*

$$\frac{a^{h'}(\widehat{X}_n^h, o)}{a^h(\widehat{X}_n^h, o)} \rightarrow 0$$

almost surely as $n \rightarrow \infty$.

Proof. We prove the corresponding statement in which a, a' are distinct extremal elements of \mathcal{P}_o . Let \widehat{X} and \widehat{X}' have laws $\widehat{\mathbf{P}}_o^a$ and $\widehat{\mathbf{P}}_o^{a'}$ respectively. One can easily verify from the definitions that

$$(Z_n)_{n \geq 1} = \left(\frac{a'(\widehat{X}_n)}{a(\widehat{X}_n)} \right)_{n \geq 1} \quad \text{and} \quad (Z'_n)_{n \geq 1} = \left(\frac{a(\widehat{X}'_n)}{a'(\widehat{X}'_n)} \right)_{n \geq 1}$$

are both non-negative martingales with respect to their natural filtrations, and hence converge almost surely to some limiting random variables Z and Z' . Since Z and Z' are measurable with respect to the invariant σ -algebras of \widehat{X} and \widehat{X}' respectively and a and a' are both extremal, there must exist constants α and α' such that $Z = \alpha$ and $Z' = \alpha'$ almost surely. We also have that $\mathbb{E}Z_n = \mathbb{E}Z'_n = 1$ for every $n \geq 1$ and hence that $\alpha, \alpha' \leq 1$. We wish to prove that $\alpha = 0$.

It follows from (10) that the conditional distributions of the initial segments $\widehat{X}[0, m]$ and $\widehat{X}'[0, m]$ are the same if we condition on $\widehat{X}_m = \widehat{X}'_m = v$ for any $v \in V$ for any $v \in V$ and $m \geq 1$ and that

$$\frac{\mathbf{P}_o^a(\widehat{X}_m = v)}{\mathbf{P}_o^{a'}(\widehat{X}'_m = v)} = \frac{a(v)}{a'(v)}$$

for every $m \geq 1$ and $v \in V$. If $\alpha > 0$ then for every $\varepsilon > 0$ there exists M such that the distribution of \widehat{X}_m puts mass at least $1 - \varepsilon$ on the set of vertices with $a'(v)/a(v) \geq (1 - \varepsilon)\alpha$ for every $m \geq M$, and it follows that for each $m \geq M$ there is a coupling of the two walks \widehat{X}' and \widehat{X} so that their initial segments of length m coincide with probability at least $(1 - \varepsilon)^2\alpha$. Taking a weak limit as $m \rightarrow \infty$ and $\varepsilon \rightarrow 0$, it follows that there exists a coupling of the two walks \widehat{X}' and \widehat{X} such that the two walks coincide forever with probability at least $\alpha > 0$. If we couple the walks in this way then on this event we must have that $Z' = 1/Z$, which can occur with positive probability only if $\alpha' = 1/\alpha$. Since $\alpha, \alpha' \leq 1$ we must have that $\alpha = \alpha' = 1$ and that we can couple the two walks to be exactly the same almost surely. This is clearly only possible if $a = a'$, and since $a \neq a'$ by assumption we must have that $\alpha = 0$. \square

3.2.5 Potential kernels and the uniform spanning tree

We now use Lemma 3.11 to show that the UST of a recurrent graph can always be sampled using a variant of Wilson's algorithm [32, 171] in which we 'root at a point in

\mathcal{H} , where again we are thinking intuitively of \mathcal{H} as a kind of boundary at infinity of the graph. Fix $h \in \text{ex}(\mathcal{H})$ and let \widehat{X}^h be the conditioned walk of the previous section. Fix some enumeration $V = \{v_1, v_2, \dots\}$ of V with $v_1 = o$. Set $E_0 = \text{LE}(\widehat{X}^h[0, \infty))$ (which is well defined because \widehat{X}^h is transient) and for each $i \geq 1$ define E_i given E_{i-1} recursively as follows:

- if $v_i \in E_{i-1}$, set $E_i = E_{i-1}$
- otherwise, set $E_i = E_{i-1} \cup \text{LE}(Y[0, \tau))$ where Y is the simple random walk started at v_i and stopped at τ , the hitting time of E_{i-1} .

Last, define $T = \bigcup_{i=0}^{\infty} E_i$. We refer to this procedure as **Wilson's algorithm rooted at h** . The random tree T generated by Wilson's algorithm rooted at h is clearly a spanning tree of G ; the next lemma shows that it is distributed as the UST of G .

Lemma 3.10 (Wilson meets Doob). *Let $G = (V, E)$ be an infinite, connected, locally finite, recurrent graph and let $h \in \text{ext}(\mathcal{H})$. The tree T generated by Wilson's algorithm rooted at h is distributed as the uniform spanning tree of G . In particular, the law of T is independent of the chosen enumeration of V and the choice of $h \in \text{ext}(\mathcal{H})$.*

Remark 3.3. It follows by taking convex combinations that the same statement also holds when h is *not* extremal.

We will deduce Lemma 3.10 from the following lemma, which allows us to think of the Doob-transformed walk \widehat{X}^h as a limit of conditioned simple random walks on G . For the purposes of this lemma we think of our walks as belonging to the space of sequences in V equipped with the product topology.

Lemma 3.11 (Local convergence). *Let $G = (V, E)$ be an infinite, connected, locally finite, recurrent graph and suppose that z_n is a sequence of vertices of G such that z_n converges to $h \in \mathcal{H}$. If X denotes the random walk on G started at o and \widehat{X}^h denotes the Doob-transformation of X as above, then the conditional law of X given that it hits z_n before first returning to o converges weakly to the law of \widehat{X}^h .*

Proof of Lemma 3.11. This is a classical result concerning Doob transforms, and can also be deduced from the limit formula (8). We give a brief proof. Let T_{z_n} be the first time the walk hits z_n , let T_o^+ be the first positive time the walk hits o , and let $\varphi = (o, \varphi_1, \dots, \varphi_m)$ be a path of length m for some $m \geq 1$ with $\varphi_i \neq o$ for every $i > 0$. By the Markov property for the simple random walk,

$$\mathbf{P}_o(X[0, m] = \varphi, T_{z_n} < T_o^+) = \mathbf{P}_o(X[0, m] = \varphi) \mathbf{P}_{\varphi_m}(T_{z_n} < T_o),$$

and it follows from (10) that

$$\mathbf{P}_o(X[0, m] = \varphi, T_{z_n} < T_o^+) = \frac{1}{\deg(o) a^h(\varphi_m, o)} \mathbf{P}_o(\widehat{X}^h[0, m] = \varphi) \mathbf{P}_{\varphi_m}(T_{z_n} < T_o).$$

The result follows once multiplying both sides by the effective resistance between o and z_n and using the representation (6) for the potential kernel. \square

Proof of Lemma 3.10. The standard implementation of Wilson's algorithm rooted at z_n allows us to sample the uniform spanning tree of G in a manner exactly analogous to above, except that we start with a walk run from o until it first hits z_n . Now, it is a combinatorial fact that the *loop erasure* of the walk run from o until it first hits z_n does not change its distribution if we condition the walk to hit z_n before returning to o : Indeed, the loop-erasure of the entire unconditioned walk is equal to the loop-erasure of the final segment of the walk between its last visit to o and its first visit to z_n , and this final segment is distributed as the conditioned walk. Thus, in the standard implementation of Wilson's algorithm, we do not change the distribution of the obtained tree if we condition the first walk to hit z_n before returning to o . The claim then follows by taking the limit as $z_n \rightarrow \infty$ and using Lemma 3.11. \square

This leads to the following connection between the ends of the UST and the extremal points of the set of harmonic measures from infinity \mathcal{H} .

Proposition 3.12. *Let $G = (V, E)$ be an infinite, connected, locally finite, recurrent graph, let T be the uniform spanning tree of T , and let H be a countable subset of $\text{ext}(\mathcal{H})$. Almost surely, for each $h \in H$ there exists an infinite simple path $\Gamma = (\Gamma_1, \Gamma_2, \dots)$ in T such that*

$$\frac{a^{h'}(\Gamma_i, x)}{a^h(\Gamma_i, x)} \rightarrow 0 \quad \text{as } i \rightarrow \infty \text{ for each } h' \in H \setminus \{h\}.$$

In particular, T almost surely has at least as many ends as there are extremal points of \mathcal{H} .

(In the last sentence of this proposition we are not distinguishing between different infinite cardinalities, but merely claiming that if \mathcal{H} has infinitely many extremal points then T has infinitely many ends almost surely.)

Proof. This is an immediate consequence of Proposition 3.9 and Lemma 3.10. \square

Orientations. Let $G = (V, E)$ be an infinite, connected, locally finite, recurrent graph and let $h \in \text{ext}(\mathcal{H})$. When we generate the UST T of G using Wilson's algorithm rooted at h , the algorithm also provides a natural *orientation* of T , where each edge is oriented in the direction that it is crossed by the loop-erased random walk that contributed that edge to the tree. When T almost surely has the same number of ends as there are extremal points in \mathcal{H} , and both numbers are finite (which will always be the case in the unimodular setting by the results of [34]), it follows from Proposition 3.12 that this

orientation is a.s. determined by the (unoriented tree): Almost surely, for each $h \in \mathcal{H}$ and $v \in V$ there is a unique infinite ray $(\Gamma_1, \Gamma_2, \dots)$ starting from v such that

$$\frac{a^{h'}(\Gamma_i, v)}{a^h(\Gamma_i, v)} \rightarrow 0 \quad \text{as } i \rightarrow \infty \text{ for each } h' \in \text{ext}(\mathcal{H}) \setminus \{h\},$$

and if we orient the tree in the direction of this ray we must recover the same orientation as if we had generated the oriented tree using Wilson's algorithm rooted at h . This fact will play a key role in the proof of our main theorem.

3.3. Proof of the main theorem

3.3.1 Reversible and unimodular graphs

We now give a very brief introduction to unimodular random rooted graphs, referring the reader to [9, 49] for detailed introductions. Let us just recall that $\mathcal{G}_{\bullet, \bullet}$ is the separable metric space of doubly rooted graphs (G, x, y) (modulo graph isomorphisms), equipped with the *local topology*, also known as *Benjamini-Schramm topology*. Similarly defined is the space \mathcal{G}_{\bullet} of rooted graphs (G, o) . A *mass transport* is a measurable function $f : \mathcal{G}_{\bullet, \bullet} \rightarrow [0, \infty]$. A measure \mathbb{P} on \mathcal{G}_{\bullet} is called *unimodular* whenever the *mass transport principle*

$$\widehat{\mathbb{E}} \left[\sum_{x \in V} f(G, o, x) \right] = \widehat{\mathbb{E}} \left[\sum_{x \in V} f(G, x, o) \right]$$

holds for all mass transports f . A probability measure \mathbb{P} on \mathcal{G}_{\bullet} is called *reversible* if $(G, o, X_1) \stackrel{d}{=} (G, X_1, o)$ where X_1 is the first step of the simple random walk. The law \mathbb{P} is called *stationary* if $(G, o) \stackrel{d}{=} (G, X_1)$ and clearly any reversible graph is stationary. For recurrent graphs, stationarity and reversibility are equivalent [27].

If \mathbb{P} is the law of a unimodular random graph, with finite expected degree, then biasing it by $\deg(o)$ gives a reversible random graph and whenever \mathbb{P} is the law of a reversible random graph, then biasing by $\deg(o)^{-1}$ gives a unimodular random graph; see for example [27].

A set $A \subseteq \mathcal{G}_{\bullet}$ is said to be **rerooting invariant** if $((g, v) \in A) \Rightarrow ((g, u) \in A)$ for every rooted graph $(g, v) \in \mathcal{G}_{\bullet}$ and every u in the vertex set of g . A unimodular random rooted graph (G, o) is said to be **ergodic** if it has probability 0 or 1 to belong to any given re-rooting invariant event in \mathcal{G}_{\bullet} . As explain in [9, Section 4], this is equivalent to the law of (G, o) being extremal in the weakly compact convex set of unimodular probability measures on \mathcal{G}_{\bullet} . As such, it follows by Choquet theory that every unimodular measure on \mathcal{G}_{\bullet} may be written as a mixture of ergodic unimodular measures. For our purposes, the upshot of this is that we may assume without loss of generality that (G, o) is ergodic when proving Theorem 3.1.

We will also rely on the following characterization of two-ended unimodular random rooted graphs due to Bowen, Kun, and Sabok [40], which builds on work of Benjamini and the second author [31]. Here, a graph G is said to have **linear volume growth** if for each vertex v of G there exists a constant C_v such that $|B(v, r)| \leq C_v r$ for every $r \geq 1$, where $B(v, r)$ denotes the graph distance ball of radius r around v .

Proposition 3.13 ([40], Proposition 2.1). *Let (G, o) be an infinite unimodular random rooted graph. Then G is two-ended almost surely if and only if it has linear volume growth almost surely.*

To prove Theorem 3.1, it will therefore suffice to prove that if (G, o) is a recurrent unimodular random rooted graph whose UST is two-ended almost surely then G has linear volume growth almost surely.

3.3.2 The effective resistance is linear on the spine

Let \mathbb{P} be the joint law of an ergodic recurrent unimodular random rooted graph (G, o) and its uniform spanning tree T , which we think of as a triple (G, o, T) . It follows by tail triviality of the UST [32, Theorem 8.3] that the number of ends of T is deterministic conditional on (G, o) , and since (G, o) is ergodic that T has some constant number of ends almost surely. Moreover, it follows from [9, Theorem 6.2 and Proposition 7.1] that this number of ends is either 1 or 2 almost surely, so that T is either one-ended almost surely or two-ended almost surely.

We wish to prove that if T is two-ended almost surely then G is two-ended almost surely also. We will rely on the following theorem of Berestycki and the first author.

Theorem 3.14 ([34], Theorem 1). *Let (G, o) be a recurrent unimodular random rooted graph with $\mathbb{E} \deg(o) < \infty$. Almost surely, the uniform spanning tree of G is one-ended if and only if the harmonic measure from infinity is uniquely defined.*

To avoid the unnecessary assumption that $\mathbb{E} \deg(o) < \infty$, we will use the following mild generalization of this theorem, whose proof is given in Appendix A.

Theorem 3.15. *Let (G, o) be a recurrent unimodular random rooted graph. Almost surely, the uniform spanning tree of G is one-ended if and only if the harmonic measure from infinity is uniquely defined.*

It follows from this theorem together with Proposition 3.12 that if T is two-ended almost surely then $|\text{ext}(\mathcal{H})| = 2$ almost surely.

Suppose that T is two-ended almost surely and let \mathcal{S} be the **spine** of T , i.e., the unique double-infinite simple path contained in T . We give T an orientation by choosing uniformly at random one of the two ends of \mathcal{S} and directing every edge towards that

end, letting the resulting oriented tree be denoted T^\rightarrow with oriented spine \mathcal{S}^\rightarrow . Since the law of T^\rightarrow is a rerooting-equivariant function of the graph (G, o) , the triple (G, T^\rightarrow, o) is unimodular. Since “everything that can happen somewhere can happen at the root” [9, Lemma 2.3] we also have that the origin belongs to \mathcal{S} with positive probability and hence that we can define a law $\mathbb{P}_{\mathcal{S}}$ on triplets (G, T^\rightarrow, o) (which we can view as a rooted network) by conditioning o to belong to \mathcal{S} . The law $\mathbb{P}_{\mathcal{S}}$ has the very useful property that it is stationary under shifts along the spine, which we now define. Each vertex $v \in \mathcal{S}$ has a unique oriented edge emanating from it in \mathcal{S}^\rightarrow , and we will write $\sigma(v)$ for the vertex on the other end of this edge. The map $v \mapsto \sigma(v)$ can be thought of as a shift, following the orientation along the spine, and there is also a well-defined backwards shift σ^{-1} mapping each $x \in \mathcal{S}$ to the unique vertex $v \in \mathcal{S}$ with $\sigma(v) = x$.

Lemma 3.16. *The law $\mathbb{P}_{\mathcal{S}}$ is invariant under the shift σ .*

Proof. Let A be any Borel set of triples (g, t^\rightarrow, v) where (g, v) is a rooted graph and t^\rightarrow is an oriented spanning tree of g , and define the mass transport

$$f(g, t^\rightarrow, v, w) := \mathbb{1}(t^\rightarrow \text{ is two-ended, } w \text{ is in the spine of } t^\rightarrow, v = \sigma(w), \text{ and } (g, t^\rightarrow, w) \in A).$$

Note that there only exists one vertex v such that $v = \sigma(w)$ and, vice-versa, for each v in the spine of t^\rightarrow there is only one w in the spine of t^\rightarrow such that $\sigma(w) = v$ and $w \in \mathcal{S}$. Therefore,

$$\begin{aligned} & \sum_{v \in V} f(G, T^\rightarrow, v, o) \\ &= \mathbb{1}(T^\rightarrow \text{ is two-ended, } o \text{ is in the spine of } T^\rightarrow, \text{ and } (G, T^\rightarrow, o) \in A) \end{aligned}$$

and

$$\begin{aligned} & \sum_{v \in V} f(G, T^\rightarrow, v, \sigma(o)) \\ &= \mathbb{1}(T^\rightarrow \text{ is two-ended, } \sigma(o) \text{ is in the spine of } T^\rightarrow, \text{ and } (G, T^\rightarrow, \sigma(o)) \in A) \end{aligned}$$

Using the mass-transport principle we thus have that

$$\begin{aligned} & \mathbb{P}(T^\rightarrow \text{ is two-ended, } o \text{ is in the spine of } T^\rightarrow, \text{ and } (G, T^\rightarrow, o) \in A) \\ &= \mathbb{P}(T^\rightarrow \text{ is two-ended, } \sigma(o) \text{ is in the spine of } T^\rightarrow, \text{ and } (G, T^\rightarrow, \sigma(o)) \in A) \end{aligned}$$

which shows the result because $\mathbb{P}(o \in \mathcal{S}) > 0$ and T is two-ended a.s. by assumption. \square

The main goal of this section is to show that along the spine of the UST, the effective resistances on the original graph must grow linearly under the assumption that the UST

has two ends (and thus a well-defined spine). Heuristically, this tells us that if a graph is unimodular and the uniform spanning tree is two-ended, then the actual graph should in some sense be “close” to the line \mathbb{Z} .

Proposition 3.17. *The limit $\lim_{n \rightarrow \infty} \frac{1}{n} \mathcal{R}_{\text{eff}}(o \leftrightarrow \sigma^n(o)) = \lim_{n \rightarrow \infty} \frac{1}{n} \mathcal{R}_{\text{eff}}(o \leftrightarrow \sigma^{-n}(o))$ exists and is positive $\mathbb{P}_{\mathcal{S}}$ -a.s.*

Note that the existence part of this proposition is an immediate consequence of the subadditive ergodic theorem; the content of the proposition is that the limit is positive.

As discussed above, it follows from Proposition 3.12 and Theorem 3.14 that, $\mathbb{P}_{\mathcal{S}}$ -almost surely, there are exactly two extremal elements of \mathcal{H} , which we call “ ℓ ” and “ r ”, which satisfy

$$\frac{a^r(\sigma^n(o), v)}{a^\ell(\sigma^n(o), v)} \rightarrow \begin{cases} \infty & \text{as } n \rightarrow +\infty \\ 0 & \text{as } n \rightarrow -\infty \end{cases} \quad (14)$$

for every $v \in V$. (In particular, the random choice of orientation of T we made when defining $\mathbb{P}_{\mathcal{S}}$ is equivalent to randomly choosing which of the two extremal elements of \mathcal{H} to call “ r ”.) Consider the function $V \rightarrow \mathbb{R}$ defined by

$$M_o(x) := a^r(x, o) - a^\ell(x, o).$$

We will show that $M_o(\sigma^n(o))$ grows linearly in n and deduce from this that the effective resistance does too. The latter fact can be seen using (6), from which it follows that

$$M_o(x) = (r_{x,o}(x) - \ell_{x,o}(x)) \mathcal{R}_{\text{eff}}(o \leftrightarrow x).$$

In the remainder we will slightly abuse notation to write $M_m(n) := M_{\sigma^m(o)}(\sigma^n(o))$ for $n, m \in \mathbb{Z}$. The first main ingredient is that $M_o(n)$ is an *additive cocycle*.

Lemma 3.18. *$M_o(n + m) = M_o(n) + M_n(n + m)$ for every $n, m \in \mathbb{Z}$.*

Proof. This is a direct consequence of Proposition 3.5 in [34], stating that

$$a^\#(x, o) - a^\#(y, o) = a^\#(x, y) - \frac{\mathbf{G}_y(x, o)}{\deg(o)}$$

for each $\# \in \{\ell, r\}$ and all $x, y \in V$. Indeed, it follows from this identity that

$$\begin{aligned} M_o(n + m) - M_o(n) &= a^r(\sigma^{n+m}(o), o) - a^\ell(\sigma^{n+m}(o), o) - a^r(\sigma^n(o), o) + a^\ell(\sigma^n(o), o) \\ &= \left[a^r(\sigma^{n+m}(o), \sigma^n(o)) - \frac{\mathbf{G}_{\sigma^n(o)}(\sigma^{n+m}(o), o)}{\deg(o)} \right] \end{aligned}$$

$$\begin{aligned}
 & - \left[a^\ell(\sigma^{n+m}(o), \sigma^n(o)) - \frac{\mathbf{G}_{\sigma^n(o)}(\sigma^{n+m}(o), o)}{\deg(o)} \right] \\
 & = a^r(\sigma^{n+m}(o), \sigma^n(o)) - a^\ell(\sigma^{n+m}(o), \sigma^n(o)) = M_n(n+m)
 \end{aligned}$$

for every $n, m \in \mathbb{Z}$ as claimed. \square

Let us also make note of the following key property of this additive cocycle.

Lemma 3.19. $\mathbb{P}_{\mathcal{S}}$ -almost surely, $M_o(n)$ is positive for all sufficiently large positive n and negative for all sufficiently large negative n . Moreover,

$$M_o(n) \sim a^r(\sigma^n(o), o) = r_{\sigma^n(o), o}(\sigma^n(o)) \mathcal{R}_{\text{eff}}(o \leftrightarrow \sigma^n(o))$$

$\mathbb{P}_{\mathcal{S}}$ -almost surely as $n \rightarrow \infty$.

Proof. This follows immediately from (14) and the definition of $M_o(n)$. \square

We will deduce Proposition 3.17 from Lemma 3.19 together with the following general fact about stationary sequences.

Proposition 3.20. Let $(Z_i)_{i \in \mathbb{Z}}$ be a stationary sequence of real-valued random variables and suppose that $\sum_{i=0}^n Z_{-i} > 0$ for all sufficiently large n almost surely. Then $\limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^n Z_i > 0$ almost surely.

Proof. For each $n \in \mathbb{Z}$ let $R_n = \inf\{m \geq 0 : \sum_{i=n}^{n+m} Z_i > 0\}$, so that $R_n = 0$ whenever $Z_n > 0$ and $(R_n)_{n \in \mathbb{Z}}$ is a stationary sequence of $\{0, 1, \dots\}$ -valued random variables. It follows from the definitions that if $n \leq m$ then either $n + R_n < m$ or $n + R_n \geq m + R_m$, so that the intervals $[n, n + R_n]$ and $[m, m + R_m]$ are either disjoint or ordered by inclusion. On the other hand, we have by stationarity and the hypotheses of the Proposition that for each $n \in \mathbb{Z}$ there almost surely exists $N_n < \infty$ such that $\sum_{i=n-m}^{n-1} Z_i > 0$ for every $m \geq N_n$ and hence that $R_{n-m} + (n - m) < n$ for every $m \geq N_n$, so that each $n \in \mathbb{Z}$ is contained in at most finitely many of the intervals $[m, m + R_m]$ almost surely. Using the fact that these intervals are either disjoint or ordered by inclusion, it follows that there is a unique decomposition of \mathbb{Z} into maximal intervals of this form

$$\mathbb{Z} = \bigcup \left\{ [k, k + R_k] : k \in \mathbb{Z}, [k, k + R_k] \not\subseteq [m, m + R_m] \text{ for every } m \in \mathbb{Z} \setminus \{k\} \right\}.$$

Thus, if we define Y_n by

$$Y_n = \begin{cases} \sum_{i=k}^n Z_i & n = k + R_k \text{ for some } k \in \mathbb{Z} \text{ such that } [k, k + R_k] \text{ maximal} \\ 0 & \text{otherwise} \end{cases}$$

then $(Y_n)_{n \in \mathbb{Z}}$ is a stationary sequence of non-negative random variables such that Y_n is positive whenever n is the right endpoint of a maximal interval. Since Y_n is non-negative

and the set of n such that $Y_n \neq 0$ is almost surely non-empty, it follows from the ergodic theorem applied to $(\min\{Y_n, 1\})_{n \in \mathbb{Z}}$ that

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^n Y_n > 0$$

almost surely. The claim follows since if $-m$ is the left endpoint of the maximal interval containing 0 then

$$\sum_{i=0}^n Y_n = \sum_{i=-m}^n Z_i$$

for every n that is the right endpoint of some maximal interval. \square

Proof. It follows from Lemma 3.16 that $(M_n(n+1))_{n \in \mathbb{Z}}$ is a stationary sequence under $\mathbb{P}_{\mathcal{S}}$ and from Lemma 3.18 that $M_o(n) = \sum_{i=0}^{n-1} M_i(i+1)$ for every $n \geq 0$ and $M_o(-n) = \sum_{i=-n}^{-1} M_i(i+1)$ for every $n \leq 0$. Thus, Lemma 3.19 implies that the stationary sequence $(M_n(n+1))_{n \in \mathbb{Z}}$ satisfies the hypotheses of Proposition 3.20 and hence that

$$\limsup_{n \rightarrow \infty} \frac{M_o(n)}{n} > 0$$

almost surely. On the other hand, the subadditive ergodic theorem implies that the limit $\lim_{n \rightarrow \infty} \frac{1}{n} \mathcal{R}_{\text{eff}}(o \leftrightarrow \sigma^n(o))$ exists $\mathbb{P}_{\mathcal{S}}$ -a.s., and since

$$M_o(n) = (r_{\sigma^n(o), o}(\sigma^n(o)) - \ell_{\sigma^n(o), o}(\sigma^n(o))) \mathcal{R}_{\text{eff}}(o \leftrightarrow \sigma^n(o)) \leq \mathcal{R}_{\text{eff}}(o \leftrightarrow \sigma^n(o))$$

we must have that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathcal{R}_{\text{eff}}(o \leftrightarrow \sigma^n(o)) > 0$$

$\mathbb{P}_{\mathcal{S}}$ -a.s. as claimed. The fact that the negative- n limit $\lim_{n \rightarrow \infty} \frac{1}{n} \mathcal{R}_{\text{eff}}(o \leftrightarrow \sigma^{-n}(o))$ also exists and is equal to the positive- n limit a.s. follows from the subadditive ergodic theorem. \square

3.3.3 Completing the proof

We now complete the proof of the main theorem.

Proof of Theorem 3.1. It suffices by Proposition 3.13 to prove that if (G, o) is a recurrent unimodular random rooted graph whose UST is two-ended almost surely then G has linear volume growth almost surely. As before, we write \mathcal{S} for the spine of the oriented UST T^{\rightarrow} , write $\mathbb{P}_{\mathcal{S}}$ for the conditional law of (G, T^{\rightarrow}, o) given that $o \in \mathcal{S}$, and write σ for the shift along the spine as in Lemma 3.16.

For each $x \in V$ let $S(x)$ be an element of \mathcal{S} of minimal graph distance to x , choosing one of the finitely many possibilities uniformly and independently at random for each x

where this point is not unique. Letting $S^{-1}(v) = \{x \in V : S(x) = v\}$ for each $v \in \mathcal{S}$, we have by the mass-transport principle that

$$\begin{aligned} \mathbb{E}_{\mathcal{S}}|S^{-1}(o)| &= \mathbb{E}[|S^{-1}(o)| \mid o \in \mathcal{S}] = \mathbb{P}(o \in \mathcal{S})^{-1} \mathbb{E} \left[\sum_{x \in V} \mathbf{1}(o = S(x)) \right] \\ &= \mathbb{P}(o \in \mathcal{S})^{-1} \mathbb{E} \left[\sum_{x \in V} \mathbf{1}(x = S(o)) \right] \\ &= \mathbb{P}(o \in \mathcal{S})^{-1} < \infty. \end{aligned}$$

We thus have a stationary sequence of random variables $(|S^{-1}(\sigma^i(o))|)_{i \in \mathbb{Z}}$ with uniformly finite mean, and the ergodic theorem implies that

$$\lim_{i \rightarrow \infty} \frac{1}{2n} \sum_{i=-n}^n |S^{-1}(\sigma^i(o))| < \infty \tag{15}$$

almost surely. On the other hand, letting $B(o, r)$ be the graph distance ball of radius r around o for each $r \geq 1$, we have by definition of S^{-1} that

$$B(o, r) \subseteq \bigcup \{S^{-1}(\sigma^n(o)) : n \in \mathbb{Z}, d(o, \sigma^n(o)) \leq 2r\} \tag{16}$$

for each $r \geq 1$. Proposition 3.17 together with the trivial inequality $\mathcal{R}_{\text{eff}}(x \leftrightarrow y) \leq d(x, y)$ imply that there exists a positive constant $c > 0$ such that $d(o, \sigma^n(o)) \geq c|n|$ for all sufficiently large (positive or negative) n almost surely, and together with (15) and (16) this implies that $\limsup_{r \rightarrow \infty} \frac{1}{r} |B(o, r)| < \infty$ almost surely. This completes the proof. \square

Part II.

**Continuous spin models and dual height
functions**

Review of height function models

The purpose of this chapter is to review some relatively recent results that will be used in the remaining three chapters. We focus on integer-valued height functions and aim to explain fundamental results due to Sheffield [157], which state that (translation invariant) ergodic Gibbs measures are extremal. (The precise definitions are given below). We will proceed by presenting a slightly different proof of the beautiful result of delocalization for integer-valued height functions due to Lammers [114]. The basic ideas are the same, although the proof we present is perhaps slightly less magical. The presentation of Section 4.2 is partially based on a forthcoming paper of the author with Marcin Lis. In that paper, the results are formalized and explained in a more general and unified approach.

4.1. Integer-valued height functions

Fix $G = (V, E)$ any finite graph. We wish to describe random, nearest neighbor height functions $h : V \rightarrow \mathbb{Z}$, penalized for having large gradients. To that end, let $\mathcal{V} : \mathbb{Z} \rightarrow \mathbb{R}$ be a symmetric potential function, which we will assume here to be convex over the integers:

$$\mathcal{V}^{(2)}(k) := \mathcal{V}(k+1) - 2\mathcal{V}(k) + \mathcal{V}(k-1) \geq 0.$$

The weight of a height configuration $h : V \rightarrow \mathbb{Z}$ is given by

$$\nu(h) \propto \prod_{e \in E} e^{-\mathcal{V}(dh_e)}, \quad (1)$$

where $dh_{xy} = h(x) - h(y)$ is the discrete gradient along the edge xy . Since \mathcal{V} is symmetric, we can indeed just take the product over the unoriented edges.

The fact that \mathcal{V} is convex corresponds to the heuristic: the larger gradients gets penalized more and more. It also implies that the partition function

$$Z_G := \sum_{h: V \rightarrow \mathbb{Z}} \prod_{e \in E} e^{-\mathcal{V}(dh_e)}$$

is finite, so the above probability measure ν is well-defined.

Extension to infinite graphs. As usual, we extend the definition to infinity graphs using the DLR formalism. Let $\Gamma = (V, E)$ a locally finite, infinite graph. Let $\varphi : \Lambda^c \rightarrow \mathbb{Z}$ be a function and define the probability measure μ_Λ^φ supported on $h : \mathcal{V} \rightarrow \mathbb{Z}$ satisfying $h|_{\Lambda^c} = \varphi$ by

$$\nu_\Lambda^\varphi(h) \propto \exp\left(-\sum_{e \in E} \mathcal{V}(dh_e)\right).$$

If μ is a probability measure on functions $h : V \rightarrow \mathbb{Z}$ which satisfies

$$\nu_\Lambda(\cdot) = \int_{\mathbb{Z}^V} \nu_\Lambda^\varphi(\cdot) d\nu(\varphi),$$

for all finite subset Λ , we call it a Gibbs measure. Here, μ_Λ denotes the restriction of μ to Λ . As such, we can view ν as an extension of the measures defined in (1).

If Γ is a vertex-transitive graph and the measure ν is invariant under shifts, it is called translation invariant.

4.2. Ergodicity and extremality

Let $\Gamma = (V, E)$ be an infinite graph invariant under some lattice action and fix a symmetric, convex potential \mathcal{V} . Denote by \mathcal{G} the collection of all Gibbs measure corresponding to this potential.

Extremal measures. The set \mathcal{G} is convex and its extreme points are called *extremal* measures. These measures are also characterized by being “tail trivial” with respect to the tail σ -algebra generated by

$$\mathcal{I} := \bigcap_{\Lambda \uparrow V} \mathcal{F}_{\Lambda^c},$$

where the intersection is taken over all finite subsets Λ of V and \mathcal{F}_A is the σ -algebra generated by $\{h_x : x \in A\}$.

An extremely useful consequence of this definition is that the backward martingale convergence theorem can be applied. For any local observable $F : V \rightarrow \mathbb{R}$,

$$\mu[F | \mathcal{F}_{\Lambda_n^c}]$$

is a martingale for any sequence $\Lambda_n \uparrow V$ and any (Gibbs) measure μ . If μ is extremal, backwards martingale convergence implies $\mu[F | \mathcal{F}_{\Lambda_n^c}] \rightarrow \mu[F]$ almost surely.

It turns out that any Gibbs measure can be written as a convex combination of extremal Gibbs measures, see for example [78, Proposition 7.22]. The proofs are general, quite standard and based on some Choquet-like theory.

Ergodic measures. Since Γ is assumed to be invariant under some lattice action, it is natural to consider the subset of Gibbs measures which are invariant under the same lattice action, denoted here by \mathcal{G}_T . We will call such measures *translation invariant* Gibbs measures.

Again, \mathcal{G}_T is a convex set and its extremal points are called *ergodic measures*. These measures are also linked to some tail σ -algebra, but this time the one induced by translations. Denote thus by \mathcal{T} the σ -algebra of events which are invariant under the translations. A measure $\mu \in \mathcal{G}_T$ is ergodic precisely when it is trivial on \mathcal{T} .

Using again relatively general machinery it follows that each translation invariant Gibbs measure is a convex combination of ergodic measures [78]. Of course, translation invariant measures are useful more or less because of this reason: they allow to apply ergodic theory.

4.2.1 The main result: ergodic measures are extremal

Any translation invariant *extremal measure* is also ergodic. Indeed, there are multiple ways to show this, the most standard perhaps is via the “mixing property”, but we will not recall such notions here. But what about ergodic Gibbs measures? Are they also extremal?

The answer turns out to be: yes (in the context of integer-valued height functions as described here). This fact is instrumental in the proof of delocalization below.

Theorem 4.1. *Fix Γ an infinite amenable graph, invariant under some translation, and \mathcal{V} a convex, symmetric potential. If μ is an ergodic Gibbs measure, then μ is extremal.*

This result was proved by Sheffield [157]. The goal of the remainder of this section is to briefly sketch a simplified version of his proof, a full version will appear later.

4.2.2 In case of the Gaussian free field

Imagine for a moment that we are looking at the Gaussian free field (GFF) on a graph Γ , and assume that there exists some ergodic Gibbs measure μ corresponding to the potential $\mathcal{V}(x) = |x|^2$, now defined on \mathbb{R} . Is μ extremal?

Although the method described here does not extend to the general setup, it nonetheless provides many of the heuristics needed. It is well known that the GFF can be extended to the *cable graph*¹, an idea first coined by Lupu [129].

Let (h, \tilde{h}) be sampled from the product measure $\mu \times \nu$ (extended to the cable graph), where both μ and ν are some translation invariant Gibbs measures. Define $\eta : E \rightarrow \{0, 1\}$

¹This is the 1-dimensional CW-complex associated to Γ , i.e. each edge $e \in E$ is replaced by a line segment $[0, 1]$ and the segments are glued together at the vertices.




