Title
Isoperimetric shapes in supercritical bond percolation

Author
Gold, Julian Thomas

Publication Date
2017

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

Isoperimetric limit shapes in supercritical bond percolation

A dissertation submitted in partial satisfaction

of the requirements for the degree

Doctor of Philosophy in Mathematics

by

Julian Thomas Gold

2017

ABSTRACT OF THE DISSERTATION

Isoperimetric limit shapes in supercritical bond percolation

by

Julian Thomas Gold
Doctor of Philosophy in Mathematics
University of California, Los Angeles, 2017
Professor Marek Biskup, Chair

This thesis is concerned with scaling limits of sequences of random isoperimetric problems. We first consider progressively larger isoperimetric subgraphs of the infinite cluster $\mathbf{C}_\infty$ of supercritical bond percolation on $\mathbb{Z}^d$ for $d \geq 3$. We prove a shape theorem for these subgraphs, showing that upon rescaling they tend almost surely to a deterministic shape, which is itself an isoperimetric set for a norm we construct. The norm represents a homogenized surface energy arising from random interfaces between subgraphs of $\mathbf{C}_\infty$. We obtain sharp asymptotics for a modification of the Cheeger constant of $\mathbf{C}_\infty \cap [-n,n]^d$, settling a conjecture of Benjamini for the version of the Cheeger constant defined here.

We also study the isoperimetric properties of the giant component in dimension two using the original definition of the Cheeger constant, taking into account the boundary of the large box $[-n,n]^d$. Analogous results are shown here, with the caveat that a more complicated continuum isoperimetric problem emerges due to the presence of the boundary.

The dissertation of Julian Thomas Gold is approved.

Georg Menz

Peter Petersen

Thomas M. Liggett

Marek Biskup, Committee Chair

University of California, Los Angeles

2017

To my parents and brothers

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGMENTS

I thank Margaret Metzger for helping me understand the importance of my education. As an undergrad at UC Davis, I developed an interest in math largely due to the enthusiasm of my friends Connor Sempek and Asad Lodhia. Being at Davis put geometry in my heart thanks to courses of Dmitry Fuchs, Albert Schwarz, Abby Thompson, Jennifer Schultens, Blake Temple, Andy Berget and Rahbar Virk. I am indebted to Andy Berget and Alex Coward for believing in my ability to do research early on.

I am fortunate to have been supported by an NSF RTG in my first year at UCLA, and by the DYF in my final year. I'm grateful for the training I've received at UCLA: I took beautiful courses from Marek Biskup, Eviatar Procaccia, Rowan Killip, Monica Vişan, Mario Bonk, Michael Hitrik, John Garnett and Terry Tao. I happily took statistical mechanics from Eric D'Hoker twice.

This lead me to work for Marek Biskup, who I thank for his guidance (mathematically and otherwise). Marek suggested the two problems making up this dissertation, and introduced me to probability with enthusiasm and caring. In addition to supporting me at UCLA, Marek made it possible for me to attend the trimester on disordered systems at the IHP. This allowed me to meet with Raphaël Cerf, who shared his expertise on the Wulff construction over several months. Without these meetings, my first problem would have taken far longer to solve. At this point, I must also thank Caroline Mini for her support and for making it possible to stay in France during this period. It is worth emphasizing the enormous amount energy Marek has put into making the probability group a lively and productive team.

I thank my collaborators Oren Louidor and Aser Cortines for teaching me about random walks and the Gaussian free field, and I thank Lincoln Chayes, Peter Petersen, John Garnett, Tom Liggett, Georg Menz, Inwon Kim, Nick Crawford and Nestor Guillen for being open to discuss research. I likewise thank David Renfrew, Martin Tassy, Pierre-Francois Rodriguez, Stephen DeSalvo, Alejandro Morales and Omer Ben Neria for their kindness.

Every math graduate student at UCLA is lucky to have the support of the math department staff. In particular, I thank Maida Basilli, Martha Contreras and Maggie Albert.

I am fortunate to have had several role models in the analysis group during my first few years of grad school: Nick Cook, Ben Krause, Will Feldman, Ben Hayes and Jeff Lin. More generally, I'm grateful for the community of grad students within and outside of the math department: I thank Stephanie Lewkiewicz, Eric Larson, Andre Pradhana, Stephen Ge, Dohyun Kwon, Adam Azzam, Ian Zemke, Jacob Rooney, Ian Charlesworth, Laure Flapan, Matt Stoffregen, Joe Hughes, Yehonatan Sella, John Susice, Ian Coley, Chris Ohrt, Sangchul Lee, Will Oakley, Bon-Soon Lin, Yacoub Kureh, Stephanie Wang, Zane Li, Angela Wu, Dimitrios Ntalampekos, Ravi Bajaj, Mike Miller, Flavien Moreau, Tom Neiser, Andrea Trivella, Allic Sivaramakrishnan and Josh Samani for making UCLA vibrant and fun.

I spend a lot of time working at coffee shops. I thank Sébastian Demesse and Luke Fleishman for making Anticafé and Espresso Profeta awesome places to work and hang out. Ben Krause, Michael Hughes, Henry Chan and Dustin Chen are largely responsible for keeping me sane in the last five years. I am so thankful for their friendship.

I'm grateful for the friend and partner I've found in Sarah Christofersen, I thank her for her love and support. Lastly, I thank my parents and brothers, who I look up to, and to whom this thesis is dedicated. Work I've been able to do has been largely due to their belief in me, their support and friendship. It's impossible to pay them back, so I just promise to pay it forward.

VITA

2012	B.S. (Mathematics), University of California, Davis
2012	Outstanding Senior in Mathematics, UCD Mathematics Department
2016-2017	Dissertation Year Fellowship, UCLA
2017	Departmental Teaching Award

PUBLICATIONS

J. Gold. A bound for orderings of Reidemeister moves, Algebraic & Geometric Topology 13 (2013), 3099-3110.

J. Gold. Isoperimetry in supercritical bond percolation in dimensions three and higher. Submitted.

J. Gold. Intrinsic isoperimetry of the giant component of supercritical bond percolation in dimension two. Submitted.

CHAPTER 1

Introduction

How does disorder at small spatial scales manifest at the scale in which we live? Such a question invokes statistical physics, a subject concerned in part with deriving large scale (*macroscopic*) properties of a physical system from information at small (*microscopic*) scales. Probability is a useful tool in rigorous statistical physics because disorder is naturally modeled by randomness. In turn, much of current research in probability is driven by ideas and models coming from statistical physics. A corresponding mathematical goal is then to try to understand how randomness at small spatial scales manifests at large spatial scales within a menagerie of probabilistic models.

In many cases, small scale randomness gets averaged out or *homogenized* at large scales, and a deterministic object emerges as a result. The study of this behavior falls in the domain of homogenization theory. We give a few examples of problems studied in this area below.

(1) *Contact angles of droplets:* Consider a large liquid droplet clinging to a rough solid surface. The precise angle at which the droplet meets the surface is determined in part by the microscopic roughness of the latter object. Upon zooming out, one begins to see a macroscopic or homogenized contact angle between droplet and surface, reflecting an *effective* surface energy between distinct phases which is accurate at large scales. We give this example first to emphasize that homogenization is not a subfield of probability: the microscopic roughness may be described by a deterministic function. We will see that despite being non-random, this example possesses extreme relevance to our work.

(2) *Spectra of random Laplacians and random PDE:* One can study the Dirichlet Laplacian associated to the random weights on $\mathbb{Z}^d$ in a large domain. These are large random matrices, and of interest is the behavior of the associated sequence of eigenvectors and eigenvalues. Moment conditions on the weights determine whether these objects scale to their counterparts for a deterministic continuum Laplacian. More generally, one can study discrete PDE with random coefficients, the goal being to discover a deterministic PDE in the continuum reflecting the behavior of the discrete equations at large spatial scales.

The rescaling of space links the phenomenon of homogenization to the notion of a *scaling limit*. To study a sequence of random objects defined over a discrete structure, one can fix the typical scale of the objects in question by rescaling the underlying discrete structure. For instance, if we are looking at random objects defined over a lattice, we may rescale the lattice by taking the spacing between neighboring sites to zero. Limiting objects obtained in this way are called scaling limits.

Scaling limits need not be deterministic: the limiting object may retain the randomness present at small scales and become richer for it. We give examples of scaling limits below. [1]

(3) *Brownian motion and super-Brownian motion:* The canonical example of a scaling limit is given by the invariance principle, which tells us that any random walk on $\mathbb{Z}^d$ with mean zero and square-integrable steps will tend, upon rescaling time and space, to a Brownian motion in $\mathbb{R}^d$. More generally, a branching random walk (with balanced birth and death rate) can be shown to scale to a measure-valued process known as super-Brownian motion. Both continuum objects provide rigorous descriptions of diffusing particles.

(4) *SLE from walks:* The loop-erased random walk was introduced to study the self-avoiding random walk, a model of a polymer in a good solvent. The scaling limit of the loop-erased random walk in $\mathbb{Z}^2$ was shown to be SLE_2, where SLE_κ is a special family of random planar curves exhibiting conformal invariance. Though the scaling limit of the self-avoiding walk in $\mathbb{Z}^2$ is unknown, it is conjectured to be $SLE_{8/3}$, and in general it is known to be distinct from the scaling limit of the loop-erased walk.

For models with several layers of randomness, it is possible for some aspects of the scaling

[1]We find it important to distinguish the notions of scaling limits and thermodynamic limits. For our purposes, the latter notion begins with probabilistic models defined over a sequence of finite discrete structures. A thermodynamic limit is then obtained by extracting a limiting probabilistic model over a corresponding limiting infinite discrete structure. Thus, when working with a probabilistic model defined over a full Euclidean lattice, one has already taken a thermodynamic limit.

limit to become homogenized and for randomness to persist in other aspects.

(5) *Random walks in a random environement:* Returning to the setup of **(2)**, we attach random edge weights to each edge of $\mathbb{Z}^d$ and form a random walk whose local movements are governed by the weights of incident edges. The discrete Laplacian in **(2)** is then the generator for this random walk in a random environment. For edge weights which are sufficiently tame, one can show this random walk in a random environment scales to a Brownian motion in $\mathbb{R}^d$ whose covariance structure is deterministic but depends on the law of the edge weights. Here, the microscopic movements of the random walk manifest as a random object (a Brownian motion), while the random environment becomes homogenized at large scales, corresponding to the deterministic covariance structure of the Brownian motion.

Of course, our goal is not to be comprehensive or to draw sharp boundaries between subfields. Rather we wish to give the reader a taste of some interesting behavior exhibited by discrete, spatial probabilistic models upon "zooming out" to large spatial scales.

Let us now turn to models closely related to our work: we consider bond percolation on the Euclidean lattice $\mathbb{Z}^d$ for $d \geq 2$. In the supercritical regime, we condition on the rare event that the origin lies within a large and finite open cluster. Such large clusters are like droplets, and it is natural to wonder if these droplets have an asymptotic *shape*. Implicit is the idea that we will rescale progressively larger droplets so that they are of the same size by taking the lattice spacing towards zero. Thus, we are interested in the scaling limits of large finite open clusters. Also implicit is that the asymptotic shape of these clusters is deterministic, which links this problem to homogenization theory.

It has taken a substantial amount of work to show that such a limit shape (in each dimension) exists. The method of proof uses a tool known as the Wulff construction, which begins by extracting a homogenized surface energy from the model. The surface energy is a function τ on the unit sphere which intuitively tells us how many bonds are used to form a large discrete interface in a given direction. We think of the boundary of the finite cluster as a full interface enveloping the droplet, and we can approximate the size of the interface by computing the surface energy (with respect to τ) of a continuum object with similar shape.

Because we are within an event of small probability, it comes as no surprise that the theory of large deviations plays a role. One shows that the probability of seeing a droplet in a certain shape is related to the τ-surface energy of this shape, so that shapes with smaller surface energy are more likely. The theory of large deviations tells us that the unlikely event that there is a large finite cluster will be achieved in the most likely way – the large finite cluster will roughly look like a shape with least possible surface energy.

It remains to classify this shape, but the solution is well-known: τ is in fact a norm on $\mathbb{R}^d$, and the optimal shape is simply the unit ball of the dual norm. This is the limit shape. We remark that there is also homogenization occurring in the bulk of the droplet: though the density is random, it concentrates around a fixed constant for large droplets.

This dissertation is devoted to studying a different class of droplets occurring in supercritical percolation. Within the infinite cluster, we look at subgraphs of the infinite cluster of a given size with minimal boundary to volume ratio. Such sets are *isoperimetric*. Our goal is, as above, to extract a deterministic limit shape as a scaling limit of these objects. We also use the Wulff construction, though there are key differences in our approach which are discussed in the introductions of the following papers:

1. "Isoperimetry in supercritical bond percolation in dimensions three and higher",
2. "Intrinsic isoperimetry of the giant component of supercritical bond percolation in dimension two".

The two papers make up the entirety of the dissertation, which is in spirit an investigation into the shapes of random crystals.

1.1 Outline of dissertation

The two papers are presented below largely unchanged, aside from new formatting. Each is self contained, with Chapters 2 through 10 making up the first paper and Chapters 11 through 16 making up the second. Both papers are similar in structure, as both essentially follow the Wulff construction program developed for studying droplets in lattice models. The structure of the argument can roughly be decomposed into three parts:

1. Constructing a surface energy,

2. Passing from continuous objects to discrete objects,

3. Passing from discrete objects to continuous objects.

We are intentionally vague, as more detailed outlines of each argument are given in the introductions found in Chapters 2 and 11 (see Section 2.4 and Section 11.4). Still, we provide a rough sketch below.