Figure 4.1: Depicted is the extension of $h - \tilde{h}$ to a single edge. Left: the signs of $h - \tilde{h}$ are different on both end-points, forcing the bridge to touch zero. Middle: bridge doesn't touch zero. Right: endpoints have the same sign, yet bridge still touches zero.

a percolation process as follows: η_e is closed if and only if $h - \tilde{h}$ touches zero somewhere on the edge.

If x is *not* connected to infinity in the percolation η , then any path from x to infinity must cross through a point where $h = \tilde{h}$. By the domain Markov property, this implies that exchanging h and \tilde{h} on the finite open clusters of η , does not change the marginals. “Switching h and \tilde{h} ” corresponds to reflecting $h - \tilde{h}$ through 0. Therefore, if all clusters of η are finite, then $\mu = \nu$.

How can we use this to say something about the extremality of μ ? We follow the argument in [45]. If we can say that $\mu \times \mu$, there is no infinite cluster almost surely, then we can conclude that μ is extremal. Indeed, write μ as some convex combination $\lambda_1 \mu_1 + \lambda_2 \mu_2$, for Gibbs measures μ_1, μ_2 . But then η would also not have any infinite cluster almost surely under $\mu_1 \times \mu_2$, so that $\mu_1 = \mu_2$, implying μ is extremal.

Thus, to prove Theorem 4.1 it suffices to rule out that η percolates if $\mu = \nu$ and μ is ergodic. For the GFF, this can be achieved using relatively elementary tools, because $h - \tilde{h}$ is again a GFF.

4.2.3 A percolation of sign clusters

The first question is thus: how to generalize the percolation of the GFF to general height functions. Of course, not every model allows for an extension to the cable graph, so it may seem hopeless. But the clusters η for the Gaussian free field can also be sampled without the cable graph representation. To that end, sample h and \tilde{h} under $\mu \times \nu$ on V as usual. Define $\eta : E \rightarrow \{0, 1\}$ by declaring it *closed* on the edge $e = (x, y)$ if $(h_x - \tilde{h}_x)(h_y - \tilde{h}_y) < 0$, and otherwise close it with probability

$$\exp(-\beta(h_x - \tilde{h}_x)(h_y - \tilde{h}_y)).$$

Since the latter is exactly the probability that a Brownian bridge between $h_x - \tilde{h}_x$ and $h_y - \tilde{h}_y$ touches 0, the percolation process η defined here is the same as the one above, see again Figure 4.1. The upshot is that this definition does not rely on the cable graph extension, and therefore may look more promising.

General definition. We go back to the setting we set out to describe at the beginning of this chapter. Let $\mathcal{V} : \mathbb{Z} \rightarrow \mathbb{R}$ a convex, symmetric potential. We will compare for any edge $e = (x, y)$ the two weights

$$e^{-\mathcal{V}(h_x - h_y) - \mathcal{V}(\tilde{h}_x - \tilde{h}_y)} \quad \text{and} \quad e^{-\mathcal{V}(\tilde{h}_x - h_y) - \mathcal{V}(h_x - \tilde{h}_y)}, \quad (2)$$

in other words: we compare the energy cost obtained from “switching the roles of h and \tilde{h} at x at the edge”. The question is: can we reduce the energy cost on the edge by applying this switch (and thus increase the weight). Note that the above definition is symmetric in exchanging the roles of x and y .

With this in mind, we will define a general percolation η in the same spirit as the one for the GFF. Again conditional on (h, \tilde{h}) , let $(m_e)_e$ and $(M_e)_e$ be the minimal respectively maximal values of (2). For all edges $e = (x, y)$, do the following. If the actual weight along an edge is the minimal one:

$$e^{-\mathcal{V}(h_x - h_y) - \mathcal{V}(\tilde{h}_x - \tilde{h}_y)} = m_{x,y},$$

then we can reduce the energy by switching h_x and \tilde{h}_x (or h_y and \tilde{h}_y). Such an edge will be called *excited*. For the percolation, we set $\eta_{x,y} = 0$ on excited edges.

If the actual weight equals the maximum in (2), then we set $\eta_{x,y} = 0$ with probability

$$\frac{m_{x,y}}{M_{x,y}},$$

which corresponds to the *residual energy* on the edge. An easy exercise for the reader is to check that this definition of η agrees with the one for the Gaussian free field given above.

We write $\mathbb{P}_{\mu,\nu}$ for the probability measures on the triplets (h, \tilde{h}, η) defined here, although sometimes we will abuse notation and only write $\mu \times \nu$.

4.2.4 Switching the clusters

Just like in the Gaussian case, this definition is chosen precisely so that the law $\mathbb{P}_{\mu,\nu}$ is invariant under switching h and \tilde{h} on any finite cluster of η . In other words, $\mathbb{P}_{\mu,\nu}$ is invariant under reflecting $h - \tilde{h}$ in 0 on the clusters of η .

To see this, we can restrict to the setting of finite graphs. Consider two configurations (h, \tilde{h}, η) and (g, \tilde{g}, η) (so with the same η), with $|h - \tilde{h}| = |g - \tilde{g}|$ everywhere, and such

that the signs of $h - \tilde{h}$ and $g - \tilde{g}$ are different on some finite clusters of η . The claim is that (h, \tilde{h}, η) has the same weight as (g, \tilde{g}, η) . This can be seen directly once we rewrite the weights in terms of excited and not excited edges:

$$\mathbb{P}_{\mu, \nu}(h, \tilde{h}, \eta) \propto \prod_{e \text{ is excited}} M_e \left(\frac{m_e}{M_e} \right)^{1-\eta_e} \left(1 - \frac{m_e}{M_e} \right)^{\eta_e} \prod_{e \text{ not excited}} m_e.$$

In other words, we deduce:

Lemma 4.2. *The measure $\mathbb{P}_{\mu, \nu}$ is invariant under switching h and \tilde{h} on any collection of finite clusters of η .*

Convexity and sign clusters. The explanation above works in a large generality, as will be explained in the forthcoming paper with Lis. In particular, so far, we have not used the convexity of the potential \mathcal{V} . At this point, a word of warning may be needed: if \mathcal{V} is *not* convex, then it is generally *not* true that the sign of $h - \tilde{h}$ is fixed on the clusters of η .

When the potential *is* convex, it turns out that the construction above percolation η correspond to a type of FK representation for the sign clusters of $h - \tilde{h}$. This is highlighted in the following result.

Lemma 4.3. *If x is connected to y in η , then $(h_x - \tilde{h}_x)(h_y - \tilde{h}_y) > 0$.*

Proof. It is enough to prove the result for $e = (x, y)$ neighbors in the graph. If $\eta_e = 1$, then we need in particular that

$$e^{-\mathcal{V}(h_x - h_y) - \mathcal{V}(\tilde{h}_x - \tilde{h}_y)} > e^{-\mathcal{V}(\tilde{h}_x - h_y) - \mathcal{V}(h_x - \tilde{h}_y)}. \quad (3)$$

Write $k = \tilde{h}_x - h_x$. It is easy to see that if $h = \tilde{h}$ at either endpoint x, y , then the above cannot hold, so we can assume without loss of generality that $k > 0$ and prove this implies $\tilde{h}_y - h_y > 0$. Equation (3) is equivalent to

$$\mathcal{V}(d\tilde{h}_e) - \mathcal{V}(d\tilde{h}_e - k) < \mathcal{V}(dh_e + k) - \mathcal{V}(dh_e).$$

Convexity of \mathcal{V} implies therefore that $d\tilde{h}_e < dh_e + k$, showing that indeed $\tilde{h}_y - h_y > 0$. \square

Thus, on the clusters of η the sign of $h - \tilde{h}$ stays constant. Moreover, on the finite clusters of η , switching the sign of $h - \tilde{h}$ leaves the law $\mathbb{P}_{\mu, \nu}$ constant. In other words, conditional on $|h - \tilde{h}|$ and η , the signs of the (finite) clusters are independent uniform random variables on $\{-1, +1\}$.

4.2.5 Finalizing the proof of Theorem 4.1

Let Γ some amenable, infinite graph that is invariant under a lattice action. Suppose that $\mu \in \mathcal{G}_T$ is a translation-invariant Gibbs measure corresponding to the potential \mathcal{V} . The aim is to show that μ is extremal. As explained for the Gaussian case, it is enough to prove that $\mathbb{P}_{\mu,\mu}$ -almost surely, η does not percolate. This is done in three steps, which we will only sketch here.

First, note that η splits naturally into two clusters: η^+ and η^- , corresponding to the signs of $h - \tilde{h}$ on them. This helps to rule out that η has more than 2 infinite connected components because both η^+ and η^- satisfy a version of the finite energy property.

The second and most difficult step is ruling out that η has two infinite clusters; one of each sign. This can be achieved using entropy arguments, which we will not explain here (but uses the fact that η has at most two infinite clusters).

The two points above thus imply

$$\mathbb{P}_{\mu,\mu}(\eta \text{ has 2 or more infinite clusters}) = 0.$$

In the third and final step, it still needs to be ruled out that η percolates at all. But since we know that there is at most one infinite cluster, this can be achieved in the same way as was done in [45].

The idea is that we can couple μ with itself in such a way that (h, \tilde{h}) is ordered on the finite clusters:

$$h_x \leq \tilde{h}_x, \quad \text{if } x \text{ in a finite cluster of } \eta.$$

This can be achieved by starting with an independent coupling (g, \tilde{g}) sampled from $\mu \times \mu$, and switching g, \tilde{g} on the finite clusters precisely when $g > \tilde{g}$.

Now suppose that $\mathbb{P}_{\mu,\mu}(\eta \text{ percolates}) > 0$. In particular we may assume that with positive probability, $h < \tilde{h}$ on the infinite clusters of η . In particular, on this event $g \leq \tilde{g}$ everywhere. But since the marginals of g and \tilde{g} are the same, this readily implies a contradiction: using translation invariance of the coupled measure (g, \tilde{g}) , the infinite cluster of η has strict positive frequency, so $g \leq \tilde{g} - 1$ has strict positive frequency. The latter is impossible. We conclude that $\mathbb{P}_{\mu,\mu}$ -almost surely, η does not percolate.

4.3. Delocalization of height functions

In this section, we will sketch a slight variation of Lammers' beautiful proof of delocalization [114]. Let us begin by stating the main result.

Theorem 4.4 (Delocalization, [114]). *Let Γ be a bi-periodic planar lattice with degrees bounded by three. Suppose $\mathcal{V} : \mathbb{Z} \rightarrow \mathbb{R}$ be a convex and symmetric potential. If*

$$\mathcal{V}(0) \leq \mathcal{V}(1) + \log(2),$$

there does not exist a translation invariant Gibbs measure for the height function.

About the proof. Our variation of Lammer’s proof uses essentially the same ideas proposed by Lammers, but differs in one way. We do not use an extension of the height function to middle of the edges, but rather, we will introduce a percolation process which corresponds to kind of sub level-set percolation, extended to the whole cable graph. We think this describes the intuition behind delocalization quite well, although it could well be a semantics discussion.

4.3.1 Extremal measures

To proof Theorem 4.4, we will argue by contradiction and suppose that an ergodic translation invariant Gibbs measure μ does exist. By Theorem 4.1 we can assume that μ is extremal and translation invariant.

This is important because it implies that the FKG property holds for the height function: for any A, B increasing events, we know that

$$\mu(\mathbb{1}_{A \cap B} \mid \mathcal{F}_{\Lambda_n^c}) \geq \mu(\mathbb{1}_A \mid \mathcal{F}_{\Lambda_n^c})\mu(\mathbb{1}_B \mid \mathcal{F}_{\Lambda_n^c})$$

because \mathcal{V} is convex. Taking $n \rightarrow \infty$ and using backwards martingale convergence, this implies $\mu(A \cap B) \geq \mu(A)\mu(B)$.

The FKG property can be used to show that the expectation of h_0 is an integer, so that we can assume without loss of generality that it equals 0. This implies, essentially, that all the extremal Gibbs measures are coming from the infinite volume limits (if we would know they existed) of wired boundary conditions, with some prescribed boundary value in \mathbb{Z} .

Lemma 4.5 (Integer-valued expectation). *Let μ be an extremal, translation invariant Gibbs measure for the potential \mathcal{V} . Then $\mu(h_0) \in \mathbb{Z}$.*

Proof. We begin by observing that the site percolations $\xi^- := \{x : h_x \leq 0\}$ and $\xi^+ := \{x : h_x \geq 1\}$ are FKG, because the height function satisfies the FKG property and μ is extremal. The fact that the height function is FKG is a consequence of convexity.

Since the underlying lattice is bi-periodic and planar, it is impossible for both ξ^- and ξ^+ to percolate by [59, 157]. If ξ^- does not percolate, this implies that there must be a circuit of ξ^+ which is blocking. Therefore, and by the FKG property, this implies that $\mu(h_0) \geq 1$. On the other hand, if ξ^+ does not percolate, the same argument implies that $\mu(h_0) \leq 0$, so that $\mu(h_0) \notin (0, 1)$. However, we can always shift h_0 by an integer, so we must have that $\mu(h_0) \in \mathbb{Z}$. \square

Supposing that $\mu(h_0) = 0$, the proof of the last lemma tells us essentially that $(\mathbb{1}\{h_x < 1\})_{x \in V}$ must percolate: if it would not, then there would be blocking circuits

of $h \geq 1$ and hence $\mu(h_0) \geq 1$. Thus, if we could show that also there must always be a circuit where the height function is larger than 1, then this implies that the expected value is actually be larger than one, which is a contradiction.

The idea to break this barrier is similar to the switching percolation described above: to use the residual energy. If there is an edge where h equals zero on both endpoints, then there may be enough energy in the system, to say that the height function will “touch” a prescribed value somewhere on the edge (which doesn’t make sense). Let us make a small detour to a setting where it does make sense: the integer-valued GFF.

4.3.2 The integer-valued Gaussian free field

Take the potential $\mathcal{V}(k) = \beta|k|^2$. This model is often called the integer-valued GFF. One way to view it is as a regular GFF on the graph, conditioned to have integer values on the vertices. As such, it allows for an extension to the cable graph, which is again given by Brownian bridges connecting the endpoints. This is similar to the discussion in Section 4.2.2 concerning the real-valued Gaussian free field.

In general, this defines a family of percolation processes ω^c indexed by $c \in \mathbb{R}$ corresponding to the “sub level-set percolation”: set $\omega_e^c = 0$ precisely when the cable graph extension of h passes through c on the edge e . If $c \in \mathbb{Z}$, then as above we can exchange the value of h with $2c - h$ on all finite clusters of ω^c , without changing the law of h . This corresponds to reflecting h around the “level line” where $h = c$.

However, if $c \in \mathbb{R}$, such as reflection is not possible because $2c - h$ is not integer valued. Thus, we must restrict to c in the half-integers and as such, the minimal $c > 0$ we can look for is $c = \frac{1}{2}$.

Assume that μ is an extremal, translation invariant Gibbs measure and $\mu(h_0) = 0$. If we can show that there is almost surely no path from 0 to infinity (on the cable system) where $h < \frac{1}{2}$, then we know that $\mu(h_0)$ is at least $\frac{1}{2}$, contradiction the assumption that $\mu(h_0) = 0$.

On an edge where the endpoints are 0, the Brownian bridge has probability $e^{-\beta}$ to touch $\frac{1}{2}$. Thus, the conditions of Theorem 4.4 imply $\beta \leq \log(2)$, so that $e^{-\beta} \geq \frac{1}{2}$.

Lammers’ idea. We will soon finally use the assumption that the underlying Γ is planar and has degrees bounded by 3.

Sample an i.i.d. family $(\sigma_e)_{e \in E}$ of uniform random variables on $\{-1, 1\}$. Define, using only the height function h sampled from μ on the vertices, two bond percolation processes ξ^+ and ξ^- as follows.

For each edge $e = (x, y)$, set ξ_e^+ equal to 1 whenever either h_x or h_y is greater than or equal to 1. If both endpoints are 0, set ξ_e^+ to one if $\sigma_e = 1$. Similarly, defined is ξ_e^- . Note that ξ^+ (and ξ^-) are translation-invariant, FKG percolation processes because h

satisfies the FKG property. Also note that since μ is ergodic², we get that the event “ ξ^+ percolates” is trivial. The same holds for ξ^- .

Lemma 4.6. *Almost surely, either $1 - \xi^-$ or $1 - \xi^+$ does not percolate.*

Proof. Because Γ has degrees bounded by three and ξ^+ satisfies the FKG property, we know that almost surely it is impossible that both ξ^+ and $1 - \xi^+$ percolate³. Suppose thus without loss of generality that $1 - \xi^+$ percolates a.s. (otherwise we are done). By construction, $1 - \xi^+$ viewed as a random subgraph of Γ is contained in ξ^- , so this implies ξ^- percolates almost surely. By the same argument as for ξ^+ , this implies $1 - \xi^-$ does not percolate, so we are done. \square

Assume, without loss of generality, that $1 - \xi^+$ does not percolate almost surely.

So how to go back to the level-set percolation? Recall that for $\beta \leq \log(2)$ the probability that the cable graph extension of h on an edge hits $\frac{1}{2}$, with two zero endpoints, is larger than $1/2$. Since the Brownian bridges are independent for each edge, conditional on $(h_x)_{x \in V}$, we can couple the sub level-set percolation $h < \frac{1}{2}$ and $1 - \xi^+$ in such a way that h touching $\frac{1}{2}$ on the edge, implies that ξ^+ is open for edges with zero end-points. For all other edges, the coupling is guaranteed by construction. We make such a coupling precise below. Since $1 - \xi^+$ does not percolate a.s., we deduce that $h < \frac{1}{2}$ does not percolate (on the cable graph), and hence $\mu(h_0) \geq \frac{1}{2}$, a contradiction.

4.3.3 Extension to general height functions

We will now finalize the proof of Theorem 4.4. Assume as Γ is some infinite, trivalent planar graph which is invariant under some lattice action.

The generalization to other models with a convex potential is similar as the one appearing in the proof of Theorem 4.1. Let \mathcal{V} be a convex, symmetric potential and take μ and ergodic Gibbs measure corresponding to the potential. Fix also $X = (X_e)_{e \in E}$ a family of independent uniform random variables on $[0, 1]$, which will be useful to obtain explicit couplings. Write \mathbb{P}_μ for the product of μ and the law of X .

For $c \in \frac{1}{2}\mathbb{Z}$, we want to study the impact of a reflection around c on an edge. Sample h from μ and for an edge $e = (x, y)$, consider the two energies:

$$e^{-\mathcal{V}(h_x - h_y)} \quad \text{and} \quad e^{-\mathcal{V}(2c - h_x - h_y)}. \quad (4)$$

²Purely technically, we would also need to deal with the extra randomness needed to sample σ , but this is just i.i.d. percolation.

³This is because a planar graph of degrees bounded above by three can be mapped to a planar graph which has the edges of the original graph as its vertices, and connectivity is preserved. But we know that for any FKG site percolation ξ , it is impossible that both ξ and $1 - \xi$ percolate [59, 157].

which correspond to the actual energy and the energy after switching h_x with $2c - h_x$. Write $m_{x,y}$ and $M_{x,y}$ for the minimum respectively maximum values of (4). Notice that $m_{x,y}$ and $M_{x,y}$ depend on h .

Define the percolation process ω^c as follows. For any edge $e = (x, y)$, set $\omega_e^c = 0$ if the actual weight equals the minimal weight: $e^{-\mathcal{V}(h_x - h_y)} = m_{x,y}$. In this case, we call the edge *excited*. Next, also set $\omega_e^c = 0$ if

$$X_e \leq \frac{m_e}{M_e}.$$

Otherwise, set $\omega_e^c = 1$

As in Section 4.2, the percolation process ω^c can be thought of as a FK representation of the sign clusters of $h + c$. Indeed, the following two important facts about the percolation and the height function establish this analogue.

Lemma 4.7. *For any $c \in \frac{1}{2}\mathbb{Z}$ and any two vertices $x, y \in V$, we have*

- (i) \mathbb{P}_μ is invariant under switching h and $2c - h$ on any of the finite clusters of ω^c .
- (ii) If x is connected to y in ω^c , then $(h_x + c)(h_y + c) > 0$.

Proof. The first fact follows immediately from rewriting the (local) Gibbs factor in terms of excited and not excited edges. The second fact is a consequence of convexity of \mathcal{V} . \square

At this point, we have achieved the generalization of “sub level-set percolation” from the integer-valued Gaussian free field to general convex, symmetric potentials. The remainder of the proof is almost the same, although this time we will make the coupling precise. Thus, fix $c = \frac{1}{2}$ (although it does not really play a special role, we could have picked *any* half-integer $c > 0$, making the result slightly weaker).

The percolation process ω^c splits naturally into two parts: $\omega^{c,+}$ and $\omega^{c,-}$ depending on the sign of $h + c$ on the cluster of ω^c . The clusters $\omega^{c,-}$ and $\omega^{c,+}$ correspond to the sub respectively super level-set percolation of the height function (extended to the edges). As for the integer-valued GFF, it is thus enough to rule out that $\omega^{\frac{1}{2},-}$ percolates:

Lemma 4.8. *If $\omega_{\frac{1}{2}}^-$ almost surely does not percolate, then $\mu(h_0) \geq \frac{1}{2}$.*

Proof. Suppose that $\omega^{\frac{1}{2},-}$ does not percolate almost surely. When the cluster of 0 is finite in $\omega_{\frac{1}{2}}$, we can flip the sign of $h + \frac{1}{2}$ on that cluster and hence

$$\mathbb{E}_\mu(h_0 \mathbf{1}\{\text{cluster of 0 finite in } \omega^{\frac{1}{2}}\}) = \mathbb{E}_\mu((1 - h_0) \mathbf{1}\{\text{cluster of 0 finite in } \omega^{\frac{1}{2}}\}),$$

which implies that

$$\mathbb{E}_\mu(h_0 \mathbf{1}\{0 \text{ in finite cluster of } \omega^{\frac{1}{2}}\}) = \frac{1}{2} \mathbb{P}_\mu(\text{cluster of 0 finite in } \omega^{\frac{1}{2}}).$$

Moreover, if the cluster of 0 is infinite in $\omega^{\frac{1}{2}}$, then $h_0 > \frac{1}{2}$ by assumption. This shows

$$\begin{aligned} \mu(h_0) &= \mathbb{E}_\mu(h_0 \mathbf{1}\{\text{cluster of 0 finite in } \omega^{\frac{1}{2}}\}) + \mathbb{E}_\mu(h_0 \mathbf{1}\{\text{cluster of 0 infinite in } \omega^{\frac{1}{2}}\}) \\ &\geq \frac{1}{2}, \end{aligned}$$

as desired. \square

The auxiliary percolation. Thus, we are left to rule out that $\omega^{\frac{1}{2},-}$ percolates. To show this, we will again use a variation of Lammers' argument and introduce auxiliary bond percolation processes ξ^+ and ξ^- as follows. Sample h from μ and let $(\sigma_e)_{e \in E}$ be a collection of independent random variables, each uniform on $\{-1, 1\}$.

Set $\xi_e^+ = 0$ for the edge $e = (x, y)$ if either h_x or h_y is greater than $\frac{1}{2}$, or if $h_x = h_y = 0$ and $\sigma_e > 0$. ξ^- is defined similarly, with the signs flipped. Since μ is extremal, the product measure of μ and the spins is ergodic, and hence $\mu(\xi^\pm \text{ percolates})$ is zero or one. Also note that ξ^+ and ξ^- satisfy the FKG property, because h does.

Lemma 4.9. *Let Γ be a trivalent planar graph, invariant under some lattice action. Either $\mu(1 - \xi^+ \text{ percolates}) = 0$ or $\mu(1 - \xi^- \text{ percolates}) = 0$.*

Proof. The proof is exactly the same as for integer-valued GFF. \square

Implications on sub level-set percolation. Of course, having that $1 - \xi^+$ or $1 - \xi^-$ does not percolate, is not quite what we want. But we will couple the one which does not percolate to the corresponding ‘‘sub level-set percolation’’ $h < \frac{1}{2}$ or $h > -\frac{1}{2}$.

Lemma 4.10. *If $\mathcal{V}(1) \leq \mathcal{V}(0) + \log(2)$, either $\omega^{\frac{1}{2},-}$ does not percolate almost surely, or $\omega^{-\frac{1}{2},+}$ does not percolate almost surely.*

Proof. We will show that

$$\mu(1 - \xi^+ \text{ percolates}) = 0$$

implies that $\omega^{\frac{1}{2},-}$ does not percolate almost surely. The result then follows from symmetry. Assume thus that $1 - \xi^+$ does not percolate almost surely.

First, we calculate the probability that an edge $e = (x, y)$ is closed in $\omega^{\frac{1}{2},-}$ given that the two endpoints satisfy $h_x = h_y = 0$. By definition, this is given by

$$\mathbb{P}_\mu(\omega_e^{\frac{1}{2}} \text{ is closed} \mid h_x = h_y = 0) = e^{-\mathcal{V}(1) + \mathcal{V}(0)} \geq \frac{1}{2},$$

where the inequality follows by assumption on \mathcal{V} . In particular, in this case $\omega_e^{\frac{1}{2}}$ is open implies $X_e > \frac{1}{2}$.

Second, we define the coupling. Set $\sigma_e = 1$ precisely when

$$X_e < \frac{1}{2}$$

and $\sigma_e = -1$ otherwise. Of course, since $(X_e)_{e \in E}$ is a family of independent random variables, uniform on $[0, 1]$, σ_e is a family of independent, uniform spins on $-1, 1$. This defines a coupling of ξ^+, ξ^- and $\omega^{\frac{1}{2}, -}$, and we claim that $\omega^{\frac{1}{2}, -}$ is less than $1 - \xi^+$.

To prove the claim, we need to show that $\omega^{\frac{1}{2}, -} = 1$ implies $\xi^+ = 0$. Let $e = (x, y)$ be any edge. For $\omega_e^{\frac{1}{2}, -}$ to be open, we need that both endpoints are less than $\frac{1}{2}$: $h_x, h_y < \frac{1}{2}$. On the one hand, if either of then endpoints does not equal 0, then $\xi_e^+ = 0$ and we are done. On the other hand, if $h_x = h_y = 0$, $\omega_e^{\frac{1}{2}, -}$ is open only only when $X_e > \frac{1}{2}$, so $\xi_e^+ = 0$ by the coupling ($\sigma_e = -1$).

Finally, recall that $1 - \xi^+$ does not percolate almost surely, so also $\omega^{\frac{1}{2}, -}$ does not percolate, finalizing this proof. \square

The proof of Theorem 4.4 is essentially finished:

Proof of Theorem 4.4. Consider the setting of Theorem 4.4. Assume translation invariant Gibbs measures exist. Let μ be an ergodic Gibbs measure. Take 0 some distinguished vertex. By Lemma 4.5, we can assume that $\mu(h_0) = 0$.

On the other hand, by Lemma 4.10 we know that either $\omega^{\frac{1}{2}, -}$ or $\omega^{\frac{1}{2}, +}$ does not percolate. This implies that $\mu(h_0) \notin (-1/2, 1/2)$, a contradiction. \square

The BKT–transition in the XY model

5.1. Introduction and main result

Let $G = (V, E)$ be a finite graph. Given a collection of nonnegative *coupling constants* $J = (J_e)_{e \in E}$, and an *inverse temperature* $\beta > 0$, the *XY model* (with free boundary conditions) is a random spin configuration $\sigma \in \mathbb{S}^V$, where $\mathbb{S} = \{z \in \mathbb{C} : |z| = 1\}$ is the complex unit circle, sampled according to the Gibbs distribution

$$d\mu_{G,\beta}(\sigma) \propto \exp\left(\frac{1}{2}\beta \sum_{vv' \in E} J_{vv'}(\sigma_v \bar{\sigma}_{v'} + \bar{\sigma}_v \sigma_{v'})\right) \prod_{v \in V} d\sigma_v, \quad (1)$$

where vv' denotes the edge $\{v, v'\}$, and $d\sigma_v$ is the uniform probability measure on \mathbb{S} . For simplicity of notation, unless stated otherwise, we will assume that $J_e = 1$ for all e . However, our results extend naturally to nonhomogeneous coupling constants. We will write $\langle \cdot \rangle_{G,\beta}$ for the expectation with respect to $\mu_{G,\beta}$. The observable of main interest for us will be the *two-point function* $\langle \sigma_a \bar{\sigma}_b \rangle_{G,\beta}$, $a, b \in V$, and its *infinite volume limit* (which is well defined by the Ginibre inequalities [80])

$$\langle \sigma_a \bar{\sigma}_b \rangle_{\Gamma,\beta} = \lim_{G \nearrow \Gamma} \langle \sigma_a \bar{\sigma}_b \rangle_{G,\beta},$$

where Γ is an infinite planar lattice.

Note that if $\sigma_v = e^{i\theta_v}$, $\theta_v \in (-\pi, \pi]$, then $\sigma_v \bar{\sigma}_{v'} + \bar{\sigma}_v \sigma_{v'} = 2 \cos(\theta_v - \theta_{v'})$. This means that the model is *ferromagnetic*, i.e., pairs of neighboring spins that are (almost) aligned have smaller energy and hence are statistically favoured. A natural question is whether varying β leads to a ferromagnetic *order–disorder* phase transition in the model. The classical theorem of Mermin and Wagner [134] excludes this possibility when the underlying lattice Γ is two-dimensional. Moreover, McBryan and Spencer showed that at any finite temperature $\langle \sigma_a \bar{\sigma}_b \rangle_{\mathbb{Z}^2,\beta}$ decays to zero at least as fast as a power of the distance between a and b . On the other hand, it is known by the work of Fröhlich, Simon and Spencer [70] that in higher dimensions the model exhibits long-range order at low temperatures and the two-point function does not decay to zero.

Even though there is no spontaneous symmetry breaking, Berezinskii [36, 37], and Kosterlitz and Thouless [108] predicted that a different type of phase transition takes place in two dimensions. It should be understood in terms of interacting topological excitations of the model, the so called *vortices* and *antivortices*. They are those faces of the graph where the XY configuration makes a full clockwise or anticlockwise turn respectively when one traverses the edges of the face in a clockwise manner. Vortices and antivortices interact through a Coulomb-like interaction, and are energetically favoured to form short-distance pairs of vortex-antivortex. The *Berezinskii–Kosterlitz–Thouless* (BKT) phase transition happens when, while decreasing the temperature, the freely spaced vortices and antivortices (high-temperature plasma) bind together into such vortex-antivortex pairs. This regime should exhibit power-law decay of the two-point functions (in contrast to exponential decay at high temperatures). A rigorous lower bound of this type for low temperatures, and therefore a proof of the BKT phase transition was first obtained in the celebrated work of Fröhlich and Spencer [71] who also derived analogous results for the Villain spin model. Their proof uses a multi-scale analysis of the Coulomb gas, and the main purpose of the present chapter is to present an alternative and less technically involved argument for the existence of phase transition in two dimensions.

To be more precise, we introduce a new loop representation for the two-point function in the XY model that can be used to transfer probabilistic information from the dual integer-valued height function model to the XY model. Along the way we also show that the height function possesses the crucial absolute-value-FKG property. This, together with a recent elementary delocalization result for general height functions obtained by Lammers [114], is used to prove existence of the BKT phase transition.

Theorem 5.1 (Berezinskii–Kosterlitz–Thouless phase transition). *There exists $\beta_c \in (0, \infty)$ such that*

(i) *for all $\beta < \beta_c$, there exists $c = c(\beta) > 0$ such that for all $v, v' \in \mathbb{Z}^2$,*

$$\langle \sigma_v \bar{\sigma}_{v'} \rangle_{\mathbb{Z}^2, \beta} \leq e^{-c|v-v'|},$$

(ii) *for all $\beta \geq \beta_c$ and all distinct $v, v' \in \mathbb{Z}^2$,*

$$\langle \sigma_v \bar{\sigma}_{v'} \rangle_{\mathbb{Z}^2, \beta} \geq \frac{1}{8|v-v'|}.$$

We note that unlike in the original proof of Fröhlich and Spencer, we do not show that the rate of decay approaches zero when so does the temperature. However, we establish a type of sharpness which says that there is no other behavior than exponential and power-law decay. The short proof of sharpness is independent of the rest of the argument. In the

first step we classically use the Lieb–Rivasseau inequality [125, 151] to establish a sharp transition between exponential decay and nonsummability of correlations (similarly to the proof for the Ising model [60]). To conclude a uniform power-law lower bound as in (ii) whenever the correlations are not summable we use the Messenger–Miracle–Sole inequality [135] on monotonicity of correlations with respect to the position of the vertex on the lattice.

We also note that our proof works (with minor modifications and a different, implicit multiplicative constant in (ii)) for other infinite planar graphs that in addition to being translation invariant possess reflection and rotation symmetries, and whose dual graph has bounded degree.

At the same time when the original article presented here appeared, an analogous result for the Villain model (without sharpness and explicit polynomial decay in the BKT phase) was given by Aizenman et al. [8]. It was later extended to also cover the XY model (including sharpness). For a more detailed overview of the XY model, we refer the reader to [69, 142], and for expositions of the argument of Fröhlich and Spencer, we refer to [76, 104].

This chapter is organized as follows.

- In Section 5.2 we introduce the dual of the planar XY model in form of an integer-valued height function defined on the faces of the graph. We also establish positive association of its absolute value (the absolute-value-FKG property), and recall the delocalization result of Lammers [114].
- In Section 5.3 we define a random collection of loops on the graph that carries probabilistic information about both the XY spins and the dual height function. Although this is a well known object that goes back to the works of Symanzik [161], and Brydges, Fröhlich and Spencer [42], the formula that relates the two-point function to the probability of two points being connected by a loop (Lemma 5.10) is new and crucial to our argument.
- In Section 5.4 we give an elementary argument which states that if the height function delocalizes at some temperature, then the spin two-point function *does not* decay exponentially.
- In Section 5.5 we use the above ingredients to show that on any translation invariant graph, there exists a finite temperature at which the two-point function does not decay exponentially. This is not immediate as the result of Lammers [114] applies only to trivalent graphs. However, a simple graph-modification argument together with the Ginibre inequality allows to change the setup from a general graph to a triangulation (a graph whose dual is trivalent).

- In Section 5.6 we finish the proof of the main theorem. We use the Lieb–Rivasseau inequality [125, 151] and the Messenger–Miracle–Sole inequality [135] to show that the absence of exponential decay implies a power-law lower bound on the two-point function.

5.2. The dual height function

To define the dual model we assume that G is planar and finite, and we introduce the notion of currents. To this end, let $\vec{E} = \{(v, v') : \{v, v'\} \in E\}$ be the set of directed edges of G , and let $\mathbb{N} = \{0, 1, \dots\}$. A function $\mathbf{n} : \vec{E} \rightarrow \mathbb{N}$ is called a *current* on G . For a current \mathbf{n} , we define $\delta\mathbf{n} : V \rightarrow \mathbb{Z}$ by

$$\delta\mathbf{n}_v = \sum_{v' \sim v} \mathbf{n}_{(v, v')} - \mathbf{n}_{(v', v)}.$$

Hence if $\delta\mathbf{n}_v$ is positive, then the amount of outgoing current is larger than the incoming current, and we think of v as a *source*. Likewise if $\delta\mathbf{n}_v$ is negative, there is more incoming current and v is a *sink*. A current is *sourceless* if $\delta\mathbf{n}_v = 0$ for all $v \in V$.

We define Ω_0 to be the set of all (sourceless) currents. Sourceless currents naturally define a height function h on the set of faces of G , denoted by U , where the height of the outer face is set to zero, and the increment of the height between two faces u and u' is equal to

$$h(u) - h(u') = \mathbf{n}_{(v, v')} - \mathbf{n}_{(v', v)},$$

where the primal directed edge (v, v') crosses the dual directed edge (u, u') from right to left. That this yields a well defined function on the faces of G follows from the fact that $\delta\mathbf{n} = 0$. We define the *XY weight* of a current by

$$w_\beta(\mathbf{n}) = \prod_{(v, v') \in \vec{E}} \frac{1}{\mathbf{n}_{(v, v')}!} \left(\frac{\beta J_{vv'}}{2} \right)^{\mathbf{n}_{(v, v')}}, \quad (2)$$

These weights appear naturally in the expansion of the partition function of the XY model into a sum over sourceless currents after one expands the exponentials in (1) into a power series in the variables $\frac{1}{2}\beta J_{vv'}\sigma_v\bar{\sigma}_{v'}$ for each directed edge $(v, v') \in \vec{E}$, and then integrates out the σ variables. They will also appear in the analogous classical expansion for spin correlations (11).

We note that using currents to define a model on the dual graph is an instance of planar duality of abelian spin systems [52], and the fact that the function is a consequence of \mathbb{Z} being the dual group of the unit circle.

Clearly, the weight (2) defines a probability measure $\mathbb{P}_{G, \beta}$ on currents and hence also on height functions. In terms of the height function it is a Gibbs measure given by

$$\mathbb{P}_{G, \beta}(h) \propto \exp\left(-\sum_{uu' \in E^\dagger} \mathcal{V}_e^\beta(h(u) - h(u'))\right), \quad (3)$$

where E^\dagger is the set of dual edges of G , and where the symmetric potentials $\mathcal{V}_e^\beta : \mathbb{Z} \rightarrow \mathbb{R}$ are given by

$$\mathcal{V}_e^\beta(k) = -\log \left(\sum_{i=0}^{\infty} \frac{1}{i!(i+|k|)!} \left(\frac{\beta J_e}{2} \right)^{2i+|k|} \right) = -\log I_k(\beta J_e) \quad (4)$$

with I_k being the modified Bessel function. We again note that we will usually set all $J_e = 1$ to simplify the notation.

A well known Turán-type inequality for modified Bessel functions [163] states that for any $k \geq 0$ and $\beta > 0$,

$$I_k^2(\beta) \geq I_{k-1}(\beta)I_{k+1}(\beta) \quad (5)$$

which means that \mathcal{V}_e^β is convex on the integers. This puts the model in the well-studied framework of height functions with a convex potential (see e.g. [157]).

5.2.1 Gibbs measures and delocalization

To state the delocalization result of Lammers we will need the notion of a Gibbs measure for height functions on infinite graphs (though we will not directly work with it in the remainder of the chapter). Let $\Gamma = (V, E)$ be an infinite planar graph and $\Gamma^\dagger = (U, E^\dagger)$ its planar dual. If ν is a measure on height functions $\varphi : \mathbb{Z}^U \rightarrow \mathbb{Z}$ and $\Lambda \subset U$ a finite subset, write ν_Λ for the measure restricted to Λ . Let $\mathcal{V} = (\mathcal{V}_e)_{e \in E^\dagger}$ be a family of convex symmetric potentials. We call ν a *Gibbs measure* for the potential \mathcal{V} if for every such Λ , it satisfies the Dobrushin–Lanford–Ruelle relation

$$\nu_\Lambda(\cdot) = \int_{\mathbb{Z}^U} \nu_\Lambda^\varphi(\cdot) d\nu(\varphi),$$

where ν_Λ^φ is the Gibbs measure on height functions $h \in \mathbb{Z}^U$ given as in (3) (but with \mathcal{V}^β replaced by \mathcal{V}) and conditioned on h being equal to φ on the boundary of Λ .

In what follows we will always assume that Γ is locally finite and invariant under the action of a \mathbb{Z}^2 -isomorphic lattice. We say that ν is translation invariant if it is invariant under the same action.

In a recent beautiful work [114] Lammers gave a condition on the potential that guarantees that there are no translation invariant Gibbs measures on graphs of degree three (trivalent graphs).

Theorem 5.2 (Lammers [114]). *Let $\Gamma^\dagger = (U, E^\dagger)$ be as above and moreover trivalent. If for every $e \in E^\dagger$,*

$$\mathcal{V}_e(\pm 1) \leq \mathcal{V}_e(0) + \log(2), \quad (6)$$

then there are no translation invariant Gibbs measures for \mathcal{V} .

This together with the dichotomy stated in Theorem 5.4 will be one of the key ingredients of the proof of the main theorem.

5.2.2 Absolute-value-FKG and dichotomy

In this section, we prove that the height function satisfies the absolute-value-FKG property, which is known to imply the dichotomy in Theorem 5.4 below [45, 117]. Here we will only work with the potential \mathcal{V}_β as defined in (4).

Proposition 5.3 (Absolute-value-FKG). *Let $G = (V, E)$ be a finite graph and U the set of its faces. Then for all $\beta > 0$, and all $\Psi, \Phi : \mathbb{N}^U \rightarrow \mathbb{R}_+$ increasing functions,*

$$\mathbb{E}_{G,\beta}[\Psi(|h|)\Phi(|h|)] \geq \mathbb{E}_{G,\beta}[\Psi(|h|)]\mathbb{E}_{G,\beta}[\Phi(|h|)].$$

We first explain briefly the dichotomy. Let $\Gamma = (V, E)$ be a translation invariant graph, and let 0 be a chosen face of Γ . Define B_n to be the subgraph of Γ induced by the vertices in V that lie on at least one face of Γ that is contained in the graph ball of radius n on Γ^\dagger . We introduce this slightly convoluted definition to guarantee the following three properties: 0 belongs to all B_n , also $B_n \nearrow \Gamma$ as $n \rightarrow \infty$, and finally, the weak dual graph of B_n (the dual graph with the vertex corresponding to the external face of B_n removed) is a subgraph of Γ^\dagger .

Theorem 5.4. *Consider the setup as above. Then for every $\beta > 0$, exactly one of the following two occurs:*

(i) (localization) *There exists a $C < \infty$ such that uniformly over all n ,*

$$\mathbb{E}_{B_n,\beta}[|h(0)|] \leq C.$$

(ii) (delocalization) *There are no translation invariant Gibbs measures for the potential (4).*

Proof. This is a consequence of the absolute-value-FKG property (Proposition 5.3) and standard arguments using monotonicity in boundary conditions. See [117, Theorem 2.7]. \square

We turn to the proof of Proposition 5.3. The first step (Lemma 5.5) consists in showing that, for β small enough, the potential satisfies an inequality known to imply the absolute-value-FKG property [117]. In the second step we use this to conclude the absolute-value-FKG property for general β .

Lemma 5.5. *The absolute-value-FKG property holds true for all $\beta \leq 1$.*

Proof. We rely on a result of Lammers and Ott [117, Theorem 2.8], stating that if

$$\mathcal{V}_e^\beta(k-1) - 2\mathcal{V}_e^\beta(k) + \mathcal{V}_e^\beta(k+1) = -\log\left(\frac{I_{k-1}(\beta)I_{k+1}(\beta)}{I_k(\beta)^2}\right)$$

is a nonincreasing function of k on $\{0, 1, \dots\}$, then $\mathbb{P}_{G,\beta}$ is absolute-value-FKG. We define $r_k = \frac{1}{\beta} \frac{I_k(\beta)}{I_{k-1}(\beta)}$, and need to show that $r_k^2 \leq r_{k-1}r_{k+1}$ for all $k \geq 0$. The well known recurrence relation

$$I_{k-1}(\beta) = \frac{2k}{\beta} I_k(\beta) + I_{k+1}(\beta) \quad \text{yields} \quad r_k = (2k + \beta^2 r_{k+1})^{-1}.$$

Hence it is enough to prove that

$$(2k + \epsilon_{k+1})(2(k+2) + \epsilon_{k+3}) \leq (2(k+1) + \epsilon_{k+2})^2,$$

where $\epsilon_k = \beta^2 r_k$. Using the Turán inequality (5), it follows that $0 \leq r_{k+1} \leq r_k$, and therefore it is sufficient to establish that

$$R_k := (2k + \epsilon_{k+1})(2k + 4 + \epsilon_{k+1}) - (2k + 2)^2 = 4(k+1)\epsilon_{k+1} + \epsilon_{k+1}^2 - 4 \leq 0.$$

At the same time, simply using the definition of r_{k+1} and comparing the Taylor expansions (4) of I_{k+1} and I_k term by term gives $\epsilon_{k+1} \leq \beta^2/(2k+2)$. Therefore, when $\beta \leq 1$, we have $R_k \leq \epsilon_{k+1}^2 - 2 \leq 0$ for all $k \geq 0$, which concludes the proof. \square

To treat general values of β , we will use a trick which consists in replacing each edge of G by $s = \lceil \beta \rceil$ consecutive edges, and reducing the parameter β by the factor s , together with the following convolution property of the modified Bessel functions.

Lemma 5.6. *For all $k, l \in \mathbb{Z}$ and all $\beta, \beta' \geq 0$,*

$$\sum_{m \in \mathbb{Z}} I_{k-m}(\beta) I_{m-l}(\beta') = I_{k-l}(\beta + \beta').$$

Proof. This is a classical identity which follows from the fact that $I_k(\beta)/e^\beta = \mathbb{P}(Z - Z' = k)$, where Z, Z' are independent Poisson random variables with mean $\beta/2$, and the fact that a sum of independent Poisson random variables is Poisson. \square

With this we can prove Proposition 5.3.

Proof of Proposition 5.3. Let $G_s = (V_s, E_s)$ be G with each edge replaced by s consecutive edges, and let h_s be the height function on G_s with law $\mu_{G_s, \beta/s}$. By Lemma 5.6 (and an induction argument) the restriction of h_s to V has the same law as h_1 . Moreover, $\beta/s \leq 1$ by definition of s , which by Lemma 5.5 implies that μ_s satisfies the absolute-value-FKG property. To finish the proof it is enough to notice that any increasing function on \mathbb{N}^V is also increasing on \mathbb{N}^{V_s} . \square

Remark 5.1. An interesting consequence of the idea above (that we will not use in this chapter) is the following. Consider the case when s from above is independent of β and diverges to infinity. In this limit, the height function becomes well defined at every

point of every dual edge. Here we think of the dual graph as the so called *cable graph*, i.e., every dual edge e is identified with a continuum interval of length $J_e\beta$. Then the distribution of the height on an edge, when conditioned on the values at the endpoints, is one of the difference of two Poisson processes with intensity $J_e\beta/2$ each, and conditioned on the value at the endpoints. One can check that the model exhibits a spatial Markov property on the full cable graph and not only on the vertices. This is in direct analogy with the cable graph representation of the discrete Gaussian free field, where the vertex-field can be extended to the edges via Brownian bridges (see e.g. [129] and the references therein).

5.3. Loop representation of currents and path reversal

The purpose of this section is mainly to develop a loop representation for the two-point function of the XY model. The important aspect of our approach is that the correlations are represented as probabilities for loop connectivities in random ensembles of closed loops. This is in contrast with most of the classical representations that write correlation functions as ratios of partition functions of loops, where in the numerator, in addition to loops, one also sums over open paths between the points of insertion in the correlator [42, 161]. We note that a similar idea to ours appears in the work of Benassi and Ueltschi [26], but due to technical differences in the framework (see Remark 5.4), the formula for the two-point function obtained in [26] is not as transparent as ours.

Let $G = (V, E)$ be a finite, not necessarily planar graph. We say that a multigraph \mathcal{M} on V is a *submultigraph* of G if after identifying the multiple copies of the same edge in \mathcal{M} it is a subgraph of G .

Definition 5.7 (Loop configurations outside S). Let \mathcal{M} be a submultigraph of G , and let $S \subseteq V$. A *loop configuration (on \mathcal{M}) outside S* is a collection of

- unrooted directed loops on \mathcal{M} avoiding S , and
- directed open paths on \mathcal{M} starting and ending in S (and not visiting S except at their start and end vertex),

such that every edge of \mathcal{M} is traversed exactly once by a loop or a path.

We write \mathcal{L}^S for the set of all loop configurations outside S , and define a weight for $\omega \in \mathcal{L}^S$ by

$$\lambda_\beta^S(\omega) = \prod_{v \in V \setminus S} \frac{1}{(\deg_{\mathcal{M}}(v)/2)!} \prod_{e \in E} \frac{1}{\mathcal{M}_e!} \left(\frac{\beta}{2}\right)^{\mathcal{M}_e}, \quad (7)$$

where \mathcal{M} is the underlying multigraph, and \mathcal{M}_e is the number of copies of e in \mathcal{M} . When $S = \emptyset$, a configuration is composed only of loops that can visit every vertex in V , and we simply call it a loop configuration.

An important feature of the weight (7) is that it depends on ω only through \mathcal{M} . Also note, that if $S' \subseteq S$, then there is a natural map $\rho : \mathcal{L}^{S'} \rightarrow \mathcal{L}^S$ that consists in forgetting (or cutting) the loop connections at the vertices in $S \setminus S'$. Under this map, each configuration in \mathcal{L}^S has $\prod_{v \in S \setminus S'} (\deg_{\mathcal{M}}(v)/2)!$ preimages, each of them having the same weight, and hence

$$\sum_{\tilde{\omega} \in \rho^{-1}[\omega]} \lambda_{\beta}^{S'}(\tilde{\omega}) = \lambda_{\beta}^S(\omega). \quad (8)$$

This consistency property will be useful later on.

For now, let $|\mathbf{n}| : E \rightarrow \mathbb{N}$ be the *amplitude* of a current \mathbf{n} , i.e.

$$|\mathbf{n}|_{vv'} := \mathbf{n}_{(v,v')} + \mathbf{n}_{(v',v)}.$$

Definition 5.8 (Multigraph of a current and consistent configurations). For a current \mathbf{n} , let $\mathcal{M}_{\mathbf{n}}$ be the submultigraph of G where each edge $e \in E$ is replaced by $|\mathbf{n}|_e$ (possibly zero) parallel copies of e . A loop configuration on $\mathcal{M}_{\mathbf{n}}$ is called *consistent with \mathbf{n}* if for every edge $(v, v') \in \vec{E}$, the number of times the loops traverse a copy of vv' in the direction of (v, v') is equal to $\mathbf{n}_{(v,v')}$. We define $\mathcal{L}_{\mathbf{n}}^S$ to be the set of all loop configurations on $\mathcal{M}_{\mathbf{n}}$ outside S that are consistent with \mathbf{n} .

For $\varphi : V \rightarrow \mathbb{Z}$, let $\Omega_{\varphi} = \{\mathbf{n} : \delta \mathbf{n} = \varphi\}$,

$$Z_{G,\beta}^{\varphi} = \sum_{\mathbf{n} \in \Omega_{\varphi}} w_{\beta}(\mathbf{n}),$$

and $\mathcal{S}(\varphi) = \{v \in V : \varphi_v \neq 0\}$. For a current \mathbf{n} , with a slight abuse of notation, we also write $\mathcal{S}(\mathbf{n}) = \mathcal{S}(\delta \mathbf{n})$. Note that $\mathcal{L}_{\mathbf{n}}^S$ can be nonempty only if $\mathcal{S}(\mathbf{n}) \subseteq S$. Indeed, each path and loop that enters a vertex in $V \setminus S$ must also leave it, and hence the total number of incoming and outgoing arrows at each such vertex must be the same. For $\varphi : V \rightarrow \mathbb{Z}$, we also define

$$\mathcal{L}_{\varphi}^S = \bigcup_{\mathbf{n} \in \Omega_{\varphi}} \mathcal{L}_{\mathbf{n}}^S.$$

Again, this is nonempty only if $\mathcal{S}(\varphi) \subseteq S$. We will write \mathcal{L}_0^S , where 0 denotes the zero function on V .

We now relate the weights of loops to those of currents. To this end, note that for each edge $vv' \in E$, there are exactly

$$\frac{|\mathbf{n}|_{vv'}!}{\mathbf{n}_{(v,v')}! \mathbf{n}_{(v',v)}!}$$

ways of assigning orientations to it so that the result is consistent with \mathbf{n} . Moreover, independently of the choices of orientations, there are exactly $(\deg_{\mathcal{M}_{\mathbf{n}}}(v)/2)!$ possible

pairings of the incoming and outgoing edges at each vertex $v \in V \setminus S$. Combining all this we arrive at a crucial loop representation for current weights: if $\mathcal{S}(\mathbf{n}) \subseteq S$, then

$$w_\beta(\mathbf{n}) = \sum_{\omega \in \mathcal{L}_\mathbf{n}^S} \lambda_\beta^S(\omega). \quad (9)$$

An important observation here is that the left-hand side is independent of S , and hence so is the right-hand side.

5.3.1 Coupling with the height function

We now apply this framework to the case of two sourceless currents and a coupling with the corresponding height function. From (9) we have

$$Z_{G,\beta}^0 = \sum_{\omega \in \mathcal{L}_0^\emptyset} \lambda_\beta^\emptyset(\omega) \quad (10)$$

where 0 denotes the zero function on V .

Remark 5.2. This loop representation of the partition function, though obtained via a different procedure, goes back to the work of Symanzik [161], and Brydges, Fröhlich and Spencer [42].

Moreover, in the case when G is planar we immediately get the following distributional identity. Define $\mathbf{P}_{G,\beta}$ to be the probability measure on $\mathcal{L}_0 := \mathcal{L}_0^\emptyset$ induced by the weights $\lambda_\beta := \lambda_\beta^\emptyset$. For each face $u \in U$ of G , and $\omega \in \mathcal{L}_0$, define $W_\omega(u)$ to be the total net winding of all the loops in ω around u .

Proposition 5.9. *The law of $(W(u))_{u \in U}$ under $\mathbf{P}_{G,\beta}$ is the same as the law of the height function $(h(u))_{u \in U}$ under $\mathbb{P}_{G,\beta}$.*

5.3.2 The two point-function and path reversal

We now turn to the loop representation of the two-point function. For reasons that will become apparent soon, we need to consider the two-point function of the squares, i.e., $\langle \sigma_a^2 \bar{\sigma}_b^2 \rangle$.

Since the resulting currents will have sources, we will need to consider nonempty S in the construction above. To this end, fix two vertices $a, b \in V$, and define $\varphi = 2(\delta_a - \delta_b)$, where $\delta_a(v) = \mathbf{1}\{a = v\}$. To lighten the notation, will write a, b instead of $\{a, b\}$ for the set S . As for the partition function, expanding the exponential in the Gibbs–Boltzmann weights (1) into a power series in $\frac{1}{2}\beta J_{vv'} \sigma_v \bar{\sigma}_{v'}$ for each directed $(v, v') \in \vec{E}$, and integrating out the σ variables, we get

$$\langle \sigma_a^2 \bar{\sigma}_b^2 \rangle_{G,\beta} = \frac{Z_{G,\beta}^\varphi}{Z_{G,\beta}^0} = \frac{\sum_{\omega \in \mathcal{L}_\varphi^{a,b}} \lambda_\beta^{a,b}(\omega)}{Z_{G,\beta}^0}, \quad (11)$$

where the first equality classically follows from the high-temperature expansion of correlation functions and the second one is a consequence of (9).

We will write $\mathcal{P}_{a,b}(\omega)$ for the set of paths in ω that start at a and end at b , and define

$$m_{a,b}(\omega) = |\mathcal{P}_{a,b}(\omega)|.$$

We now want to “erase the sources” at a and b from the currents underlying $\mathcal{L}_\varphi^{a,b}$, and hence rewrite the numerator as a sum over $\mathcal{L}_0^{a,b}$. We will then ultimately connect the open paths at a and b in all possible ways, and hence get a sum over \mathcal{L}_0^\emptyset (see Figure 5.1 for an example). To this end note that in each $\omega \in \mathcal{L}_\varphi^{a,b}$ there are exactly two more paths going from a to b , than those going from b to a , i.e., $m_{a,b}(\omega) = m_{b,a}(\omega) + 2$. The elementary operation that we will perform on the former paths is reversal. To this end, denote by $r(\gamma)$ the path γ with the orientation of all the visited edges reversed. Obviously this does not change the underlying multigraph, and hence also the weight of the loop configuration. The crucial observation now is that it maps $\omega \in \mathcal{L}_\varphi^{a,b}$ to a configuration $\omega' \in \mathcal{L}_0^{a,b}$, and hence erases the sources of the underlying currents. Indeed one can easily check that after reversing a path, the number of incoming minus the number of outgoing edges at every vertex $v \notin \{a, b\}$ in ω' is the same as in ω , whereas at a (resp. b) this number is decreased (resp. increased) by two. More precisely, our transformation maps bijectively a pair (ω, γ) where $\omega \in \mathcal{L}_\varphi^{a,b}$ and $\gamma \in \mathcal{P}_{a,b}(\omega)$ to the pair $(\omega', r(\gamma))$ where $\omega' \in \mathcal{L}_0^{a,b}$ and $r(\gamma) \in \mathcal{P}_{b,a}(\omega')$. Moreover, $m_{b,a}(\omega') = m_{b,a}(\omega) + 1$, which in particular means that $m(\omega') > 0$. Since path reversal does not change the weight of a loop configuration, we obtain

$$\begin{aligned} \sum_{\omega \in \mathcal{L}_\varphi^{a,b}} \lambda_\beta^{a,b}(\omega) &= \sum_{\omega \in \mathcal{L}_\varphi^{a,b}, \gamma \in \mathcal{P}_{a,b}(\omega)} \frac{1}{m_{b,a}(\omega) + 2} \lambda_\beta^{a,b}(\omega) \\ &= \sum_{\omega' \in \mathcal{L}_0^{a,b}, \gamma' \in \mathcal{P}_{b,a}(\omega')} \frac{1}{m_{b,a}(\omega') + 1} \lambda_\beta^{a,b}(\omega') \mathbb{1}\{m_{a,b}(\omega') > 0\} \\ &= \sum_{\omega' \in \mathcal{L}_0^{a,b}} \frac{m_{b,a}(\omega')}{m_{b,a}(\omega') + 1} \lambda_\beta^{a,b}(\omega') \mathbb{1}\{m_{b,a}(\omega') > 0\} \\ &= \sum_{\omega'' \in \mathcal{L}_0^\emptyset} \frac{m_{b,a}(\omega'')}{m_{b,a}(\omega'') + 1} \lambda_\beta^\emptyset(\omega'') \mathbb{1}\{m_{b,a}(\omega'') > 0\}, \end{aligned}$$

where in the second equality we used path reversal, the last equality follows from (8) with $S' = \emptyset$, and where, with a slight abuse of notation, for $\omega'' \in \mathcal{L}_0^\emptyset$, $m_{b,a}(\omega'')$ is the number of pieces of loops going from b to a and not visiting b nor a except for the start and end vertex. Recall that $\mathbf{P}_{G,\beta}$ is the probability measure on \mathcal{L}_0^\emptyset induced by the weights λ_β^\emptyset , and note that $m_{b,a}$ has the same distribution as $m_{a,b}$ under $\mathbf{P}_{G,\beta}$ (the law on loops is invariant under a global orientation reversal). We therefore obtain from (10)




Figure 5.1: Left to right: an Eulerian multigraph \mathcal{M} ; a loop configuration $\omega \in \mathcal{L}_{2(\delta_a - \delta_b)}^{a,b}$ on \mathcal{M} (a is the top left and b the bottom right vertex) together with a path from a to b marked red; a loop configuration $\omega' \in \mathcal{L}_0^{a,b}$ with the path reversed; and one of the final loop configurations $\omega'' \in \mathcal{L}_0^\emptyset$ corresponding to ω' , i.e., such that $\rho(\omega'') = \omega'$. Here $m_{a,b}(\omega) = 3$, $m_{b,a}(\omega) = 1$, and $m_{a,b}(\omega') = m_{b,a}(\omega') = 2$

and (11) the following loop representation of the two-point function.

Lemma 5.10. *Let $a, b \in V$ be distinct. Then*

$$\langle \sigma_a^2 \bar{\sigma}_b^2 \rangle_{G,\beta} = \mathbf{E}_{G,\beta} \left[\frac{m_{a,b}}{m_{a,b} + 1} \right],$$

and in particular

$$\frac{1}{2} \mathbf{P}_{G,\beta}(m_{a,b} > 0) \leq \langle \sigma_a^2 \bar{\sigma}_b^2 \rangle_{G,\beta} \leq \mathbf{P}_{G,\beta}(m_{a,b} > 0).$$

Let us finish with a number of remarks.

Remark 5.3. We stress again that the crucial property of this loop representation is that the measure $\mathbf{P}_{G,\beta}$ is supported on collections of closed loops, and is independent of the choice of a and b . A similar idea was used by Lees and Taggi [123] to study spin $O(n)$ models with an external magnetic field. Moreover, by Proposition 5.9 and Lemma 5.10, the random loops under $\mathbf{P}_{G,\beta}$ carry probabilistic information about *both* the spin XY model (in terms of correlation functions) and its dual height function (as an exact coupling). An analogous role for the Ising and Ashkin–Teller model is played by the (double) random current measure that encodes both an integer-valued height function and the spin correlations [54, 126, 127]. The difference is that for the XY model, the correlations are determined by loop connectivities instead of percolation connectivities. This comparison offers an alternative explanation for the different types of phase transition in discrete and continuous spin systems.

Remark 5.4. The approach above is different from [26, 42, 123, 161] in that in the loop configurations, we never make connections at vertices with sources. This leads to different combinatorics than in [26], and in particular a more transparent formula for the two-point function.

Remark 5.5. We call a multigraph \mathcal{M} *Eulerian* if its degree is even at every vertex. Another way to sample the loop configuration that easily follows from the above definitions is the following procedure:

- First sample an Eulerian submultigraph \mathcal{M} of G with probability proportional to

$$\mathcal{E}(\mathcal{M}) \prod_{e \in E} \frac{1}{\mathcal{M}_e!} \left(\frac{\beta}{2}\right)^{\mathcal{M}_e},$$

where $\mathcal{E}(\mathcal{M})$ is the number of *Eulerian orientations* of \mathcal{M} , i.e., assignments of orientations to every edge of \mathcal{M} with an equal number of incoming and outgoing edges at every vertex.

- Then choose uniformly at random an Eulerian orientation of \mathcal{M} .
- Finally, at each vertex, independently of other vertices, connect the incoming edges with the outgoing edges uniformly at random.

Remark 5.6. Using the same argument as above one obtains the following formula for higher power two-point functions. For $k \geq 1$, we have

$$\langle \sigma_a^{2k} \bar{\sigma}_b^{2k} \rangle_{G,\beta} = \mathbf{E}_{G,\beta} \left[\frac{(m_{a,b})_k}{(m_{a,b} + k)_k} \right],$$

where $(m)_k = m(m-1) \cdots (m-k+1)$ is the falling factorial. One can also consider multi-point functions and get more complicated loop representation formulas.

Remark 5.7. The isomorphism theorem of Le Jan [119] says that the discrete complex Gaussian free field can be coupled with a Poissonian collection of random walk loops, the so called *random walk loop soup*, in such a way that one half of the square of the absolute value of the field is equal to the total occupation time of the random walk loops. On the other hand, it is immediate that conditioned on the absolute value of the field, its complex phase is distributed like the XY model with coupling constants depending on this absolute value. With some work, e.g. using [120], one can show that under this conditioning the random walk loops have the same distribution as the loops described above.

5.4. Delocalization implies no exponential decay

In this section we prove that if the height function delocalizes, then the spin correlations are not summable along certain sets of vertices. In the next section, we will show how to apply this together with the delocalization results of Lammers [114] to deduce a BKT-type phase transition in a wide range of periodic planar graphs.

Suppose $\Gamma = (V, E)$ is a translation invariant planar graph, and write

$$\langle \sigma_a \bar{\sigma}_b \rangle_{\Gamma, \beta} = \lim_{G \nearrow \Gamma} \langle \sigma_a \bar{\sigma}_b \rangle_{G, \beta} \quad (12)$$

for the infinite volume two-point function, where the limit is taken along any increasing sequence of subgraphs G exhausting Γ . That this is well defined is guaranteed by the fact that the sequence is nondecreasing, i.e., $\langle \sigma_a \bar{\sigma}_b \rangle_{G, \beta} \leq \langle \sigma_a \bar{\sigma}_b \rangle_{G', \beta}$ if G is a subgraph of G' , which in turn is a classical consequence of the Ginibre inequality [80].

Definition 5.11. Let 0 be a distinguished face of Γ . A bi-infinite self-avoiding path in Γ that goes through at least one edge incident to 0 is called a *cut (at 0)*. Note that a cut L naturally splits into two connected infinite sets of vertices L_+ and L_- with the property that any cycle in Γ that surrounds 0 must intersect both L_+ and L_- .

The main quantity of interest for us will be the sum of correlations along cuts. To be more precise for $\varepsilon > 0$, let

$$\chi_{\Gamma, \beta}^\varepsilon(L) = \sum_{a \in L_+, b \in L_-} (\langle \sigma_a \bar{\sigma}_b \rangle_{\Gamma, \beta})^{2-\varepsilon}. \quad (13)$$

Proposition 5.12. *For every $\varepsilon > 0$, there exists $C = C(\varepsilon, \beta, \Gamma) < \infty$ such that for all finite subgraphs G of Γ containing 0 , we have*

$$\mathbb{E}_{G, \beta}[|h(0)|] \leq C \inf_L \chi_{\Gamma, \beta}^\varepsilon(L),$$

where the infimum is over all cuts at 0 .

Before presenting the proof, let us mention that a direct corollary of this proposition is the following. A natural example of a cut is any path that stays at a constant distance from a straight line going through 0 . In this case it is easy to see that $\chi_{\Gamma, \beta}^\varepsilon(L)$ is finite whenever there is exponential decay of spin correlations. We can now state the main conclusion of this section.

Corollary 5.13. *If the height function delocalizes in the sense of Theorem 5.4, then*

$$\chi_{\Gamma, \beta}^\varepsilon(L) = \infty$$

for all $\varepsilon > 0$ and all cuts L at 0 . In particular the two-point function does not decay exponentially fast with the distance between the vertices.

Proof. We know that situation (i) from Theorem 5.4 does not happen. This means that $\sup_n \mathbb{E}_{B_n, \beta}[|h(0)|] = \infty$, and the claim follows directly from Proposition 5.12. \square

Remark 5.8. One naturally expects that the localization-delocalization phase transition for the height function happens at the same temperature as the BKT transition for the XY model. The remaining part of this prediction is therefore to show that if the spin correlations do not decay exponentially, then the height function delocalizes. We do not do this in this chapter.

Recall that $m_{a,b}$ is the number of paths (pieces of loops) in a loop configuration that go from a to b . We will need the following lemma.

Lemma 5.14. *For all $\beta > 0$ and $p > 1$, there exists a $C_p < \infty$ such that for all finite graphs $G = (V, E)$ and all $a, b \in V$,*

$$\mathbf{E}_{G,\beta}[m_{a,b}] \leq C_p \deg_G(a) (\mathbf{P}_{G,\beta}(m_{a,b} > 0))^{\frac{1}{p}}.$$

Proof. Fix $\beta > 0$, $G = (V, E)$ and $a, b \in V$, and let $\omega \in \mathcal{L}_0$ be a loop configuration on G . Denote by ω_e , the number of visits of all loops in ω to an undirected edge $e \in E$. If there are $m \geq 1$ paths going from a to b in ω , then in particular $\sum_{c \sim a} \omega_{\{a,c\}} \geq m$. This implies that

$$\mathbf{E}_{G,\beta}[m_{a,b}] \leq \mathbf{E}_{G,\beta} \left[\sum_{c \sim a} \omega_{\{a,c\}} \mathbf{1}\{m_{a,b} > 0\} \right] \leq \deg_G(a) \max_{c \sim a} \mathbf{E}_{G,\beta}[\omega_{\{a,c\}} \mathbf{1}\{m_{a,b} > 0\}].$$

Applying Hölder's inequality gives

$$\mathbf{E}_{G,\beta}[\omega_{\{a,c\}} \mathbf{1}\{m_{a,b} > 0\}] \leq (\mathbf{E}_{G,\beta}[\omega_{\{a,c\}}^q])^{1/q} \mathbf{P}_{G,\beta}(m_{a,b} > 0)^{1/p},$$

where $1/p + 1/q = 1$. We now notice that by definition, ω_e under $\mathbf{P}_{G,\beta}$ has the same distribution as the amplitude $|\mathbf{n}|_e$ under $\mathbb{P}_{G,\beta}$. Therefore, to finish the proof it is enough to show that for all $q > 1$, there exists $C_q < \infty$ depending on β but independent of G such that

$$\mathbb{E}_{G,\beta}[|\mathbf{n}|_e^q] \leq C_q. \tag{14}$$

We postpone the proof of this bound to Lemma 5.16 and Lemma 5.17. \square

The last ingredient that we will need is the following inequality

Lemma 5.15. *For any $a, b \in V$, we have*

$$\langle \sigma_a^2 \bar{\sigma}_b^2 \rangle_{G,\beta} \leq 2 \langle \sigma_a \bar{\sigma}_b \rangle_{G,\beta}^2.$$

Proof. A version of the Ginibre inequality (see e.g. [25]) says that

$$\langle \Im(\sigma_a) \Im(\sigma_b) \Re(\sigma_a) \Re(\sigma_b) \rangle_{G,\beta} \leq \langle \Im(\sigma_a) \Im(\sigma_b) \rangle_{G,\beta} \langle \Re(\sigma_a) \Re(\sigma_b) \rangle_{G,\beta},$$

which after rearrangement gives the desired inequality. \square

We are now ready to prove the main theorem.

Proof of Proposition 5.12. Fix a finite subgraph G and a cut L . By Proposition 5.9 the height function $h(0)$ under $\mathbb{P}_{G,\beta}$ has the sam law as $W(0)$ – the total net winding around 0 of all loops in a loop configuration – drawn according to $\mathbf{P}_{G,\beta}$. Moreover, any piece of a loop that adds to the winding (in any orientation) must intersect both L_+ and L_- by definition of a cut. Therefore, taking $p = 2/(2 - \varepsilon)$, we have

$$\begin{aligned} \mathbb{E}_{G,\beta}[|h(0)|] &= \mathbf{E}_{G,\beta}[|W(0)|] \leq \sum_{a \in L_+, b \in L_-} \mathbf{E}_{G,\beta}[m_{a,b}] \\ &\leq \tilde{C} \sum_{a \in L_+, b \in L_-} (\mathbf{P}_{G,\beta}(m_{a,b} > 0))^{1/p} \\ &\leq 2\tilde{C} \sum_{a \in L_+, b \in L_-} (\langle \sigma_a^2 \bar{\sigma}_b^2 \rangle_{G,\beta})^{1-\varepsilon/2} \\ &\leq 4\tilde{C} \sum_{a \in L_+, b \in L_-} (\langle \sigma_a \bar{\sigma}_b \rangle_{G,\beta})^{2-\varepsilon} \leq C\chi_{\Gamma,\beta}^\varepsilon(L). \end{aligned}$$

where the third line follows from Lemma 5.14, the forth one from Lemma 5.10, the fifth one from Lemma 5.15, and the last one from (12). This completes the proof. \square

It therefore remains to show (14), which will directly follow from Lemma 5.16 and Lemma 5.17 below. To that end, define for $k \in \mathbb{N}$ and $\beta > 0$, a random variable Y_k by

$$\mathbb{P}_\beta(Y_k = i) \propto \frac{1}{i!(i+k)!} \left(\frac{\beta}{2}\right)^{2i+k},$$

so that the normalizing constant is $I_k(\beta)$. For $e = vv'$, let

$$|\nabla h|_e = |\mathbf{n}_{(v,v')} - \mathbf{n}_{(v',v)}|$$

be the absolute value of the gradient of the height function across the dual edge e^\dagger . Note that the random variables $(X_e = X_e(\mathbf{n}))_{e \in E}$ defined through

$$X_e = \frac{|\mathbf{n}|_e - |\nabla h|_e}{2}$$

have the same distribution as $Y_{|\nabla h|_e}$. Moreover, conditionally on $|\nabla h|_e$, they are an independent family. To show (14) it is enough to bound the moments of $|\nabla h|_e$ and X_e separately, which we will now do.

Lemma 5.16. *For all $\beta > 0$ and all $r \in \mathbb{N}$, there exists a $C_r < \infty$ such that for all finite planar graphs $G = (V, E)$ and all $e \in E$,*

$$\mathbb{E}_{G,\beta}[|\nabla h|_e^r] \leq C_r.$$

Proof. Fix a finite planar graph G , and let $e = vv' \in E$. Write $\mathbb{P}_{e,\beta}$ for the law of the height function on the graph consisting of just one edge e , say with $h(v) = 0$. We claim first that there exists some absolute constant C not depending on G , e or r such that

$$\mathbb{E}_{G,\beta} |\nabla h|_e^r \leq C \mathbb{E}_{e,\beta} |\nabla h|_e^r. \quad (15)$$

This implies the result because (as \mathcal{V}_β is convex and symmetric) the law of ∇h_e is log-concave and symmetric under $\mathbb{P}_{e,\beta}$ so that it has all moments.