Chapter 3 sets up some notation for the first paper. Chapters 4 and 5 construct the surface energy and develop concentration estimates. Chapter 6 passes from continuous objects to discrete objects, while Chapters 7 and 8 accomplish the considerably harder task of going in the other direction. Chapter 9 puts everything together, yielding the main results of the paper, and Chapter 10 is an appendix.

As far as the second paper is concerned, Chapters 12, 14 and 15 respectively handle the three items above. As mentioned in the abstract, the second paper differs from the first in part because the continuum isoperimetric problem is more complicated, so Chapter 13 is devoted to an analysis of this problem, in which we deduce necessary stability estimates. Chapter 16 is a much smaller appendix for the second paper.

For the convenience of the reader, we remark that the main results of the first paper are Theorem 2.2.1 and Theorem 2.2.2, and that the main results of the second paper are Theorem 11.3.1 and Theorem 11.3.3.

CHAPTER 2

Isoperimetry in supercritical bond percolation in dimensions three and higher

2.1 Motivation

Isoperimetric problems, namely the problem of finding a set of given size and minimal boundary measure, have been studied for millennia [Bl05]. In the continuum, such problems are the subject of geometric measure theory and the calculus of variations. Isoperimetric inequalities give a lower bound on the boundary measure of a set in terms of the volume measure of the set. Their applications in mathematics range from concentration of measure to PDE theory.

Isoperimetric problems are also well-studied in the discrete setting. One can encode isoperimetric inequalities for graphs in the *Cheeger constant*, or modifications thereof. For a graph G, define the Cheeger constant of G to be

$$\Phi_G := \min \left\{ \frac{|\partial_G H|}{|H|} : H \subset G, 0 < |H| \leq |G|/2 \right\} , \tag{2.1}$$

where $\partial_G H$ is the edge boundary of H in the graph G and where $|H|$ and $|G|$ respectively denote cardinalities of the vertex sets of H and G. This constant was originally introduced for manifolds in Cheeger's thesis [Che70], in which the Cheeger constant was used to give a lower bound on the smallest positive eigenvalue of the negative Laplacian. Its discrete analogue, introduced by Alon [Alo86], plays a similar role in spectral graph theory (see for instance Chapter 2 of [Chu97]). Indeed, Cheeger's inequality and its variants are used to study mixing times of random walks and Markov chains. Ultimately, the Cheeger constant provides one of many ways to study the geometry of a graph.

Broadly, the goal of this paper is to explore the geometry of random graphs arising from bond percolation on $\mathbb{Z}^d$. Specifically, we view $\mathbb{Z}^d$ as a graph, with edge set $\mathrm{E}(\mathbb{Z}^d)$ determined by nearest-neighbor pairs, and we form the probability space $(\{0,1\}^{\mathrm{E}(\mathbb{Z}^d)}, \mathscr{F}, \mathbb{P}_p)$, where $\mathscr{F}$ denotes the product σ-algebra on $\{0,1\}^{\mathrm{E}(\mathbb{Z}^d)}$ and where $\mathbb{P}_p$ is the product Bernoulli measure associated to the *percolation parameter* $p \in [0,1]$. Elements $\omega = (\omega_e)_{e\in\mathrm{E}(\mathbb{Z}^d)}$ of our probability space are referred to as *percolation configurations*. We say that an edge $e \in \mathrm{E}(\mathbb{Z}^d)$ is *open* in the configuration ω if $\omega_e = 1$; we say that an edge is *closed* otherwise. The collection of open edges determine a random subgraph of $\mathbb{Z}^d$; the connected components of this subgraph are referred to as *open clusters*. It is well-known (see Grimmett [Gri99] for details) that for $d \geq 2$, bond percolation exhibits a phase transition. That is, there exists a $p_c(d) \in (0,1)$ such that whenever $p > p_c(d)$, there exists a unique infinite open cluster $\mathbb{P}_p$-almost surely, and whenever $p < p_c(d)$, there is no infinite open cluster $\mathbb{P}_p$-almost surely. We work in the supercritical regime, and we denote the unique infinite (open) cluster by $\mathbf{C}_\infty$.

We may now be more specific: our goal is to explore the geometry of $\mathbf{C}_\infty$. There are many ways to do this, for example, one can study the asymptotic graph distance in $\mathbf{C}_\infty$ (e.g. Antal and Pisztora [AP96]), the asymptotic shapes of balls in the graph distance metric of $\mathbf{C}_\infty$ (e.g. Cox and Durrett [CD81]), or the effective resistance of $\mathbf{C}_\infty$ within a large box (e.g. Grimmett and Kesten [GK84]). We study the isoperimetry of $\mathbf{C}_\infty$ through the Cheeger constant.

By definition, $\Phi_G = 0$ for any amenable graph, and one can show that $\Phi_{\mathbf{C}_\infty} = 0$ almost surely. We will instead study the Cheeger constant of $\mathbf{C}_n := \mathbf{C}_\infty \cap [-n,n]^d$. Let $\widetilde{\mathbf{C}}_n$ be the largest connected component of $\mathbf{C}_n$. It is known (Benjamini and Mossel [BM03], Mathieu and Remy [MR04], Rau [Rau07], Berger, Biskup, Hoffman and Kozma [BBH08] and Pete [Pet08]) that $\Phi_{\widetilde{\mathbf{C}}_n} \asymp n^{-1}$ as $n \to \infty$, prompting the following conjecture of Benjamini.

Conjecture 2.1.1. For $p > p_c(d)$ and $d \geq 2$, the limit

$$\lim_{n\to\infty} n\Phi_{\widetilde{\mathbf{C}}_n} \tag{2.2}$$

exists $\mathbb{P}_p$-almost surely and is a positive deterministic constant.

Procaccia and Rosenthal [PR12] made progress towards resolving this conjecture: they proved upper bounds on the variance of the Cheeger constant, showing $\mathrm{Var}(n\Phi_{\widetilde{\mathbf{C}}_n}) \leq cn^{2-d}$

for some positive $c = c(p,d)$. Recently, Biskup, Louidor, Procaccia and Rosenthal [BLP15] settled this conjecture positively for a natural modification of $\Phi_{\widetilde{\mathbf{C}}_n}$ in the case $d = 2$. We define the *modified Cheeger constant* $\widehat{\Phi}_n$ of $\mathbf{C}_n$ for dimension $d \geq 2$ to be

$$\widehat{\Phi}_n := \min\left\{ \frac{|\partial_{\mathbf{C}_\infty} H|}{|H|} : H \subset \mathbf{C}_n, 0 < |H| \leq |\mathbf{C}_n|/d! \right\}, \tag{2.3}$$

where $\partial_{\mathbf{C}_\infty} H$ denotes the open edge boundary of H within *all of* $\mathbf{C}_\infty$ as opposed to $\mathbf{C}_n$. Thanks to (for instance) Proposition 1.2 of [BM03], the asymptotics of $\widehat{\Phi}_n$ are unchanged whether we use $\mathbf{C}_n$ or $\widetilde{\mathbf{C}}_n$ in (2.3).

This modification is natural in the sense that a candidate subgraph H is treated as living within $\mathbf{C}_\infty$, and the $d!$ in the upper volume bound ensures that H need not touch the boundary of the box. Both $\Phi_{\widetilde{\mathbf{C}}_n}$ and $\widehat{\Phi}_n$ are closely related to the so-called *anchored isoperimetric profile*, defined in the context of the infinite cluster as

$$\Phi_{\mathbf{C}_\infty,0}(n) := \inf\left\{ \frac{|\partial_{\mathbf{C}_\infty} H|}{|H|} : 0 \in H \subset \mathbf{C}_\infty, H \text{ connected}, 0 < |H| \leq n \right\}, \tag{2.4}$$

where of course we must condition on the positive probability event $\{0 \in \mathbf{C}_\infty\}$. In [BLP15], the analogue of Benjamini's conjecture for the anchored isoperimetric profile was also established in dimension two. Moreover, the subgraphs of $\mathbf{C}_n$ and $\mathbf{C}_\infty$ achieving each minimum were studied in both cases, and in fact were shown to scale uniformly to the same deterministic limit shape. This latter *shape theorem* implies the existence of the limit in Conjecture 2.1.1 for (2.3), indeed, the perimeter of this limit shape appears in the limiting value of the modified Cheeger constant.

2.2 Results

We extend the work of [BLP15] to the setting $d \geq 3$, settling Benjamini's conjecture for the modified Cheeger constant and proving a shape theorem for isoperimetric subgraphs of $\mathbf{C}_n$. As the arguments in [BLP15] rely heavily on planar geometry and graph duality, a much different approach is needed. Nevertheless, we share a common starting point with the Wulff construction, described below, and there are similarities between the overall structure of the argument presented here and the argument of [BLP15]. We state the main theorem of the paper first. For each n, let $\mathcal{G}_n$ be the (random) collection of subgraphs of $\mathbf{C}_n$ which realize the minimum $\widehat{\Phi}_n$. For $A \subset \mathbb{R}^d$, $r > 0$ and $x \in \mathbb{R}^d$ the sets rA and $x + A$ are defined as usual by

$$rA := \Big\{ra \ : a \in A\Big\}, \quad x + A := \Big\{x + a \ : a \in A\Big\}, \tag{2.5}$$

and we write $|| \cdot ||_{\ell^1}$ to denote the ℓ^1-norm of a function on $\mathbb{Z}^d$. Here is our main result:

Theorem 2.2.1. Let $d \geq 3$ and $p > p_c(d)$. There exists a convex set $W_{p,d} \subset [-1,1]^d$ such that

$$\max_{G_n \in \mathcal{G}_n} \inf_{x \in \mathbb{R}^d} n^{-d} \Big\| \mathbf{1}_{G_n} - \mathbf{1}_{\mathbf{C}_n \cap (x + nW_{p,d})} \Big\|_{\ell^1} \xrightarrow[n \to \infty]{} 0 \tag{2.6}$$

holds $\mathbb{P}_p$-almost surely.

Following [BLP15], we build the limit shape $W_{p,d}$ through what is known as the Wulff construction. This is a method for solving anisotropic isoperimetric problems, first introduced by Wulff [Wul01] in 1901. Given a norm τ on $\mathbb{R}^d$, one can form the associated isoperimetric problem, which we state in the Lipschitz setting:

$$\text{minimize } \frac{\mathcal{I}_\tau(E)}{\mathcal{L}^d(E)} \quad \text{subject to } \mathcal{L}^d(E) \leq 1, \tag{2.7}$$

where the minimum runs over $E \subset \mathbb{R}^d$ with Lipschitz boundary, where $\mathcal{L}^d$ denotes d-dimensional Lebesgue measure, and where $\mathcal{I}_\tau(E)$ is defined as

$$\mathcal{I}_\tau(E) := \int_{\partial E} \tau(v_E(x)) \mathcal{H}^{d-1}(dx) \,. \tag{2.8}$$

Here $\mathcal{H}^{d-1}$ is the $(d-1)$-dimensional Hausdorff measure on ∂E and $v_E(x)$ the unit exterior normal to E at the point $x \in \partial E$, which is defined for $\mathcal{H}^{d-1}$-almost every point of ∂E. Wulff's candidate isoperimetric set is constructed as the following intersection of half-spaces:

$$\widehat{W}_\tau := \bigcap_{v \in \mathbb{S}^{d-1}} \Big\{ x \in \mathbb{R}^d \ : x \cdot v \leq \tau(v) \Big\}, \tag{2.9}$$

where $\cdot$ denotes the standard dot product. We call $\widehat{W}_\tau$ the *unit Wulff crystal* associated to τ; this object is the unit ball in the norm τ' dual to τ (recall that τ' is defined on $y \in \mathbb{R}^d$ by

$\tau'(y) = \sup\{x \cdot y : x \in \mathbb{R}^d, \tau(x) \leq 1\})$. When $\widehat{W}_\tau$ is scaled to have unit volume, it becomes a candidate for (2.7). It was Taylor [Tay74] who ultimately proved that this rescaled shape is optimal within a wide class of Borel sets, and moreover (in [Tay75]) that this rescaled shape is the unique optimizer up to translations and modifications on a null set.

In the following section, and at the beginning of Chapter 4, we will observe that a norm naturally emerges when our problem viewed in the correct context. This norm, denoted $\beta_{p,d}$, is first defined on the unit sphere in $\mathbb{R}^d$: in a given direction $v \in \mathbb{S}^{d-1}$, we first rotate a large cube so that its top and bottom faces are normal to v. We intersect this cube with $\mathbb{Z}^d$, and the percolation configuration on $\mathbb{Z}^d$ restricts naturally to the discretization of the cube. We then consider the minimum size of a cutset separating the top and bottom faces of the cube in this percolated graph. We require these cutsets to be anchored near the middle of the cube, so that after taking expectations and dividing by the area of a face of the cube, we may employ a subadditivity argument to extract a limit as the diameter of the cube tends to infinity. This limit is the value of $\beta_{p,d}$ in the direction v.