Let $G \setminus e$ be the graph without the edge e . For $l \in \mathbb{Z}$, we define $\Omega_l(G) = \{\mathbf{n} \text{ on } G : \delta \mathbf{n} = l(\delta_v - \delta_{v'})\}$, and

$$Z_G^l = \sum_{\mathbf{n} \in \Omega_l(G)} w_\beta(\mathbf{n}),$$

and analogously $Z_{G \setminus e}^l$. Similarly to (11), we get from the high-temperature expansion of correlation functions that

$$\langle \sigma_v^l \bar{\sigma}_{v'}^l \rangle_{G \setminus e, \beta} = \frac{Z_{G \setminus e}^l}{Z_{G \setminus e}^0}.$$

By the definition of the height function and currents, we therefore have

$$\begin{aligned} \mathbb{P}_{G,\beta}(|\nabla h|_e = l) &= I_l(\beta) \frac{(Z_{G \setminus e}^l + Z_{G \setminus e}^{-l})}{Z_G^0} \\ &= 2I_l(\beta) \frac{Z_{G \setminus e}^l}{Z_{G \setminus e}^0} \frac{Z_{G \setminus e}^0}{Z_G^0} \leq 2I_l(\beta) = \mathbb{P}_{e,\beta}(|\nabla h|_e = l) Z_e^0, \end{aligned}$$

where we used the obvious bounds $\langle \sigma_v^l \bar{\sigma}_{v'}^l \rangle_{G \setminus e, \beta} \leq 1$, and $Z_{G \setminus e}^0 / Z_G^0 \leq 1$. Setting $C = Z_e^0$ we establish (15). \square

Lemma 5.17. *For all $\beta > 0$ and all $r \in \mathbb{N}$, there exists a $\tilde{C}_r < \infty$ such that for all finite planar graphs $G = (V, E)$ and $e \in E$,*

$$\mathbb{E}_{G,\beta}[|X_e|^r] \leq \tilde{C}_r.$$

Proof. For two nonnegative integers i, r , let $(i)_r = i(i-1)\cdots(i-r+1)$ be the falling factorial with the convention that $(i)_0 = 1$. Note that $(i)_r = 0$ whenever $i < r$. It will be convenient to look at the falling factorial moments. First note that by definition of Y_k ,

$$\begin{aligned} \mathbb{E}_\beta[(Y_k)_r] &= \frac{1}{I_k(\beta)} \sum_{i \geq 0} \frac{(i)_r}{i!(i+k)!} \left(\frac{\beta}{2}\right)^{2i+k} \\ &= \frac{\left(\frac{\beta}{2}\right)^r}{I_k(\beta)} \sum_{i \geq 0} \frac{1}{i!(i+k+r)!} \left(\frac{\beta}{2}\right)^{2i+k+r} = \left(\frac{\beta}{2}\right)^r \frac{I_{k+r}(\beta)}{I_k(\beta)}. \end{aligned}$$

By the Turán inequality (5), the map $k \mapsto I_{k+1}(\beta)/I_k(\beta)$ is decreasing and hence

$$\mathbb{E}_\beta[(Y_k)_r] = \left(\frac{\beta}{2}\right)^r \frac{I_{k+r}(\beta)}{I_k(\beta)} \leq \left(\frac{\beta}{2}\right)^r \frac{I_r(\beta)}{I_0(\beta)} =: C.$$

Now note that $(i)_r \geq |i-r|^r$ when $i \geq r$, and hence $i^r \leq 2^{r-1}(|i-r|^r + r^r) \leq 2^r((i)_r + r^r)$. Finally

$$\mathbb{E}_\beta[|X_e|^r \mid |\nabla h|_e = k] = \mathbb{E}_\beta[|Y_k|^r] \leq 2^r(C + r^r) := \tilde{C}_r,$$

where the last bound does not depend on k . Integrating over the possible values of $|\nabla h|_e$ concludes the proof. \square

5.5. Existence of phase transition in the XY model

In this section, we prove that for all translation invariant planar graphs $\Gamma = (V, E)$, the XY model undergoes a non-trivial phase transition in terms of the quantity $\chi_\beta^\varepsilon(L)$. As before, let 0 denote an arbitrary distinguished face of Γ . We define

$$\beta_0 = \inf\{\beta > 0 : \text{for all } \varepsilon > 0 \text{ and all cuts } L \text{ at } 0, \chi_\beta^\varepsilon(L) = \infty\}.$$

Theorem 5.18. *Let Γ be as above. Then $\beta_0 < \infty$.*

By Corollary 5.13 it is enough to show that for any such Γ , there exists a finite $\beta_0 > 0$ such that the associated height function delocalizes in the sense that there are no translation invariant Gibbs measures on the dual Γ^\dagger . We first implement this strategy for triangulations, where delocalization can be shown directly using the general result of Lammers [114] (Theorem 5.2).

Proof of Theorem 5.18 for triangulations. Let Γ be a translation invariant triangulation. Note that condition (6) in our case is equivalent to $I_1(\beta)/I_0(\beta) \geq \frac{1}{2}$. It is known that this fraction converges to 1 as $\beta \rightarrow \infty$ (see for example [155]), and therefore in light of Theorem 5.2, there are no translation invariant Gibbs measures for β large enough. \square

To extend beyond triangulations, we will use a different approach. We stress that in particular, we will not show delocalization of the height function on graphs that are not triangulations. Instead, we exploit monotonicity in coupling constants to bound from below the spin correlations on an arbitrary translation invariant graph by correlations on a modified graph that is a triangulation. We explain this procedure in detail for the square lattice, and briefly mention the extension to other lattices at the end.

In what follows, we will need the following well known monotonicity of spin correlations that is a classical consequence of the Ginibre inequality [80].




Figure 5.2: The transformation to a triangulation. The red edge on the right is the edge with different potential.

Lemma 5.19. *For each (infinite or finite) graph $G = (V, E)$, $\beta > 0$, $e \in E$, and $a, b \in V$, the function*

$$J_e \mapsto \langle \sigma_a \bar{\sigma}_b \rangle_{G, \beta}$$

is nondecreasing.

Proof of Theorem 5.18 for the square lattice. Let $\Gamma = (V, E)$ denote the square lattice.

In order to use (6), we need to transform Γ into a triangulation. See Figure 5.2 for guidance. Fix a square and double the bottom and left edge and put coupling constants $\beta/2$ on the doubled edges instead of β . Next, double the common vertex of the left and bottom edge and add an additional edge e , on which we set the coupling constant to infinity. This does not change the distribution of the spins. Finally, set the coupling constant on the edge e to 0, which is equivalent to removing the edge from the square, and repeat the procedure for all other squares. In this way, we obtain a new lattice Γ' , which consists of squares with a diagonal on which there is an additional vertex. Note that all coupling constants are now equal to $\beta/2$. By Lemma 5.19,

$$\langle \sigma_a \bar{\sigma}_b \rangle_{\Gamma, \beta} \geq \langle \sigma_a \bar{\sigma}_b \rangle_{\Gamma', \beta/2} \quad (16)$$

for all pairs of vertices a, b in Γ , using the natural embedding of Γ on Γ' .

Since Γ' is a translation invariant graph, the dichotomy statement of Theorem 5.4 holds. To show that there are no translation invariant Gibbs measures for the associated height function, notice that the dual $(\Gamma')^\dagger$ of Γ' (after collapsing the doubled edges to a single edge) is trivalent. Moreover, the height function on any finite subgraph of $(\Gamma')^\dagger$ has a potential given by $\mathcal{V}'_e = \mathcal{V}_e^{\beta/2}$ for the nondiagonal edges and $\mathcal{V}'_e = 2\mathcal{V}_e^{\beta/2}$ otherwise, and the potential \mathcal{V}' satisfies Lammers' condition (6) precisely when $(I_1(\beta/2)/I_0(\beta/2))^2 \geq \frac{1}{2}$. Since the fraction on the left-hand side tends to 1 as $\beta \rightarrow \infty$, we can choose β large enough so that there are no translation invariant Gibbs measures for the height function on $(\Gamma')^\dagger$.




Figure 5.3: The transformation of a general graph to a triangulation (after identifying the resulting multiple edges). The dashed edges are such that the coupling constant is set to infinity first, and then to zero (which is equivalent to removing the edges) and hence the spin correlations in the final graph are smaller than in the original graph.

Note that every cut on Γ embeds naturally as a cut on Γ' . Therefore, by Proposition 5.12 together with (16), we have that for each cut L on Γ and each $\epsilon > 0$,

$$\chi_{\Gamma,\beta}^{\epsilon}(L) \geq \chi_{\Gamma',\beta/2}^{\epsilon}(L) = \infty.$$

This finishes the proof. \square

To extend this proof to general graphs, we make each face into a triangulation by “zig-zagging” (see Figure 5.3).

5.6. No exponential decay implies a power-law lower bound

In this section we finish the proof of the main theorem by showing that the absence of exponential decay implies a power-law lower bound on the two-point function when $\Gamma = \mathbb{Z}^2$. Similar arguments can be applied to other graphs that in addition to being translation invariant possess reflection and rotation symmetries.

We will use the following two classical inequalities.

Lemma 5.20 (Lieb–Rivasseau inequality [125, 151]). *Let $G = (V, E)$ be any graph. Let $a, b \in V$ be distinct, and let H be a finite subgraph of G containing a and not containing b , and let ∂H be the set of vertices of H adjacent to at least one vertex outside H . Then*

$$\langle \sigma_a \bar{\sigma}_b \rangle_{G,\beta} \leq \sum_{c \in \partial H} \langle \sigma_a \bar{\sigma}_c \rangle_{H,\beta} \langle \sigma_c \bar{\sigma}_b \rangle_{G,\beta}.$$

Lemma 5.21 (Messenger–Miracle-Sole inequality [135]). *For any $n \in \mathbb{Z}$, the two sequences $\langle \sigma_0 \bar{\sigma}_{(n,k)} \rangle_{\mathbb{Z}^2,\beta}$ and $\langle \sigma_0 \bar{\sigma}_{(n+k,n-k)} \rangle_{\mathbb{Z}^2,\beta}$ are nonincreasing in k for $k \geq 0$.*

Proof of Theorem 5.1. Let 0 denote the vertex at the origin. For a finite subgraph G of \mathbb{Z}^2 containing 0, let

$$\varphi_{G,\beta} = \sum_{w \in \partial G} \langle \sigma_0 \bar{\sigma}_w \rangle_{G,\beta},$$




Figure 5.4: The $[-n, n]^2$ box Λ_n shaded in grey and the L^1 ball Λ'_n of radius $2n$.

where ∂G is the set of vertices of G adjacent to at least one vertex outside G . Define

$$\beta_c = \sup\{\beta : \text{there exists finite } G \text{ with } \varphi_{G,\beta} < 1\}. \quad (17)$$

We will show that β_c satisfies the properties listed in Theorem 5.1. To this end first fix $\beta < \beta_c$. By Lemma 5.19, there exists a finite graph G with $\varphi_{G,\beta} < 1$. Take m such that $G \subset \Lambda_m$ and let $x \in V$. Fix n so that $(n+1)m \geq |x|_1 \geq nm$. Iteratively applying the Lieb–Rivasseau inequality [125, 151] to translates of G gives

$$\langle \sigma_0 \bar{\sigma}_x \rangle_{\mathbb{Z}^2, \beta} \leq \sum_{w \in \partial G} \langle \sigma_0 \bar{\sigma}_w \rangle_{G, \beta} \sum_{w' \in \partial(G+w)} \langle \sigma_w \bar{\sigma}_{w'} \rangle_{G+w, \beta} \langle \sigma_{w'} \bar{\sigma}_x \rangle_{\mathbb{Z}^2, \beta} \leq \cdots \leq (\varphi_{G,\beta}(0))^n,$$

hence (i) holds true if $\beta < \beta_c$.

To conclude (ii), note that for each finite G , $\varphi_{G,\beta}$ is a continuous function of β , and hence the set in (17) is open. This means that for every $\beta \geq \beta_c$, we have $\varphi_{G,\beta} \geq 1$ for all finite subgraphs G .

Now let Λ_n be the box $[-n, n]^2$, and let Λ'_n be the ball in L^1 of radius $2n$ (see Figure 5.4). We write $x_n := (n, n) \in \partial\Lambda_n \cap \partial\Lambda'_n$ and $a_n = \langle \sigma_0 \bar{\sigma}_{x_n} \rangle_{\mathbb{Z}^2, \beta}$. By rotation symmetry and the Messager–Miracle–Sole [135] inequality, we have

$$a_n = \min_{v \in \partial\Lambda_n} \langle \sigma_0 \bar{\sigma}_v \rangle_{\mathbb{Z}^2, \beta} = \max_{v \in \partial\Lambda'_n} \langle \sigma_0 \bar{\sigma}_v \rangle_{\mathbb{Z}^2, \beta}.$$

For $\beta \geq \beta_c$, we moreover have

$$\sum_{w \in \partial\Lambda'_n} \langle \sigma_0 \bar{\sigma}_w \rangle_{\mathbb{Z}^2, \beta} \geq \varphi_{\Lambda'_n, \beta} \geq 1.$$

These two observations together imply that for any $v \in \partial\Lambda_n$,

$$\langle \sigma_0 \bar{\sigma}_v \rangle_{\mathbb{Z}^2, \beta} \geq a_n \geq \frac{1}{|\partial\Lambda'_n|} = \frac{1}{8n} \geq \frac{1}{8|v|}$$

which implies (ii).

Finally by Theorem 5.18 we know that there exists a finite β at which there is no exponential decay, and by classical expansions there exists a nonzero β at which there is exponential decay (see e.g. [4]). We conclude that $0 < \beta_c < \infty$. \square

Correlation Inequalities

The main purpose of this chapter is to present a new technique that may be applied to a further study of the XY model (possibly in higher dimensions). We develop a loop representation for *squares and products* of correlation functions. This is a generalization of the construction from Section 5.3, and to the best of our knowledge has not yet been described in the literature. It is also analogous to the double random current representation of the Ising model [1, 6, 82] but is more subtle as one has to deal with path switching rather than connection switching in a percolation model. We stress the fact that we do not use any of the results from this chapter in the remainder of this thesis, except for the well known inequalities of Lieb and Rivasseau, and Messenger and Miracle-Sole.

There will be two major differences in the definition of a loop configuration compared to Section 5.3: the edges will come in two colors, red and blue, corresponding to two currents \mathbf{r} and \mathbf{b} respectively, and we will allow the paths to enter vertices v at which the number of incoming and outgoing edges is not the same, i.e., $\delta(\mathbf{r} + \mathbf{b})_v \neq 0$. To be more precise, consider the following definition.

Definition 6.1 (Colored loop configurations outside S with sources φ). Let \mathcal{M} be a multigraph on V , let $S \subseteq V$, and $\varphi : V \rightarrow \mathbb{Z}$ with $\sum_{v \in V} \varphi_v = 0$. A *colored loop configuration (on \mathcal{M}) outside S with sources φ* is

- an assignment of a red or blue color to each edge of \mathcal{M} , together with
- a collection of
 - unrooted directed loops on \mathcal{M} avoiding S , and
 - directed open paths on \mathcal{M} not visiting S except possibly at their start and end vertex,

such that

- every edge of \mathcal{M} is traversed exactly once by a loop or a path, and

- at each vertex $v \in V \setminus S$, there are exactly $\varphi_v \mathbb{1}\{\varphi_v > 0\}$ outgoing and $-\varphi_v \mathbb{1}\{\varphi_v < 0\}$ incoming paths.

We write $\tilde{\mathcal{L}}_\varphi^S$ for the set of all colored loop configurations outside S with sources φ , and define a weight on $\tilde{\mathcal{L}}_\varphi^S$ by

$$\tilde{\lambda}_\beta^S(\omega) = \prod_{v \in V \setminus S} \frac{|\varphi_v|!}{((\deg_{\mathcal{M}}(v) + |\varphi_v|)/2)!} \prod_{e \in E} \frac{1}{\mathcal{M}_e!} \left(\frac{\beta}{2}\right)^{\mathcal{M}_e}, \quad (1)$$

where \mathcal{M} is the underlying multigraph, and \mathcal{M}_e is the number of copies of e in \mathcal{M} .

Note that this weight no longer only depends on the multigraph M and on S , but also on $|\varphi(\omega)|$, where $\varphi(\omega)$ are the sources of ω . Also note that in the above definition S and φ can be chosen independently. S denotes the set of vertices where we do not resolve any connections between paths and loops, and φ prescribes where the sources and sinks are (vertices with nonzero value of φ). At any such vertex v , we resolve as many connections as possible leaving only $|\varphi_v|$ incoming or outgoing arrows unmatched, depending on the sign of φ_v . This is the reason why φ appears in the above weight, which was not the case in Section 5.3.

As before if $S' \subseteq S$, then there is a natural map $\rho : \mathcal{L}^{S'} \rightarrow \mathcal{L}^S$ that consists in forgetting (or cutting) the loop and path connections at the vertices in $S \setminus S'$, and

$$\sum_{\tilde{\omega} \in \rho^{-1}[\omega]} \tilde{\lambda}_\beta^{S'}(\tilde{\omega}) = \tilde{\lambda}_\beta^S(\omega). \quad (2)$$

Definition 6.2 (Colored currents and consistent configurations). We will consider a pair of currents \mathbf{r}, \mathbf{b} that we think of as *red* and *blue* respectively. A colored loop configuration ω on $\mathcal{M}_{\mathbf{r}+\mathbf{b}}$ is called *consistent with \mathbf{r} and \mathbf{b}* if for every edge $vv' \in E$, the number of times the loops and paths traverse a red (resp. blue) copy of vv' in the direction of (v, v') is equal to $\mathbf{r}_{(v, v')}$ (resp. $\mathbf{b}_{(v, v')}$). In particular ω has sources $\delta(\mathbf{r} + \mathbf{b})$. We define $\tilde{\mathcal{L}}_{\mathbf{r}, \mathbf{b}}^S$ to be the set of all colored loop configurations on $\mathcal{M}_{\mathbf{r}+\mathbf{b}}$ outside S that are consistent with \mathbf{r} and \mathbf{b} .

For $\varphi, \psi : V \rightarrow \mathbb{Z}$, we also define

$$\tilde{\mathcal{L}}_{\varphi, \psi}^S = \bigcup_{\mathbf{r} \in \Omega_\varphi, \mathbf{b} \in \Omega_\psi} \tilde{\mathcal{L}}_{\mathbf{r}, \mathbf{b}}^S \subseteq \tilde{\mathcal{L}}_{\varphi+\psi}^S.$$

where the union is clearly disjoint. For brevity, we will write $\tilde{\mathcal{L}}_0^S$ instead of $\tilde{\mathcal{L}}_{0,0}^S$, where 0 denotes the zero function on V .

We now relate the weights of loops to those of pairs of currents. To this end, note that for each edge $vv' \in E$, there are exactly

$$\frac{|\mathbf{r} + \mathbf{b}|_{vv'}!}{\mathbf{r}_{(v, v')}! \mathbf{r}_{(v', v)}! \mathbf{b}_{(v, v')}! \mathbf{b}_{(v', v)}!}$$




Figure 6.1: Path switching behavior at a vertex v that is (resp. is not) the start or end-point of the switched path. Left: the values of *both* $\delta \mathbf{r}_v$ and $\delta \mathbf{b}_v$ are increased by one after switching. Right: the values are not changed.

ways of assigning color to the copies of vv' in $\mathcal{M}_{\mathbf{r}+\mathbf{b}}$, and to orient them in the two possible ways so that the result is consistent with \mathbf{r} and \mathbf{b} . Moreover, independently of the choices of colors and orientations, there are exactly

$$\frac{((\deg_{\mathcal{M}_{\mathbf{r}+\mathbf{b}}}(v) + |\varphi_v|)/2)!}{|\varphi_v|!}$$

possible pairings of the incoming and outgoing edges at each vertex $v \in V \setminus S$ such that there are exactly $\varphi_v \mathbf{1}\{\varphi_v > 0\}$ outgoing and $-\varphi_v \mathbf{1}\{\varphi_v < 0\}$ incoming edges unpaired. This is equivalent to choosing the possible steps that all the loops and paths in the configuration make at v . Combining all this, we get the following identity:

$$\sum_{\omega \in \tilde{\mathcal{L}}_{\mathbf{r},\mathbf{b}}^S} \tilde{\lambda}_{\beta}^S(\omega) = w_{\beta}(\mathbf{r})w_{\beta}(\mathbf{b}). \quad (3)$$

An important observation again is that the right-hand side is independent of S , and hence so is the left-hand side.

In particular, for two sourceless currents, we have

$$\sum_{\omega \in \tilde{\mathcal{L}}_0^{\emptyset}} \tilde{\lambda}_{\beta}^{\emptyset}(\omega) = \left(\sum_{\mathbf{n} \in \Omega_0} w_{\beta}(\mathbf{n}) \right)^2 = (Z_{G,\beta}^0)^2. \quad (4)$$

Again, in the case when G is planar we get the following distributional identity. Let $\tilde{\mathbf{P}}_{G,\beta}$ to be the probability measure on $\tilde{\mathcal{L}}_0 := \tilde{\mathcal{L}}_0^{\emptyset}$ induced by the weights $\tilde{\lambda}_{\beta} := \tilde{\lambda}_{\beta}^{\emptyset}$. For each face $u \in U$ of G , and $\omega \in \tilde{\mathcal{L}}_0$, define $W_{\omega}(u)$ to be the total net winding of all the loops in ω around u .

Proposition 6.3. *The law of $(W(u))_{u \in U}$ under $\tilde{\mathbf{P}}_{G,\beta}$ is the same as the law of the sum of two independent height functions $(h(u) + h'(u))_{u \in U}$ under $\mathbb{P}_{G,\beta}$.*

6.1. The two point-function and path switching

We now turn to the loop representation of the square of the two-point function. To this end, write $\varphi = \delta_a - \delta_b$. Similar to (11), we get

$$\langle \sigma_a \bar{\sigma}_b \rangle_{G,\beta}^2 = \left(\frac{\sum_{\mathbf{n} \in \Omega_\varphi} w_\beta(\mathbf{n})}{Z_{G,\beta}^0} \right)^2 = \frac{\sum_{\mathbf{r}, \mathbf{b} \in \Omega_\varphi} w_\beta(\mathbf{r}) w_\beta(\mathbf{b})}{(Z_{G,\beta}^0)^2} = \frac{\sum_{\omega \in \tilde{\mathcal{L}}_{\varphi,\varphi}^{a,b}} \tilde{\lambda}_\beta^{a,b}(\omega)}{(Z_{G,\beta}^0)^2}, \quad (5)$$

where the last equality follows from (3).

As before, we now want to reverse some of the paths. However, this time we also need to take care of the colors of the edges visited by a path. This motivates the following definition.

Definition 6.4 (Path switching). For a path γ in a colored loop configuration ω , we define $s(\gamma)$ to be the path obtained from γ by

- reversing the orientation of γ , and
- swapping the colors of the edges visited by γ .

We also define ω' to be the configuration where γ is replaced by $s(\gamma)$. This operation does not change the underlying multigraph. Moreover if γ starts at a and ends at b , then for any $\varphi, \psi : V \rightarrow \mathbb{Z}$, path switching maps

$$\omega \in \tilde{\mathcal{L}}_{\varphi,\psi}^S \quad \text{to} \quad \omega' \in \tilde{\mathcal{L}}_{\varphi-\delta_a+\delta_b,\psi-\delta_a+\delta_b}^S$$

(see Figure 6.1).

We note that there are two important cases in which path switching does not change the weight $\tilde{\lambda}_\beta^S$. The first one is when $\{a, b\} \subset S$, and the second one is when $\varphi_a + \psi_a = 1$, and $\varphi_b + \psi_b = -1$, since then the absolute value of the sources of the configuration does not change.

Again the crucial observation now is that switching a path going from a to b maps $\omega \in \tilde{\mathcal{L}}_{\varphi,\varphi}^{a,b}$ to $\omega' \in \tilde{\mathcal{L}}_0^{a,b}$, and hence erases the sources and sinks of the underlying currents. Indeed one can easily check (see Figure 6.1) that after reversing a path and swapping the colors, the number of incoming minus the number of outgoing red and blue edges at every vertex $v \notin \{a, b\}$ in ω' is the same as in ω , whereas at a and b this number is decreased by one. Since we did not change the sources outside $\{a, b\}$, we do not change the weight of a loop configuration, and hence obtain in the same way as in Section 5.3.2 that

$$\sum_{\omega \in \tilde{\mathcal{L}}_{\varphi,\varphi}^{a,b}} \tilde{\lambda}_\beta^{a,b}(\omega) = \sum_{\omega \in \tilde{\mathcal{L}}_0^{a,b}} \frac{m_{b,a}(\omega)}{m_{b,a}(\omega) + 1} \tilde{\lambda}_\beta^\theta(\omega) \mathbf{1}\{m_{b,a}(\omega) > 0\}.$$

Together with (5) this implies the following loop representation of the square of the two-point function.

Proposition 6.5. *Let $a, b \in V$ be distinct. Then*

$$\langle \sigma_a \bar{\sigma}_b \rangle_{G,\beta}^2 = \tilde{\mathbf{E}}_{G,\beta} \left[\frac{m_{a,b}}{m_{a,b} + 1} \right],$$

and in particular

$$\frac{1}{2} \tilde{\mathbf{P}}_{G,\beta}(m_{a,b} > 0) \leq \langle \sigma_a \bar{\sigma}_b \rangle_{G,\beta}^2 \leq \tilde{\mathbf{P}}_{G,\beta}(m_{a,b} > 0).$$

Remark 6.1. The constant in the inequality of Lemma 5.15 can be improved to 1 using the same method as above but starting from colored loop configurations in $\tilde{\mathcal{L}}_{0,2\varphi}^{a,b}$ instead of $\tilde{\mathcal{L}}_{\varphi,\varphi}^{a,b}$.

6.2. Application to some inequalities

As a further application we now prove an inequality that is related to but independent of the Ginibre inequality.

Lemma 6.6. *Let $a, b, c \in V$. Then*

$$\langle \sigma_a \bar{\sigma}_b \rangle_{G,\beta} \geq \langle \sigma_a \bar{\sigma}_c \rangle_{G,\beta} \langle \sigma_c \bar{\sigma}_b \rangle_{G,\beta} \geq \langle \sigma_a \sigma_b \bar{\sigma}_c^2 \rangle_{G,\beta}.$$

Proof. The two inequalities have, maybe quite surprisingly, almost the same proof. We only prove the first and leave the second to the reader. We set $S = \{c\}$ and will write c instead of $\{c\}$ in our notation. We also define $\varphi = \delta_a - \delta_c$, $\psi = \delta_b - \delta_c$, and note that $\psi - \varphi = \delta_b - \delta_a$. Also note that for each $\omega \in \tilde{\mathcal{L}}_{\varphi,\psi}^c$, the unique path starting at a must end at c . Consider the map $\omega \mapsto \omega'$ that switches this path. Clearly this is a bijection between $\tilde{\mathcal{L}}_{\varphi,\psi}^c$ and

$$\{\omega \in \tilde{\mathcal{L}}_{0,\psi-\varphi}^c : \text{the unique path ending at } a \text{ starts at } c\}.$$

Moreover, we have $|\varphi_v(\omega)| = |\varphi_v(\omega')|$ for all $v \neq c$, and hence the weights $\tilde{\lambda}_\beta^c$ are preserved. This means that

$$\begin{aligned} \langle \bar{\sigma}_a \sigma_c \rangle_{G,\beta} \langle \bar{\sigma}_b \sigma_c \rangle_{G,\beta} (Z_{G,\beta}^0)^2 &= \sum_{\omega \in \tilde{\mathcal{L}}_{\varphi,\psi}^c} \tilde{\lambda}_\beta^c(\omega) \\ &= \sum_{\omega \in \tilde{\mathcal{L}}_{0,\psi-\varphi}^c} \mathbb{1}\{c \rightarrow a \text{ in } \omega\} \tilde{\lambda}_\beta^c(\omega) \leq \langle \sigma_a \bar{\sigma}_b \rangle_{G,\beta} (Z_{G,\beta}^0)^2, \end{aligned}$$

where we used (3) twice. This finishes the proof. \square

Remark 6.2. Note that the Ginibre inequality in e.g. [142, Theorem 2.3] is equivalent to a relation between increments given by

$$\langle \sigma_a \bar{\sigma}_b \rangle_{G,\beta} - \langle \sigma_a \bar{\sigma}_c \rangle_{G,\beta} \langle \sigma_c \bar{\sigma}_b \rangle_{G,\beta} \geq \langle \sigma_a \bar{\sigma}_c \rangle_{G,\beta} \langle \sigma_c \bar{\sigma}_b \rangle_{G,\beta} - \langle \sigma_a \sigma_b \bar{\sigma}_c^2 \rangle_{G,\beta}.$$

Comparing this with the statement of Lemma 6.6, the latter proves nonnegativity of the increments. Hence, the Ginibre inequality and Lemma 6.6 do not imply one another.

The purpose of the remainder of this section is to give more applications of the representation introduced above. We start with two new bijective proofs of the classical inequalities that we used in the proof of our main theorem.

Lemma 6.7 (Lieb–Rivasseau inequality [125, 151]). *Let $G = (V, E)$ be any graph. Let $a, b \in V$ be distinct, and let H be a finite subgraph of G containing a and not containing b , and let ∂H be the set of vertices of H adjacent to at least one vertex outside H . Then*

$$\langle \sigma_a \bar{\sigma}_b \rangle_{G, \beta} \leq \sum_{c \in \partial H} \langle \sigma_a \bar{\sigma}_c \rangle_{H, \beta} \langle \sigma_c \bar{\sigma}_b \rangle_{G, \beta}.$$

Proof. It is enough to assume that G is finite, and then approximate an infinite graph by finite subgraphs. The proof is similar to the previous one. Assume $a \notin \partial H$. Otherwise, there is nothing to prove. Fix $c \in \partial H$ and $S = \{c\}$. We will write c instead of $\{c\}$ in our notation. Let $\varphi = \delta_c - \delta_a$, $\psi = \delta_c - \delta_b$, and note that $\psi - \varphi = \delta_a - \delta_b$.

Write $\tilde{\mathcal{L}}_c$ for the collection of colored loop configurations $\omega \in \tilde{\mathcal{L}}_{0, \psi - \varphi}^c$ with the property that the unique path starting at a exits $H \setminus \partial H$ at c , and ω has no red edges outside of H . For $\omega \in \tilde{\mathcal{L}}_c$ consider a colored loop configuration where this path is switched. Clearly this is a bijection between $\tilde{\mathcal{L}}_c$ and the set of configurations $\omega' \in \tilde{\mathcal{L}}_{\varphi, \psi}^c$ that have no red edges outside H , and for which the unique path ending at a stays within $H \setminus \partial H$ until it hits c . Denote this collection of configurations by $\tilde{\mathcal{L}}'_c$. Moreover, we have $|\varphi_v(\omega)| = |\varphi_v(\omega')|$ for all $v \neq c$, and hence the weights $\tilde{\lambda}_\beta^c$ are preserved.

Let $\tilde{\mathcal{E}}_c$ be the collection of $\omega \in \tilde{\mathcal{L}}_{0, \psi - \varphi}^c$ with the property that the unique path from a to b exits $H \setminus \partial H$ in c , and ω does not have red edges outside of H . Clearly, the subset of $\tilde{\mathcal{L}}_{0, \psi - \varphi}^c$ consisting of configurations with no red edges outside of H equals the disjoint union $\cup_{c \in \partial H} \tilde{\mathcal{E}}_c$ and cutting $\omega \in \tilde{\mathcal{E}}_c$ at c gives an element of $\tilde{\mathcal{L}}_c$. In light of (2), we therefore have

$$\begin{aligned} \langle \sigma_a \bar{\sigma}_b \rangle_{G, \beta} Z_{G, \beta}^0 Z_{H, \beta}^0 &= \sum_{c \in \partial H} \sum_{\omega \in \tilde{\mathcal{L}}_c} \tilde{\lambda}_\beta^c(\omega) \\ &= \sum_{c \in \partial H} \sum_{\omega' \in \tilde{\mathcal{L}}'_c} \tilde{\lambda}_\beta^c(\omega') \leq \sum_{c \in \partial H} \langle \sigma_c \bar{\sigma}_a \rangle_{H, \beta} \langle \sigma_c \bar{\sigma}_b \rangle_{G, \beta} Z_{G, \beta}^0 Z_{H, \beta}^0, \end{aligned}$$

which completes the proof. \square

We are also able to use the colored loop representation to prove the Messenger–Miracle–Sole inequality.

Lemma 6.8 (Messenger–Miracle–Sole inequality [135]). *For any $n \in \mathbb{Z}$, the two sequences $\langle \sigma_0 \bar{\sigma}_{(n, k)} \rangle_{\mathbb{Z}^2, \beta}$ and $\langle \sigma_0 \bar{\sigma}_{(n+k, n-k)} \rangle_{\mathbb{Z}^2, \beta}$ are nonincreasing in k for $k \geq 0$.*

Geometrically, this in particular implies that the largest correlation with the spin at 0 on any vertical, horizontal or diagonal straight line is attained by the vertex closest to

0. This will follow from the following lemma after taking $G \nearrow \mathbb{Z}^2$. The proof is inspired by the one from [7] for the Ising model. The idea is to fold a graph across a line and think of the parts of the current coming from both sides of the line as the red and blue current in the colored loop representation.

Lemma 6.9. *Let $G = (V, E)$ be a subgraph of \mathbb{Z}^2 symmetric under reflection across a line L . Let $a, b \in V$ lie on the same side of L , and let $L(b) \in V$ be the reflection of b . Then*

$$\langle \sigma_a \bar{\sigma}_b \rangle_{G, \beta} \geq \langle \sigma_a \bar{\sigma}_{L(b)} \rangle_{G, \beta}.$$

Proof. We only consider the easier case when L passes through vertices. This means that it is either a diagonal, or a horizontal (vertical) line at integer height. The more involved case when L passes only through the edges (this case implies Lemma 5.21 for horizontal and vertical lines) we leave to the interested reader.

If L is horizontal or vertical, then split the edges that lie on L into two parallel edges with coupling constants $\beta/2$, and think of the resulting graph as a new graph G . Write Z for the set of vertices on L , and $G_- = (V_-, E_-)$ and $G_+ = (V_+, E_+)$ for the two isomorphic parts of G separated by L where G_- contains a and b (each of them also containing Z).

We can decompose a current \mathbf{n} on G into two parts: \mathbf{r} and \mathbf{b} on G_- and G_+ . In what follows, we identify G_- with G_+ under the obvious isomorphism, and all currents are considered on G_- unless stated otherwise. Let \mathcal{C}_k , for $k = 0, 1$, be the set of functions $\varphi : V_- \rightarrow \mathbb{Z}$ such that $\varphi_v = 0$ for $v \in V_- \setminus Z$, and $\sum_{v \in Z} \varphi_v = k$. Since every current in $\Omega_{\delta_a - \delta_{L(b)}}(G)$ must have a total flux of $+1$ across L , we can write

$$\begin{aligned} \langle \sigma_a \bar{\sigma}_{L(b)} \rangle_{G, \beta} Z_{G, \beta}^0 &= \sum_{\varphi \in \mathcal{C}_1} \sum_{\mathbf{r} \in \Omega_{\delta_a - \varphi}, \mathbf{b} \in \Omega_{-\delta_b + \varphi}} w_\beta(\mathbf{r}) w_\beta(\mathbf{b}) \\ &= \sum_{\varphi \in \mathcal{C}_1} \sum_{\mathbf{r} \in \Omega_{\delta_a - \varphi}, \mathbf{b} \in \Omega_{\delta_b - \varphi}} w_\beta(\mathbf{r}) w_\beta(\mathbf{b}) \\ &= \sum_{\varphi \in \mathcal{C}_1} \sum_{\omega \in \tilde{\mathcal{L}}_{\delta_a - \varphi, \delta_b - \varphi}^Z} \tilde{\lambda}_\beta^Z(\omega), \end{aligned}$$

where the second inequality holds true as the weight w_β is invariant under reversal of the current, and the last equality is a consequence of (3). Now, for each $\omega \in \tilde{\mathcal{L}}_{\delta_a - \varphi, \delta_b - \varphi}^Z$ switch the unique path γ starting at b . This transformation preserves weights and results in a configuration $\omega' \in \tilde{\mathcal{L}}_{\delta_a - \delta_b - \varphi', -\varphi'}^Z$, where $\varphi' = \varphi - \delta_z \in \mathcal{C}_0$ and $z \in Z$ is the vertex at which γ ends. Reversing the order of the steps above we therefore get

$$\langle \sigma_a \bar{\sigma}_{L(b)} \rangle_{G, \beta} Z_{G, \beta}^0 = \sum_{\varphi \in \mathcal{C}_1} \sum_{z \in Z} \sum_{\omega' \in \tilde{\mathcal{L}}_{\delta_a - \delta_b - \varphi', -\varphi'}^Z} \tilde{\lambda}_\beta^Z(\omega') \mathbf{1}\{\text{the path ending at } b \text{ starts at } z\}$$

$$\begin{aligned}
 &= \sum_{\varphi' \in \mathcal{C}_0} \sum_{\omega' \in \tilde{\mathcal{L}}_{\delta_a - \delta_b - \varphi', -\varphi'}^Z} \tilde{\lambda}_\beta^Z(\omega') \mathbb{1}\{\text{the path ending at } b \text{ starts in } Z\} \\
 &\leq \sum_{\varphi' \in \mathcal{C}_0} \sum_{\mathbf{r} \in \Omega_{\delta_a - \delta_b - \varphi'}, \mathbf{b} \in \Omega_{\varphi'}} w_\beta(\mathbf{r}) w_\beta(\mathbf{b}) \\
 &= \langle \sigma_a \bar{\sigma}_b \rangle_{G, \beta} Z_{G, \beta}^0,
 \end{aligned}$$

where the last equality follows since the total flux of a current in $\Omega_{\delta_a - \delta_b}(G)$ across L is zero. \square

6.3. Limitations of the colored loop representation

A natural idea is to try to prove the Ginibre inequality in form of Lemma 5.19 using our representation. One would like to show that the derivative of the two-point function with respect to one coupling constant J_e is nonnegative. Using colored loop configurations we can write

$$\begin{aligned}
 (Z_{G, \beta}^0)^2 \frac{\partial}{\partial J_e} \langle \sigma_a \bar{\sigma}_b \rangle_{G, \beta} &= Z_{G, \beta}^0 \frac{\partial}{\partial J_e} Z_{G, \beta}^{\delta_a - \delta_b} - Z_{G, \beta}^{\delta_a - \delta_b} \frac{\partial}{\partial J_e} Z_{G, \beta}^0 \\
 &= J_e^{-1} \sum_{\omega \in \tilde{\mathcal{L}}_{\delta_a - \delta_b, 0}^\emptyset} \tilde{\lambda}_\beta^\emptyset(\omega) (R_e(\omega) - B_e(\omega)),
 \end{aligned}$$

where $R_e(\omega)$ and $B_e(\omega)$ is respectively the number of red and blue copies of e in the multigraph visited by the unique path from a to b in ω . Without going into too many details, to justify the second equality we make the following observations. First, taking the derivative with respect to J_e is equivalent to dividing by J_e and marking one of the copies of e of the right color (here the currents in $\Omega_{\delta_a - \delta_b}$ are red and those in Ω_0 are blue). Then, if the marked edge is not on the path from a to b , we switch the corresponding loop (reverse it and swap the colors). This does not change the weight of the configuration. Such terms hence cancel out from the expression above as the loops going through a marked blue copy of e are counted with a minus sign. The remaining terms are those whose marked edge lies on the distinguished path. This gives the final formula.

Clearly, the final result is not evidently nonnegative and we would need additional arguments to conclude the Ginibre inequality. On the other hand, the Ginibre inequality implies the distinguished path visits red edges more often than blue edges on average.

Duality between height functions and spin models

7.1. Introduction

The phenomenon of duality in statistical mechanics goes back to the famous work of Kramers and Wannier who discovered an exact identity between the partition functions of the Ising model on a finite planar graph and an Ising model (at a different temperature) on its dual graph [110]. They used it to identify the self-dual temperature (that stays invariant under the duality transformation) as the point of phase transition in the model on the square lattice (that is itself a self-dual graph). In an extended version of this correspondence, spin correlation functions are mapped to correlators of dual disorder variables introduced by Kadanoff and Ceva [101]. This construction has been very fruitful in the study of the Ising model. A notable example are the works of Smirnov [159], and Chelkak and Smirnov [46], who derived scaling limits of certain variants of order-disorder correlations (the so called fermionic observables), providing the first proofs of conformal invariance of the critical Ising model.

It is by now classical that analogous duality transformations exist for more general spin models whose state space is a locally compact abelian group [71, 154, 173] (we refer to [52] for an introductory account). In such a setting Fourier transforms can be used to map one model with values in a group \mathbb{G} to another model with values in the Pontryagin-dual group $\widehat{\mathbb{G}}$. For example, for $\mathbb{G} = \mathbb{Z}/q\mathbb{Z}$, $q \in \mathbb{N}$, duality was a crucial tool in the study of the planar q -state Potts model, the associated random cluster model, and the Ashkin–Teller model (see e.g. [15, 23, 81, 126, 128, 146]). These groups and models are self-dual in the sense that $\mathbb{G} \cong \widehat{\widehat{\mathbb{G}}}$, and moreover dual to the same model on the dual graph, but with possibly different coupling constants. Another famous self-dual example is $\mathbb{G} = \mathbb{R}$ together with the discrete Gaussian free field, where duality exchanges electric and magnetic operators of the field (see e.g. [52]).

In this chapter we go beyond the self-dual domain and consider the mutually dual but distinct groups of the integers \mathbb{Z} and the circle \mathbb{S} . This results in two very different objects

facing each other on the opposite sides of the duality relation: one is a discrete *random height function* with an unbounded set of values, and the other is a *continuous spin model* with spins in the circle. Duality can be then used to transfer probabilistic information from one side to the other. A landmark application of this relation appeared in the work of Fröhlich and Spencer [71] who rigorously established the Berezinskii–Kosterlitz–Thouless (BKT) phase transition in the classical XY model on the square lattice (see Section 7.4 for more background). They first showed delocalization of the associated height function, and then used duality to conclude that spin correlations decay at most polynomially fast. New proofs of the latter implication appeared recently in [8, 64] which together with a novel approach to delocalization introduced by Lammers [114] yields alternative proofs of the BKT transition. All three proofs of [8, 64, 71] use duality “in the same direction” in that they study spin correlations via disorder correlations in the dual height function. This leads to technical complications as disorders are non-local functions of the height field. In [8, 64] these issues were taken care of by considering different graphical representations of the models. Here we argue that following duality “in the opposite direction” leads to an even more concise proof (that only uses duality itself) of the implication that delocalization of the height function implies the BKT transition in the spin model. Indeed, when the disorders appear on the spin model side, one can “localize” them by simply using the Taylor expansion to the second order, which is clearly impossible when disorders come as discrete excitations of the heights.

Duality is an exact correspondence, and hence one expects that the critical point of the localization-delocalization phase transition is dual to the BKT critical point. This was recently proved for the XY and the Villain model by Lammers [116]. Here we also provide a result in the same direction for a larger class of models that includes the XY model. To be more precise, we establish an equivalence between the delocalization of the height function and the divergence of a certain series (a type of susceptibility) of correlation functions in the spin model.

Another contribution of this chapter is a universal upper bound on the variance of the height function in terms of the variance of the discrete GFF. This holds true for all height function models with positive definite potentials, and moreover irrespective of the graph being planar or not. There are two main applications. In the planar case, e.g., on \mathbb{Z}^2 , this leads to a conjecturally optimal (up to a constant) logarithmic in the size of the system upper bound when the height function is delocalized. On the other hand, it shows that on transient graphs, e.g. on \mathbb{Z}^d , $d \geq 3$, the variance is always uniformly bounded and the height function is localized.

For a special class of potentials, we also establish monotonicity of the variance of the height function in a natural temperature parameter. To the best of our knowledge, this is the first results of this type. We achieve this by transporting, through duality,

the (appropriately generalized) Ginibre inequalities. One consequence is a direct proof of the fact that for the XY height function there is only one point of phase transition from a localized to a delocalized regime. A (more involved) proof of this fact was first given in [116]. Together with the dichotomy of Lammers [115], this also shows that for the XY height function on \mathbb{Z}^2 , the transition is sharp. Another application is that for two-dimensional Euclidean lattices with non nearest neighbor interactions, the height function undergoes a localization-delocalization phase transition.

We note that we only consider height functions with positive definite potentials, i.e., those that have well defined dual spin models, and vice versa. The chapter is organized as follows:

- In Section 7.2 we recall the notion of duality, and state in Lemma 7.3 its consequence for the covariances of the gradient of the height function and gradient of the spin model. This is the stepping stone to the remaining results in this chapter.
- In Section 7.3 we establish an upper bound on the variance of the height function in terms of the Green's function of the underlying simple random walk. The bound is valid on any, not necessarily planar graph, and in particular implies localization of the height function on graphs on which simple random walk is transient.
- In Section 7.4 in Theorem 7.7 we give a direct proof of the fact that in two dimensions, delocalization of the height function implies a BKT phase transition in the spin model in the sense that certain spin correlation functions are not summable. In Corollary 7.8 using classical correlation inequalities we translate this to an analogous statement for the standard two-point functions, recovering the main result of [64]. Finally, for a subclass of spin models (that includes the classical XY model) we show that the above mentioned implication is actually an equivalence.
- In Section 7.5 in Theorem 7.10 we show that for a certain class of height functions, the variance is increasing in the inverse temperature. We use this to prove that a phase transition occurs for these random height functions, when the underlying graph is planar or “almost planar”.
- In Section 7.6, using only duality, we show that a certain (non-local) observable of a classical spin model has, up to multiplicative constants, the covariance of the discrete Gaussian free field. Remarkably, this holds for all graphs and does not depend on the temperature.
- In Section 7.7 we show a central limit theorem in the planar spin model that holds irrespectively of the temperature.

- In Appendix B we recall and give concise proofs of the main results needed for duality.
- In Appendix B.1 we extend the Ginibre inequalities to the setting that we need in Section 7.5.
- In Appendix B.2 we review the notion of reflection positivity that we use in Section 7.4.

7.2. General duality

7.2.1 Discrete calculus

We first give a basic background on discrete calculus on graphs, staying close to the language of [131]. Let $G = (V, E)$ be a locally finite graph and let \mathbb{G} be a group (we will consider $\mathbb{G} = \mathbb{R}, \mathbb{Z}$ with addition and $\mathbb{G} = \mathbb{S} := \{z \in \mathbb{C} : |z| = 1\}$ with multiplication). To keep the exposition homogenous, we will use the additive notation for all considered groups.

A 1-form ω taking values in \mathbb{G} is an antisymmetric function defined on the directed edges \vec{E} of G , i.e., such that $\omega_{vv'} = -\omega_{v'v}$, where vv' denotes the directed edge (v, v') . The set of 1-forms will be denoted by $\Omega^1(\mathbb{G}) = \Omega^1(G, \mathbb{G})$, and the set \mathbb{G}^V by $\Omega^0(\mathbb{G}) = \Omega^0(G, \mathbb{G})$. We will identify the space of 1-forms with \mathbb{G}^E by fixing, once and for all, one of the two orientations for each edge in E . Define the *boundary* operator $d^* : \Omega^1(\mathbb{G}) \rightarrow \Omega^0(\mathbb{G})$ by

$$d^* \omega_x = \sum_{y \sim x} \omega_{yx},$$

where $y \sim x$ indicates that y and x are adjacent in G , and the *co-boundary* operator $d : \Omega^0(\mathbb{G}) \rightarrow \Omega^1(\mathbb{G})$ by

$$df_{xy} = f_y - f_x.$$

Note that d^* and d are homomorphisms between groups \mathbb{G}^E and \mathbb{G}^V , and hence we can define groups

$$H_\star(\mathbb{G}) = H_\star(G, \mathbb{G}) := \text{Im}(d) \cong \mathbb{G}^V / \ker(d) \quad \text{and} \quad H_\diamond(\mathbb{G}) = H_\diamond(G, \mathbb{G}) := \ker(d^*).$$

For $\mathbb{G} = \mathbb{S}$, we will write dJ to be the Haar probability measure on the induced (compact) groups $H_\star(\mathbb{S})$ and $H_\diamond(\mathbb{S})$. If \mathbb{G} is only locally compact, the Haar measure is defined up to a multiplicative constant and we fix some normalisation. For a more concrete definition, we refer to Appendix B. We make the convention that the space over which we integrate determines the measure.

Notation. In what follows we will use the letters ϵ, ω (resp. f, g, τ) to denote deterministic elements of $\Omega^1(\mathbb{G})$ (resp. $\Omega^0(\mathbb{G})$) when $\mathbb{G} = \mathbb{R}$ or when \mathbb{G} is not specified. We will write \mathbf{n} and h for (mostly random) elements of $\Omega^1(\mathbb{Z})$ and $\Omega^0(\mathbb{Z})$ respectively, and J and θ for (mostly random) elements of $\Omega^1(\mathbb{S})$ and $\Omega^0(\mathbb{S})$ respectively.

We will also often abuse notation in the following sense: through the identification $\exp(i\theta) \leftrightarrow \theta$, we have $\mathbb{S} \cong (-\pi, \pi]$, and we will view $J \in \Omega^1(\mathbb{S})$ as belonging to $\Omega^1(\mathbb{R})$. On the other hand, by considering real numbers modulo 2π , we will map $\Omega^1(\mathbb{R})$ to $\Omega^1(\mathbb{S}^1)$. One has to be careful when going from one space to the other: the embedding does not map $H_{\#}(\mathbb{S})$ to $H_{\#}(\mathbb{R})$, because for example a 1-form $\omega \in \Omega^1(\mathbb{S})$ which satisfies $d^*\omega = 0$ in \mathbb{S} , only satisfies $d^*\omega = 0$ modulo 2π when viewed as a 1-form in $\Omega^1(\mathbb{R})$. We will also think of $H_{\#}(\mathbb{Z})$ as a subset of $H_{\#}(\mathbb{R})$ in the obvious way.

7.2.2 Spin models and random height functions

In this section we consider a finite graph $G = (V, E)$. We will study random *spin* and *height* 1-forms taking values in the spaces $H_{\#}(\mathbb{S})$ and $H_{\#}(\mathbb{Z})$ respectively for $\# \in \{\diamond, \star\}$.

Definition 7.1 (Height function and spin potentials). Let $\mathcal{V} : \mathbb{Z} \rightarrow \mathbb{R} \cup \{+\infty\}$ be symmetric, i.e., $\mathcal{V}(n) = \mathcal{V}(-n)$, such that

$$\sum_{n \in \mathbb{Z}} n^2 \exp(-\mathcal{V}(n)) < \infty, \tag{1}$$

and moreover such that $\exp(-\mathcal{V})$ is *positive definite*: for all $\alpha \in \mathbb{R}$,

$$w(\alpha) := \exp(-\mathcal{V}(0)) + \sum_{n=1}^{\infty} \exp(-\mathcal{V}(n)) 2 \cos(n\alpha) > 0. \tag{2}$$

We call \mathcal{V} the *height function potential*, and $\mathcal{U}(\alpha) := -\log w(\alpha)$ the *spin potential*.

We will always assume that the considered potentials satisfy the conditions of Definition 7.1. Note that condition (1) implies that the series in (2) is absolutely summable, and moreover that w , as well as \mathcal{U} , is twice continuously differentiable in α . Also note that even though this is not the classical definition of positive definiteness, it is equivalent to it by Bochner's theorem.

Example 7.1. The following potentials satisfy the conditions of Definition 7.1:

- $\mathcal{V}(n) = -\log(I_n(\beta))$, where $I_n(\beta)$ is the modified Bessel function of the first kind, and $\mathcal{U}(t) = -\beta \cos(t)$ for all $\beta > 0$ is the potential of the *classical XY model*,
- $\mathcal{V}(n) = \beta n^2$ for all $\beta > 0$ is the potential of the *integer-valued Gaussian free field* and the corresponding \mathcal{U} defined through the series in (2) is the potential of the *Villain spin model*,

- $\mathcal{V}(n) = \beta \mathbf{1}\{n = \pm 1\} + \infty \mathbf{1}\{|n| > 1\}$ for $\exp(-\beta) < 1/2$ is a model of random (nonuniform) *Lipschitz functions*.
- Any *annealed Gaussian* potential \mathcal{V} meaning that there exists a finite Borel measure λ on $[0, \infty)$ such that

$$e^{-\mathcal{V}(n)} = \int_{[0, \infty)} e^{-\frac{\gamma}{2}n^2} \lambda(d\gamma)$$

for all n . It satisfies Definition 7.1 because the function $n \mapsto \frac{\gamma}{2}n^2$ does and because by dominated convergence, we can exchange the integral and the summation in (2). This class includes the potentials $\mathcal{V}(n) = \beta|n|^a$ for any $a \in (0, 2]$ (see [8]).

Let ω be as in Definition 7.1. Fix $\# \in \{\diamond, \star\}$, and consider a probability measure on *spin 1-forms* $J \in H_{\#}(\mathbb{S})$ defined by

$$d\mu_{\#}(J) = d\mu_{G, \#}(J) = \frac{1}{Z_{\#}} \left(\prod_{e \in E} w(J_e) \right) dJ, \quad (3)$$

where $Z_{\#}$ is the *partition function*, and dJ denotes the Haar probability measure on the group $H_{\#}(\mathbb{S})$. For a 1-form $\epsilon \in \Omega^1(\mathbb{R})$, we define the *twisted partition function*

$$Z_{\#}(\epsilon) = \int_{H_{\#}(\mathbb{S})} \prod_{e \in E} w(J_e + \epsilon_e) dJ,$$

and note that $Z_{\#}(0) = Z_{\#}$. We also define a probability measure on *height 1-forms* $\mathbf{n} \in H_{\#}(\mathbb{Z})$ by

$$\nu_{\#}(\mathbf{n}) = \nu_{G, \#}(\mathbf{n}) \propto \exp \left(- \sum_{vv' \in E} \mathcal{V}(\mathbf{n}_{vv'}) \right). \quad (4)$$

Note that this is well defined as the normalisation constant is finite by assumption (1).

For $f, g \in \Omega^1(\mathbb{R})$ and $\epsilon, \omega \in \Omega^1(\mathbb{R})$, we will write

$$(f, g)_{\Omega^0} = \sum_{v \in V} f_v g_v, \quad \text{and} \quad (\epsilon, \omega)_{\Omega^1} = \frac{1}{2} \sum_{\vec{e} \in \vec{E}} \epsilon_{\vec{e}} \omega_{\vec{e}} = \sum_{e \in E} \epsilon_e \omega_e$$

for the standard inner products. We will usually drop the subscripts and simply write (\cdot, \cdot) in case the space is clear from the context.

The central result that we will use is the following duality formula. Even though it is classical (see e.g. Appendix A in [71]), we will provide its derivation in Appendix B.

Lemma 7.2 (Fourier–Pontryagin duality). *Let $\# \in \{\diamond, \star\}$ and let $-\#$ denote the other element of $\{\diamond, \star\}$. Then for any $\epsilon \in \Omega^1(\mathbb{R})$, we have*

$$\nu_{-\#}[\exp(i(\mathbf{n}, \epsilon))] = \frac{Z_{\#}(\epsilon)}{Z_{\#}} = \mu_{\#} \left[\prod_{e \in E} \frac{w(J_e + \epsilon_e)}{w(J_e)} \right].$$

Clearly there are two intertwined random objects in the statement of Lemma 7.2: the height and spin 1-forms \mathbf{n} and J respectively. We will mostly apply the duality to analyse one of these two models whose values are the *exact* 1-forms $H_\star(\mathbb{G})$, since then for each $\omega \in H_\star(\mathbb{G})$, there exists a unique $\tau \in \mathbb{G}^V$ such that

$$d\tau = \omega \quad \text{and} \quad \tau_\partial = 0,$$

where $\partial \in V$ is a fixed *boundary* vertex of G , and 0 is the identity element of \mathbb{G} . The random configuration τ is then distributed as a classical spin system with spins assigned to vertices with 0 boundary conditions at ∂ , and that interact through edges.

Remark 7.1. In two dimensions there is a special form of duality where H_\diamond on the planar graph G can be seen as H_\star on the *planar dual* graph G^* by simply rotating all directed edges by $\pi/2$ to the left. Therefore if ω is a 1-form such that $d^*\omega = 0$, there exists a function τ on the vertices of the dual graph G^* (faces of G) which has ω as its gradient, i.e.

$$\omega_{vv'} = \tau_u - \tau_{u'} = d\tau_{uu'},$$

where u, u' are the two faces adjacent to vv' from the right and left respectively. In this case, both models in Lemma 7.2 can be seen as classical spin and height function models.

Remark 7.2. As mentioned in the introduction, the Fourier–Pontryagin duality is usually applied in the opposite direction to Lemma 7.2, i.e., to compute the characteristic function of the spin model rather than the height function. On the height function side this results in expectations of nonlocal observables (disorders) which are in general difficult to analyse. In our case however the disorder appears on the spin model side, and can be removed from the picture by taking derivatives at zero of the characteristic function. This is the main point of view which allows to obtain most of the results in this article using comparatively elementary arguments.

One of the main tools in this article is the following identity. Even though it is a rather direct consequence of duality, we were unable to find this formulation in the literature.

Lemma 7.3 (Covariance duality). *Let $\# \in \{\diamond, \star\}$ and let $-\#$ be the other element of $\{\diamond, \star\}$. For any $\epsilon, \omega \in \Omega^1(\mathbb{R})$, we have*

$$\nu_\#[(\mathbf{n}, \epsilon)(\mathbf{n}, \omega)] + \mu_{-\#}[(\mathcal{U}'(J), \epsilon)(\mathcal{U}'(J), \omega)] = \sum_{e \in E} \mu_{-\#}[\mathcal{U}''(J_e)] \epsilon_e \omega_e.$$

Proof. It is enough to compute

$$\left. \frac{\partial}{\partial s} \frac{\partial}{\partial t} \left(\nu_\#[\exp(i(h, s\epsilon + t\omega))] \right) \right|_{s=t=0}$$

by differentiating under the sign of integration on the right-hand side of the formula from Lemma 7.2. \square

Remark 7.3. In the case of measures ν_\star on true height functions h , the quantity

$$\nu_\star[(\mathbf{n}, \epsilon)(\mathbf{n}, \omega)] = \nu_\star[(h, d^*\epsilon)(h, d^*\omega)]$$

explicitly depends only on $d^*\epsilon$ and $d^*\omega$, and hence the rest of the equation above does so implicitly.

Choosing $\epsilon = \mathbf{1}_{xy} - \mathbf{1}_{yx}$ for some edge xy and $\tau = \mathbf{1}_{uv} - \mathbf{1}_{vu}$ for another edge uv , as a corollary we immediately get the following identities:

$$\nu_\#[\mathbf{n}_{xy}^2] = \mu_{-\#}[\mathcal{U}''(J_{xy}) - \mathcal{U}'(J_{xy})^2] \quad (5)$$

and

$$\nu_\#[\mathbf{n}_{xy}\mathbf{n}_{uv}] = -\mu_{-\#}[\mathcal{U}'(J_{xy})\mathcal{U}'(J_{uv})]. \quad (6)$$

Remark 7.4. Note that Lemma 7.3 implies that the sum of the covariance matrices of two mutually dual edge fields is diagonal, i.e., equals the covariance matrix of (possibly inhomogeneous) white noise. This was known for the discrete Gaussian Free Field ($\mathbb{G} = \mathbb{R}$), see e.g. [52] and Remark 7.5, in which case the independent sum of the mutually dual edge fields is a collection of independent normal random variables.

Having established the covariance duality formula in Lemma 7.3, we will now discuss several of its rather direct consequences. Unless stated otherwise, we study the models on a finite graph $G = (V, E)$ with a prescribed boundary vertex $\partial \in V$. We will write $\Omega_0^0(\mathbb{G})$ for the set of functions $f \in \Omega^0(\mathbb{G})$ with $f_\partial = 0$.

7.3. Upper bound on the variance of the height function

In this section we consider random exact 1-forms $\mathbf{n} \in H_\star(\mathbb{Z})$ distributed according to ν_\star . As mentioned before, for each such 1-form $\mathbf{n} \in H_\star(\mathbb{Z})$, there exists exactly one *height function* $h \in \Omega_0^0(\mathbb{Z})$ such that $dh = \mathbf{n}$. Note that in the case of $\mathbb{G} = \mathbb{R}$, d and d^* are adjoint as linear operators, i.e., for all $f \in \Omega^0(\mathbb{R})$ and $\omega \in \Omega^1(\mathbb{R})$, we have

$$(f, d^*\omega)_{\Omega^0} = (df, \omega)_{\Omega^1}. \quad (7)$$

Also note that the operator

$$\Delta := d^*d : \Omega^0(\mathbb{R}) \rightarrow \Omega^0(\mathbb{R})$$

is the *graph Laplacian* on G , and it has a well defined inverse Δ^{-1} on $\Omega_0^0(\mathbb{R})$. Moreover, as matrices,

$$\Delta^{-1} = \mathbf{G}D^{-1},$$

where $D = \text{Diag}(\deg(v))_{v \in V \setminus \{\partial\}}$ and \mathbf{G} is the Green's function of simple random walk on G killed upon hitting ∂ .

Let $f \in \Omega_0^0(\mathbb{R})$ and $\epsilon := d\Delta^{-1}f$ so that $d^*\epsilon = f$. Discarding the explicitly nonnegative term $\mu_\diamond[(\mathcal{U}'(J), \epsilon)^2]$ in Lemma 7.3 applied to $\epsilon = \omega$, and using (7) we get

$$\nu_\star[(h, f)_{\Omega^0}^2] = \nu_\star[(\mathbf{n}, \epsilon)_{\Omega^1}^2] \leq \sum_{e \in E} \epsilon_e^2 |\mu_\diamond[\mathcal{U}''(J_e)]| \leq C(\epsilon, \epsilon)_{\Omega^1}, \quad (8)$$

where

$$C = \sup_{e \in E} |\mu_\diamond[\mathcal{U}''(J_e)]| \leq \sup_{J \in \mathbb{S}} |\mathcal{U}''(J)| < \infty.$$

On the other hand, by (7) again $(\epsilon, \epsilon)_{\Omega^1} = (d\Delta^{-1}f, d\Delta^{-1}f)_{\Omega^1} = (\Delta^{-1}f, f)_{\Omega^0}$.

Corollary 7.4 (GFF upper bound on variance). *For any $f \in \Omega_0^0(\mathbb{R})$,*

$$\nu_\star[(h, f)^2] \leq C(\Delta^{-1}f, f)_{\Omega^0} = C \sum_{v, v' \in V} f_v f_{v'} \frac{\mathbf{G}(v, v')}{\deg(v')},$$

where C is as above.

Remark 7.5. One can also apply duality to the discrete Gaussian free field (GFF) (in this case both the primal and dual fields are real-valued as \mathbb{R} is self-dual as a locally compact abelian group). The GFF is defined similarly to the integer-valued GFF with potential $\mathcal{V}(t) = t^2$ with the difference that the reference measure in (4) is the Lebesgue measure on \mathbb{R} and not the counting measure on \mathbb{Z} . The model is self dual in that $\mathcal{U}(t) = \mathcal{V}(t) = t^2$, and in the analog of the corollary above actually get an equality since $(\epsilon, \mathcal{U}'(J_e)) = 0$ since $\mathcal{U}'(J_e) = 2J_e \in H_\diamond$, and $d^*\epsilon \in H_\star$ by definition. This agrees with the fact that the covariance of the GFF is given *exactly* by the inverse Laplacian.

Remark 7.6. Consider an infinite countable graph $\Gamma = (\mathcal{V}, \mathcal{E})$ and a sequence of increasing finite subgraphs exhausting Γ , i.e., $G_N \nearrow \Gamma$ as $N \rightarrow \infty$. If $f : \mathcal{V} \rightarrow \mathbb{R}$ has bounded support and mean zero, i.e., $\sum_{v \in \mathcal{V}} f_v = 0$, where this sum is actually taken over a finite set of vertices, then we can find a 1-form ϵ on \mathcal{E} with *bounded* support such that $d^*\epsilon = f$, and hence

$$\prod_{e \in E} \frac{w(J_e + \epsilon_e)}{w(J_e)} = \prod_{e \in \text{supp}(\epsilon)} \frac{w(J_e + \epsilon_e)}{w(J_e)} \quad (9)$$

is a local bounded continuous function of J (in the product topology on $\mathbb{S}^\mathcal{E}$) whenever \mathcal{V} and w are as in Definition 7.1. Moreover, since \mathbb{S} is compact metrizable so is $\mathbb{S}^\mathcal{E}$ with the product topology by Tychonoff's theorem, and hence the edge spin models $\mu_{G_N, \#}$ always form a tight sequence of measures on $\mathbb{S}^\mathcal{E}$ as $N \rightarrow \infty$. This in particular implies that $\mu_{G_N, \diamond}$ converges weakly along a subsequence. Therefore Lemma 7.2 together with (9) and the fact that

$$\nu_{G_N, \star}[\exp(i(\mathbf{n}, \epsilon))] = \nu_{G_N, \star}[\exp(i(f, h))]$$

for N large enough so that G_N contains $\text{supp}(\epsilon)$, imply that the random height 1-forms \mathbf{n} under $\nu_{G_N, \star}$, and hence also the *differences* of the associated height function h , converge weakly along the same subsequence.

One has to be careful as this is in general no longer true if f does not have zero mean, e.g., $f = \delta_v$. Then ϵ with $d^*\epsilon = f$ cannot be taken with bounded support (there always has to be an infinite path with nonzero values of ϵ). In this case tightness may fail when *delocalization* of the height function arises, i.e., $\nu_{G_N, \star}[(h, f)^2] = \nu_{G_N, \star}[h_v^2] \rightarrow \infty$ as $N \rightarrow \infty$ (e.g. if Γ is planar, see Section 7.4).

We also immediately deduce that delocalization of the height function does not happen on transient graphs for potentials as in Definition 7.1. We note that our result, despite its simple proof, seems new in this generality, and that such behavior is expected for a larger class of potentials. We also note that the special case of the integer valued GFF follows from a stronger estimate proved by Fröhlich and Park [73] (see also [104]).

To state the result, we briefly recall the notion of Gibbs measures and gradient Gibbs measures (we do it for height functions only, and the definition for spin models used later in the chapter is completely analogous). From now on we assume that $\Gamma = (\mathcal{V}, \mathcal{E})$ is a locally finite, infinite graph. For a finite set $\Lambda \subset \mathcal{V}$ write $E(\Lambda)$ for the set of edges with at least one vertex in Λ . Let $\varphi : \Lambda^c \rightarrow \mathbb{Z}$ be a function and define the probability measure μ_Λ^φ supported on $h : \mathcal{V} \rightarrow \mathbb{Z}$ satisfying $h|_{\Lambda^c} = \varphi$ by

$$\nu_\Lambda^\varphi(h) \propto \exp\left(-\sum_{e \in E(\Lambda)} \mathcal{V}(dh_e)\right).$$

In other words, ν_Λ^φ is the measure ν_\star from (4) with φ -boundary conditions outside Λ . A probability measure ν supported on height functions $h : \mathcal{V} \rightarrow \mathbb{Z}$ is called a *Gibbs measure* (on Γ with respect to the potential \mathcal{V}) if it satisfies the Dobrushin–Lanford–Ruelle (DLR) relations: for all finite sets $\Lambda \subset \mathcal{V}$,

$$\nu_\Lambda(\cdot) = \int_{\mathbb{Z}^{\mathcal{V}}} \nu_\Lambda^\varphi(\cdot) d\nu(\varphi),$$

where ν_Λ denotes the restriction of ν to Λ . If Γ is a Cayley graph and the measure ν is invariant under shifts, it is called translation invariant. In terms of Gibbs measures, delocalization corresponds to non-existence of translation invariant Gibbs measures.

A *gradient* Gibbs measure is a slight variation of the above, where we consider measures supported only on gradients. Fix this time a finite set of edges $\Lambda \subset \mathcal{E}$. Let ω be an exact 1-form (thus taking value in $H_\star(\mathcal{E}, \mathbb{Z})$). Define the probability measure μ_Λ^ω supported on 1-forms $\mathbf{n} \in H_\star(\mathcal{E}, \mathbb{Z})$ satisfying $h|_{\Lambda^c} = \omega|_{\Lambda^c}$ as

$$\mu_\Lambda^\omega(\mathbf{n}) \propto \exp\left(-\sum_{e \in \Lambda} \mathcal{V}(\mathbf{n}_e)\right).$$

A probability measure supported on 1-forms $\mathbf{n} \in H_\star(\mathcal{E}, \mathbb{Z})$ will be called a gradient Gibbs measure if it satisfies the analog of the DLR equation above in this setup.

Theorem 7.5. *Let $\Gamma = (\mathcal{V}, \mathcal{E})$ be a transient graph and \mathcal{V} a height function potential as in Definition 7.1. Then, there exists an infinite volume Gibbs measure on Γ with respect to \mathcal{V} . If Γ is moreover an amenable Cayley graph, there exist translation invariant Gibbs measures.*

Proof. Let now $G_N \nearrow \Gamma$, as $N \rightarrow \infty$ be an exhaustion of Γ by finite subgraphs G_N . Define the boundary $\partial_N := \partial G_N$ to be the set of vertices in G_N adjacent to a vertex from outside of G_N . Let $\nu_{G_N, \star}[\cdot]$ be the expectation associated with the height function on G_N with 0-boundary conditions. Fix any $\Lambda \subset \mathcal{V}$ finite and let $f : \Lambda \rightarrow \mathbb{R}$ be any function. It follows from Corollary 7.4 that

$$\nu_{G_N, \star}[(h, f)^2] \leq C \max_{v \in \Lambda} \frac{\mathbf{G}_N(v, v)}{\deg(v)} (f, f)_\Lambda.$$

where $C < \infty$, and \mathbf{G}_N is the Green's function of simple random walk on G_N killed on hitting ∂_N . Since Γ is transient, the right-hand side is uniformly bounded in N . Therefore, the sequence $\nu_{G_N, \star}(h|_\Lambda)$ is tight and subsequential limits exist by Prokhorov's theorem. By a diagonal argument, we can extract a further sub-sequence N_K so that $\nu_{G_{N_K}, \star}(h|_\Lambda)$ converges for each Λ finite. Any such subsequential limit is a Gibbs measure as it satisfies the DLR relations. This proves the first part of the theorem.

For the second part, suppose that Γ is an amenable Cayley graph, so that $\nu_{G_N, \star}[h_v^2] \leq C'$, for some $C' < \infty$ which is independent of v and the exhaustion $(G_N)_{N \geq 1}$. Let μ be a subsequential limit (which exists by the argument above, and is a Gibbs measure). Let $o \in \mathcal{V}$. Since Γ is amenable, there is some Følner sequence (also called Van Hove sequence) $(F_N)_{N \geq 1}$ of sets of vertices containing o . This means $F_N \nearrow \mathcal{V}$ and $|\partial F_N|/|F_N| \rightarrow 0$ as $N \rightarrow \infty$. Let

$$\nu_N := \frac{1}{|F_N|} \sum_{x \in F_N} \mu \circ \theta_x,$$

where θ_x is the shift towards x (since Γ is a Cayley graph, this is the same as left multiplication in the group). This is a Gibbs measure because the set of Gibbs measures is closed under translations and convex combinations. Moreover, $\nu_N[h_v^2] \leq C'$ for each N and $v \in \mathcal{V}$. Therefore, $(\nu_N)_{N \geq 1}$ is tight. Let ν be any subsequential limit, which is again a Gibbs measure. By construction and since $|\partial F_N|/|F_N| \rightarrow 0$, we have $\nu \circ \theta_x = \nu$ for each x , and hence ν is translation invariant. \square

7.4. Delocalization implies the BKT phase transition

7.4.1 Background

In this section we consider the spin and height function models on the square lattice \mathbb{Z}^2 and we show that delocalization of the height function (defined below) is equivalent to

the divergence of a certain series of two-point functions in the dual spin model. One of the conclusions is that delocalization implies the Berezinskii–Kosterlitz–Thouless (BKT) phase transition in the dual spin model [36, 108].

This implication for the classical XY and the Villain spin models, together with a proof of delocalization of the associated height functions, was first obtained by Fröhlich and Spencer in their seminal work establishing the BKT transition [71] (also see [104] for an exhaustive account). Recently alternative proofs were provided by Aizenman et al. [8] (first for the Villain and later also for the XY model) and by the authors [64] for the XY model. Together with the new conceptual approach to delocalization introduced by Lammers [114], these works improve our mathematical understanding of the BKT transition. These results can be thought of as an inequality between the critical points of the mutually dual spin and height function models. A natural conjecture is that these critical points always coincide. In the case of the XY and Villain model this was confirmed in a recent work Lammers [116].