We construct $\beta_{p,d}$ in Chapter 4, and we define the *Wulff crystal* $W_{p,d}$ to be the dilate of the unit Wulff crystal $\widehat{W}_{p,d}$ associated via (2.9) to $\beta_{p,d}$ so that $\mathcal{L}^d(W_{p,d}) = 2^d/d!$. The Wulff crystal is then the limit shape from Theorem 2.2.1, and we note that the norm $\beta_{p,d}$ gives rise to a functional of the form (2.8) which we write as $\mathcal{I}_{p,d}$ and which we refer to as the *surface energy*. As in [BLP15], the shape theorem we present is intimately linked with the limiting value of the Cheeger constant. Let $\theta_p(d) := \mathbb{P}_p(0 \in \mathbf{C}_\infty)$ be the density of the infinite cluster within $\mathbb{Z}^d$.

Theorem 2.2.2. Let $d \geq 3$, $p > p_c(d)$ and let $\beta_{p,d}$ be the norm defined in Proposition 4.2.3. Let $W_{p,d}$ be the Wulff crystal for this norm, that is, the ball in the dual norm $\beta'_{p,d}$ such that $\mathcal{L}^d(W_{p,d}) = 2^d/d!$. Then,

$$\lim_{n\to\infty} n\widehat{\Phi}_n = \frac{\mathcal{I}_{p,d}(W_{p,d})}{\theta_p(d)\mathcal{L}^d(W_{p,d})} \tag{2.10}$$

holds $\mathbb{P}_p$-almost surely.

2.3 History and discussion

Within the last thirty years, the Wulff construction has grown into an important tool in the rigorous analysis of equilibrium crystal shapes. Such problems are concerned with understanding the macroscopic behavior of one phase of matter immersed within another.

The present work fits into this paradigm in that we may regard each Cheeger optimizer G_n as a large droplet of a *crystalline* phase within $\mathbf{C}_\infty \setminus G_n$, regarded as the *ambient* phase. The value of the norm $\beta_{p,d}$ in a given direction represents the energy required to form a flat interface between the two phases in this direction, and gives rise to a surface energy functional of the form (2.8). It was Gibbs [Gib78] who postulated that, in general, the asymptotic shape of the crystalline phase should minimize this surface energy. The Wulff construction then furnishes this minimal shape.

The spirit of Theorem 2.2.1 can be traced back to the work of Milnos and Sinai [MS67, MS68] from the 1960s, in which the geometric properties of phase separation in a material are rigorously studied. The first rigorous characterizations of phase separation via the Wulff construction are due independently to Dobrushin, Kotecký and Shlosman [DKS92] in the context of the two-dimensional Ising model and to Alexander, Chayes and Chayes [ACC90] in the context of two-dimensional bond percolation. The results of [DKS92], valid in the low-temperature regime, were extended up to the critical temperature thanks to the work of Ioffe [Iof95] and Ioffe and Schonmann [IS98].

The first rigorous derivation of the Wulff construction for a genuine short-range model in three dimensions was achieved by Cerf in the context of bond percolation [Cer00]. Analogous results for the Ising model and in higher dimensions were achieved in several substantial works of Bodineau [Bod99, Bod02] and Cerf and Pisztora [CP00, CP01]. The coarse graining results of Pisztora [Pis96] played an integral role in this study of the Ising model, FK percolation and bond percolation in higher dimensions. A comprehensive survey of these results and of others can be found in Section 5.5 of Cerf's monograph [Cer06] and in the review article of Bodineau, Ioffe and Velenik [BIV00].

In all cases, the jump to dimensions strictly larger than two has, at least so far, necessitated a shift from the uniform topology to the ℓ^1 topology on the space of shapes (we are intentionally vague about which space we consider). Indeed, the variational problem (2.7) is not stable in $d \geq 3$ when the space of shapes is equipped with the uniform topology. That is, in $d \geq 3$, it is possible to construct sequence of shapes which are bounded away from the optimal shape in the uniform topology, but whose surface energies tend to the optimal surface energy. This has implications at the microscopic level; if one desired to prove a uniform shape theorem in $d \geq 3$ for the Cheeger optimizers, one would first have to rule out the existence of long but thin filaments (as in Figure 2.1) in these discrete objects with high probability.

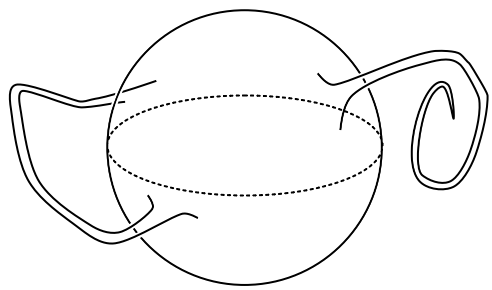

Figure 2.1: In $d = 3$, filaments added to the optimal shape for the Euclidean isoperimetric problem produce a set which is almost optimal and and which has large uniform distance to the sphere.

This lack of regularity at the microscopic level requires that we consider the variational problem over a wider class of shapes, and it is here that geometric measure theory emerges as a valuable tool, as first realized by Alberti, Bellettini, Cassandro and Presutti [ABC96].

2.4 Outline

Our goals may be summarized as follows: we wish to show that the sequence of discrete, random isoperimetric problems (2.3) scale to a continuous, deterministic isoperimetric problem (2.7) corresponding to some norm $\beta_{p,d}$ on $\mathbb{R}^d$. This could be phrased in terms of Γ-convergence, and indeed, this language has been used in some of the results outlined in our historical discussion, as well as in recent related work of Braides and Piatnitski [BP12, BP13].

The first task is to construct a suitable norm $\beta_{p,d}$ on $\mathbb{R}^d$. This is done in Chapter 4 after introducing some definitions and notation in Chapter 3. The key to the existence of $\beta_{p,d}$ is a spatial subadditivity argument applied to the geometric setting described briefly before Theorem 2.2.2.

The resulting norm $\beta_{p,d}$ gives rise to a surface energy $\mathcal{I}_{p,d}$, and the remainder of the paper is concerned with demonstrating that the unique optimizer of the isoperimetric problem associated to $\mathcal{I}_{p,d}$ faithfully describes the macroscopic shape of each large $G_n \in \mathcal{G}_n$. We must show a strong correspondence between discrete objects (the various subgraphs of $\mathbf{C}_n$) and continuous objects (Borel subsets of $[-1,1]^d$ for which isoperimetric problems can be defined). Specifically, this correspondence should link the isoperimetric ratio of subgraphs of $\mathbf{C}_n$ to the corresponding ratio for continuous objects, as in the limiting value of Theorem 2.2.2.

In order to pass from continuous objects to discrete objects, we must first prove concentration estimates for the random variables used to define $\beta_{p,d}$. This is in line with [BLP15], and is in contrast to large deviation methods used in some of the works mentioned in Section 2.3, where the nature of these earlier problems requires working within events of small probability. We prove concentration estimates in Chapter 5.

Consequences of these concentration estimates are presented in Chapter 6, where we use the results of Chapter 5 to give a high probability upper bound on $n\widehat{\Phi}_n$. To obtain this upper bound, we intersect a large polytope P with $\mathbf{C}_n$ to produce a subgraph H_n of $\mathbf{C}_n$. We then exhibit control on the volume of H_n which ensures that H_n meets the criteria of (2.3). Our concentration estimates relate the size of the open edge boundary of H_n to the surface energy of the original polytope. At the end of Chapter 6, we will have proved half of Theorem 2.2.2. All of the arguments presented up to this point work in the setting $d \geq 2$.

Passing from discrete objects to continuous objects is more delicate. The main difficulty is to construct a suitable continuum object P_n (for instance, a polytope in $[-1,1]^d$) for each $G_n \in \mathcal{G}_n$ so that the dilate nP_n has a similar isoperimetric ratio to that of G_n. In particular, we need the perimeters of these P_n to stay uniformly bounded in n. A natural first guess for P_n to take the union of unit cubes centered at all vertices of G_n and to scale this set by a factor of n^{-1}. However, due to percolation of closed edges near $p_c(d)$, we do not have control on the perimeters of such P_n unless p is very close to one.

This suggests a renormalization argument, which we introduce in Chapter 7. We base our argument on a construction due to Zhang from [Zha07], but we must modify this construction and study it carefully in order to apply it to our situation. It is here that, for reasons which will be made clear in Chapter 7, we must restrict ourselves to the setting $d \geq 3$. This is no loss as the case $d = 2$ is covered by the results in [BLP15].

In Chapter 8, we reap the efforts of Chapter 7, passing from $G_n \in \mathcal{G}_n$ to sets of finite perimeter (defined in Chapter 3). Such sets have just enough regularity that we may work locally on their boundaries. We use this feature in Chapter 9 to show that whenever a G_n is close (in the appropriate sense) to a set of finite perimeter, the surface energy of this set is roughly a lower bound on the open edge boundary of G_n. Our notion of closeness also allows us to relate the volumes of these discrete and continuous objects; we may then deduce that whenever G_n is close to a set of finite perimeter, the isoperimetric ratio of G_n (hence the Cheeger constant) is controlled from below by the isoperimetric ratio of the given continuum set.

We then invoke the results of Chapter 6 and the work of Taylor [Tay74, Tay75] to see that with high probability, each G_n must be close to the Wulff crystal. This gives Theorem 2.2.1 and Theorem 2.2.2 in quick succession.

2.5 Open problems

We now pose several open questions, some of which were originally stated in [BLP15].

(1) *Free boundary conditions and more general domains:* Conjecture 2.1.1 is still open for the unmodified Cheeger constant. We expect that it is possible to adapt the approach of [BLP15] to resolve this in $d = 2$. It is not obvious what the limit shape should be in this case, or even whether it is unique. One conjecture is that an optimal shape will be a rescaled quarter-Wulff crystal in one of the corners of the square. This conjecture is motivated by the Winterbottom construction introduced in [Win67], which is an analogue of the Wulff construction for crystals in the presence of a wall. This construction has been used successfully in the two-dimensional Ising model by Pfister and Velenik [PV97, PV96] and in higher dimensions by Bodineau, Ioffe and Velenik [BIV01].

One can generalize Benajmini's conjecture in the two-dimensional setting to domains other than boxes; given a nice bounded open set $\Omega \subset \mathbb{R}^2$, one can study the asymptotics of the unmodified Cheeger constant as well as the shapes of the Cheeger optimizers for the largest connected component of $\mathbf{C}_\infty \cap n\Omega$. Results characterizing the limiting value of the Cheeger constant or the limiting shapes of the optimizers would be isoperimetric analogues of the work of Cerf and Théret [CT12] on minimal cutsets (in $d \geq 2$).

***(2)** More information on the Wulff crystal:* Little is known about the geometric properties of the Wulff crystal. One recent result of Garet, Marchand, Procaccia and Théret [GMP15] is that, in two dimensions, the Wulff crystal varies continuously with respect to the uniform metric on compact sets as a function of the percolation parameter $p \in (p_c(2), 1]$. It was conjectured in [BLP15] that the two-dimensional Wulff crystal tends to a Euclidean ball as $p \downarrow p_c(2)$; this is still widely open. It is natural to ask whether the Wulff crystal has facets (open portions of the boundary with zero curvature) or corners, and how such questions depend on the percolation parameter.

***(3)** Uniform convergence for $d \geq 3$:* An interesting and challenging question is whether a form of Theorem 2.2.1 holds in $d \geq 3$ when we replace ℓ^1 convergence by uniform convergence. Despite the complications we have described, Dobrushin, Kotecký and Shlosman [DKS92] were optimistic that filaments could be removed in the context of the Ising model in dimensions greater than two.

2.6 Acknowledgements

I thank my advisor Marek Biskup for suggesting this problem, for his guidance and his support. I am deeply indebted to Raphaël Cerf for patiently sharing his expertise and insight during my time in Paris. I likewise thank Eviatar Procaccia for his guidance. I am grateful to Vincent Vargas, Claire Berenger and Shannon Starr for making it possible for me to attend the IHP Disordered Systems trimester. Finally, I thank Yoshihiro Abe, Ian Charlesworth, Arko Chatterjee, Hugo Duminil-Copin, Aukosh Jagannath, Ben Krause, Sangchul Lee, Tom Liggett, Jeff Lin, Peter Petersen, Jacob Rooney and Ian Zemke for helpful conversations. This research has been partially supported by the NSF grant DMS-1407558.

CHAPTER 3

Definitions and notation

3.1 Graphs

Throughout this paper, we work within the graph $\mathbb{Z}^d$. The vertex set $\mathrm{V}(\mathbb{Z}^d)$ of $\mathbb{Z}^d$ is the set of integer d-tuples. There is an edge between two d-tuples if, when viewed as vectors in $\mathbb{R}^d$, they have Euclidean distance one, and we denote the edge set of $\mathbb{Z}^d$ as $\mathrm{E}(\mathbb{Z}^d)$. Edges have no orientation. If x is a vertex in the graph $\mathbb{Z}^d$, we will often write $x \in \mathbb{Z}^d$ in place of $x \in \mathrm{V}(\mathbb{Z}^d)$, and similarly for edges. The same convention will be used for subgraphs of $\mathbb{Z}^d$.

If vertices $x, y \in \mathbb{Z}^d$ share an edge, we say they are adjacent and we write $x \sim y$. Given two vertices $x, y \in \mathbb{Z}^d$, a *path from x to y* is a finite, alternating sequence of vertices and edges $x_0, e_1, x_1, \ldots, e_m, x_m$ such that e_i is the edge shared by x_{i-1} and x_i for $i = 1, \ldots, m$, and such that $x_0 = x$ and $x_m = y$. We say the path *joins* the vertices x and y, and the *length* of the path is the number of edges m in this sequence. A subgraph $G \subset \mathbb{Z}^d$ is *connected* if for any two vertices $x, y \in G$, there is a path using only vertices and edges in G which joins x and y.