In this section we provide yet another, and arguably the simplest so far, proof of the fact that delocalization of the height function implies that correlations functions of certain observables in the spin model do not decay exponentially fast in the distance. For reflection positive models (which is the case when $-\mathcal{U}$ is itself positive definite, i.e., has nonnegative coefficients in the Fourier series), we moreover obtain an equivalence between delocalization and nonsummability of spin correlations. Our approach, unlike the previous ones, is based solely on duality, and does not invoke any additional (e.g. graphical) representations of the models at hand.

7.4.2 Notions of delocalization

It is now well established that integer-valued height functions on \mathbb{Z}^2 (or in general on periodic planar lattices) undergo a phase transition between a *localized (smooth)* and a *delocalized (rough)* regime [45, 57, 58, 71, 114–117, 128]. We say that a potential \mathcal{V} is localized (on \mathbb{Z}^2) if it admits a translation-invariant Gibbs measure on height functions $h : \mathbb{Z}^2 \rightarrow \mathbb{Z}$. Otherwise we say that \mathcal{V} delocalizes. It is known that if \mathcal{V} is convex on the integers, and moreover its second discrete derivative is nonincreasing, i.e., \mathcal{V} is a so-called *supergaussian* potential [115, 117], then delocalization in this sense is equivalent to the fact that

$$\sup_{N \geq 1} \nu_{\Lambda_N, \star}[h_{\mathbf{0}}^2] = \infty, \tag{10}$$

where $\mathbf{0}$ is the origin of \mathbb{Z}^2 , and $\Lambda_N = [-N, N]^2 \cap \mathbb{Z}^2$ where we identify all vertices in Λ_N that are adjacent to $\Lambda_N^c := \mathbb{Z}^2 \setminus \Lambda_N$ as one boundary vertex ∂ (wired boundary) and set $h_{\partial} = 0$. Moreover for such potentials, the sequence in (10) is nondecreasing in N [117], and it grows up to a multiplicative constant at least like $\log N$ [115] (which is consistent

with the general conjecture stating that delocalized height functions should behave like the GFF at large scales).

Yet another approach to delocalization is to work with infinite volume *gradient* measures and study the variance of the increment of the height between two distant points. This was e.g. studied in [126, 128] in the context of the six-vertex model and it will be convenient for us to follow the same route here, as we already know by Remark 7.6 that translation invariant gradient Gibbs measures always exist for potentials as in Definition 7.1. We say that a potential \mathcal{V} is ∇ -delocalized (on \mathbb{Z}^2) if for any translation-invariant gradient Gibbs measure ν (with expectation ν), we have

$$\sup_{v \in \mathbb{Z}^2} \nu[(h_v - h_{\mathbf{0}})^2] = \infty. \tag{11}$$

Lemma 7.6. *If a potential \mathcal{V} is delocalized, then it is also ∇ -delocalized.*

Proof. Suppose otherwise that there exists a translation-invariant gradient Gibbs measure ν for which the supremum in (11) is finite. Then, as in Theorem 7.5, by considering convex combinations of translations of ν thought of as a measure on height functions \tilde{h} given by $\tilde{h}_v = h_v - h_{\mathbf{0}}$ we can construct a translation invariant Gibbs measure on height functions which is a contradiction. We leave the details to the reader. \square

We note that the opposite implication is also true e.g. for potentials \mathcal{V} that are convex on the integers. Indeed, in this case it is known from the foundational work of Sheffield [157] that each Gibbs measure for height functions has a finite second moment (since the height at every point has a log-concave distribution).

7.4.3 Setup

Let us fix mutually dual potentials \mathcal{V} and \mathcal{U} as in Definition 7.1. It will be convenient to consider the spin and height function models on finite, exponentially growing tori $\mathbb{T}_N = (\mathbb{Z}/2^N\mathbb{Z})^2$. This way we achieve three properties by construction:

- we work with measures that are translation invariant and invariant under $\pi/2$ -rotations,
- we can apply the duality of Lemma 7.2 first in the finite volume \mathbb{T}_N , and then take simultaneous (subsequential) infinite-volume limits, $\mathbb{T}_N \rightarrow \mathbb{Z}^2$ as $N \rightarrow \infty$, on both sides of the duality relation,
- we get an explicit monotonicity in N for the Green's function of the random walk on \mathbb{T}_N (see below).

Let $\mu = \mu_{\mathbb{Z}^2, \diamond}$ be any subsequential limit of $\mu_{\mathbb{T}_N, \diamond}$, and let $\nu = \nu_{\mathbb{Z}^2, \star}$ denote the limit of $\nu_{\mathbb{T}_N, \star}$ taken along the same subsequence (it exists by Remark 7.6). One can think of

ν as a probability measure on height functions $h : \mathbb{Z}^2 \rightarrow \mathbb{Z}$ satisfying $h(\mathbf{0}) = 0$. By weak convergence, the duality of Lemma 7.2 holds also for μ and ν whenever $\epsilon \in \Omega^1(\mathbb{Z}^2, \mathbb{R})$ is of bounded support. The same is true for Corollary 7.3 and Corollary 7.4, where we choose $\partial = \mathbf{0}$ and consider the Green's function of a random walk on \mathbb{Z}^2 killed at $\mathbf{0}$.

Note that by planar duality, we have $H_\diamond(\mathbb{Z}^2, \mathbb{S}) \cong H_\star((\mathbb{Z}^2)^*, \mathbb{S})$, where $(\mathbb{Z}^2)^* \cong \mathbf{0}^* + \mathbb{Z}^2$ with $\mathbf{0}^* := (1/2, 1/2)$, is the *dual* square lattice. Since $H_\star((\mathbb{Z}^2)^*, \mathbb{S}) \cong \mathbb{S}^{(\mathbb{Z}^2)^* \setminus \{\mathbf{0}^*\}}$, we can think of μ as a Gibbs measure on spin configurations θ on $(\mathbb{Z}^2)^*$ where the spin at $\mathbf{0}^*$ is fixed to be the identity element of \mathbb{S} .

Finally, let $v_n = (n, 0) \in \mathbb{Z}^2$, and let $p_n = (e_0, e_1, \dots, e_{n-1})$ be the directed horizontal path from v_0 to v_n . We identify p_n with the 1-form that assigns 1 to each directed edge in p_n , and 0 to the directed edges of \mathbb{Z}^2 that are not in p_n . For compactness of notation, we write $J_i = J_{e_i}$ and $h_i = h_{u_i}$.

7.4.4 The implication

Applying Lemma 7.3 and Corollary 7.4 in finite volume, and then taking the subsequential limit as in Section 7.4.3, we have

$$\begin{aligned} 0 &\leq \sum_{i=0}^{n-1} \mu[\mathcal{U}''(J_i)] - \mu\left[\left(\sum_{i=0}^{n-1} \mathcal{U}'(J_i)\right)^2\right] \\ &\leq \limsup_{N \rightarrow \infty} \nu_{\mathbb{T}_N, \star}[(h_0 - h_n)^2] \leq \limsup_{N \rightarrow \infty} \mathbf{G}_N(v_n, v_n), \end{aligned} \quad (12)$$

where \mathbf{G}_N is the Green's function of simple random walk on \mathbb{T}_N killed at $\mathbf{0}$. This is equivalent to a random walk on \mathbb{Z}^2 killed at all points in $2^N \mathbb{Z}^2$. Hence, $\mathbf{G}_N \nearrow \mathbf{G}$ as $N \rightarrow \infty$, where \mathbf{G} is the Green's function of a random walk on \mathbb{Z}^2 killed at $\mathbf{0}$. Classically we have $\mathbf{G}(v_n, v_n) \leq \text{const} \times \log n$ (see e.g. [131]). Plugging this bound into (12), dividing both sides by n , letting $n \rightarrow \infty$, and finally using translation invariance of μ , we get

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^{n-1} u_k = \frac{1}{2} \mu[\mathcal{U}''(J_0) - \mathcal{U}'(J_0)^2], \quad \text{where } u_k = \sum_{i=1}^k \mu[\mathcal{U}'(J_0)\mathcal{U}'(J_i)]. \quad (13)$$

In particular u_k converges in the Cesàro sense as $k \rightarrow \infty$.

Theorem 7.7 (Delocalization implies the BKT phase transition). *Consider the setup from Section 7.4.3. If the height function delocalizes in the sense that (11) holds true for ν , then*

$$\sum_{i=1}^{\infty} i |\mu[\mathcal{U}'(J_0)\mathcal{U}'(J_i)]| = \infty.$$

In particular, there is no exponential decay of the two-point function $\mu[\mathcal{U}'(J_0)\mathcal{U}'(J_i)]$ as $i \rightarrow \infty$.

Proof. We can assume that $\sum_{i=1}^{\infty} |\mu[\mathcal{U}'(J_0)\mathcal{U}'(J_i)]| < \infty$ since otherwise we are done. This means that u_k converges in the classical sense to its Cesàro limit from (13). Hence,

$$\mu[\mathcal{U}''(J_0) - \mathcal{U}'(J_0)^2] = \lim_{k \rightarrow \infty} 2u_k = 2 \sum_{i=1}^{\infty} \mu[\mathcal{U}'(J_0)\mathcal{U}'(J_i)]. \quad (14)$$

By Lemma 7.3 applied in the infinite volume (p_n has bounded support) and translation invariance of μ , we have

$$\begin{aligned} \nu[(h_n - h_0)^2] &= \sum_{i=0}^{n-1} \mu[\mathcal{U}''(J_i) - \mathcal{U}'(J_i)^2] - 2 \sum_{i=1}^{n-1} (n-i) \mu[\mathcal{U}'(J_0)\mathcal{U}'(J_i)] \\ &= 2n \sum_{i=1}^{\infty} \mu[\mathcal{U}'(J_0)\mathcal{U}'(J_i)] - 2 \sum_{i=1}^{n-1} (n-i) \mu[\mathcal{U}'(J_0)\mathcal{U}'(J_i)] \\ &= 2 \sum_{i=1}^{n-1} i \mu[\mathcal{U}'(J_0)\mathcal{U}'(J_i)] + 2n \sum_{i=n}^{\infty} \mu[\mathcal{U}'(J_0)\mathcal{U}'(J_i)] \\ &\leq 2 \sum_{i=1}^{\infty} i |\mu[\mathcal{U}'(J_0)\mathcal{U}'(J_i)]|. \end{aligned}$$

By the assumption, and translation and $\pi/2$ -rotation invariance of ν , we have

$$\infty = \sup_{v \in \mathbb{Z}^2} \nu[(h_v - h_0)^2] \leq 2 \sup_{n \geq 1} \nu[(h_n - h_0)^2],$$

which together with the inequality above finishes the proof. \square

It is classical that spin correlation functions decay exponentially fast at high temperatures (here the temperature is incorporated in the definition of \mathcal{U}). This in particular implies that $\sum_{i=1}^{\infty} i |\mu[\mathcal{U}'(J_0)\mathcal{U}'(J_i)]| < \infty$. From this point of view Theorem 7.7 says that if the height function delocalises, then the associated spin model undergoes a BKT phase transition from a regime with exponential decay to a regime with slow decay of correlations.

7.4.5 The case of the XY-model

The change of behavior of the two-point functions $\mu[\mathcal{U}'(J_0)\mathcal{U}'(J_i)]$ as $i \rightarrow \infty$ clearly indicates a phase transition in the spin model. However it is more common to look at correlations of the type $\mu[\mathcal{F}(\theta_u - \theta_{u'})]$ when u and u' are far apart, where θ is the underlying spin field on $(\mathbb{Z}^2)^*$ (the faces of \mathbb{Z}^2), and where \mathcal{F} is some chosen function, e.g. $\mathcal{F} = \mathcal{U}$.

For general spin models, it is not clear how to compare these two types of correlation functions. Here we present an approach based on correlation inequalities in the case of the classical XY model, i.e., when $\mathcal{U}(t) = -\beta \cos(t)$, where $\beta > 0$ is the inverse temperature in the spin model.

To this end, consider the setup as in Theorem 7.7. If $\{u_i, u'_i\}$ is the dual edge of e_i , writing $\theta_i = \theta_{u_i}$ and $\theta'_i = \theta'_{u_i}$, we have

$$\begin{aligned} \frac{2}{\beta^2} \mu[\mathcal{U}'(e_0)\mathcal{U}'(e_i)] &= 2\mu[\sin(\theta_0 - \theta'_0) \sin(\theta_i - \theta'_i)] \\ &= \mu[\cos(\theta_0 - \theta'_0 - \theta_i + \theta'_i)] - \mu[\cos(\theta_0 - \theta'_0 + \theta_i - \theta'_i)] \end{aligned} \quad (15)$$

A version of the classical Ginibre inequality for the XY model [80] (see also [25]) says that

$$\mu[\sin(\theta_0) \cos(\theta'_0) \sin(\theta_i) \cos(\theta'_i)] \leq \mu[\sin(\theta_0) \sin(\theta_i)] \mu[\cos(\theta'_0) \cos(\theta'_i)],$$

which after expanding into cosines of sums of angles and disregarding terms that are not invariant under global rotation (shift of angles mod 2π) whose expectations vanish, we obtain

$$\begin{aligned} \mu[\cos(\theta_0 + \theta'_0 - \theta_i - \theta'_i)] + \mu[\cos(\theta_0 - \theta'_0 - \theta_i + \theta'_i)] - \mu[\cos(\theta_0 - \theta'_0 + \theta_i - \theta'_i)] \\ \leq 2\mu[\cos(\theta_0 - \theta_i)] \mu[\cos(\theta'_0 - \theta'_i)] \\ = 2\mu[\cos(\theta_0 - \theta_i)]^2, \end{aligned}$$

where the last identity follows by reflection invariance of μ across the real line. Analogous inequality follows by exchanging the roles of θ_i and θ'_i , which results in swapping the signs of the second and third term in the first line. Using that the first term is positive by the first Griffiths inequality, and combining with (15), we get that

$$\begin{aligned} \frac{1}{2} \mu[\cos(\theta_0 + \theta'_0 - \theta_i - \theta'_i)] + |\mu[\mathcal{U}'(e_0)\mathcal{U}'(e_i)]| &\leq \max\{\mu[\cos(\theta_0 - \theta_i)]^2, \mu[\cos(\theta_0 - \theta'_i)]^2\} \\ &= \mu[\cos(\theta_0 - \theta_i)]^2, \end{aligned} \quad (16)$$

where the last identity follows from the Messager–Miracle–Sole inequality [135] by applying the reflection across the real line.

These considerations, together with Lemma 7.6, lead us to the following corollary that recovers the main result of [64].

Corollary 7.8. *If the height function associated with the XY model on \mathbb{Z}^2 delocalizes, then*

$$\sum_{i=1}^{\infty} i \mu[\cos(\theta_0 - \theta_i)]^2 \geq \sum_{i=1}^{\infty} i (\mu[\cos(\theta_0 - \theta_i)]^2 - \frac{1}{2} \mu[\cos(\theta_0 + \theta'_0 - \theta_i - \theta'_i)]) = \infty.$$

Remark 7.7. For the XY model it is known that there exists only one translation-invariant Gibbs measure μ on \mathbb{Z}^2 [145], and hence regular, instead of subsequential, limits may be taken Section 7.4.3.

7.4.6 An equivalence

When $-\mathcal{U}$ is itself positive definite, i.e., all its Fourier coefficients are nonnegative, we can actually conclude more than in the above discussion. Indeed, in this case μ is reflection positive (see Appendix B.2). This implies that

$$\mu[\mathcal{U}'(J_0)\mathcal{U}'(J_i)] \geq 0$$

as this holds true on \mathbb{T}_N for every $N \geq i$ by reflection positivity. Therefore the Cesàro convergence from (13) implies classical convergence, and (14) always holds true. The same argument as in the proof of Theorem 7.7 yields the following corollary.

Corollary 7.9. *Consider the setup from Section 7.4.3, and moreover assume that $-\mathcal{U}$ is positive definite. Then (11) holds true for ν if and only if*

$$\sum_{i=1}^{\infty} i\mu[\mathcal{U}'(J_0)\mathcal{U}'(J_i)] = \infty. \quad (17)$$

Remark 7.8. The identity from (14) can be rewritten in a more symmetric form as

$$\sum_{i \in \mathbb{Z}} \mu[\mathcal{U}'(J_0)\mathcal{U}'(J_i)] = \mu[\mathcal{U}''(J_0)], \quad (18)$$

where now the sum is over a bi-infinite path of edges. Curiously, this is an exact (but nonlocal) identity for correlation functions that is independent of the (hidden in the definition of \mathcal{U}) temperature parameter. In particular the series in (18) is always convergent, independently of the temperature. This is in contrast with the behavior of the series in (17) that does undergo a phase transition. This, together with the relation to the gradient of the height function (6), is consistent with the conjecture that in delocalized regime the discrete GFF describes the large-scale fluctuations of the height function. Indeed, the two-point function of the gradient of the discrete GFF is known to decay like the inverse square of the distance (see e.g. [22]).

7.5. Monotonicity of variance of the height function

In this section we will show that the variance of the height function is monotone in a natural temperature parameter under some further assumptions on the potential. To the best of our knowledge the result is new, even in the case of planar graphs. Together with the dichotomy of Lammers [115], this directly implies that the height function of the XY model undergoes a sharp phase transition on the square lattice.

7.5.1 Setup

Let $G = (V, E)$ be a finite graph. To each edge $e \in E$, associate

- a twice continuously differentiable spin potential $\mathcal{U}_e : \mathbb{S} \rightarrow \mathbb{R}$ such that $-\mathcal{U}_e$ is positive definite,
- a non-negative real β_e (thought of as the inverse temperature in the spin model),
- the dual potential $\mathcal{V}_{\beta_e} := \mathcal{V}_{\beta_e, e}$ of $\beta_e \mathcal{U}_e$ as in Definition 7.1.

We will consider the family of measures $\nu_{\beta, \star}$ for height functions and their dual measures $\mu_{\beta, \diamond}$, indexed by $\beta = (\beta_e)_{e \in E}$.

We wish to point out that the above requirements on the spin potential \mathcal{U} can also be described purely in terms of the height function potential: if \mathcal{V} satisfies the conditions of Definition 7.1 and $e^{-\mathcal{V}}$ is infinitely divisible (in the sense that each division satisfies Definition 7.1), then the corresponding spin potential satisfies the above conditions, see Appendix B.3. This equivalence is not important in the remainder of this section.

7.5.2 Increasing variance

The main result of this section is the following fact.

Theorem 7.10. *Consider the setup as in Section 7.5.1. For each $x, y \in V$ and $e \in E$ the function*

$$\beta_e \mapsto \nu_{\beta, \star}[(h_x - h_y)^2]$$

is non-decreasing.

To prove the theorem, let us begin by slightly extending Ginibre's inequalities [80] to spin models on $H_\diamond(\mathbb{S})$ (the original inequality deals with $H_\star(\mathbb{S})$).

Lemma 7.11 (Ginibre). *Consider the setup as in Section 7.5.1. For all positive definite functions $F : \mathbb{S} \rightarrow \mathbb{R}$ and all $e, f \in E$, we have*

$$\frac{\partial}{\partial \beta_e} \mu_{\beta, \diamond}[F(J_f)] \geq 0.$$

Proof. This is proved in Appendix B.1. □

Proof of Theorem 7.10. We first add an additional edge g connecting x and y (even if there was already such an edge present). On this edge, we put the potential $-\mathcal{U}_g(t) = \cos(t)$ and parameter $\beta_g = \lambda \geq 0$. Thus, we remain in the setup of Section 7.5.1. We write $\mu_{\beta, \lambda, \diamond}$ and $\nu_{\beta, \lambda, \star}$ for the corresponding spin and height-function measure respectively, and note that $\mu_{\beta, \lambda, \diamond} \rightarrow \mu_{\beta, \diamond}$ as $\lambda \rightarrow \infty$.

Let ϵ be any 1-form vanishing on g and such that $d^* \epsilon = \delta_x - \delta_y$, and let ϵ' be the 1-form vanishing outside of g and such that $d^* \epsilon' = \delta_x - \delta_y$. By Lemma 7.3 applied first to ϵ' and then to ϵ , we have that

$$\nu_{\beta, \lambda, \star}[(h_x - h_y)^2] = \beta_g \mu_{\beta, \lambda, \diamond}[\cos(J_g)] + \beta_g^2 \mu_{\beta, \lambda, \diamond}[\cos(J_g)^2] - \beta_g^2$$

$$= \sum_{e \in E} (\mu_{\beta, \lambda, \diamond}[\mathcal{U}_e''(J_e)]\epsilon_e^2 - \mu_{\beta, \lambda, \diamond}[(\mathcal{U}_e'(J), \epsilon)^2]). \quad (19)$$

Since $2 \cos^2 t = 1 + \cos 2t$ is positive definite we can apply Lemma 7.11 to the first line above and conclude that (19) is nondecreasing in β_e for any $e \neq g$. By weak convergence, the same holds for

$$\sum_{e \in E} (\mu_{\beta, \diamond}[\mathcal{U}_e''(J_e)]\epsilon_e^2 - \mu_{\beta, \diamond}[(\mathcal{U}_e'(J), \epsilon)^2]) = \nu_{\beta, \star}[(h_x - h_y)^2],$$

where the last equality again follows from Lemma 7.3. This ends the proof. \square

7.5.3 Delocalization of roughly planar height function models

In this section, we will use Theorem 7.10 to deduce that on many planar graphs, the height function delocalizes. Consider here an infinite lattice $\Gamma = (\mathcal{V}, \mathcal{E})$ embedded in the plane, but not necessarily planar. We will assume throughout that Γ (under this embedding) invariant under a bi-periodic lattice action, and that it has finite degrees. An example of such Γ is \mathbb{Z}^2 where all vertices are connected if they are within distance $M < \infty$ from each other. Given Γ , recall that $(G_N)_{N \geq 1}$ is an exhaustion of Γ if G_N is a finite subgraph of Γ for each N , $G_N \subset G_{N+1}$ and $G_N \nearrow \Gamma$. We will also consider the *wiring* of G_N by identifying G_N^c in Γ to a single vertex ∂ and removing all the self-loops created in this process. The obtained graph will be denoted by G_N^* . On such graphs, we will take the measures $\nu_{N, \beta, \star}$ as in Section 7.5.1, and we identify the space of 1-forms \mathbf{n} in $H_\star(\mathbb{Z})$ with functions h in $\Omega_0^0(\mathbb{Z})$.

Theorem 7.12 (Delocalization). *Let Γ be as above and consider the setup as in Section 7.5.1 where we assume that \mathcal{U}_e is the same for all edges. There exists a $\beta_c < \infty$ such that for all $\beta \geq \beta_c$ and all wired exhaustions $G_N^* \nearrow \Gamma$,*

$$\nu_{N, \beta, \star}[h_\partial^2] \rightarrow \infty.$$

To prove this theorem, we rely on a beautiful result of Lammers [114]:

Theorem 7.13 (Theorem 2.7 [114]). *Let $\Gamma' = (V, E)$ be an infinite graph with degree at most three, that is invariant under some lattice action. If \mathcal{V} is a convex potential for the height function with*

$$\mathcal{V}(\pm 1) \leq \mathcal{V}(0) + \log(2),$$

then the height function delocalizes in the sense that there are no translation invariant Gibbs measures.

In general, the potentials \mathcal{V} as in the setup of Section 7.5.1 need *not* be convex. However, in some special cases they are, as we will show next. This will be crucial for what follows: in Section 7.5.3 it will be shown that we can always reduce to this case.

Lemma 7.14. *If $-\mathcal{U}(J) = \cos(iJ)$ for some $i \in \mathbb{N}$, then \mathcal{V}_β is convex over $i\mathbb{Z}$ for all β . Moreover, translation invariant Gibbs measures exist if and only if $\nu_{N,\beta,\star}[h_0^2]$ is bounded uniformly in N .*

Proof. In the case $-\mathcal{U}(J) = \cos(J)$, convexity of \mathcal{V}_β over the integers is an easy consequence of the Turán inequality, see e.g. [64]. The extension to $-\mathcal{U}(J) = \cos(iJ)$ follows from a change of variables. The second statement of the lemma was proved in the case of the XY model in [64, Theorem 4]. It follows from a standard dichotomy (see e.g. [117]) in the case where the height function satisfies the so called “absolute value FKG” property, meaning that $|h|$ is FKG, see also (10). \square

Remark 7.9. We wish to point out that the result of Lammers does *not* depend on the potential \mathcal{V} being the same on each edge, just that it satisfies the condition of Theorem 7.13 for all edges, and that the potentials are invariant under some lattice action.

We will first modify the potentials $-\mathcal{U}$ so they will fit the framework of Theorem 7.13 and Lemma 7.14. Next, we modify the graph Γ to obtain a graph Γ' to which we can apply Theorem 7.13 in such a way that Γ' embeds into Γ and the variance of the height function in Γ' is smaller.

Reduction to convex potentials

We will apply here a simplification that allows us to only consider potentials of the form $-\mathcal{U}(J) = \alpha_i \cos(iJ)$. Since $-\mathcal{U}$ is positive definite, it can be written as

$$-\mathcal{U}(J) = \alpha_0 + \sum_{i=1}^{\infty} \alpha_i \cos(iJ),$$

with $\alpha_i \geq 0$. Now let $i \geq 1$ be the first mode where $\alpha_i > 0$. Write $-\mathcal{U}' = \alpha_i \cos(iJ)$ and $\nu'_{G,\beta,\star}$ for the corresponding height function measure.

Lemma 7.15. *For any finite graph $G = (V, E)$ with boundary ∂ and any $x \in V \setminus \{\partial\}$, we have*

$$\nu'_{G,\beta,\star}[h_x^2] \leq \nu_{G,\beta,\star}[h_x^2].$$

Proof. Take $-\mathcal{U}'' = -\mathcal{U} + \mathcal{U}'$ which is positive definite. Write for any $\alpha \geq 0$

$$\mathcal{U}_\alpha(J) = \mathcal{U}'(J) + \alpha \mathcal{U}''(J),$$

so that $\mathcal{U}_1 = \mathcal{U}$, $\mathcal{U}_0 = \mathcal{U}'$ and $-\mathcal{U}_\alpha$ is positive definite for each α . Let $\nu_{G,\beta,\alpha,\star}$ be the corresponding height function measure. Theorem 7.10 implies that for any $x \in V$

$$\frac{\partial}{\partial \alpha} \nu_{G,\beta,\alpha,\star}[h_x^2] \geq 0,$$

so that the variance is minimized at $\alpha = 0$. This shows the result. \square

Graph Modifications.

Fix Γ an infinite graph and $G_N \nearrow \Gamma$ an exhaustion as above. We wish to perform two operations:

- (a) splitting edges into multiple sub-edges and
- (b) gluing vertices together,

in such a way that the variance of the height function does not increase.

Operation (a) is the easiest: to add $k-1$ “evenly spaced” vertices to an edge without changing the height function on the original graph, we wish to find a potential $\mathcal{V}_\beta^{(k)}$ such that

$$e^{-\mathcal{V}_\beta} = (e^{-\mathcal{V}_\beta^{(k)}})^{*k},$$

where by *k we mean k -fold convolution.

Using basic properties of the Fourier transform, we can take the potential $\mathcal{V}_\beta^{(k)} = \mathcal{V}_{\beta/k}$ which is dual to $-(\beta/k)\mathcal{U}$. This offers the following lemma.

Lemma 7.16 (Splitting edges). *Suppose \mathcal{V} corresponds to a spin potential \mathcal{U} , such that $-\mathcal{U}$ is positive definite. For each $k \in \mathbb{N}$, we have*

$$e^{-\mathcal{V}_\beta} = (e^{-\mathcal{V}_{\beta/k}})^{*k}.$$

Operation (b) will make use of Theorem 7.10. Let v_1, v_2 be two vertices in the graph, with or without an edge between them and add to the graph the edge $g = \{v_1, v_2\}$ with the XY potential $\mathcal{V}_\lambda(k) = -\log(I_k(\lambda))$ with parameter λ . Write $\nu_{N,\beta,\lambda,\star}$ for the corresponding height function measure on G_N^* . We will show now that gluing the vertices v_1, v_2 corresponds to sending λ to 0 in this setting. Indeed, as $\lambda \rightarrow 0$, we have

$$e^{-\mathcal{V}_\lambda(k)} = I_k(\lambda) \rightarrow \begin{cases} 1, & \text{if } k = 0 \\ 0, & \text{else} \end{cases}$$

which means that the height function measure $\nu_{N,\beta,0,\star}$ is supported on height functions with $h_{v_1} = h_{v_2}$. Moreover, Theorem 7.10 implies that for any vertex x of G_N^* ,

$$\frac{\partial}{\partial \lambda} \nu_{N,\beta,\lambda,\star}(h_x^2) \geq 0,$$

so that we find the following result.

Lemma 7.17 (Gluing vertices). *Let $x, v_1, v_2 \in V$ and H_N^* be obtained from G_N^* by gluing together v_1 and v_2 . Then $\nu_{H_N,\beta,\star}[h_x^2] \leq \nu_{G_N,\beta,\star}[h_x^2]$.*

To summarize, we have established that gluing two vertices together reduces the variance of the height function, and splitting edges as in Lemma 7.16 does not change the model. These two facts together imply that we can modify Γ to obtain a planar graph Γ' of degree at most three as we will explain now. We first show how to go from any planar graph to a planar graph of degree at most three.

Star-tree transform

There are many ways to transform a planar graph into a planar graph with degree at most three. We follow here the elegant idea presented in [86], where it was (implicitly) stated for the Gaussian free field. Suppose that $G = (V, E)$ is a planar graph with boundary $\partial \in V$ and take the setup of Section 7.5.1. Fix a vertex $v_0 \in V$. It will be slightly more convenient to make a distinction between the number of neighboring vertices of v_0 and its degree in multigraphs.

Degree reduction algorithm at v_0 .

1. If the the number of neighbors of v_0 is strictly less than 4, do nothing.
2. Label all neighbors of v_0 by v_1, \dots, v_{2d} by starting somewhere and going clockwise around v_0 , where we *do not* include the last vertex if the number of neighbors is odd.
3. Add to each edge v_0v_i an intermediate vertex x_i (note that if there are multiple edges between v_0 and v_i , then we have created many new vertices).
4. put the potential $\mathcal{V}_{\beta v_0 v_i/2}$ on the edges v_0x_i and x_iv_i , for each i .
5. Glue together each pair x_{2i-1} and x_{2i} (this includes gluing together multiple vertices x_i if they exist).

Note that this algorithm reduces the *number of neighbors of v_0* by a factor 2 if this number is even. Also note that it creates a multigraph. From the splitting and gluing lemmas, we obtain the next result.

Lemma 7.18. *Applying the degree reduction algorithm at v_0 does not increase the variance of h_x for any $x \in V$.*

Thus, to reduce the number of neighbors of v_0 to 3 or less, we are left to apply the reduction algorithm inductively, and to get a graph of degree three we apply it to all vertices in G other than the boundary vertex.

To finalize the star-tree transform, we still need to transform the multi-graph into a simple graph. Of course, we need to do so without changing the height function model. If e_1, e_2 are two edges with the same end-points x and y then

$$\nu_{N,\beta,\star}(h_x - h_y = k) \propto e^{-\mathcal{V}_\beta(k)} e^{-\mathcal{V}_\beta(k)} = e^{-(\mathcal{V}_\beta + \mathcal{V}_\beta)(k)}. \quad (20)$$

This observation implies that applying inductively the reduction algorithm and then applying the above observation does result in a graph where:




Figure 7.1: Left: original graph with v_0 in the center. Middle two: first step of the degree-reduction algorithm, dotted lines correspond to vertices to be glued together. Right: Final graph after “star-tree” transform, with vertices glued together and all vertices have three or less neighbors.

- (i) the number of neighbors of each vertex is less than or equal to 3,
- (ii) the variance of the height function is not increased,
- (iii) the potentials are of the form $DV_{\beta/k}$ for some D and k that can depend on the edges.

Proof of Theorem 7.12

Before we finish the proof of Theorem 7.12, let us briefly mention how to go from a “roughly planar” graph to a planar graph, see also Figure 7.2. We will do so for \mathbb{Z}^2 where $x \sim y$ if $|x - y|_2 \leq 2$. Add to an edge connecting two vertices x and y that are at distance 2 from each other a new vertex. Glue it to the unique vertex between x and y that is at distance 1 of each. Apply this algorithm to all edges. The obtained graph is planar and the variance of the height function is not increased by Lemma 7.17.

Proof of Theorem 7.10. Consider the setup of Section 7.5.1. By Lemma 7.15, we can assume without loss of generality that $-\mathcal{U}(J) = \cos(iJ)$. Write \mathcal{V}_β for the corresponding height function potential.

Let Γ' be the planar graph obtained from Γ as in Section 7.5.3, with the corresponding potentials $D_e \mathcal{V}_{\beta/k_e}$, $D_e \in (0, \infty)$ and $k_e \in \mathbb{N}$. Although D_e, k_e may be different on distinct edges, they are uniformly bounded because Γ (and hence Γ') is invariant a bi-periodic lattice action.

Let $(G_N)_{N \geq 1}$ be any exhaustion of Γ and let $(G'_N)_{N \geq 1}$ be the induced exhaustion of Γ' , obtained from applying the degree reduction algorithm to all of G_N (but not the boundary vertex). Write $\nu'_{N, \beta, \star}$ for the corresponding height function measure on




Figure 7.2: Left: an example of \mathbb{Z}^2 with long-range interactions; only the edges of the (red) origin are drawn. Middle: gluing. The square (gray) vertices are added, together with the green edges where the gluing will happen. Right: the final (planar) graph.

G'_N . It follows from Section 7.5.3 that it suffices to prove that for all β large enough, $\nu'_{N,\beta,\star}[h_o^2] \rightarrow \infty$. Indeed, in this case we also have $\nu_{N,\beta,\star}[h_o^2] \rightarrow \infty$.

Note that since $\mathcal{V}_{\beta/k}$ is convex, so is any multiple. Moreover, for each $D \in (0, \infty)$ and $k \in \mathbb{N}$ we have that for all β large enough,

$$D\mathcal{V}_{\beta/k}(0) \leq D\mathcal{V}_{\beta/k}(1) + \log(2).$$

Indeed, this follows from the fact that the modified Bessel function satisfies the ratio $I_m(\beta)/I_{m'}(\beta)$ tends to 1 as $\beta \rightarrow \infty$ (see e.g. [64]). Therefore, we can apply Theorem 7.13 and Lemma 7.14 to deduce that for β large enough,

$$\nu'_{N,\beta,\star}[h_o^2] \rightarrow \infty$$

as $N \rightarrow \infty$. This proves the theorem. \square

7.6. GFF covariance for a projection of the spin model

Let $G = (V, E)$ be a finite graph. Recall that $\Omega^1(\mathbb{R})$ equipped with the l^2 -inner product $(\epsilon, \omega)_{\Omega^1}$ is a Hilbert space. In this setting, the linear operators d and d^* are mutually adjoint, and hence the spaces $H_\star(\mathbb{R})$ and $H_\diamond(\mathbb{R})$ are orthogonal in $\Omega^1(\mathbb{R})$ and span the whole space, i.e.,

$$\Omega^1(\mathbb{R}) = H_\star(\mathbb{R}) \oplus H_\diamond(\mathbb{R}).$$

We denote by P_\diamond and P_\star the orthogonal projection onto $H_\diamond(\mathbb{R})$ and $H_\star(\mathbb{R})$ respectively.

We focus on finite graphs $G = (V, E)$ with boundary vertex $\partial \in V$ and take mutually dual potentials \mathcal{V} and \mathcal{U} as in Definition 7.1.

Since \mathcal{U} is symmetric around 0 by assumption, the derivative \mathcal{U}' of \mathcal{U} is odd and hence $\mathcal{U}'(J)$ is a 1-form in $\Omega^1(\mathbb{R})$. It thus makes sense to look at the orthogonal decomposition

of $\mathcal{U}'(J)$ in the space $H_\star(\mathbb{R}) \oplus H_\diamond(\mathbb{R})$. Define τ to be the unique element of $\Omega_0^0(\mathbb{R})$ such that

$$d\tau = P_\star \mathcal{U}'(J).$$

We will next obtain the – in our eyes somewhat remarkable – result that τ has the covariance of a Gaussian free field irrespective of \mathcal{U} .