Given $x \in \mathbb{Z}^d$, a *path from x to ∞* is an infinite alternating sequence of vertices and edges $x_0, e_1, x_1, \ldots$ such that $x = x_0$ and such that no finite box contains all edges in this path. A path is said to be *simple* if it does not use any vertex more than once. In either case, we often regard paths as sequences of edges out of convenience. We now define several useful notions of graph boundaries. Given G a finite subgraph of $\mathbb{Z}^d$, we define the *edge boundary* and *outer edge boundary* of G to respectively be

$$\partial G := \Big\{ e \in \mathrm{E}(\mathbb{Z}^d) \ : \ \text{exactly one endpoint of } e \text{ lies in } G \Big\}, \tag{3.1}$$

$$\partial_o G := \left\{ e \in \partial G \ : \ \begin{array}{c} \text{the endpoint of } e \text{ in } G \text{ is connected to } \infty \\ \text{via a path which uses no other vertices of } G \end{array} \right\}. \tag{3.2}$$

Note that ∂G and $\partial_o G$ are sets of edges. It will also be necessary to work with the *vertex boundary* of G, defined as

$$\partial_* G := \Big\{ v \in \mathrm{V}(G) \ : \ v \text{ is an endpoint of an edge in } \partial_o G \Big\}. \tag{3.3}$$

Given a finite subgraph $G \subset \mathbb{Z}^d$, a *cutset separating G from ∞* is a finite collection of edges $S \subset \mathrm{E}(\mathbb{Z}^d)$ such any path from a vertex of G to ∞ must use an edge in the set S. If $A, B \subset \mathrm{V}(G)$ are disjoint vertex sets, a *cutset separating A and B* is a finite collection of edges $S \subset \mathrm{E}(G)$ such that any path from a vertex of A to a vertex of B must use an edge of S. In either case, a cutset S is said to be *minimal* if it is no longer a cutset upon removing any edge in S.

We define $\mathbb{L}^d$ to be the graph with vertex set $\mathrm{V}(\mathbb{Z}^d)$ and edge set consisting of pairs of vertices $x, y \in \mathbb{Z}^d$ which, when viewed as vectors in $\mathbb{R}^d$, have ℓ^∞-distance one. If $x, y \in \mathbb{Z}^d$ are joined by an edge in $\mathbb{L}^d$, we say the two vertices are $\mathbb{L}^d$-*adjacent* and write $x \sim_{\mathbb{L}} y$. We define $\mathbb{L}^d$-*paths* analogously to paths in $\mathbb{Z}^d$, and we say that a subgraph $G \subset \mathbb{Z}^d$ is $\mathbb{L}^d$-*connected* if any two vertices $x, y \in G$ are joined by an $\mathbb{L}^d$-path. The following proposition, which is standard in the literature, provides a link between cutsets in $\mathbb{Z}^d$ and the notion of $\mathbb{L}^d$-connectivity. A proof of this proposition may be found in Deuschel and Pisztora [DP96], Lemma 2.1. More recently, Timár [Tim13] has given a concise and more combinatorial proof of a stronger statement.

Proposition 3.1.1. Let $G \subset \mathbb{Z}^d$ be a finite, connected subgraph of $\mathbb{Z}^d$. Then the vertex boundary $\partial_* G$ is $\mathbb{L}^d$-connected, as is the set of vertices which are endpoints of edges in $\partial_o G$.

This result is fundamental to the execution of Peierls estimates (used in conjunction with bounds as in Proposition 10.3.2) which appear frequently in the study of lattice models.

For G a subgraph of $\mathbb{Z}^d$ and $K \subset \mathbb{R}^d$ compact, we will often write $G \cap K$ as shorthand for $\mathrm{V}(G) \cap K$. The vertex set $\mathrm{V}(G) \cap K$ inherits a graph structure from G, so that when referring to $e \in G \cap K$ for an edge e, it is understood that both endpoints of e lie in $\mathrm{V}(G) \cap K$. Finally, if G is a finite set, we use $|G|$ to denote the cardinality of G. If G is a finite subgraph of $\mathbb{Z}^d$, we write $|G|$ in place of the cardinality of $\mathrm{V}(G)$.

3.2 Percolation

The probabilistic setting of this paper is bond percolation on $\mathbb{Z}^d$ with $d \geq 2$. We have already defined this model in the introduction; here we introduce a bit more terminology. Percolation gives rise to another notion of graph boundary: if G is a finite subgraph of $\mathbf{C}_\infty$, we define the *open edge boundary* of G as

$$\partial^\omega G := \Big\{ e \in \partial G \ : \omega(e) = 1 \Big\} . \tag{3.4}$$

Recall that $\mathbf{C}_n$ was defined to be $\mathbf{C}_\infty \cap [-n, n]^d$. Given a subgraph G of $\mathbf{C}_n$, the ratio $|\partial^\omega G|/|G|$ shall be called the *conductance* of G and written as φ_G. This is consistent with the terminology used in [MP05], for instance. We may then rewrite the definition of the modified Cheeger constant as

$$\widehat{\Phi}_n := \min \left\{ \varphi_G \ : G \subset \mathbf{C}_n, \, 0 < |G| \leq \frac{|\mathbf{C}_n|}{d!} \right\} . \tag{3.5}$$

We say a subgraph G of $\mathbf{C}_n$ is *valid* if it satisfies $0 < |G| \leq (|\mathbf{C}_n|/d!)$. A valid subgraph $G \subset \mathbf{C}_n$ is *optimal* if $\varphi_G = \widehat{\Phi}_n$, and we let $\mathcal{G}_n$ denote the collection of all optimal subgraphs of $\mathbf{C}_n$.

Remark 3.2.1. We observe that each element of $\mathcal{G}_n$ is determined by its vertex set, in the sense that each $G_n \in \mathcal{G}_n$ inherits its graph structure from $\mathbf{C}_n$. If this were not the case for some G_n, we could strictly reduce its open edge boundary.

3.3 A preliminary setting of convergence

To prove Theorem 2.2.1, we will first encode each optimizer G_n as a measure and prove convergence to a set of limiting measures. Given $G_n \in \mathcal{G}_n$, we define the *empirical measure* of G_n as

$$\mu_n := \frac{1}{n^d} \sum_{x \in G_n} \delta_{x/n} \,, \tag{3.6}$$

where the sum ranges over all vertices of G_n, so that each μ_n is a random, non-negative Borel measure on $[-1,1]^d$. Given a Borel set $E \subset [-1,1]^d$, we define ν_E as the measure on $[-1,1]^d$ having density $\theta_p(d)\mathbf{1}_E$ with respect to Lebesgue measure, and we say that ν_E *represents* the set E. The collection of signed Borel measures on $[-1,1]^d$ shall be denoted as $\mathcal{M}([-1,1]^d)$, and the closed ball (with respect to total variation norm) of radius 3^d about the zero measure in this space shall be written as $\mathcal{B}_d$. For each percolation configuration ω, the empirical measures μ_n lie within $\mathcal{B}_d$, as do the representative measures of every Borel set $E \subset [-1,1]^d$. We equip $\mathcal{B}_d$ with a metric $\mathfrak{d}$ defined as follows.

For $k \in \{0, 1, 2, \dots\}$, let Δ^k denote the collection of closed dyadic cubes in $[-1,1]^d$ at scale k, so that each cube is a translate of $[-2^{-k}, 2^{-k}]^d$. Given $\mu, \nu \in \mathcal{B}_d$, we define

$$\mathfrak{d}(\mu, \nu) := \sum_{k=0}^{\infty} \frac{1}{2^k} \sum_{Q \in \Delta^k} \frac{1}{|\Delta^k|} |\mu(Q) - \nu(Q)| \,. \tag{3.7}$$

The metric $\mathfrak{d}$ is useful for comparing discrete objects to continuous objects (provided both are suitably encoded as measures); it figures prominently in the final section of the paper.

3.4 Some geometric measure theory

Throughout the paper, we write $\mathcal{L}^d$ for d-dimensional Lebesgue measure and $\mathcal{H}^d$ for d-dimensional Hausdorff measure. We now introduce sets of finite perimeter; the following definitions are taken from Sections 13.3 and 14.1 of [Cer06].

For $E \subset \mathbb{R}^d$ a Borel set and $O \subset \mathbb{R}^d$ open, we define the *perimeter of E in O* to be

$$\mathrm{per}(E, O) := \sup \left\{ \int_E \mathrm{div} f(x) \mathcal{L}^d(dx) \;:\; f \in C_c^\infty(O, B(0,1)) \right\} , \tag{3.8}$$

where $B(0,1)$ denotes the Euclidean unit ball in $\mathbb{R}^d$. We write $\mathrm{per}(E, \mathbb{R}^d)$ as $\mathrm{per}(E)$ and say that $\mathrm{per}(E)$ is the *perimeter* of E. A Borel set $E \subset \mathbb{R}^d$ has *finite perimeter* if $\mathrm{per}(E) < \infty$, and has *locally finite perimeter* if $\mathrm{per}(E, O) < \infty$ for each bounded open O. Sets of locally finite perimeter are also known as *Caccioppoli sets*.

We can generalize the definition (3.8) to other norms τ on $\mathbb{R}^d$ by using the unit Wulff crystal (2.9) for τ in place of $B(0,1)$. This extends the definition of the surface energy (2.8) to Borel sets. For a Borel set $E \subset \mathbb{R}^d$ and O open, we define the *surface energy of E in O with respect to τ* as

$$\mathcal{I}_\tau(E, O) = \sup \left\{ \int_E \mathrm{div} f(x) \mathcal{L}^d(dx) \ : f \in C_c^\infty(O, \widehat{W}_\tau) \right\} , \tag{3.9}$$

and we write $\mathcal{I}_\tau(E, \mathbb{R}^d)$ as $\mathcal{I}_\tau(E)$. It is a consequence of the divergence theorem that (3.9) is consistent with (2.8). We now formally state the theorem of Taylor [Tay74, Tay75, Tay78] mentioned in the introduction, which is vital to the proof of Theorem 2.2.1.

Theorem 3.4.1. Let τ be a norm on $\mathbb{R}^d$ and consider the variational problem for Borel sets $E \subset \mathbb{R}^d$:

$$\text{minimize } \mathcal{I}_\tau(E) \qquad \text{subject to } \mathcal{L}^d(E) \geq \mathcal{L}^d(\widehat{W}_\tau) \, , \tag{3.10}$$

A set E is a solution to this variational problem if and only if there exists $x \in \mathbb{R}^d$ such that the symmetric difference of $\widehat{W}_\tau$ and $E + x$ has Lebesgue measure zero.

3.5 Common notation

We collect miscellaneous notation used throughout the paper. Given $E \subset \mathbb{R}^d$ and $a > 0$, we let $\mathcal{N}_a(E)$ denote the closed Euclidean a-neighborhood of E:

$$\mathcal{N}_a(E) := \left\{ x \in \mathbb{R}^d \ : \exists y \in E \text{ with } |x - y|_2 \leq a \right\} , \tag{3.11}$$

where $|\cdot|_2$ denotes the Euclidean norm on $\mathbb{R}^d$. More generally, for $p \in [1, \infty]$, we write $|\cdot|_p$ to denote the ℓ^p norm on $\mathbb{R}^d$. For $E \subset \mathbb{R}^d$ and $a > 0$, define the closed ℓ^1 a-neighborhood of E similarly:

$$\mathcal{N}_a^{(1)}(E) := \left\{ x \in \mathbb{R}^d \ : \exists y \in E \text{ with } |x - y|_1 \leq a \right\} . \tag{3.12}$$

Define the ℓ^∞-Hausdorff metric on compact subsets of $\mathbb{R}^d$ via

$$\mathrm{d}_H(A,B) := \max\left(\sup_{x\in A}\inf_{y\in B}\left|x-y\right|_\infty, \sup_{y\in B}\inf_{x\in A}\left|x-y\right|_\infty\right). \tag{3.13}$$

Finally, we write α_d for the volume of the d-dimensional Euclidean unit ball.

CHAPTER 4

The norm $\beta_{p,d}$ and the Wulff crystal

We now introduce objects fundamental to defining the norm $\beta_{p,d}$ and hence the Wulff crystal $W_{p,d}$. To motivate our construction, we appeal to the following heuristic: we regard an optimizer $G_n \in \mathcal{G}_n$ as a droplet in $\mathbf{C}_\infty$, and we look at a small but macroscopic (diameter on the order of n) box intersecting the boundary of G_n, as depicted in Figure 4.1.

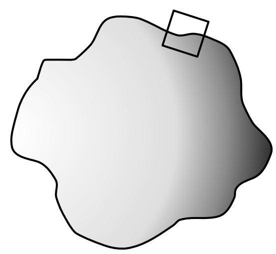

Figure 4.1: A small macroscopic box on the boundary of G_n.

Let us first discuss some of the assumptions implicitly made by Figure 4.1: the G_n are discrete objects, and the representation of G_n above treats ∂G_n as though it is a continuous object. This is justified by appealing to Proposition 3.1.1, in which $\mathbb{L}^d$-connectivity of $\partial_o G_n$ may be thought of as a discrete substitute for continuity. Accepting that the boundary of G_n behaves like a continuous object macroscopically, we note that G_n may not even be connected, simply connected or otherwise may have various holes.