Proposition 7.19 (GFF covariance). *Let τ be as above and $f, g \in \Omega_0^0(\mathbb{R})$. Then*

$$\inf_{e \in E} \mu_{G,\star}[\mathcal{U}''(J_e)](f, \mathbf{G}g) \leq \mu_{G,\star}[(\tau, f)(\tau, g)] \leq \sup_{e \in E} \mu_{G,\star}[\mathcal{U}''(J_e)](f, \mathbf{G}g).$$

We begin with an easy consequence of the duality lemma for covariance 7.3.

Lemma 7.20. *For any $f, g \in \Omega_0^0(\mathbb{R})$, we have*

$$\mu_{G,\star}[(\mathcal{U}'(J), dg)(\mathcal{U}'(J), df)] = \sum_{e \in E} \mu_{G,\star}[\mathcal{U}''(J_e)]df_e dg_e$$

Proof. Let f, g be as in the statement so that $df, dg \in H_\star$. Lemma 7.3 implies

$$\mu_{G,\star}[(\mathcal{U}'(J), df)(\mathcal{U}'(J), dg)] = \sum_{e \in E} \mu_{G,\star}[\mathcal{U}''(J_e)]df_e dg_e - \nu_{G,\diamond}[(\mathbf{n}, df)(\mathbf{n}, dg)]$$

Since H_\star and H_\diamond are orthogonal, and the dual height 1-form \mathbf{n} takes value in H_\diamond , the right-most term vanishes and the result follows. \square

Proof of Proposition 7.19. For any $g \in \Omega^0(\mathbb{R})$, we have

$$(\mathcal{U}'(J), dg) = (P_\star \mathcal{U}'(J), dg) = (\tau, \Delta g)$$

where the first equality holds because $dg \in H_\star$ and the second by definition of τ , the duality of d and d^* and because $\Delta = d^*d$. As before, \mathbf{G} denotes the inverse of Δ (so defined that functions take value 0 at the boundary) and take now $g = \mathbf{G}m$, $h = \mathbf{G}f$. It follows from this and the previous lemma that

$$\mu_{G,\star}[(\tau, m)(\tau, f)] = \sum_{e \in E} \mu_{G,\star}[\mathcal{U}''(J_e)]df_e dg_e,$$

implying the desired result. \square

7.7. A central limit theorem

Here we consider the same setup as in Section 7.4.3, and we establish a central limit theorem for $\mathcal{U}'(J)$ summed over the path p_n . The main conclusion of this section is that even though the decay of the correlations of $\mathcal{U}'(J)$ changes if the height function delocalizes, $\mathcal{U}'(J)$ always satisfies a central limit theorem as shown below.

Let $(N_k)_{k \geq 1}$ be a sequence along which $\mu_{\mathbb{T}_N, \diamond}$ converges weakly to a measure $\mu = \mu_{\mathbb{T}_N}$. As usual, by duality, μ can be thought of as a Gibbs measure on $H_\star(\mathbb{Z}^2, \mathbb{S})$. By Remark 7.6, for any fixed n , the difference $h_{v_0} - h_{v_n}$ converges weakly under $\nu_{\mathbb{T}_N, \star}$, as $N \rightarrow \infty$ (as long as $v_0, v_n \in \mathbb{T}_N$). Moreover by Corollary 7.4,

$$\lim_{n \rightarrow \infty} \limsup_{k \rightarrow \infty} \nu_{\mathbb{T}_{N_k}, \star} [(h_{v_0} - h_{v_n})^2] / n \leq \lim_{n \rightarrow \infty} c \frac{\log n}{n} = 0.$$

and hence by Lemma 7.2, for all $t \in \mathbb{R}$,

$$1 = \lim_{n \rightarrow \infty} \lim_{k \rightarrow \infty} \nu_{\mathbb{T}_{N_k}, \star} \left[\exp \left(i \frac{t}{\sqrt{n}} (h_{v_0} - h_{v_n}) \right) \right] = \lim_{n \rightarrow \infty} \lim_{k \rightarrow \infty} Z_{\mathbb{T}_{N_k}, \diamond} \left(\frac{-t}{\sqrt{n}} p_n \right) / Z_{\mathbb{T}_{N_k}, \diamond} (0), \quad (21)$$

where again we identify the path p_n with the associated 1-form. Using that

$$\mathcal{U}(J + \varepsilon) - \mathcal{U}(J) = \mathcal{U}'(J)\varepsilon + \frac{1}{2}\mathcal{U}''(J)\varepsilon^2 + o(\varepsilon^2)$$

we can write

$$\begin{aligned} & Z_{\mathbb{T}_{N_k}, \diamond} \left(\frac{-t}{\sqrt{n}} p_n \right) / Z_{\mathbb{T}_{N_k}, \diamond} (0) \\ &= \mu_{\mathbb{T}_{N_k}, \diamond} \left[\exp \left(-t \frac{1}{\sqrt{n}} \sum_{i=1}^n \mathcal{U}'(J_i) - t^2 \frac{1}{2n} \sum_{i=1}^n \mathcal{U}''(J_i) + o(t^2) \right) \right], \end{aligned}$$

where the error is uniform in k . We first note that by weak convergence, as $k \rightarrow \infty$, the right-hand side approaches to the same expectation but with respect to μ (the infinite volume limit of $\mu_{\mathbb{T}_N, \star}$). Moreover, the error term vanishes in the limit $n \rightarrow \infty$. Assuming that the spin measure μ is ergodic, we also have that

$$\frac{1}{n} \sum_{i=1}^n \mathcal{U}''(J_i) \rightarrow \mu[\mathcal{U}''(J_0)] \quad \mu\text{-a.s. as } n \rightarrow \infty$$

by Birkhoff's pointwise ergodic theorem (since J is invariant under the shift along the path, and \mathcal{U}'' is bounded). Note that in the case of the XY model, there is only one translation invariant Gibbs measure in two dimensions [136] which must therefore be ergodic. By the dominated convergence theorem and (21), we conclude the following central limit theorem.

Theorem 7.21. *If μ is ergodic, then for any $t \in \mathbb{R}$,*

$$\lim_{n \rightarrow \infty} \mu \left[\exp \left(- \frac{t}{\sqrt{n}} \sum_{i=1}^n \mathcal{U}'(J_i) \right) \right] = \exp \left(\frac{t^2}{2} \mu[\mathcal{U}''(J_0)] \right).$$

In particular,

$$\frac{1}{\sqrt{n}} \sum_{i=1}^n \mathcal{U}'(J_i) \rightarrow \mathcal{N}(0, \mu[\mathcal{U}''(J_0)])$$

in distribution as $n \rightarrow \infty$.

Uniqueness of the potential kernel implies one-endedness of the UST, without finite expected degree

In this appendix we prove Theorem 3.15, which generalizes the theorem of Berestycki and the first author concerning the equivalence of the UST being one-ended and uniqueness of the harmonic measure from infinity to the case that the unimodular random rooted graph does not necessarily have finite expected degree. A secondary purpose of this appendix is to give a brief and self-contained account of those results of [34] that are needed for our main results. Since recurrent graphs whose USTs are one-ended always have unique harmonic measure from infinity [32, Theorem 14.2], it suffices to prove that the converse holds under the additional assumption of unimodularity. Moreover, it suffices as usual to consider the case that (G, o) is ergodic.

Suppose that (G, o) is an ergodic recurrent unimodular random rooted graph for which \mathcal{H} is a singleton almost surely. We write h for the unique element of \mathcal{H} and a for the corresponding potential kernel. For each $c > 0$ consider the event $A_c = \{\limsup_{x \rightarrow \infty} h_{x,o}(x) \geq c\} = \{\text{for each } \varepsilon > 0 \text{ there exist infinitely many vertices } x \text{ with } h_{x,o}(x) \geq c - \varepsilon\}$. As explained in detail in [34, Lemma 5.3] (which concerns deterministic recurrent graphs), we have that

$$h_{x,o}(x) \sim h_{x,w}(x) \quad \text{as } x \rightarrow \infty \text{ for each fixed } w \in V, \quad (1)$$

which implies that A_c is re-rooting invariant. Since (G, o) was assumed to be ergodic we deduce the following.

Lemma A.1. *Let (G, o) be an ergodic unimodular random rooted graph. If G is almost surely recurrent with a uniquely defined harmonic measure from infinity then the event A_c has probability 0 or 1 for each $c \in (0, 1)$.*

The next lemma is proven in [34] using an argument that relies on reversibility (and hence on the assumption $\mathbb{E} \deg(o) < \infty$). We give an alternative proof using Følner sequences that works without this assumption.

Lemma A.2. *Let (G, o) be an ergodic unimodular random rooted graph. If G is almost surely recurrent with a uniquely defined harmonic measure from infinity then the event $A_{1/2}$ holds almost surely.*

Proof. It suffices to prove that A_c holds with positive probability for every $c < 1/2$. Since (G, o) is recurrent, it follows from the results of [9, §8] that (G, o) is *hyperfinite*, meaning that there exists a sequence of random subsets $(\omega_n)_{n \geq 1}$ of E such that

1. Every component of the subgraph spanned by ω_n is finite almost surely for each $n \geq 1$.
2. $\omega_n \subseteq \omega_{n+1}$ for each $n \geq 1$ and $\bigcup_{n \geq 1} \omega_n = E$.
3. The random rooted edge-labelled graph $(G, o, (\omega_n)_{n \geq 1})$ is unimodular.

Let $n \geq 1$ and let K_n be the component of o in ω_n . Then we have by the mass-transport principle that

$$\mathbb{E} \left[\frac{1}{|K_n|} \sum_{x \in K_n} \mathbf{1} \left(h_{x,o}(x) \geq \frac{1}{2} \right) \right] = \mathbb{E} \left[\frac{1}{|K_n|} \sum_{x \in K_n} \mathbf{1} \left(h_{x,o}(o) \geq \frac{1}{2} \right) \right],$$

and since the sum of the two sides is at least 1 it follows that

$$\mathbb{E} \left[\frac{1}{|K_n|} \sum_{x \in K_n} \mathbf{1} \left(h_{x,o}(x) \geq \frac{1}{2} \right) \right] \geq \frac{1}{2}$$

and hence by Markov's inequality that

$$\mathbb{P} \left(|\{x \in K_n : h_{x,o}(x) \geq \frac{1}{2}\}| \geq \frac{1}{4}|K_n| \right) \geq 1 - \frac{4}{3} \mathbb{E} \left[\frac{1}{|K_n|} \sum_{x \in K_n} \mathbf{1} \left(h_{x,o}(x) < \frac{1}{2} \right) \right] \geq \frac{1}{3}.$$

Since $|K_n| \rightarrow \infty$ almost surely as $n \rightarrow \infty$, it follows from this and Fatou's lemma that

$$\mathbb{P}(A_{1/2}) \geq \mathbb{P} \left(|\{x \in K_n : h_{x,o}(x) \geq \frac{1}{2}\}| \geq \frac{1}{4}|K_n| \text{ for infinitely many } n \right) \geq \frac{1}{3}$$

and hence by ergodicity that $\mathbb{P}(A_{1/2}) = 1$ as claimed. \square

Lemma A.3. *Let $G = (V, E)$ be an infinite, connected, locally finite recurrent graph with uniquely defined harmonic measure from infinity h , let $o \in V$ and let a be the associated potential kernel. If A is any infinite set of vertices with $\inf_{x \in A} h_{x,o}(x) > 0$, the Doob-transformed walk \widehat{X} visits A infinitely often almost surely.*

Proof. We have by (11) that $\widehat{\mathbf{P}}_o(\widehat{X} \text{ hits } x) = h_{o,x}(x)$ for every $x \in V$, and it follows by Fatou's lemma that $\widehat{\mathbf{P}}(\text{hit } A \text{ infinitely often}) \geq \inf_{x \in A} h_{x,o}(x) > 0$. On the other hand, we have by Theorem 3.8 and the assumption that h is unique that \widehat{X} has trivial tail σ -algebra, so that $\widehat{\mathbf{P}}(\text{hit } A \text{ infinitely often}) = 1$ as claimed. \square

Proposition A.4. *Let $G = (V, E)$ be an infinite, connected, locally finite recurrent graph with uniquely defined harmonic measure from infinity h , let $o \in V$, let a be the associated potential kernel, and suppose that $\liminf_{x \rightarrow \infty} h_{x,o}(x) > 0$. If \widehat{X} and \widehat{Y} are independent copies of the Doob-transformed walk started at some vertices x and y , then $\{\widehat{X}_n : n \geq 0\} \cap \{\widehat{Y}_n : n \geq 0\}$ is infinite almost surely.*

Proof. Let $\delta > 0$ be such that $A = \{x \in V : h_{x,o}(x) \geq \delta\}$ is infinite. Applying Lemma A.3 yields that $\widehat{X} \cap A$ is infinite almost surely, and applying Lemma A.3 a second time yields that $\widehat{Y} \cap \widehat{X} \cap A$ is infinite almost surely. \square

Applying this proposition together with the results of [132], which imply that an independent Markov process and loop-erased Markov process intersect infinitely almost surely whenever the corresponding two independent Markov processes do, we deduce the following immediate corollary.

Corollary A.5. *Let $G = (V, E)$ be an infinite, connected, locally finite recurrent graph with uniquely defined harmonic measure from infinity h , let $o \in V$, let a be the associated potential kernel, and suppose that $\liminf_{x \rightarrow \infty} h_{x,o}(x) > 0$. If \widehat{X} and \widehat{Y} are independent copies of the Doob-transformed walk started at some vertices x and y , then $\{\widehat{X}_n : n \geq 0\} \cap \{\text{LE}(\widehat{Y})_n : n \geq 0\}$ is infinite almost surely.*

Proposition A.6. *Let $G = (V, E)$ be an infinite, connected, locally finite recurrent graph with uniquely defined harmonic measure from infinity h , let $o \in V$, let a be the associated potential kernel, and suppose that $\liminf_{x \rightarrow \infty} h_{x,o}(x) > 0$. For each $x \in V$, let X be a random walk started at x and let \widehat{Y} be a Doob-transformed walk started at o . Then*

$$\lim_{x \rightarrow \infty} \mathbb{P} \left(\{X_n : 0 \leq n \leq T_o\} \cap \{\text{LE}(\widehat{Y})_m : m \geq 0\} = \{o\} \right) = 0.$$

Proof. As $x \rightarrow \infty$, the law of the time-reversed final segment $(X_{T_o}, X_{T_o-1}, \dots, X_{T_o-k})$ converges to that of $(\widehat{X}_0, \dots, \widehat{X}_k)$ for each $k \geq 1$, and the claim follows from Corollary A.5. \square

Proof of Theorem 3.15. The fact that G has a unique harmonic measure from infinity means that we can endow the uniform spanning tree of G with an orientation in a canonical way: Suppose that we exhaust V by finite sets $V = \bigcup V_n$ and let G_n^* be defined by contracting $V \setminus V_n$ into a single boundary vertex ∂_n , so that the UST of G can be expressed as the weak limit of the USTs of the graphs G_n^* . If for each $n \geq 1$ we orient the UST of G_n^* towards the boundary vertex ∂_n to obtain an oriented tree T_n^\rightarrow , then the uniqueness of the harmonic measure from infinity on G implies that the law of T_n^\rightarrow converges weakly to the law of an oriented spanning tree T^\rightarrow of G , which can be thought of as a canonical (but potentially random) orientation of the UST of G .

Indeed, if we fix an enumeration v_1, v_2, \dots of V with $v_1 = o$ we can sample T_n^{\rightarrow} using Wilson's algorithm rooted at ∂_n , starting with the vertices in the order they appear in the enumeration of V , and orienting the edges of the tree in the direction they are crossed by the loop-erased walk that contributed them to the tree. In the infinite-volume limit (since only the part of the first walk after its final visit to o contributes to its loop erasure), this corresponds to doing Wilson's algorithm where the first walk started at o is Doob-transformed and the remaining walks are ordinary simply random walks.

An important consequence of this discussion is that if we sample the oriented uniform spanning tree using Wilson's algorithm rooted at infinity, where the first random walk is a Doob-transformed walk started at o and the remaining walks are ordinary simple random walks, the distribution of the resulting oriented tree T^{\rightarrow} does not depend on the choice of the root vertex o , since it is given by the limit of the USTs of G_n^* oriented towards ∂_n independently of the choice of exhaustion. Given the oriented tree T^{\rightarrow} , we say that a vertex u is in the **future** of a vertex v if the unique infinite oriented path emanating from v passes through u , and say that u is in the **past** of v if v is in the future of u .

Let $(\omega_n)_{n \geq 1}$ be a sequence witnessing the fact that (G, o) is hyperfinite as in the proof of Lemma A.2 and let K_n be the cluster of o in ω_n for each $n \geq 1$. We have by the mass-transport principle that

$$\mathbb{E} \left[\frac{1}{|K_n|} \sum_{x \in K_n} \mathbf{1}(x \text{ in past of } o) \right] = \mathbb{E} \left[\frac{1}{|K_n|} \sum_{x \in K_n} \mathbf{1}(x \text{ in future of } o) \right].$$

On the other hand, letting \mathcal{S} be the set of vertices belonging to a doubly infinite path in T , we also have that

$$\mathbb{E} \left[\frac{1}{|K_n|} \sum_{x \in K_n} \mathbf{1}(x \text{ in past or future of } o) \right] \geq \mathbb{E} \left[\frac{1}{|K_n|} \sum_{x \in K_n} \mathbf{1}(o, x \in \mathcal{S}) \right]$$

and we can use the mass-transport principle again to bound

$$\begin{aligned} \mathbb{E} \left[\frac{1}{|K_n|} \sum_{x \in K_n} \mathbf{1}(o, x \in \mathcal{S}) \right] &= \mathbb{E} \left[\frac{1}{|K_n|^2} \sum_{x, y \in K_n} \mathbf{1}(o, x \in \mathcal{S}) \right] \\ &= \mathbb{E} \left[\frac{1}{|K_n|^2} \sum_{x, y \in K_n} \mathbf{1}(x, y \in \mathcal{S}) \right] \\ &= \mathbb{E} \left[\left(\frac{|K_n \cap \mathcal{S}|}{|K_n|} \right)^2 \right] \geq \mathbb{E} \left[\frac{|K_n \cap \mathcal{S}|}{|K_n|} \right]^2 = \mathbb{P}(o \in \mathcal{S})^2. \end{aligned}$$

Putting these two estimates together, it follows that

$$\mathbb{E} \left[\frac{1}{|K_n|} \sum_{x \in K_n} \mathbf{1}(x \text{ in past of } o) \right] \geq \frac{1}{2} \mathbb{P}(o \in \mathcal{S})^2. \quad (2)$$

On the other hand, if we sample T^{\rightarrow} using Wilson's algorithm rooted at infinity, starting with a Doob-transformed \widehat{Y} started at o followed by an ordinary random walk X started at x , the vertex x belongs to the past of o if and only if the walk X first hits the loop-erasure of \widehat{Y} at the vertex o . Proposition A.6 implies that this probability tends to zero as $x \rightarrow \infty$, and it follows by bounded convergence that

$$\mathbb{E} \left[\frac{1}{|K_n|} \sum_{x \in K_n} \mathbf{1}(x \text{ in past of } o) \right] \rightarrow 0 \tag{3}$$

as $n \rightarrow \infty$. Putting together (2) and (3) yields that $\mathbb{P}(o \in \mathcal{S}) = 0$. Since “everything that can happen somewhere can happen at the root” [9, Lemma 2.3], it follows that $\mathcal{S} = \emptyset$ almost surely and hence that T is one-ended almost surely as claimed. \square

Duality between height functions and spin models

Let $G = (V, E)$ be a finite graph and write \vec{E} for its oriented edges. Here and in what follows, we will always fix an implicit embedding $E \hookrightarrow \vec{E}$, which fixes for each edge a prescribed orientation. The particular embedding chosen is of no importance.

Recall that the two Hilbert spaces $\Omega^0(\mathbb{R}), \Omega^1(\mathbb{R})$ are equipped with the natural inner products:

$$(f, g)_{\Omega^0} := \sum_{x \in V} f_x g_x \quad \text{and} \quad (\epsilon, \omega)_{\Omega^1} := \frac{1}{2} \sum_{\vec{e} \in \vec{E}} \epsilon_{\vec{e}} \omega_{\vec{e}}$$

respectively, and it is standard to see that d and d^* are adjoints: for all $f \in \Omega^0(\mathbb{R})$ and $\omega \in \Omega^1(\mathbb{R})$

$$(df, \omega)_{\Omega^1} = (f, d^*\omega)_{\Omega^0}.$$

We will use the embedding $\mathbb{S} \hookrightarrow \mathbb{R}$ uniquely defined by identifying $e^{i\theta} \leftrightarrow \theta$ in such a way that $\theta \in (-\pi, \pi]$.

Recall the definition for any 1-form $\epsilon \in \Omega^1(\mathbb{R})$ of the twisted partition function

$$Z_{\#}(\epsilon) = \int_{H_{\#}(G, \mathbb{S})} \prod_{e \in E} w(J_e + \epsilon_e) dJ.$$

We wish to emphasize again that when $\# = \diamond$, unlike in the case of planar graphs, this does not generally correspond to a spin model on vertices, but rather to a measure on 1-forms taking value in the group \mathbb{S} and satisfying $d^*J = 0$.

We will need an appropriate description of the Haar measure on $H_{\#}(G, \mathbb{S})$. Let us start with the case $\# = \star$.

Lemma B.1. *For any function $F : \Omega^1(G, \mathbb{S}) \rightarrow \mathbb{R}$*

$$\int_{H_{\star}(G, \mathbb{S})} F(J) dJ = \int_{\mathbb{S}^V} F(d\theta) d\theta$$

where $d\theta$ is the product uniform measure on \mathbb{S}^V .

Proof. We first note that the measure ν on $H_\star(E, \mathbb{S})$ defined via

$$\nu(A) := \int_{\mathbb{S}^V} \mathbb{1}_A(d\theta) d\theta$$

is a Radon probability measure. We are left to argue that it is invariant under the group action, since then it is the unique Haar measure. To that end, let $J' \in H_\star$ and recall that we use additive notation for abelian groups. Note that $J' = d\theta'$ for some $\theta' \in \mathbb{S}^V$, so that $J' + d\theta = d\theta' + d\theta = d(\theta' + \theta)$. In particular,

$$\nu(A - J') = \int_{\mathbb{S}^V} \mathbb{1}_{A - J'}(d\theta) d\theta = \int_{\mathbb{S}^V} \mathbb{1}_A(d(\theta' + \theta)) d\theta = \nu(A),$$

where the last equality follows as $d\theta$ is the product Lebesgue measure and hence invariant under rotations (translations) of each of the coordinates. \square

For the case $\# = \diamond$, we will need a different argument. An element $J \in H_\diamond(\mathbb{S})$ satisfies $d^*J \equiv 0$ by definition. Therefore, knowing the value of J at all edges containing a vertex x but one, uniquely determines the value of J on the last edge. Let $T \subset E$ be a spanning tree of $G = (V, E)$ (the exact choice does not matter). Let $G_T = (V, E \setminus T)$. If $\partial \in V$ is a chosen root, then T can be oriented towards the root and as such, each vertex in $V \setminus \{\partial\}$ satisfies that there is exactly one edge in the oriented tree pointing out of x . Therefore, for each $J \in \Omega^1(G_T, \mathbb{S})$, there is a unique way to extend J to E in such a way that $J \in H_\diamond(G, \mathbb{S})$, and we will write \bar{J} for this extension.

Lemma B.2. *For any function $F : \Omega^1(G, \mathbb{S}) \rightarrow \mathbb{R}$, we have*

$$\int_{H_\diamond(\mathbb{S})} F(J) dJ = \int_{\mathbb{S}^{E \setminus T}} F(\bar{J}) dJ,$$

where the measure on the right-hand side is the product uniform measure on $\mathbb{S}^{E \setminus T}$. In particular, the right-hand side does not depend on T .

Proof. Define the measure ν on $H_\diamond(G, \mathbb{S}) = \ker(d^*)$ through $\nu(A) := \int_{\mathbb{S}^{E \setminus T}} \mathbb{1}_A(\bar{J}) dJ$. It is enough to show that ν is invariant under the group action. Indeed, for any $\tau \in \mathbb{S}^{E \setminus T}$, we have

$$\nu(A - \bar{\tau}) = \int_{\mathbb{S}^{E \setminus T}} \mathbb{1}_A(\overline{J + \tau}) dJ = \int_{\mathbb{S}^{E \setminus T}} \mathbb{1}_A(\bar{J}) dJ = \nu(A),$$

since the product uniform measure is invariant under the group action. This ends the proof. \square

We will also need the following classical results from Fourier series theory. For proofs, see e.g. [20] or [170, Theorem IV.2.9] (in German).

Lemma B.3. *Let $f : \mathbb{S} \rightarrow \mathbb{R}$ be continuously differentiable. Then $f(\theta) = \lim_{K \rightarrow \infty} f_K(\theta)$, with*

$$f_K(\theta) = a_0 + \sum_{k=1}^K (a_k e^{ik\theta} + a_{-k} e^{-ik\theta}), \quad \text{where} \quad a_k = \int_{\mathbb{S}} e^{-ik\theta} f(\theta) d\theta.$$

Moreover, the convergence is uniform on \mathbb{S} . Finally

$$f_K(\theta) = \int_{\mathbb{S}} \left(\sum_{k=-K}^K e^{ik(\theta' - \theta)} \right) f(\theta') d\theta'.$$

Proof of Lemma 7.2. Case I: ν_\diamond . It follows from condition 1 and from the dominated convergence theorem that for any $\epsilon \in \Omega^1(\mathbb{R})$, we have

$$\begin{aligned} Z_\star(\epsilon) &= \int_{H_\star(\mathbb{S})} \prod_{e \in E} w(J_e + \epsilon_e) dJ \\ &= \int_{H_\star(\mathbb{S})} \sum_{\mathbf{n}: E \rightarrow \mathbb{Z}} \prod_{e \in E} e^{i\mathbf{n}_e(J_e + \epsilon_e)} \exp(-\mathcal{V}(\mathbf{n}_e)) dJ \\ &= \sum_{\mathbf{n}: E \rightarrow \mathbb{Z}} e^{i(\mathbf{n}, \epsilon)} \prod_{e \in E} \exp(-\mathcal{V}(\mathbf{n}_e)) \int_{H_\star(\mathbb{S})} e^{i(\mathbf{n}, J)} dJ. \end{aligned}$$

Moreover, by Lemma B.1 we have

$$\begin{aligned} \int_{H_\star(\mathbb{S})} e^{i(\mathbf{n}, J)} dJ &= \int_{\mathbb{S}^V} e^{i(\mathbf{n}, d\theta)} d\theta = \int_{\mathbb{S}^V} e^{i(d^*\mathbf{n}, \theta)} d\theta \\ &= \prod_{x \in V} \int_{\mathbb{S}} e^{i d^*\mathbf{n}_x \theta_x} d\theta_x = \mathbf{1}\{\mathbf{n} \in H_\diamond(\mathbb{Z})\}, \end{aligned}$$

which ends the proof of case I.

Case II: ν_\star . We will show equality of partition functions with $\epsilon = 0$, and the general case follows the same steps. By Lemma B.3 we have

$$\begin{aligned} \prod_{e \in E} \exp(-\mathcal{V}(dh_e)) &= \prod_{e \in E} \int_{\mathbb{S}} e^{i dh_e \theta_e} w(\theta_e) d\theta_e \\ &= \int_{\mathbb{S}^E} e^{i(dh, \theta)_{\Omega_1}} \prod_{e \in E} w(\theta_e) \prod_{e \in E} d\theta_e \\ &= \int_{\mathbb{S}^E} e^{i(h, d^*\theta)_{\Omega_0}} \prod_{e \in E} w(\theta_e) \prod_{e \in E} d\theta_e, \end{aligned}$$

and hence

$$\begin{aligned} \sum_{\substack{h: V \rightarrow \mathbb{Z} \\ h_\theta = 0}} \prod_{e \in E} \exp(-\mathcal{V}(dh_e)) &= \lim_{K_{v_n} \rightarrow \infty} \cdots \lim_{K_{v_1} \rightarrow \infty} \sum_{h_{v_n} \in I_{K_{v_n}}} \cdots \sum_{h_{v_1} \in I_{K_{v_1}}} \prod_{e \in E} \exp(-\mathcal{V}(dh_e)) \\ &= \lim_{K_{v_n} \rightarrow \infty} \cdots \lim_{K_{v_1} \rightarrow \infty} \sum_{h_{v_n} \in I_{K_{v_n}}} \cdots \sum_{h_{v_1} \in I_{K_{v_1}}} \int_{\mathbb{S}^E} e^{i(h, d^*\theta)_{\Omega_0}} \prod_{e \in E} w(\theta_e) \prod_{e \in E} d\theta_e, \end{aligned} \quad (1)$$

where $I_K = \{-K, \dots, K\}$, and v_1, v_2, \dots, v_n is any ordering of $V \setminus \{\partial\}$ that explores the tree T from the leaves towards the root ∂ .

We will now evaluate the expression above with the use of Lemma B.3 by iteratively (over i) exchanging the order of summation of h_{v_i} with the integration over $\theta_{e_{v_i}}$, and then taking the $K_{v_i} \rightarrow \infty$ limit. To this end, we orient each edge in T towards the root vertex ∂ , and to each vertex $v \neq \partial$ we assign the unique outgoing edge e_v from v .

In the first step we choose the leaf vertex $v = v_1$, and write

$$d^*\theta_v = \sum_{w \sim v} \theta_{wv} = -\theta_{e_v} + \theta_{e_1} + \dots + \theta_{e_l},$$

where $l+1$ is the degree of v in G , and e_1, \dots, e_l are the remaining edges in E incident to v and pointing at v . Let x be the other endpoint of the edge e_v , so that $e_v = (v, x)$. Given $h_x \in I_{K_x}$ and $(\theta_e)_{e \in E \setminus \{e_v\}}$ apply Lemma B.3 (separately to the imaginary and real part) with $f(\theta_{e_v}) := w(\theta_{e_v})e^{ih_x d^*\theta_x}$ to get

$$\int_{\mathbb{S}} \left(\sum_{h_v \in I_{K_v}} e^{ih_v d^*\theta_v} \right) w(\theta_{e_v}) e^{ih_x d^*\theta_x} d\theta_{e_v} = f_{K_v}(\theta_{e_1} + \dots + \theta_{e_l}) \rightarrow f(\theta_{e_1} + \dots + \theta_{e_l}),$$

as $K_v \rightarrow \infty$ uniformly in $\theta_{e_1} + \dots + \theta_{e_l}$. This means that we can take $K_v \rightarrow \infty$ inside the integral over $\mathbb{S}^{E \setminus \{e_v\}}$. All in all this removes the variables h_v, K_v from (1), and θ_{e_v} is replaced by $\theta_{e_1} + \dots + \theta_{e_l}$. Define now $\theta_e^{(1)} = \theta_e$ for $e \in E \setminus \{e_v\}$ and $\theta_{e_v}^{(1)} = \theta_{e_1} + \dots + \theta_{e_l}$. In other words, after step one, (1) becomes

$$\lim_{K_{v_n} \rightarrow \infty} \dots \lim_{K_{v_2} \rightarrow \infty} \sum_{h_{v_n} \in I_{K_{v_n}}} \dots \sum_{h_{v_2} \in I_{K_{v_2}}} \int_{\mathbb{S}^{E \setminus \{e_{v_1}\}}} \prod_{w \in V \setminus \{v_1\}} e^{ih_w (d^*\theta^1)_w} \prod_{e \in E} w(\theta_e^1) \prod_{e \in E \setminus \{e_{v_1}\}} d\theta_e.$$

We continue this procedure for edge e_{v_2} where we take the corresponding f (x is replaced by the other endpoint of e_{v_2}). In this step we remove the variables h_{v_2}, K_{v_2} and $\theta_{e_{v_2}}^{(1)}$, and replace the latter by $\theta_{e_1}^{(1)} + \dots + \theta_{e_l}^{(1)}$ (where l depends on v_2 now). Define then $(\theta_e^{(2)})_{e \in E}$ through $\theta_e^{(2)} = \theta_e^{(1)}$ on $e \in E \setminus \{e_{v_2}\}$ and $\theta_{e_{v_2}}^{(2)} = \theta_{e_1}^{(1)} + \dots + \theta_{e_l}^{(1)}$. We iterate the procedure until we have done so for all vertices of $V \setminus \{\partial\}$ and arrive at $\theta^{(n)}$. It is clear that $(d^*\theta^{(n)})_x = 0$ for all $x \in V \setminus \{\partial\}$, and therefore

$$(d^*\theta^{(n)})_{\partial} = (d^*\theta^{(n)}, 1)_{\Omega^0} = (\theta^{(n)}, d1)_{\Omega^1} = 0,$$

so that $d^*\theta^{(n)}$ vanishes on all of V .

Now let $(J_e)_{e \in E \setminus T} = (\theta_e)_{e \in E \setminus T}$ and define \bar{J} the unique extension to $H_{\diamond}(G, \mathbb{S})$ as before. It is easy to check that \bar{J} equals $\theta^{(n)}$. Therefore, at the end of the iterative procedure, we have that (1) becomes

$$\int_{\mathbb{S}^{E \setminus T}} \prod_{e \in E} w(\bar{J}_e) \prod_{e \in E \setminus T} dJ_e = Z_{\diamond}(0),$$

where the equality follows from Lemma B.2. This ends the proof of case II. \square

B.1. Proof of the Ginibre inequality

We focus here only on the case where J takes values in H_\diamond , as the other case is just the classical Ginibre inequality [80]. To be precise, we will prove the following fact, from which Lemma 7.11 follows immediately.

Lemma B.4. *Consider the setup as in Section 7.5.1. Let $F, F' : \mathbb{S} \rightarrow \mathbb{R}$ be two positive definite functions. Then*

$$\mu_{\beta, \diamond}(FF') - \mu_{\beta, \diamond}(F)\mu_{\beta, \diamond}(F') \geq 0.$$

As in the classical proof by Ginibre [80], we will rely on the following result.

Lemma B.5. *For any $n \in \mathbb{N}$ and $(m_i)_{i=1}^n \in \mathbb{Z}^n$, we have*

$$\int_{H_\diamond(\mathbb{S})^2} \prod_{i=1}^n (\cos(m_i J_{e_i}) \pm \cos(m_i J'_{e_i})) dJ dJ' \geq 0,$$

where the signs \pm might be different for each i .

Proof. We begin by noticing that for any linear $M : \mathbb{R}^E \rightarrow \mathbb{R}$

$$\begin{aligned} \cos(MJ) + \cos(MJ') &= 2 \cos\left(M \frac{J - J'}{2}\right) \cos\left(M \frac{J + J'}{2}\right) \quad \text{and} \\ \cos(MJ) - \cos(MJ') &= 2 \sin\left(M \frac{J - J'}{2}\right) \sin\left(M \frac{J + J'}{2}\right). \end{aligned}$$

Let $T \subset E$ be a spanning tree of $G = (V, E)$ and recall for $J \in \Omega^1(G_T, \mathbb{S})$ the definition of \bar{J} as in Lemma B.2. By Lemma B.2 we have

$$\int_{H_\diamond(\mathbb{S})^2} \prod_{i=1}^n (\cos(m_i J_{e_i}) \pm \cos(m_i J'_{e_i})) dJ dJ' = \int_{(\mathbb{S}^{E \setminus T})^2} \prod_{i=1}^n (\cos(m_i \bar{J}_{e_i}) \pm \cos(m_i \bar{J}'_{e_i})) dJ dJ'.$$

Now consider J in $\Omega^1(G_T, \mathbb{R})$ (via the usual identification of \mathbb{S} with $(-\pi, \pi]$) and define by $A_T J$ the unique extension of J to $\Omega^1(G, \mathbb{R})$ so that $J \in H_\diamond(G, \mathbb{R})$, i.e. so that $d^*(A_T J) = 0$ in \mathbb{R} . Notice that $A_T J$ and \bar{J} (seen in \mathbb{R}) are equal on all edges in $E \setminus T$, while on an edge $e_i \in T$, the difference is of the form $2\pi k_i$ for some integer k_i . Since the cosine is 2π -periodic, we notice that each factor where the edge e_i is in T is of the form $\cos(m_i (A_T J)_{e_i}) \pm \cos(m_i (A_T J')_{e_i})$. All together, we can thus write

$$\int_{H_\diamond(\mathbb{S})^2} \prod_{i=1}^n (\cos(m_i J_{e_i}) \pm \cos(m_i J'_{e_i})) dJ dJ' = \int_{(\mathbb{S}^{E \setminus T})^2} F\left(\frac{J + J'}{2}\right) F\left(\frac{J - J'}{2}\right) dJ dJ'$$

for some function $F : \mathbb{S}^{E \setminus T} \rightarrow \mathbb{R}$. Next, make the change of variables via $\tau_e := \frac{J_e - J'_e}{2}$ and $\tau'_e = \frac{J_e + J'_e}{2}$, so that

$$\int_{(\mathbb{S}^{E \setminus T})^2} F\left(\frac{J + J'}{2}\right) F\left(\frac{J - J'}{2}\right) dJ dJ' = \int_{(\mathbb{S}^{E \setminus T})^2} F(\tau) F(\tau') d\tau d\tau'$$

$$= \left(\int_{(\mathbb{S}^{E \setminus T})^2} F(\tau) d\tau \right)^2 \geq 0.$$

This ends the proof. \square

Since the collection of functions $t \mapsto \cos(mt)$, $m \in \mathbb{Z}$, generates the positive cone of positive definite functions, for any collection $\{F_i\}$ of positive definite functions $\mathbb{S} \rightarrow \mathbb{R}$, we have

$$\int_{H_\diamond(\mathbb{S})^2} \prod_{i=1}^n (F_i(J_{e_i}) \pm F_i(J'_{e_i})) dJ dJ' \geq 0.$$

From this, Lemma B.4 can be proved in exactly the same way as in [80].