Nonetheless, let us proceed with our heuristic: the small box B captures a piece of ∂G_n, and one can imagine this portion of ∂G_n separating the top and bottom faces of B. Even this makes assumptions on the "regularity" of ∂G_n, as it may be that ∂G_n has thin spikes which shoot out of B, preventing $\partial G_n \cap B$ from truly separating the top and bottom faces of B. We will see in Chapter 9 that for boxes (or other objects) chosen well relative to G_n, we can arrange that $\partial G_n \cap B$ almost separates the top and bottom faces of B, meaning that we can produce a cutset from $\partial G_n \cap B$ by adding only a few more edges.

As G_n is a Cheeger optimizer, if $\partial G_n \cap B$ separates the top and bottom faces of B, this cutset should use the fewest possible open edges in order to minimize $\partial^\omega G_n$. The choice of the position of this cutset should not greatly affect the enclosed volume $|G_n|$ because B is so small relative to G_n, so we imagine that optimizing the number of open edges used by this cut is most important to minimizing the conductance φ_{G_n} of G_n. Thus the minimal number of open edges used by a cut separating the top and bottom faces of B acts as a microscopic surface energy in the direction normal to these faces. We expect this energy to grow as n^{d-1} with n, regardless of the normal direction. Our strategy is to construct $\beta_{p,d}$ as a limit of these microscopic surface energies, properly normalized. To implement a subadditivity argument, it will be important that the cuts considered are "anchored" to the equator of the box. We will be more explicit in Section 4.1.

In two dimensions, the dual edges to any cutset form a path, so studying the minimal random weight along all such cutsets falls under the umbrella of first passage percolation. This was essentially the perspective which motivated the definition of the norm in [BLP15]. We are fortunate that minimal randomly weighted cutsets in boxes for dimensions $d \geq 3$ are also well-studied objects. They were first examined by Kesten in [Kes87] as a means of studying a higher dimensional version of first passage percolation (where the dual squares to each edge in a cut form a surface with random weights). Since this time, Théret [The14], Rossignol and Théret [RT10a, RT10b], Zhang [Zha07] and Garet [Gar09] have all studied variants of this problem. For a detailed list of these results, see Section 3.1 of [CT12]. As mentioned in Section 2.5, Cerf and Théret [CT12] have obtained a law of large numbers for the randomly weighted cuts separating pieces of the boundary of a very general domain in d dimensions.

The norm $\beta_{p,d}$ which we will soon construct has been used in most of the work just mentioned, so we emphasize that the results presented in this section and in Chapter 5 are not new or even the best possible. Nevertheless, we find it important to present a relatively self-contained argument, and the notation introduced in Section 4.1 will be used heavily throughout the paper.

4.1 Discrete cylinders, cutsets and connectivity

We take much of our notation from the work of Cerf [Cer06] and of Cerf and Théret [CT12]. Let $F \subset \mathbb{R}^d$ be the isometric image of either a non-degenerate polytope in $\mathbb{R}^{d-1}$ or a Euclidean ball in $\mathbb{R}^{d-1}$. We write $\mathsf{hyp}(F)$ to denote the hyperplane spanned by F, and we let $v(F)$ denote one of the two unit vectors in $\mathbb{S}^{d-1}$ normal to $\mathsf{hyp}(F)$; the choice does not matter for our definitions. We will define exactly what we mean by polytope at the beginning of Section 6.2; in the present section we will only ever need F to be a square.

For $\rho > 0$, we define $\mathsf{cyl}(F, \rho)$ to be the closed cylinder in $\mathbb{R}^d$ whose top and bottom faces are respectively $F_\rho^+ := F + \rho v(F)$ and $F_\rho^- := F - \rho v(F)$. The choice of $v(F)$ creates some ambiguity over which face of the cylinder is the top, but as mentioned, this ambiguity will be unimportant throughout the paper, and will play no role in the definition of the norm. We also define

$$\mathsf{slab}(F, \rho) := \Big\{ y \in \mathsf{hyp}(F) + av(F) \ : a \in [-\rho, \rho] \Big\} . \tag{4.1}$$

Figure 4.2 depicts the geometric objects introduced so far.

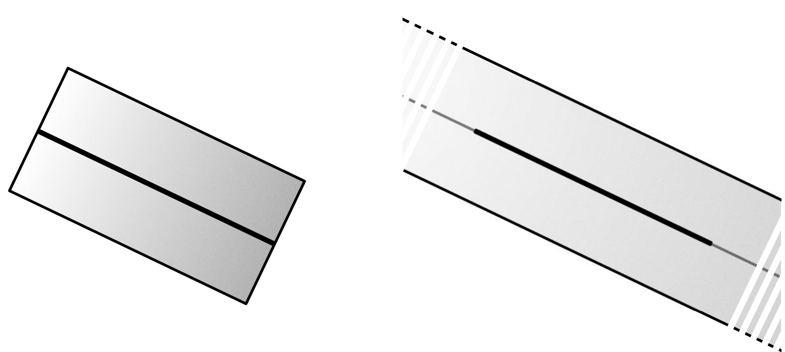

Figure 4.2: In both graphics, the bold line is the set F. The set $\mathsf{cyl}(F, \rho)$ is depicted as a box on the left. The top and bottom faces of this box are F_ρ^+ and F_ρ^- respectively. The set $\mathsf{slab}(F, \rho)$ is on the right, and the pale line running through the center of this set is $\mathsf{hyp}(F)$.

We now dilate and then discretize these objects. For $r > 0$, which we think of as large, define the discrete cylinder $\mathsf{d\text{-}cyl}(F,\rho,r)$ as

$$\mathsf{d\text{-}cyl}(F,\rho,r) := \left\{x \in \mathbb{Z}^d \ : x \in r\mathsf{cyl}(F,\rho)\right\}, \tag{4.2}$$

so that the parameter ρ controls the aspect ratio of these discrete cylinders. Towards defining $\beta_{p,d}$, the discrete cylinders will play the role of the small but macroscopic box in Figure 4.1. In order to discuss cutsets within these cylinders, we will also need to identify pairs of disjoint subsets of $\partial_* \mathsf{d\text{-}cyl}(F,\rho,r)$ to be separated, which we do now.

Note that $\mathsf{cyl}(F,\rho) \setminus \mathsf{hyp}(F)$ consists of two connected components. We will denote the top component by $\mathsf{cyl}^+(F,\rho)$, this is the component containing F_ρ^+. Likewise, the bottom component is the one containing F_ρ^- and shall be denoted $\mathsf{cyl}^-(F,\rho)$. The following sets of vertices are the top (corresponding to "+") and bottom ("−") *hemispheres* of $\mathsf{d\text{-}cyl}(F,\rho,r)$:

$$\mathsf{d\text{-}hemi}^\pm(F,\rho,r) := \left\{x \in \partial_* \mathsf{d\text{-}cyl}(F,\rho,r) \ : x \in r\mathsf{cyl}^\pm(F,\rho)\right\}. \tag{4.3}$$

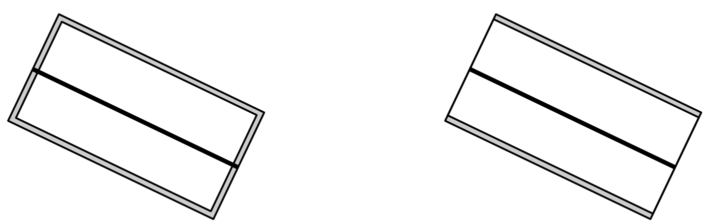

Figure 4.3: On the left, the vertex set $\mathsf{d\text{-}hemi}^+(F,\rho,r)$ (respectively $\mathsf{d\text{-}hemi}^-(F,\rho,r)$) is represented by the shaded region above (respectively below) the bold line. On the right, the vertex sets $\mathsf{d\text{-}face}^\pm(F,\rho,r)$ are represented the shaded regions above (+) and below (−) the bold line.

We also define the top and bottom *faces* of $\mathsf{d\text{-}cyl}(F,\rho,r)$:

$$\mathsf{d\text{-}face}^\pm(F,\rho,r) := \left\{x \in \partial_*[(r\mathsf{slab}(F,\rho)) \cap \mathbb{Z}^d] \ : x \in r\mathsf{cyl}^\pm(F,\rho,r)\right\}. \tag{4.4}$$

This definition looks complicated, but conceptually this vertex set is even simpler than its hemisphere counterpart, it is depicted on the right side of Figure 4.3.

Remark 4.1.1. The vertex sets $\mathsf{d\text{-}hemi}^{\pm}(F,\rho,r)$ and $\mathsf{d\text{-}face}^{\pm}(F,\rho,r)$ are contained in $\mathsf{d\text{-}cyl}(F,\rho,r)$. As $\mathsf{d\text{-}cyl}(F,\rho,r)$ inherits a natural graph structure from $\mathbb{Z}^d$, we may thus consider cutsets within $\mathsf{d\text{-}cyl}(F,\rho,r)$ which separate opposite hemispheres or opposite faces. Given such a cutset S, we write $|S|$ for the number of edges in S. Bond percolation on $\mathbb{Z}^d$ induces bond percolation within $\mathsf{d\text{-}cyl}(F,\rho,r)$, and this gives us a more relevant method for assigning a weight to these cutsets. For any fixed cutset S, let $|S|_\omega$ denote the number of open edges in S, so that $|S|_\omega$ is a random variable.

We define the random variable $\Xi_{\text{hemi}}(F,\rho,r)$ as

$$\Xi_{\text{hemi}}(F,\rho,r) := \min\Big(|S|_\omega \ : S \text{ separates } \mathsf{d\text{-}hemi}^{\pm}(F,\rho,r) \text{ within } \mathsf{d\text{-}cyl}(F,\rho,r)\Big)\,, \tag{4.5}$$

and we likewise define $\Xi_{\text{face}}(F,\rho,r)$ as

$$\Xi_{\text{face}}(F,\rho,r) := \min\Big(|S|_\omega \ : S \text{ separates } \mathsf{d\text{-}face}^{\pm}(F,\rho,r) \text{ within } \mathsf{d\text{-}cyl}(F,\rho,r)\Big)\,. \tag{4.6}$$

In either case, we may restrict the minimum to one taken over all minimal cutsets. The difference between these random variables is that $\Xi_{\text{hemi}}(F,\rho,r)$ is a minimum over cutsets which are in some sense "anchored" at the equator of the cylinder $\mathsf{cyl}(F,\rho,r)$, whereas the cutsets involved in the definition of $\Xi_{\text{face}}(F,\rho,r)$ are allowed to meet the sides of $\mathsf{cyl}(F,\rho,r)$ at any height relative to the equator.

Remark 4.1.2. Whenever r or ρ are too small relative to F, the random variables $\Xi_{\text{hemi}}(F,\rho,r)$ or $\Xi_{\text{face}}(F,\rho,r)$ may not be well-defined. We say that the parameters r and ρ are *suitable* for F if the vertex sets $\mathsf{d\text{-}hemi}^{\pm}(F,\rho,r)$ and $\mathsf{d\text{-}face}^{\pm}(F,\rho,r)$ are non-empty, and if the vertex sets $\mathsf{d\text{-}face}^{\pm}(F,\rho,r)$ are a Euclidean distance of at least $5d$. When ρ and r are suitable for F, we define $\Xi_{\text{hemi}}(F,\rho,r)$ and $\Xi_{\text{face}}(F,\rho,r)$ as in (4.5) and (4.6) respectively. Otherwise we define these random variables to be zero.

4.2 Defining the norm

A *square* in $\mathbb{R}^d$ shall be any isometric image of $[-1,1]^{d-1}\times\{0\}$. For $v\in\mathbb{S}^{d-1}$, we consider a square in $\mathbb{R}^d$ centered at 0 whose spanning hyperplane is normal to v. In dimensions at least three, this constraint does not uniquely determine the square. Such squares will be used to define $\beta_{p,d}$, so it will be important to assign to each direction $v \in \mathbb{S}^{d-1}$ a unique square. Let S be such an assignment; that is for each $v \in \mathbb{S}^{d-1}$, $\mathsf{S}(v)$ is a square in $\mathbb{R}^d$ centered at 0 with $\mathsf{hyp}(\mathsf{S}(v))$ normal to v. We refer to S as the *chosen orientation*.

In the next section, we will be more explicit about the properties this chosen orientation should have. For now we simply assume we have such an S, and use this to construct $\beta_{p,d}$.

Remark 4.2.1. The value of $\beta_{p,d}$ in a given direction will not depend on S (as we will show in Proposition 4.2.3). However, we will facilitate later proofs by building $\beta_{p,d}$ from an S which varies nicely over the sphere.

Throughout this section, treat S as given. The random variable used to define $\beta_{p,d}$ is

$$\mathfrak{X}(x, v, r) := \Xi_{\text{hemi}}(\mathsf{S}(v) + x, 1, r)\,. \tag{4.7}$$

As was mentioned at the beginning of this section, the $\mathfrak{X}(x, v, r)$ are well-studied objects. In particular, precise large deviation estimates are well known for the sequence $\mathfrak{X}(x, v, n)/(2n)^{d-1}$, giving rise to a law of large numbers.

To define $\beta_{p,d}$ as quickly as possible, we use a subadditivity argument on the expectations of these random variables. Our argument is essentially the one given by Rossignol and Théret in Section 4.3 of [RT10b]. Before carrying out this argument, we make one observation about the $\mathfrak{X}(x, v, r)$.

Lemma 4.2.2. Let $d \geq 2$. There is a positive constant $c(d)$ so that for all $p \in [0, 1]$, $x \in \mathbb{R}^d$, $v \in \mathbb{S}^{d-1}$ and $r > 0$,

$$\mathbb{E}_p \mathfrak{X}\big(x, v, r + d^{1/2}\big) \leq \mathbb{E}_p \mathfrak{X}(0, v, r) + c(d) r^{d-2}\,. \tag{4.8}$$

Proof. Let $x \in \mathbb{R}^d$ and $v \in \mathbb{S}^{d-1}$ be given. Choose $x' \in \mathbb{Z}^d$ so that $|x - x'|_\infty \leq 1$. For notational ease, make the following abbreviations within this proof:

$$\mathsf{cyl}(x') := r\mathsf{cyl}(\mathsf{S}(v) + x', 1)\,, \tag{4.9}$$

$$\mathsf{cyl}(x) := \left(r + d^{1/2}\right) \mathsf{cyl}(\mathsf{S}(v) + x, 1)\,, \tag{4.10}$$

$$\mathsf{hyp}(x) := \mathsf{hyp}\left(\left(r + d^{1/2}\right)(\mathsf{S}(v) + x)\right)\,. \tag{4.11}$$

For any $r > 0$, it follows that $\mathsf{cyl}(x') \subset \mathsf{cyl}(x)$. Let A be the collection of edges in $\mathbb{Z}^d$ having non-empty intersection with the neighborhood

$$\mathcal{N}_{5d}\left((\mathsf{cyl}(x) \setminus \mathsf{cyl}(x')) \cap \mathsf{hyp}(x)\right)\,. \tag{4.12}$$

The construction of the set A is illustrated in Figure 4.5.

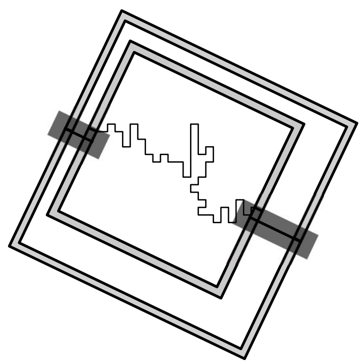

Figure 4.4: The inner box is $\mathsf{cyl}(x')$, the outer box is $\mathsf{cyl}(x)$, and the darker shaded region is the neighborhood (4.12) used to define the set A. The thin interface depicts the minimal cutset E.

It follows from the construction of A that there is some positive constant $c(d)$ with $|A| \leq c(d)r^{d-2}$. We now choose a cutset for the smaller cylinder and augment these edges with A to form a cutset for the larger cylinder. Let us again make some abbreviations for the sake of clarity.

$$\mathsf{d\text{-}cyl}(x') := \mathsf{d\text{-}cyl}(\mathsf{S}(v) + x', 1, r)\,, \tag{4.13}$$

$$\mathsf{d\text{-}hemi}^{\pm}(x') := \mathsf{d\text{-}hemi}^{\pm}(\mathsf{S}(v) + x', 1, r)\,, \tag{4.14}$$

$$\mathsf{d\text{-}cyl}(x) := \mathsf{d\text{-}cyl}(\mathsf{S}(v) + x, 1, r + d^{1/2})\,, \tag{4.15}$$

$$\mathsf{d\text{-}hemi}^{\pm}(x') := \mathsf{d\text{-}hemi}^{\pm}(\mathsf{S}(v) + x, 1, r + d^{1/2})\,. \tag{4.16}$$

Let E be a minimal cutset separating the hemispheres $\mathsf{d\text{-}hemi}^{\pm}(x')$ within $\mathsf{d\text{-}cyl}(x')$. We claim the collection of edges in $A \cup E$ which lie in $\mathsf{d\text{-}cyl}(x)$ separate the hemispheres $\mathsf{d\text{-}hemi}^{\pm}(x)$ in $\mathsf{d\text{-}cyl}(x)$. Assuming this for now, it follows that

$$\mathfrak{X}(x, v, r + d^{1/2}) \leq \mathfrak{X}(x', v, r) + c(d) r^{d-2}\,, \tag{4.17}$$

and the Lemma is proved upon taking expectations as $x' \in \mathbb{Z}^d$.

To complete the proof, it suffices to show that any $\mathbb{Z}^d$ path joining the hemispheres $\mathsf{d\text{-}hemi}^{\pm}(x)$ within $\mathsf{d\text{-}cyl}(x)$ must use an edge of $A \cup E(\omega)$. We do this carefully here, as we will appeal to this argument in other proofs without repeating the details. Let $y^{\pm} \in \mathsf{d\text{-}hemi}^{\pm}(x)$, and let γ be a simple path from y^- to y^+ using only edges of $\mathsf{d\text{-}cyl}(x)$. If γ does not pass through any vertex of $\partial_* \mathsf{d\text{-}cyl}(x')$, it must be that γ lies entirely within $\mathsf{cyl}(x) \setminus \mathsf{cyl}(x')$, in which case γ must use an edge of A. We may then suppose that γ passes through a vertex of $\partial_* \mathsf{d\text{-}cyl}(x')$. We now consider several cases.

Case (i): Suppose that the last vertex z^+ of $\mathsf{d\text{-}cyl}(x')$ used by γ lies within the bottom hemisphere $\mathsf{d\text{-}hemi}^{-}(x')$. Let γ' denote the subpath of γ connecting z^+ to y^+, and observe that γ' is contained within $\mathsf{cyl}(x) \setminus \mathsf{cyl}(x')$. As γ' starts either in the bottom half of $\mathsf{cyl}(x)$ or in the neighborhood defined in (4.12), we see γ' must use an edge in A.

Case (ii): Suppose that the first vertex z^- of $\mathsf{d\text{-}cyl}(x')$ used by γ lies in $\mathsf{d\text{-}hemi}^{+}(x')$. Using the same reasoning as in Case (i), we see that γ must use an edge in A between y^- and z^-.

Case (iii): We may now suppose that $z^{\pm} \in \mathsf{d\text{-}hemi}^{\pm}(x')$. Let z be the vertex of $\mathsf{d\text{-}hemi}^{-}(x')$ used last by γ, and consider the subpath γ' of γ joining z to z^+. If γ' is contained completely within $\mathsf{d\text{-}cyl}(x')$, it must be that γ' uses an edge of $E(\omega)$. On the other hand, if γ' is not contained in $\mathsf{d\text{-}cyl}(x')$, we may assume that the vertex following z in the path γ' lies outside of $\mathsf{d\text{-}cyl}(x')$, else γ' would either use an edge of $E(\omega)$, or would not use z last among all vertices of $\mathsf{d\text{-}hemi}^{-}(x')$. With this assumption made, we see that γ' leaves $\mathsf{d\text{-}cyl}(x')$ at the vertex z, and only returns to $\mathsf{d\text{-}cyl}(x')$ at some vertex $z' \in \mathsf{d\text{-}hemi}^{+}(x')$. Thus along the subpath γ'' of γ' joining z with z', we see that all intermediate vertices lie in $\mathsf{cyl}(x) \setminus \mathsf{cyl}(x')$, so that γ'' uses an edge of A. □

With this lemma in place, we now define $\beta_{p,d}$ as a function on $\mathbb{S}^{d-1}$.

Proposition 4.2.3. Let $d \geq 2$. For all $v \in \mathbb{S}^{d-1}$, the limit

$$\beta_{p,d}(v) := \lim_{n\to\infty} \frac{\mathbb{E}_p \mathfrak{X}(0,v,n)}{(2n)^{d-1}} \tag{4.18}$$

exists and is finite. Moreover, this limit is independent of the chosen orientation S.

Proof. Let $n, m \in \mathbb{N}$ with $n > m$ and both numbers larger than d. We write $n = km + r$ for $k, r \in \mathbb{N} \cup \{0\}$ and $r < m$. Let S be the chosen orientation, and let $\widetilde{\mathsf{S}}$ be another assignment of unit vectors $v \in \mathbb{S}^{d-1}$ to squares $\widetilde{\mathsf{S}}(v)$ so that v is normal to $\mathsf{hyp}(\widetilde{\mathsf{S}}(v))$. We define $\widetilde{\mathfrak{X}}(x,v,r)$ analogously to $\mathfrak{X}(x,v,r)$ using $\widetilde{\mathsf{S}}$ in place of S:

$$\widetilde{\mathfrak{X}}(x,v,r) := \Xi_{\text{hemi}}\left(\widetilde{\mathsf{S}}(v) + x, 1, r\right) . \tag{4.19}$$

Choose a finite collection $\{\widetilde{S}_i\}_{i=1}^{\ell}$ of translates of $(m+d^{1/2})\widetilde{\mathsf{S}}(v)$, each contained within $n\mathsf{S}(v)$, so that:

(i) The translates $\{\widetilde{S}_i\}_{i=1}^{\ell}$ are disjoint.

(ii) There is a positive constant $c(d)$ so that $\mathcal{H}^{d-1}\left(n\mathsf{S}(v) \setminus \bigcup_{i=1}^{\ell} \widetilde{S}_i\right) \leq c(d)mn^{d-2}$.

(iii) $\ell \leq (k+1)^{d-1}$.

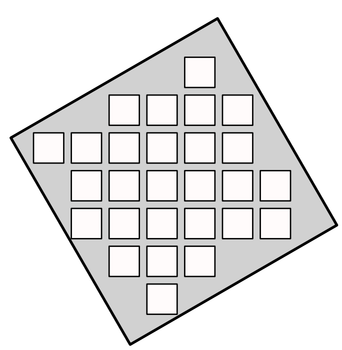

Figure 4.5: The small white squares are the collection $\{\widetilde{S}_i\}_{i=1}^{\ell}$, which are disjoint and nearly exhaust the large square $n\mathsf{S}(v)$. Note that in this diagram, we are representing squares as two-dimensional objects, whereas in all previous diagrams they were represented as one-dimensional objects.

Let us make the abbreviations

$$\mathsf{d\text{-}cyl}(i) := \mathsf{d\text{-}cyl}(\widetilde{S}_i, m + d^{1/2}, 1)\,, \tag{4.20}$$

$$\mathsf{d\text{-}hemi}^{\pm}(i) := \mathsf{d\text{-}hemi}^{\pm}(\widetilde{S}_i, m + d^{1/2}, 1)\,, \tag{4.21}$$

$$\mathsf{d\text{-}cyl} := \mathsf{d\text{-}cyl}(\mathsf{S}(v), 1, n)\,, \tag{4.22}$$

$$\mathsf{d\text{-}hemi}^{\pm} := \mathsf{d\text{-}hemi}^{\pm}(\mathsf{S}(v), 1, n)\,. \tag{4.23}$$

For each $\widetilde{S}_i$, let E_i be a minimal cutset in separating the hemispheres $\mathsf{d\text{-}hemi}^{\pm}(i)$ within $\mathsf{d\text{-}cyl}(i)$. Let A be the collection of edges in $\mathbb{Z}^d$ having non-empty intersection with the neighborhood

$$\mathcal{N}_{5d}\left(n\mathsf{S}(v) \setminus \bigcup_{i=1}^{\ell} \widetilde{S}_i\right)\,, \tag{4.24}$$

and note that by *(ii)* above, there is a positive constant $c(d)$ so that $|A| \leq c(d)mn^{d-2}$. We will soon take n to infinity, thus we lose no generality supposing n is large enough so that each $\mathsf{d\text{-}cyl}(i)$ is contained in $\mathsf{d\text{-}cyl}$, so that in particular, each E_i is contained in the edge set of $\mathsf{d\text{-}cyl}$ across all configurations ω.