B.2. Reflection positivity.

We recall briefly a condition for potentials to be reflection positive. For further reference, see e.g. [38] and [69]. Fix the torus $\mathbb{T}_n = (\mathbb{Z}/n\mathbb{Z})^d$ and let Θ be any reflection (either through edges or through vertices). This naturally splits the torus into two parts \mathbb{T}_n^+ and \mathbb{T}_n^- . Let \mathcal{U}^\pm be the set of real-valued functions on \mathbb{T}_n depending only on \mathbb{T}_n^\pm . Then Θ induces a map $\Theta : \mathcal{U}^\pm \rightarrow \mathcal{U}^\mp$. We will say that a probability measure μ on $\mathbb{S}^{\mathbb{T}_n}$ is reflection positive with respect to Θ if

- (a) $\mu(g\Theta f) = \mu(f\Theta g)$ for all $f, g \in \mathcal{U}^+$,
- (b) $\mu(g\Theta g) \geq 0$.

Although the property (a) is not important for us, it is also the easier part and it holds precisely for all measures that are invariant under the reflection Θ . It is not hard to see that *all* measures we consider in this text thus satisfy (a). We recall the following lemma.

Lemma B.6. *Let $\mathcal{H}_n : \mathbb{S}^{\mathbb{T}_n} \rightarrow \mathbb{R}$ be the Hamiltonian of a spin-system on the torus satisfying*

$$-\mathcal{H}_n = A + \Theta A + \sum_i C_i \Theta C_i$$

for some functions $A, C_i \in \mathcal{U}^+$. Then $\mu_n \propto e^{-\mathcal{H}_n}$ is reflection positive w.r.t. Θ .

For a proof we refer to e.g. [38] or [69, Lemma 10.8]. We point out already that for reflections going through vertices, the decomposition of Lemma B.6 is trivial as we consider only nearest-neighbor interactions.

For reflection through edges, we need that we can decompose

$$-\mathcal{U}(t_x - t_y) = \sum_i F_i(t_x) F_i(t_y),$$

for some collection of functions $\{F_i\}$. By classical trigonometric identities, this can be easily deduced whenever $-\mathcal{U}$ is positive definite and regular enough so that

$$-\mathcal{U}(t_x - t_y) = \sum_{i=0}^{\infty} \alpha_i \cos(i(t_x - t_y)) = \sum_{i=0}^{\infty} \alpha_i (\cos(it_x) \cos(it_y) + \sin(it_x) \sin(it_y)),$$

with $\alpha_i \geq 0$.

B.3. Positive definite functions

We will call an even function $F : \mathbb{S} \rightarrow \mathbb{R}$ *conditionally positive definite* if for any vector $\xi = (\xi_1, \dots, \xi_n) \in \mathbb{R}^n$ with mean 0 and all $t_1, \dots, t_n \in \mathbb{S}$ it holds that

$$\sum_{i,j} \xi_i \xi_j F(t_i - t_j) \geq 0$$

Lemma B.7. *A function F is conditionally positive definite if and only if e^{cF} is positive definite for each $c > 0$.*

Proof. Assume that e^{cF} is p.d. for each $c > 0$. Then

$$\frac{1}{c} \sum_{i,j} \xi_i \xi_j (e^{cF(t_i - t_j)} - 1) = \frac{1}{c} \sum_{i,j} \xi_i \xi_j e^{cF(t_i - t_j)} \geq 0,$$

and taking $c \rightarrow 0$ shows one implication, since the derivative at zero of e^{cF} is F . The other implication follows from expanding the exponential and using that the space of conditional positive definite functions is closed under addition, multiplication by nonnegative reals and multiplication. \square

Without proof, we will also state the following result.

Lemma B.8. *If $F : \mathbb{S} \rightarrow \mathbb{R}$ is conditionally positive definite, then there exists a positive definite function $\varphi : \mathbb{S} \rightarrow \mathbb{R}$ and a constant c such that $F = \varphi + c$.*

Proof. See e.g. Corollary 2.10.3 in [24]. \square

Since Gibbs measures are invariant under adding constants to the potential, this lemma implies that taking $-\mathcal{U}$ positive definite is the same as requiring $e^{-\mathcal{V}}$ to be infinitely divisible in a way so that Definition 7.1 remains satisfied for each division.

Bibliography

- [1] M. Aizenman, *Geometric analysis of φ^4 fields and Ising models. I, II*, Comm. Math. Phys. **86** (1982), no. 1, 1–48.
- [2] M. Aizenman and H. Duminil-Copin, *Marginal triviality of the scaling limits of critical 4D Ising and φ_4^4 models*, Annals of Mathematics **194** (2021), no. 1, 163–235.
- [3] M. Aizenman and R. Fernández, *On the critical behavior of the magnetization in high-dimensional Ising models*, Journal of Statistical Physics **44** (1986), no. 3, 393–454.
- [4] M. Aizenman and B. Simon, *A comparison of plane rotor and Ising models*, Physics Letters A **76** (1980), no. 3, 281–282.
- [5] M. Aizenman, *On the slow decay of $O(2)$ correlations in the absence of topological excitations: remark on the Patrascioiu-Seiler model*, J. Statist. Phys. **77** (1994), no. 1-2, 351–359. MR1300539
- [6] M. Aizenman, H. Duminil-Copin, and V. Sidoravicius, *Random Currents and Continuity of Ising Model’s Spontaneous Magnetization*, Communications in Mathematical Physics **334** (2015), no. 2, 719–742.
- [7] M. Aizenman, H. Duminil-Copin, V. Tassion, and S. Warzel, *Emergent planarity in two-dimensional Ising models with finite-range Interactions*, Inventiones mathematicae **216** (2019), no. 3, 661–743.
- [8] M. Aizenman, M. Harel, R. Peled, and J. Shapiro, *Depinning in the integer-valued Gaussian Field and the BKT phase of the 2D Villain model*, arXiv preprint arXiv:2110.09498 (2021).
- [9] D. Aldous and R. Lyons, *Processes on unimodular random networks*, Electr. J. Probab. **54** (2007), 1454–1508.
- [10] D. J. Aldous, *The random walk construction of uniform spanning trees and uniform labelled trees*, SIAM J. Discrete Math. **3** (1990), no. 4, 450–465. MR1069105
- [11] A. Ancona, *Harnack’s principle at the frontier and fatou’s theory for an elliptic operator in a lipschitzian domain*, Annales of the Fourier Institute **28** (1978), no. 4, 169–213 (fr). MR80d:31006
- [12] O. Angel, D. A. Croydon, S. Hernandez-Torres, and D. Shiraishi, *Scaling limits of the three-dimensional uniform spanning tree and associated random walk*, Ann. Probab. **49** (2021), no. 6, 3032–3105. MR4348685
- [13] O. Angel, T. Hutchcroft, A. Nachmias, and G. Ray, *Hyperbolic and parabolic unimodular random maps*, Geometric and Functional Analysis **28** (2018), no. 4, 879–942.
- [14] O. Angel and O. Schramm, *Uniform infinite planar triangulations*, Comm. Math. Phys. **241** (2003), no. 2-3, 191–213. MR2013797
- [15] Y. Aoun, M. Dober, and A. Glazman, *Phase diagram of the Ashkin-Teller model*, arXiv preprint arXiv:2301.10609 (2023).

-
- [16] J. Aru, *The geometry of the Gaussian free field combined with SLE processes and the KPZ relation*, Theses, 2015.
- [17] M. T. Barlow, *Random walks and heat kernels on graphs*, London Mathematical Society Lecture Note Series, Cambridge University Press, 2017.
- [18] M. T. Barlow, A. A. Járai, T. Kumagai, and G. Slade, *Random walk on the incipient infinite cluster for oriented percolation in high dimensions*, Communications in mathematical physics **278** (2008), no. 2, 385–431.
- [19] M. T. Barlow and M. Murugan, *Stability of the elliptic Harnack inequality*, Annals of Mathematics **187** (2018), no. 3, 777–823.
- [20] A. Bashirov, *Mathematical analysis fundamentals*, Academic Press, 2014.
- [21] R. Bauerschmidt, *Ferromagnetic spin systems*, 2016. Available at <http://www.statslab.cam.ac.uk/~rb812/teaching/math253x/index.html>.
- [22] R. Bauerschmidt and C. Webb, *The Coleman correspondence at the free Fermion point*, arXiv preprint arXiv:2010.07096 (2020).
- [23] V. Beffara and H. Duminil-Copin, *The self-dual point of the two-dimensional random-cluster model is critical for $q \geq 1$* , Probability Theory and Related Fields **153** (2012), no. 3-4, 511–542.
- [24] B. Bekka, P. de la Harpe, and A. Valette, *Kazhdan’s property (T)*, New Mathematical Monographs, vol. 11, Cambridge University Press, Cambridge, 2008. MR2415834
- [25] C. Benassi, B. Lees, and D. Ueltschi, *Correlation Inequalities for Classical and Quantum XY Models* (A. Michelangeli and G. Dell’Antonio, eds.), Springer International Publishing, 2017.
- [26] C. Benassi and D. Ueltschi, *Loop correlations in random wire models*, Communications in Mathematical Physics **374** (2020), no. 2, 525–547.
- [27] I. Benjamini and N. Curien, *Ergodic theory on stationary random graphs*, Electronic Journal of Probability **17** (2012), 1–20.
- [28] ———, *Simple random walk on the uniform infinite planar quadrangulation: subdiffusivity via pioneer points*, Geometric and Functional Analysis **23** (2013), no. 2, 501–531.
- [29] I. Benjamini, N. Curien, and A. Georgakopoulos, *The Liouville and the intersection properties are equivalent for planar graphs*, Electron. Commun. Probab. **17** (2012), 5 pp.
- [30] I. Benjamini, H. Duminil-Copin, G. Kozma, and A. Yadin, *Disorder, entropy and harmonic functions*, Ann. Probab. **43** (2015), no. 5, 2332–2373.
- [31] I. Benjamini and T. Hutchcroft, *Large, lengthy graphs look locally like lines*, Bulletin of the London Mathematical Society **53** (2021), no. 2, 482–492.
- [32] I. Benjamini, R. Lyons, Y. Peres, and O. Schramm, *Uniform spanning forests*, The Annals of Probability **29** (2001), no. 1, 1–65.
- [33] I. Benjamini and O. Schramm, *Recurrence of distributional limits of finite planar graphs*, Electron. J. Probab. **6** (2001), no. 23, 13. MR1873300
- [34] N. Berestycki and D. van Engelenburg, *Harnack inequality and one-endedness of UST on reversible random graphs*, arXiv preprint arXiv:2110.05100 (2021).
- [35] N. Berestycki and E. Gwynne, *Random walks on mated-CRT maps and Liouville Brownian motion*, arXiv (2020). <https://arxiv.org/abs/2003.10320>.
- [36] V. L. Berezinskii, *Destruction of long range order in one-dimensional and two-dimensional systems having a continuous symmetry group. I. Classical systems*, Sov. Phys. JETP **32** (1971), 493–500.

-
- [37] ———, *Destruction of Long-range Order in One-dimensional and Two-dimensional Systems Possessing a Continuous Symmetry Group. II. Quantum Systems.*, Sov. Phys. JETP **34** (1972), no. 3, 610.
- [38] M. Biskup, *Reflection positivity and phase transitions in lattice spin models*, Methods of contemporary mathematical statistical physics, 2009, pp. 1–86.
- [39] L. Boltzmann, *Lectures on gas theory*, University of California Press, Berkeley-Los Angeles, Calif., 1964. Translated by Stephen G. Brush. MR0158708
- [40] M. Bowen, G. Kun, and M. Sabok, *Perfect matchings in hyperfinite graphings*, arXiv preprint arXiv:2106.01988 (2021).
- [41] A. Z Broder, *Generating random spanning trees*, Focs, 1989, pp. 442–447.
- [42] D. Brydges, J. Fröhlich, and T. Spencer, *The random walk representation of classical spin systems and correlation inequalities*, Communications in Mathematical Physics **83** (1982), no. 1, 123–150.
- [43] R. M. Burton and M. Keane, *Density and uniqueness in percolation*, Comm. Math. Phys. **121** (1989), 501–505.
- [44] R. Burton and R. Pemantle, *Local characteristics, entropy and limit theorems for spanning trees and domino tilings via transfer-impedances*, Ann. Probab. **21** (1993), no. 3, 1329–1371. MR1235419
- [45] N. Chandgotia, R. Peled, S. Sheffield, and M. Tassy, *Delocalization of uniform graph homomorphisms from \mathbb{Z}^2 to \mathbb{Z}* , Communications in Mathematical Physics (2021).
- [46] D. Chelkak and S. Smirnov, *Universality in the 2D Ising model and conformal invariance of fermionic observables*, Invent. Math. **189** (2012), no. 3, 515–580. MR2957303
- [47] F. Comets, S. Popov, and M. Vachkovskaia, *Two-dimensional random interacements and late points for random walks*, Communications in Mathematical Physics **343** (2015), no. 1, 129–164.
- [48] T. Coulhon and L. Saloff-Coste, *Variétés riemanniennes isométriques à l’infini*, Revista Matemática Iberoamericana **11** (1995), no. 3, 687–726.
- [49] N. Curien, *Stationary random graphs*, 2021. Available at <https://www.imo.universite-paris-saclay.fr/~curien/enseignement.html>.
- [50] R. Diestel and I. Leader, *A conjecture concerning a limit of non-Cayley graphs*, J. Algebraic Combin. **14** (2001), no. 1, 17–25. MR1856226
- [51] J. Dubédat, *Exact bosonization of the Ising model*, arXiv preprint arXiv:1112.4399 (2011).
- [52] ———, *Topics on abelian spin models and related problems*, Probability Surveys **8** (2011), 374–402.
- [53] H. Duminil-Copin, *Random currents expansion of the Ising model*, 2016. arXiv:1607.06933.
- [54] H. Duminil-Copin and M. Lis, *On the double random current nesting field*, Probability Theory and Related Fields **175** (2019), no. 3, 937–955.
- [55] H. Duminil-Copin, *Parafermionic observables and their applications to planar statistical physics models*, Ensaios Matemáticos [Mathematical Surveys], vol. 25, Sociedade Brasileira de Matemática, Rio de Janeiro, 2013. MR3184487
- [56] ———, *Lectures on the Ising and Potts models on the hypercubic lattice*, Random graphs, phase transitions, and the Gaussian free field, [2020] ©2020, pp. 35–161. MR4043224
- [57] H. Duminil-Copin, M. Harel, B. Laslier, A. Raoufi, and G. Ray, *Logarithmic variance for the height function of square-ice*, Communications in Mathematical Physics **396** (2022), no. 2, 867–902.
- [58] H. Duminil-Copin, A. Karrila, I. Manolescu, and M. Oulamara, *Delocalization of the height function of the six-vertex model*, arXiv preprint arXiv:2012.13750 (2020).

-
- [59] H. Duminil-Copin, A. Raoufi, and V. Tassion, *Sharp phase transition for the random-cluster and Potts models via decision trees*, Ann. of Math. (2) **189** (2019), no. 1, 75–99. MR3898174
- [60] H. Duminil-Copin and V. Tassion, *A New Proof of the Sharpness of the Phase Transition for Bernoulli Percolation and the Ising Model*, Communications in Mathematical Physics **343** (2016), no. 2, 725–745.
- [61] E. B. Dynkin, *Boundary theory of markov processes (the discrete case)*, Russian Mathematical Surveys **24** (1969), no. 2, 1.
- [62] R. G. Edwards and A. D. Sokal, *Generalization of the Fortuin-Kasteleyn-Swendsen-Wang representation and Monte Carlo algorithm*, Phys. Rev. D (3) **38** (1988), no. 6, 2009–2012. MR965465
- [63] D. van Engelenburg and T. Hutchcroft, *The number of ends in the uniform spanning tree for recurrent unimodular random graphs*, arXiv preprint arXiv:2301.03875 (2023).
- [64] D. van Engelenburg and M. Lis, *An elementary proof of phase transition in the planar XY model*, Comm. Math. Phys. **399** (2023), no. 1, 85–104. MR4567369
- [65] ———, *On the duality between height functions and continuous spin models*, arXiv preprint arXiv:2303.08596 (2023).
- [66] A. Eskin, D. Fisher, and K. Whyte, *Coarse differentiation of quasi-isometries I: Spaces not quasi-isometric to Cayley graphs*, Ann. of Math. (2) **176** (2012), no. 1, 221–260. MR2925383
- [67] R. Fitzner and R. van der Hofstad, *Mean-field behavior for nearest-neighbor percolation in $d > 10$* , Electron. J. Probab. **22** (2017), Paper No. 43, 65. MR3646069
- [68] C. M. Fortuin and P. W. Kasteleyn, *On the random-cluster model. I. Introduction and relation to other models*, Physica **57** (1972), 536–564. MR359655
- [69] S. Friedli and Y. Velenik, *Statistical Mechanics of Lattice Systems: A Concrete Mathematical Introduction*, Cambridge University Press, 2017.
- [70] J. Fröhlich, B. Simon, and T. Spencer, *Infrared bounds, phase transitions and continuous symmetry breaking*, Communications in Mathematical Physics **50** (1976), no. 1, 79–95.
- [71] J. Fröhlich and T. Spencer, *The Kosterlitz-Thouless transition in two-dimensional Abelian spin systems and the Coulomb gas*, Communications in Mathematical Physics **81** (1981), no. 4, 527–602.
- [72] J. Fröhlich, *On the triviality of $\lambda\varphi_d^4$ theories and the approach to the critical point in $d > 4$ dimensions*, Nuclear Phys. B **200** (1982), no. 2, 281–296. MR643591
- [73] J. Fröhlich and Y. M. Park, *Remarks on exponential interactions and the quantum sine-gordon equation in two space-time dimensions*, Helv. Phys. Acta;(Switzerland) **50** (1977), no. 3.
- [74] J. Fröhlich and T. Spencer, *Massless phases and symmetry restoration in abelian gauge theories and spin systems*, Comm. Math. Phys. **83** (1982), no. 3, 411–454. MR649811
- [75] N. Gantert, S. Popov, and M. Vachkovskaia, *On the range of a two-dimensional conditioned simple random walk*, Annales Henri Lebesgue **2** (2019), 349–368.
- [76] C. Garban and A. Sepúlveda, *Statistical reconstruction of the Gaussian free field and KT transition*, arXiv preprint arXiv:2002.12284 (2020).
- [77] ———, *Quantitative bounds on vortex fluctuations in 2d Coulomb gas and maximum of the integer-valued Gaussian free field*, arXiv preprint arXiv:2012.01400 (2021).
- [78] H.-O. Georgii, *Gibbs measures and phase transitions*, de Gruyter, 2011.
- [79] J. W. Gibbs, *Elementary principles in statistical mechanics: developed with especial reference to the rational foundations of thermodynamics*, C. Scribner’s sons, 1902.

-
- [80] J. Ginibre, *General formulation of Griffiths' inequalities*, Communications in Mathematical Physics **16** (1970), no. 4, 310–328.
- [81] A. Glazman and R. Peled, *On the transition between the disordered and antiferroelectric phases of the 6-vertex model*, arXiv preprint arXiv:1909.03436 (2019).
- [82] R. B. Griffiths, C. A. Hurst, and S. Sherman, *Concavity of Magnetization of an Ising Ferromagnet in a Positive External Field*, Journal of Mathematical Physics **11** (1970), no. 3, 790–795.
- [83] A. A. Grigor'yan, *The heat equation on noncompact Riemannian manifolds (in russian)*, Matematicheskii Sbornik **182** (1991), no. 1, 55–87.
- [84] M. Gromov, *Groups of polynomial growth and expanding maps*, Inst. Hautes Études Sci. Publ. Math. **53** (1981), 53–73. MR623534
- [85] T. S. Gunaratnam, C. Panagiotis, R. Panis, and F. Severo, *Random tangled currents for φ^4 : translation invariant Gibbs measures and continuity of the phase transition*, arXiv preprint arXiv:2211.00319 (2022).
- [86] O. Gurel-Gurevich and A. Nachmias, *Recurrence of planar graph limits*, Annals of Mathematics (2013), 761–781.
- [87] E. Gwynne and J. Miller, *Random walk on random planar maps: spectral dimension, resistance, and displacement*, Ann. Probab. **49** (2021), no. 3, 1097–1128.
- [88] E. Gwynne, J. Miller, and S. Sheffield, *The Tutte embedding of the mated-CRT map converges to Liouville quantum gravity*, Ann. Probab. **49** (2021), no. 4, 1677–1717.
- [89] T. Hara and G. Slade, *Mean-field critical behaviour for percolation in high dimensions*, Comm. Math. Phys. **128** (1990), no. 2, 333–391. MR1043524
- [90] C. Hongler and S. Smirnov, *The energy density in the planar ising model*, Acta Mathematica **211** (2013), no. 2, 191–225.
- [91] T. Hutchcroft, *Wired cycle-breaking dynamics for uniform spanning forests*, The Annals of Probability **44** (2016), no. 6, 3879–3892.
- [92] ———, *Interlacements and the wired uniform spanning forest*, The Annals of Probability **46** (2018), no. 2, 1170–1200.
- [93] ———, *Universality of high-dimensional spanning forests and sandpiles*, Probability Theory and Related Fields **176** (2020), no. 1, 533–597.
- [94] T. Hutchcroft and A. Nachmias, *Indistinguishability of trees in uniform spanning forests*, Probab. Theory Related Fields **168** (2017), no. 1-2, 113–152. MR3651050
- [95] ———, *Uniform spanning forests of planar graphs*, Forum Math. Sigma **7** (2019), Paper No. e29, 55. MR4010561
- [96] T. Hutchcroft and Y. Peres, *Collisions of random walks in reversible random graphs*, Electronic Communications in Probability **20** (2015), 1–6.
- [97] T. Hutchcroft and P. Sousi, *Logarithmic corrections to scaling in the four-dimensional uniform spanning tree*, arXiv preprint arXiv:2010.15830 (2020).
- [98] E. Ising, *Beitrag zur Theorie des Ferromagnetismus*, Z. Physik **31** (1925FEB), 253–258.
- [99] A. A. Járai, W. M. Ruszel, and E. Saada, *Mean-field avalanche size exponent for sandpiles on Galton-Watson trees*, Probab. Theory Related Fields **177** (2020), no. 1-2, 369–396. MR4095018
- [100] J. V. José, L. P. Kadanoff, S. Kirkpatrick, and D. R. Nelson, *Renormalization, vortices, and symmetry-breaking perturbations in the two-dimensional planar model*, Physical Review B **16** (1977), no. 3, 1217.

-
- [101] L. P. Kadanoff and H. Ceva, *Determination of an Operator Algebra for the Two-Dimensional Ising Model*, Phys. Rev. B **3** (1971Jun), 3918–3939.
- [102] J. G. Kemeny, J. L. Snell, and A. W. Knapp, *Denumerable markov chains: with a chapter of Markov random fields by David Griffeath*, Vol. 40, Springer, 1976.
- [103] H. Kesten, *The critical probability of bond percolation on the square lattice equals $\frac{1}{2}$* , Communications in Mathematical Physics **74** (1980), no. 1, 41–59.
- [104] V. Kharash and R. Peled, *The Fröhlich-Spencer Proof of the Berezinskii-Kosterlitz-Thouless Transition*, arXiv preprint arXiv:1711.04720 (2017).
- [105] G. Kirchhoff, *Über die auflösung der gleichungen, auf welche man bei der untersuchung der linearen vertheilung galvanischer ströme geführt wird*, Annalen der Physik **148** (1847), no. 12, 497–508.
- [106] B. Kleiner, *A new proof of Gromov’s theorem on groups of polynomial growth*, J. Amer. Math. Soc. **23** (2010), no. 3, 815–829. MR2629989
- [107] J M Kosterlitz, *The critical properties of the two-dimensional xy model*, Journal of Physics C: Solid State Physics **7** (1974mar), no. 6, 1046.
- [108] J. M. Kosterlitz and D. J. Thouless, *Ordering, metastability and phase transitions in two-dimensional systems*, Journal of Physics C: Solid State Physics **6** (1973apr), no. 7, 1181–1203.
- [109] G. Kozma and A. Nachmias, *The alexander-orbach conjecture holds in high dimensions*, Inventiones mathematicae **178** (2009), no. 3, 635.
- [110] H. A. Kramers and G. H. Wannier, *Statistics of the two-dimensional ferromagnet. part i*, Physical Review **60** (1941), no. 3, 252.
- [111] M. Krikum, *Local structure of random quadrangulations*, arXiv (2005). <https://arxiv.org/abs/0512304>.
- [112] T. Kumagai, *Random walks on disordered media and their scaling limits*, École d’Été de probabilités de Saint-Flour XL - 2010, Springer, 2014.
- [113] A. J. Kupiainen, *On the $1/n$ expansion*, Comm. Math. Phys. **73** (1980), no. 3, 273–294. MR574175
- [114] P. Lammers, *Height function delocalisation on cubic planar graphs*, Probability Theory and Related Fields (2021).
- [115] P. Lammers, *A dichotomy theory for height functions*, arXiv preprint arXiv:2211.14365 (2022).
- [116] ———, *Bijecting the BKT transition*, arXiv preprint arXiv:2301.06905 (2023).
- [117] P. Lammers and S. Ott, *Delocalisation and absolute-value-FKG in the solid-on-solid model*, arXiv preprint arXiv:2101.05139 (2021).
- [118] G. F. Lawler and V. Limic, *Random walk: a modern introduction*, Cambridge Studies in Advanced Mathematics, vol. 123, Cambridge University Press, Cambridge, 2010. MR2677157 (2012a:60132)
- [119] Y. Le Jan, *Markov loops and renormalization*, The Annals of Probability **38** (2010/5/1), no. 3, 1280–1319.
- [120] ———, *Markov loops, free field and Eulerian networks*, Journal of the Mathematical Society of Japan **67** (2015), no. 4, 1671–1680.
- [121] J. R. Lee, *Conformal growth rates and spectral geometry on distributional limits of graphs*, arXiv (2017). <https://arxiv.org/abs/1701.01598>.

-
- [122] J. R. Lee, *Relations between scaling exponents in unimodular random graphs*, arXiv preprint arXiv:2007.06548 (2020).
- [123] B. Lees and L. Taggi, *Exponential decay of transverse correlations for $O(N)$ spin systems and related models*, Probability Theory and Related Fields **180** (2021), no. 3, 1099–1133.
- [124] W. Lenz, *Beitrag zum verständnis der magnetischen erscheinungen in festen körpern*, Z. Phys. **21** (1920), 613–615.
- [125] E. H. Lieb, *A refinement of Simon’s correlation inequality*, Comm. Math. Phys. **77** (1980), no. 2, 127–135.
- [126] M. Lis, *Spins, percolation and height functions*, 2019. arXiv:1909.07351.
- [127] ———, *On Boundary Correlations in Planar Ashkin–Teller Models*, International Mathematics Research Notices (202102). rnaa380.
- [128] ———, *On delocalization in the six-vertex model*, Communications in Mathematical Physics **383** (2021), 1181–1205.
- [129] T. Lupu, *From loop clusters and random interacements to the free field*, The Annals of Probability **44** (2016), no. 3, 2117–2146.
- [130] R. Lyons, B. Morris, and O. Schramm, *Ends in uniform spanning forests*, Electronic Journal of Probability **13** (2008), 1702–1725.
- [131] R. Lyons and Y. Peres, *Probability on trees and networks*, Vol. 42, Cambridge University Press, 2017.
- [132] R. Lyons, Y. Peres, and O. Schramm, *Markov chain intersections and the loop-erased walk*, Annales de l’Institut Henri Poincaré (B) Probability and Statistics **39** (2003), no. 5, 779–791.
- [133] O. A. McBryan and T. Spencer, *On the decay of correlations in $so(n)$ -symmetric ferromagnets*, Communications in Mathematical Physics **53** (1977/02/01), no. 3, 299–302.
- [134] N. D. Mermin and H. Wagner, *Absence of Ferromagnetism or Antiferromagnetism in One- or Two-Dimensional Isotropic Heisenberg Models*, Phys. Rev. Lett. **17** (1966Nov), 1133–1136.
- [135] A. Messenger and S. Miracle-Sole, *Correlation functions and boundary conditions in the Ising ferromagnet*, Journal of Statistical Physics **17** (1977), no. 4, 245–262.
- [136] A. Messenger, S. Miracle-Sole, and Ch. Pfister, *Correlation inequalities and uniqueness of the equilibrium state for the plane rotator ferromagnetic model*, Communications in Mathematical Physics **58** (1978), no. 1, 19–29.
- [137] J. Norris, *Markov chains*, Cambridge university press, 1998.
- [138] L. Onsager, *Crystal statistics. I. A two-dimensional model with an order-disorder transition*, Phys. Rev. (2) **65** (1944), 117–149.
- [139] C. Panagiotis and F. Severo, *Gap at 1 for the percolation threshold of Cayley graphs*, arXiv preprint arXiv:2111.00555 (2021).
- [140] A. Patrascioiu and E. Seiler, *Phase structure of two-dimensional spin models and percolation*, J. Statist. Phys. **69** (1992), no. 3-4, 573–595. MR1193851
- [141] R. Peierls, *On Ising’s model of ferromagnetism*, Proc. Cambridge Phil. Soc. **32** (1936), 477–481.
- [142] R. Peled and Y. Spinka, *Lectures on the spin and loop $O(n)$ models*, Sojourns in probability theory and statistical physics-i, 2019, pp. 246–320.
- [143] R. Pemantle, *Choosing a spanning tree for the integer lattice uniformly*, The Annals of Probability (1991), 1559–1574.

-
- [144] G. Pete and Á. Timár, *The free uniform spanning forest is disconnected in some virtually free groups, depending on the generator set*, *The Annals of Probability* **50** (2022), no. 6, 2218–2243.
- [145] C. E. Pfister, *Translation invariant equilibrium states of ferromagnetic abelian lattice systems*, *Communications in Mathematical Physics* **86** (1982), 375–390.
- [146] C. E. Pfister and Y. Velenik, *Random-cluster representation of the Ashkin-Teller model*, *Journal of Statistical Physics* **88** (1997Sep), no. 5, 1295–1331.
- [147] A. M. Polyakov, *Interaction of goldstone particles in two dimensions. applications to ferromagnets and massive yang-mills fields*, *Physics Letters B* **59** (1975), no. 1, 79–81.
- [148] S. Popov, *Two-dimensional random walk: From path counting to random interlacements*, *Institute of Mathematical Statistics Textbooks*, Cambridge University Press, 2021.
- [149] S. Popov, L. T. Rolla, and D. Ungaretti, *Transience of conditioned walks on the plane: encounters and speed of escape*, *Electronic Journal of Probability* **25** (2020), no. none, 1–23.
- [150] A. Raoufi, *Translation-invariant Gibbs states of the Ising model: general setting*, *Ann. Probab.* **48** (2020), no. 2, 760–777. MR4089494
- [151] V. Rivasseau, *Lieb’s correlation inequality for plane rotors*, *Communications in Mathematical Physics* **77** (1980), no. 2, 145–147.
- [152] L. Saloff-Coste, *A note on Poincaré, Sobolev and Harnack inequalities*, *Inter. Math. Res. Notices* **2** (1992), 27–38.
- [153] M. Sami, A. Bouaziz, and M. Sifi, *Discrete harmonic functions on an orthant in \mathbb{Z}^d* , *Electron. Commun. Probab.* **20** (2015), 13 pp.
- [154] R. Savit, *Duality transformations for general abelian systems*, *Nuclear Physics B* **200** (1982), no. 2, 233–248.
- [155] J. Segura, *Bounds for ratios of modified Bessel functions and associated Turán-type inequalities*, *Journal of Mathematical Analysis and Applications* **374** (2011), no. 2, 516–528.
- [156] Y. Shalom and T. Tao, *A finitary version of Gromov’s polynomial growth theorem*, *Geom. Funct. Anal.* **20** (2010), no. 6, 1502–1547. MR2739001
- [157] S. Sheffield, *Random surfaces*, *Astérisque*, Société mathématique de France, 2005 (en). MR2251117
- [158] B. Simon, *Correlation inequalities and the decay of correlations in ferromagnets*, *Comm. Math. Phys.* **77** (1980), no. 2, 111–126.
- [159] S. Smirnov, *Conformal invariance in random cluster models. I. Holomorphic fermions in the Ising model*, *Ann. of Math. (2)* **172** (2010), no. 2, 1435–1467. MR2680496 (2011m:60302)
- [160] P. M. Soardi and W. Woess, *Amenability, unimodularity, and the spectral radius of random walks on infinite graphs*, *Dipartimento di Matematica “F. Enriques”*, 1988.
- [161] K. Symanzik, *Euclidean quantum field theory*, *New York Univ.*, NY, 1969.
- [162] P. Tang, *Weights of uniform spanning forests on nonunimodular transitive graphs*, *Electronic Journal of Probability* **26** (2021), 1–62.
- [163] V. R. Thiruvengatachar and T. S. Nanjundiah, *Inequalities concerning Bessel functions and orthogonal polynomials*, *Proceedings of the Indian Academy of Sciences-Section A*, 1951, pp. 373.
- [164] Á. Timár, *Indistinguishability of the components of random spanning forests*, *Ann. Probab.* **46** (2018), no. 4, 2221–2242. MR3813990
- [165] V. I. Trofimov, *Graphs with polynomial growth*, *Mathematics of the USSR-Sbornik* **51** (1985), no. 2, 405.

-
- [166] R. Van der Hofstad and A. A. Járai, *The incipient infinite cluster for high-dimensional unoriented percolation*, Journal of statistical physics **114** (2004), no. 3, 625–663.
- [167] B. L. van der Waerden, *Die lange reichweite der regelmässigen atomanordnung in mischkristallen*, Zeitschrift für Physik **118** (1941), no. 7-8, 473–488.
- [168] N. Varopoulos, *Analysis and geometry on groups*, Proceedings of the international congress of mathematicians 1990 kyoto, 1991, pp. 951–957.
- [169] N. T. Varopoulos, L. Saloff-Coste, and T. Coulhon, *Analysis and geometry on groups*, Cambridge university press, 2008.
- [170] D. Werner, *Funktionalanalysis*, Springer-Verlag, 2006.
- [171] D. B. Wilson, *Generating random spanning trees more quickly than the cover time*, Proceedings of the twenty-eighth annual acm symposium on theory of computing, 1996, pp. 296–303.
- [172] W. Woess, *Random walks on infinite graphs and groups*, Cambridge university press, 2000.
- [173] F. Y. Wu and Y. K. Wang, *Duality transformation in a many-component spin model*, Journal of Mathematical Physics **17** (1976), no. 3, 439–440.
- [174] A. Yadin, *Harmonic functions on groups*, 2023. Available at <https://www.math.bgu.ac.il/~yadina/>.