One can repeat the argument at the end of the proof of Lemma 4.2.2 to show that the collection of edges $A \cup \left(\bigcup_{i=1}^{\ell} E_i\right)$ which lie in $\mathsf{d\text{-}cyl}$ separate the vertex sets $\mathsf{d\text{-}hemi}^{\pm}$ in $\mathsf{d\text{-}cyl}$. Though we are dealing with many more boxes in this case, the complexity of the argument does not go up: we can always reduce to the case that our simple path γ last uses any vertex of $\mathsf{d\text{-}hemi}^{-}(i)$ for all i, and we may also assume γ uses a vertex within some $\mathsf{d\text{-}hemi}^{+}(j)$ at a later point. Between these two points, we find that we must either use an edge in A, or an edge in one of the E_i. Thus, we may conclude

$$\mathfrak{X}(0, v, n) \leq \sum_{i=1}^{\ell} \Xi_{\text{hemi}}(\widetilde{S}_i, m + d^{1/2}, 1) + c(d)mn^{d-2}\,. \tag{4.25}$$

The chosen orientation thus far has been arbitrary, so the preceding lemma applies to $\widetilde{\mathfrak{X}}(0, v, n)$ as well as to $\mathfrak{X}(0, v, n)$. We may take expectations of both sides and apply Lemma 4.2.2 to each term in the above sum, while also using our bound $\ell \leq (k+1)^{d-1}$ from *(iii)*.

$$\mathbb{E}_p\mathfrak{X}(0, v, n) \leq \ell\mathbb{E}_p\widetilde{\mathfrak{X}}(0, v, m) + \ell c(d)m^{d-2} + c(d)mn^{d-2}\,, \tag{4.26}$$

$$\leq (k+1)^{d-1}\mathbb{E}_p\widetilde{\mathfrak{X}}(0, v, m) + (k+1)^{d-1}c(d)m^{d-2} + c(d)mn^{d-2}\,. \tag{4.27}$$

We divide through by n^{d-1}:

$$\frac{\mathbb{E}_p\mathfrak{X}(0,v,n)}{n^{d-1}} \leq (k+1)^{d-1}\frac{\mathbb{E}_p\widetilde{\mathfrak{X}}(0,v,m)}{n^{d-1}} + \frac{(k+1)^{d-1}m^{d-2}c(d)}{n^{d-1}} + \frac{c(d)m}{n}\,, \tag{4.28}$$

$$\leq \left(\frac{k+1}{k}\right)^{d-1} k^{d-1}\cdot\frac{\mathbb{E}_p\widetilde{\mathfrak{X}}(0,v,m)}{n^{d-1}} + \left(\frac{k+1}{k}\right)^{d-1}\left(\frac{k}{n}\right)^{d-1} m^{d-2}c(d) + \frac{c(d)m}{n}\,, \tag{4.29}$$

$$\leq \left(\frac{k+1}{k}\right)^{d-1}\frac{\mathbb{E}_p\widetilde{\mathfrak{X}}(0,v,m)}{m^{d-1}} + \left(\frac{k+1}{k}\right)^{d-1}\frac{c(d)}{m} + \frac{c(d)m}{n}\,. \tag{4.30}$$

We first take the lim sup of both sides in n, and then the lim inf of both sides in m:

$$\limsup_{n\to\infty}\frac{\mathbb{E}_p\mathfrak{X}(0,v,n)}{n^{d-1}} \leq \frac{\mathbb{E}_p\widetilde{\mathfrak{X}}(0,v,m)}{m^{d-1}} + \frac{c(d)}{m}\,, \tag{4.31}$$

$$\leq \liminf_{m\to\infty}\frac{\mathbb{E}_p\widetilde{\mathfrak{X}}(0,v,m)}{m^{d-1}}\,, \tag{4.32}$$

and the proof is complete upon dividing both sides by 2^{d-1}: setting $\widetilde{\mathsf{S}} \equiv \mathsf{S}$ gives us the existence of the limit in question, and interchanging $\widetilde{\mathsf{S}}$ and S in the above argument tells us this limit does not depend on the chosen orientation. The finiteness of this limit can be seen, for instance, in the following way: given a direction $v \in \mathbb{S}^{d-1}$, the collection of edges intersecting the neighborhood $\mathcal{N}_{5d}(n\mathsf{S}(v))$ forms a cutset in $\mathsf{d\text{-}cyl}(\mathsf{S}(v),1,n)$ separating $\mathsf{d\text{-}hemi}^{\pm}(\mathsf{S}(v),1,n)$ and this cutset has cardinality bounded above by $c(d)n^{d-1}$ for some positive constant $c(d)$ which does not depend on the direction. □

The above proposition defines $\beta_{p,d}$ as a function on $\mathbb{S}^{d-1}$. We can immediately deduce that $\beta_{p,d}$ inherits the symmetries of $\mathbb{Z}^d$.

Corollary 4.2.4. Let $d \geq 2$. For all $v \in \mathbb{S}^{d-1}$ and for all linear transformations $L : \mathbb{R}^d \to \mathbb{R}^d$ such that $L(\mathbb{Z}^d) = \mathbb{Z}^d$, we have $\beta_{p,d}(Lv) = \beta_{p,d}(v)$.

Proof. Let $v \in \mathbb{S}^{d-1}$, and let S be the chosen orientation. Then

$$\widetilde{\mathsf{S}}(v) := L^{-1}\mathsf{S}(Lv) \tag{4.33}$$

is a rotation of $\mathsf{S}(v)$ contained in $\mathsf{hyp}(\mathsf{S}(v))$. From the preceding Proposition 4.2.3, we know

$$\lim_{n\to\infty}\frac{\mathbb{E}_p\Xi_{\text{hemi}}(\widetilde{\mathsf{S}}(v),1,n)}{(2n)^{d-1}} = \lim_{n\to\infty}\frac{\mathbb{E}_p\mathfrak{X}(0,v,n)}{(2n)^{d-1}}\,. \tag{4.34}$$

Moreover, because L induces a graph automorphism of $\mathbb{Z}^d$, we have also that $\mathbb{E}_p \Xi_{\text{hemi}}(\widetilde{\mathsf{S}}(v), 1, n) = \mathbb{E}_p \mathfrak{X}(0, Lv, n)$, so that

$$\lim_{n\to\infty} \frac{\mathbb{E}_p \mathfrak{X}(0, Lv, n)}{(2n)^{d-1}} = \lim_{n\to\infty} \frac{\mathbb{E}_p \mathfrak{X}(0, v, n)}{(2n)^{d-1}}\,, \tag{4.35}$$

which gives the desired result. □

4.3 The chosen orientation and properties of $\beta_{p,d}$

The function $\beta_{p,d}$ could have been defined using cylinders based at discs instead of squares, but this would have made it harder to execute the above subadditivity argument (and similar arguments in Chapter 6). There is a tradeoff between the tidiness of these arguments and the artificial nature of the chosen orientation; we feel we have taken the route which is ultimately cleanest.

Part of this tradeoff is that we need the chosen orientation S to vary over most of the sphere in a Lipschitz way; for instance, we would like there to be a positive constant $M = M(d)$ so that S to satisfies

$$\mathrm{d}_H(\mathsf{S}(v), \mathsf{S}(w)) \leq M\epsilon \tag{4.36}$$

whenever $|v - w|_2 < \epsilon$, where d_H is defined in (3.13). For S with this property, it is easy to show the sequence of functions $v \mapsto \mathbb{E}_p \mathfrak{X}(0, v, n)/(2n)^{d-1}$ on $\mathbb{S}^{d-1}$ converge uniformly to $\beta_{p,d}$. This will in turn allow us to prove concentration estimates in Chapter 5.

For topological reasons, it is not possible to have S vary as in (4.36) over the entire sphere, so we must first work over the closed upper hemisphere. Let $\mathbb{S}^{d-1}_+$ be the closed upper hemisphere of $\mathbb{S}^{d-1}$, that is,

$$\mathbb{S}^{d-1}_+ := \mathbb{S}^{d-1} \cap \left\{ x \in \mathbb{R}^d \,:\, x_d \geq 0 \right\}. \tag{4.37}$$

A corollary of Proposition 10.4.4 is that we may first define S over $\mathbb{S}^{d-1}_+$ so that (4.36) holds for some $M(d) > 0$ whenever $v, w \in \mathbb{S}^{d-1}_+$ and $|v - w|_2 < \epsilon$.

With such S defined on the upper hemisphere, we extend the definition of S to the rest of $\mathbb{S}^{d-1}$ in a natural way. Let $e_1, \dots, e_d$ be the standard basis for $\mathbb{R}^d$, and let $A : \mathbb{R}^d \to \mathbb{R}^d$

be the reflection which preserves the first $d-1$ basis elements and takes e_d to $-e_d$. Given $v \in \mathbb{S}^{d-1} \setminus \mathbb{S}^{d-1}_+$, we define $\mathsf{S}(v)$ to be the square $A\mathsf{S}(Av)$.

Remark 4.3.1. This defines a chosen orientation S on the whole sphere such that the above property (4.36) holds whenever the vectors v and w are sufficiently close and lie in the same hemisphere, and from this point forward we assume that S has this property and is fixed. This "piecewise-Lipschitz" property is sufficient to show the aforementioned uniform convergence of the sequence of functions $v \mapsto \mathbb{E}_p\mathfrak{X}(0,v,n)/(2n)^{d-1}$ to $\beta_{p,d}$.

Proposition 4.3.2. Let $d \geq 2$. The sequence of functions on $\mathbb{S}^{d-1}$ defined by $v \mapsto \mathbb{E}_p\mathfrak{X}(0,v,n)/(2n)^{d-1}$ converges uniformly to $\beta_{p,d}$.

Proof. Let $\epsilon > 0$, Let $v, w \in \mathbb{S}^{d-1}_+$ be such that $|v-w|_2 < \epsilon$. Let us fix some notation:

$$\mathsf{cyl}(v) := n\mathsf{cyl}(\mathsf{S}(v), 1)\,, \tag{4.38}$$

$$\textsf{d-cyl}(v) := \textsf{d-cyl}(\mathsf{S}(v), 1, n)\,, \tag{4.39}$$

$$\textsf{d-hemi}^{\pm}(v) := \textsf{d-hemi}^{\pm}(\mathsf{S}(v), 1, n)\,, \tag{4.40}$$

$$\mathsf{cyl}(w) := \lceil n(1+M\epsilon)\rceil \mathsf{cyl}(\mathsf{S}(w), 1)\,, \tag{4.41}$$

$$\textsf{d-cyl}(w) := \textsf{d-cyl}(\mathsf{S}(w), 1, \lceil n(1+M\epsilon)\rceil)\,, \tag{4.42}$$

$$\textsf{d-hemi}^{\pm}(w) := \textsf{d-hemi}^{\pm}(\mathsf{S}(w), 1, \lceil n(1+M\epsilon)\rceil)\,. \tag{4.43}$$

Let E be a minimal cutset which separates the hemispheres $\textsf{d-hemi}^{\pm}(v)$ within $\textsf{d-cyl}(v)$. By the hypothesis (4.36) on the chosen orientation, we have $\mathsf{cyl}(v) \subset \mathsf{cyl}(w)$, so that E is contained in the edge set of $\textsf{d-cyl}(w)$. As we have done before, we would like to use E in conjunction with a small collection of edges to produce a cut separating the hemispheres of $\textsf{d-cyl}(w)$. We will actually use two other collections of edges to do this.

Writing $\mathsf{hyp}(w)$ for $\mathsf{hyp}(\lceil n(1+M\epsilon)\rceil \mathsf{S}(w))$, we define the edge set A as in Lemma 4.2.2 to be the edges in $\mathbb{Z}^d$ having non-empty intersection with the neighborhood

$$\mathcal{N}_{5d}\left((\mathsf{cyl}(w) \setminus \mathsf{cyl}(v)) \cap \mathsf{hyp}(w)\right)\,. \tag{4.44}$$

Let B be the collection of edges having non-empty intersection with

$$\mathcal{N}_{5d}\left(\partial \mathsf{cyl}(v) \cap \mathsf{slab}(\mathsf{S}(v), nM\epsilon)\right)\,. \tag{4.45}$$

Here we suppose that ϵ is small enough so that $\mathsf{slab}(\mathsf{S}(v), nM\epsilon)$ does not contain the top and bottom faces of the cube $\mathsf{cyl}(v)$. The neighborhood (4.45) is thus slight thickening of an equatorial band of height $nM\epsilon$ in $\partial\mathsf{cyl}(v)$, and it follows that $|B| \leq c(d)M\epsilon n^{d-1}$ for some positive constant $c(d)$.

The above neighborhood (4.44) forms a bridge between the neighborhood (4.45) defining B and the equator of the larger cube $\mathsf{cyl}(w)$. By construction, we also have $|A| \leq c(d)M\epsilon n^{d-1}$ for some positive constant $c(d)$. Figure 4.6 is an illustration of the cut E with the edge sets A and B.

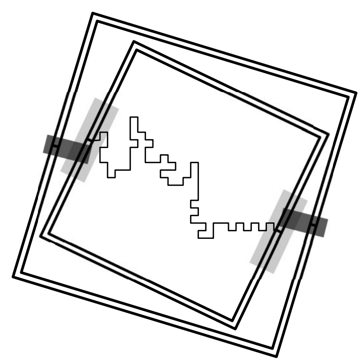

Figure 4.6: The cut E in the smaller cube, $\mathsf{cyl}(v)$, is central. At the equator of the smaller cube, this cut meets with the edge set B, which is represented by the lightly shaded regions. The edge set B is joined to the equator of the larger cube, $\mathsf{cyl}(w)$, by the edge set A, depicted as the darker shaded regions.

The edges of the union $E \cup A \cup B$ contained in $\mathsf{d\text{-}cyl}(w)$ form a cutset separating the hemispheres $\mathsf{d\text{-}hemi}^{\pm}(w)$. The argument for this is nearly identical to the one given at the end of the proof of Lemma 4.2.2. Indeed, we are looking at nested cubes, the only difference being that one is tilted slightly relative to the other. This tilt necessitates our use of the edge set B in our case. With this established, we conclude

$$\mathfrak{X}(0, w, \lceil n(1 + M\epsilon)\rceil) \leq \mathfrak{X}(0, v, n) + c(d)M\epsilon n^{d-1}, \tag{4.46}$$

so that by taking expectations, we have

$$\frac{\mathbb{E}_p\mathfrak{X}(0, w, \lceil n(1 + M\epsilon)\rceil)}{(2\lceil n(1 + M\epsilon)\rceil)^{d-1}} \leq \frac{\mathbb{E}_p\mathfrak{X}(0, v, n)}{(2n)^{d-1}} + c(d)M\epsilon\,. \tag{4.47}$$

Taking $n \to \infty$, and using the symmetry in v and w, we have shown that when $v, w \in \mathbb{S}^{d-1}_+$ satisfy $|v - w|_2 < \epsilon$, we have

$$|\beta_{p,d}(v) - \beta_{p,d}(w)| < c(d)M\epsilon\,. \tag{4.48}$$

A symmetric argument shows that we have the same bounds whenever $v, w \in \mathbb{S}^{d-1} \setminus \mathbb{S}^{d-1}_+$ and $|v - w|_2 < \epsilon$.

Choose a finite collection of unit vectors $\{v_i\}_{i=1}^m$ (with $m = m(\epsilon)$), arranged so that if $v \in \mathbb{S}^{d-1}_+$, we can find $v_i \in \mathbb{S}^{d-1}_+$ so that $|v - v_i|_2 < \epsilon$, and if $v \in \mathbb{S}^{d-1} \setminus \mathbb{S}^{d-1}_+$, there is $v_i \in \mathbb{S}^{d-1} \setminus \mathbb{S}^{d-1}_+$ with $|v - v_i|_2 < \epsilon$. Take N large enough so that whenever $n \geq N$, we have for each i,

$$\left| \frac{\mathbb{E}_p \mathfrak{X}(0, v_i, n)}{(2n)^{d-1}} - \beta_{p,d}(v_i) \right| < \epsilon\,. \tag{4.49}$$

Let $v \in \mathbb{S}^{d-1}$ and take v_i in the same hemisphere so that $|v - v_i|_2 < \epsilon$. Two applications of (4.47) show

$$\frac{\mathbb{E}_p \mathfrak{X}(0, v_i, \lceil \lceil n(1+M\epsilon) \rceil (1+M\epsilon) \rceil)}{(2\lceil \lceil n(1+M\epsilon) \rceil (1+M\epsilon) \rceil)^{d-1}} - c(d)M\epsilon \leq \frac{\mathbb{E}_p \mathfrak{X}(0, v, \lceil n(1+M\epsilon) \rceil)}{(2\lceil n(1+M\epsilon) \rceil)^{d-1}} \tag{4.50}$$

$$\leq \frac{\mathbb{E}_p \mathfrak{X}(0, v_i, n)}{(2n)^{d-1}} + c(d)M\epsilon\,. \tag{4.51}$$

By (4.48), we have

$$\beta_p(v) - \epsilon - 2c(d)M\epsilon \leq \frac{\mathbb{E}_p \mathfrak{X}(0, v, \lceil n(1+M\epsilon) \rceil)}{(2\lceil n(1+M\epsilon) \rceil)^{d-1}} \leq \beta_p(v) + \epsilon + 2c(d)M\epsilon\,, \tag{4.52}$$

which establishes the desired uniform convergence. □

We now extend $\beta_{p,d}$ to a function on all of $\mathbb{R}^d$ via homogeneity; for $x \in \mathbb{R}^d$ define

$$\beta_{p,d}(x) := \begin{cases} |x|_2 \beta_{p,d}(x/|x|_2) & |x|_2 > 0 \\ 0 & |x|_2 = 0 \end{cases}. \tag{4.53}$$

To show $\beta_{p,d}$ defines a norm on $\mathbb{R}^d$ when $p > p_c(d)$, we first appeal to the fact that $\beta_{p,d}$ satisfies the so-called weak triangle inequality. For two distinct points $a, b \in \mathbb{R}^d$, we let $[ab]$ denote the closed line segment in $\mathbb{R}^d$ connecting a and b, and for $a, b, c \in \mathbb{R}^d$ not co-linear, we let $[abc]$ denote the closed triangle in $\mathbb{R}^d$ with vertices a, b, c. For a triangle $[abc]$, we define v_a to be the outward pointing unit normal to the side $[bc]$ within the plane spanned by $[abc]$. We define v_b and v_c analogously.

Proposition 4.3.3. For $d \geq 2$ and $p > p_c(d)$, the function $\beta_{p,d} : \mathbb{S}^{d-1} \to [0, \infty)$ satisfies the weak triangle inequality. That is, for $a, b, c \in \mathbb{R}^d$ not co-linear,

$$\mathcal{H}^1([bc])\beta_{p,d}(v_a) \leq \mathcal{H}^1([ac])\beta_{p,d}(v_b) + \mathcal{H}^1([ab])\beta_{p,d}(v_c)\,. \tag{4.54}$$

Proof. The proof here is identical to Propostion 11.6 in [Cer06], see also Proposition 4.5 of [RT10b]. □

The following is a consequence of the weak triangle inequality.

Proposition 4.3.4. For $d \geq 2$ and $p > p_c(d)$, the function $\beta_{p,d} : \mathbb{R}^d \to [0, \infty)$ defines a norm on $\mathbb{R}^d$.

Proof. We combine Proposition 4.3.3 with Corollary 11.7 of [Cer06] to conclude that $\beta_{p,d}$ is a convex function on $\mathbb{R}^d$. To show non-degeneracy of $\beta_{p,d}$ then, it suffices to show non-degeneracy in the cardinal directions. Thanks to the symmetries of $\beta_{p,d}$ (Corollary 4.2.4) it suffices to show non-degeneracy in a single cardinal direction.

This non-degeneracy is a consequence of Theorem 7.68 in [Gri99], for instance, which says that within a large axis-parallel cube, with high probability, there are at least cn^{d-1} edge-disjoint open paths between the top and bottom faces, for some $c = c(p, d) > 0$. Menger's theorem allows us to convert this into a high probability lower bound on the size of a minimal cut separating the top and bottom faces of this cube, and the proof is complete. □

Remark 4.3.5. That $\beta_{p,d}$ is a norm allows us to define the associated surface energy $\mathcal{I}_{p,d}$, as in Chapter 3, as well as the unit Wulff crystal $\widehat{W}_{p,d}$, which is the unit ball in the norm dual to $\beta_{p,d}$. We define the *Wulff crystal* $W_{p,d}$ to be the dilate of $\widehat{W}_{p,d}$ about the origin so that $\mathcal{L}^d(W_{p,d}) = 2^d/d!$. The Wulff crystal $W_{p,d}$ is the limit shape which shows up in Theorem 2.2.1. So that this theorem makes sense, we must know that $W_{p,d}$ is contained in $[-1, 1]^d$.

Lemma 4.3.6. For $d \geq 2$ and $p > p_c(d)$, the Wulff crystal $W_{p,d}$ is contained in $[-1, 1]^d$.

Proof. From Corollary 4.2.4 we use the symmetries of $\beta_{p,d}$ to deduce that the unit volume Wulff crystal $\widehat{W}_{p,d}$ satisfies

$$cB_1 \subset \widehat{W}_{p,d} \subset cB_\infty \tag{4.55}$$

for some $c > 0$, where B_1 and B_∞ respectively denote unit ℓ^1- and unit ℓ^∞-balls in $\mathbb{R}^d$ centered at the origin. The claim follows from the fact that $\mathcal{L}^d(B_1) = 2^d/d!$. □

Remark 4.3.7. We may use $\mathcal{I}_{p,d}$ to define an analogous notion of conductance in the continuum: for $E \subset \mathbb{R}^d$ a set of finite perimeter, we define the *conductance* of E as $\mathcal{I}_{p,d}(E)/\theta_p(d)\mathcal{L}^d(E)$. The central theorem of the paper, Theorem 2.2.1, is a statement that the discrete notion of conductance (defined in Section 3.2) scales to the continuum notion just introduced.

CHAPTER 5

Concentration estimates for $\beta_{p,d}$

We now derive concentration estimates for the random variables used to define $\beta_{p,d}$, following an argument of Zhang in Section 9 of [Zha07]. We use results from his paper in conjunction with the following concentration estimate due to Talagrand (we find it aesthetically pleasing to use a result powered by isoperimetry towards the proof of Theorem 2.2.1).

5.1 Key result and application

Theorem 5.1.1. Let (V, E) be a finite graph with $\{X_e\}_{e \in \mathrm{E}}$ a collection of iid Bernoulli(p) random variables. Let $\mathcal{S}$ denote a family of sets of edges and for $S \in \mathcal{S}$, let $X_S = \sum_{e \in S} X_e$. Let $Z_{\mathcal{S}} = \inf_{S \in \mathcal{S}} X_S$, and let $M_{\mathcal{S}}$ be a median of $Z_{\mathcal{S}}$. There is a constant $c = c(p) > 0$ so that for all $u > 0$,

$$\mathbb{P}_p(|Z_{\mathcal{S}} - M_{\mathcal{S}}| \geq u) \leq 4 \exp\left(-c \min\left(\frac{u^2}{\alpha}, u\right)\right), \tag{5.1}$$

where $\alpha = \sup_{S \in \mathcal{S}} |S|$.

Remark 5.1.2. The above theorem of Talagrand was stated in the context of first passage percolation in Section 8.3 of [Tal95]; we have specialized it to Bernoulli percolation. It is clear how our random variables Ξ_{hemi} and Ξ_{face} could be expressed as $Z_{\mathcal{S}}$ for some family of edge sets $\mathcal{S}$. However, use of Theorem 5.1.1 requires control over the size of the largest edge set in $\mathcal{S}$ through the term α. In our case, we must control the size of the largest minimal cut separating opposing hemispheres (or faces) of a cube.

Remark 5.1.3. The chosen orientation S, introduced in the previous section, has been fixed since the beginning of Section 4.3. We also treat $v \in \mathbb{S}^{d-1}$ as fixed for now; to simplify our notation further, we write $\mathfrak{X}(0,v,n)$ as $\mathfrak{X}_n$. Following Zhang in [Zha07], we first use Theorem 5.1.1 to prove concentration for a variant of the $\mathfrak{X}_n$ in which we restrict our attention to cutsets using $O(n^{d-1})$-many edges.

Let $\gamma > 0$, and let $\mathcal{S}_n(\gamma)$ be the family of cutsets in $\mathsf{d\text{-}cyl}(\mathsf{S}(v),1,n)$ which satisfy $|S| \leq \gamma(2n)^{d-1}$, and which separate the hemispheres $\mathsf{d\text{-}hemi}^{\pm}(\mathsf{S}(v),1,n)$. Define

$$Z_n^{(\gamma)}(\omega) := \inf_{S\in\mathcal{S}_n(\gamma)} |S|_\omega\,. \tag{5.2}$$

We now apply Theorem 5.1.1 to $Z_n^{(\gamma)}$, using the bound $\alpha \leq \gamma(2n)^{d-1}$.

Proposition 5.1.4. Let $\epsilon, \gamma > 0$. There are positive constants $c_1(p,\gamma,\epsilon)$ and $c_2(p,\gamma,\epsilon)$ so that for $n \geq 1$,

$$\mathbb{P}_p\left(\frac{|Z_n^{(\gamma)} - \mathbb{E}_p Z_n^{(\gamma)}|}{(2n)^{d-1}} \geq \epsilon\right) \leq c_1 \exp\left(-c_2 n^{(d-1)/3}\right)\,. \tag{5.3}$$

Proof. We follow the argument at the beginning of Section 9 in [Zha07]. Write $A = A(n) := (2n)^{d-1}$ for the $\mathcal{H}^{d-1}$-measure (or "area") of the square $n\mathsf{S}(v)$. Let $M_n^{(\gamma)}$ be the median of $Z_n^{(\gamma)}$. Then,

$$|\mathbb{E}_p Z_n^{(\gamma)} - M_n^{(\gamma)}| \leq \mathbb{E}_p|Z_n^{(\gamma)} - M_n^{(\gamma)}|\,, \tag{5.4}$$

$$\leq \sum_{j=1}^{\lfloor A^{2/3}\rfloor} \mathbb{P}_p(|Z_n^{(\gamma)} - M_n^{(\gamma)}| \geq j) + \sum_{j=\lceil A^{2/3}\rceil}^{\infty} \mathbb{P}_p(|Z_n^{(\gamma)} - M_n^{(\gamma)}| \geq j)\,. \tag{5.5}$$

Apply Theorem 5.1.1 with $\alpha \leq \gamma A$ to the right-most sum above:

$$|\mathbb{E}_p Z_n^{(\gamma)} - M_n^{(\gamma)}| \leq A^{2/3} + 4\left(\sum_{j=\lceil A^{4/3}\rceil}^{\infty} \exp\left(-c\frac{j}{\gamma A}\right) + \sum_{j=\lceil A^{2/3}\rceil}^{\infty} \exp(-cj)\right)\,, \tag{5.6}$$

$$\leq A^{2/3} + \frac{4}{1-\exp(-c/\gamma A)}\exp(-cA^{1/3}/\gamma) + \frac{4}{1-\exp(-c)}\exp(-cA^{2/3})\,. \tag{5.7}$$

Thus for all n sufficiently large, in a way which depends on p and γ, we have $|\mathbb{E}_p Z_n^{(\gamma)} - M_n^{(\gamma)}| \leq (3/2)A^{2/3}$. We may thus use the triangle inequality to conclude that, for such n,

$$\mathbb{P}_p(|Z_n^{(\gamma)} - \mathbb{E}_p Z_n^{(\gamma)}| \geq 4A^{2/3}) \leq \mathbb{P}_p(|Z_n^{(\gamma)} - M_n^{(\gamma)}| + |M_n^{(\gamma)} - \mathbb{E}_p Z_n^{(\gamma)}| \geq 4A^{2/3})\,, \tag{5.8}$$

$$\leq \mathbb{P}_p(|Z_n^{(\gamma)} - M_n^{(\gamma)}| \geq 2A^{2/3})\,. \tag{5.9}$$

Another application of Theorem 5.1.1 gives that, when n is sufficiently large depending on p and γ,

$$\mathbb{P}_p\left(|Z_n^{(\gamma)} - \mathbb{E}_p Z_n^{(\gamma)}| \geq 4A^{2/3}\right) \leq 4\exp\left(-c\min\left(\frac{4}{\gamma}A^{1/3}, 2A^{2/3}\right)\right), \tag{5.10}$$

and the proof is complete. □

Remark 5.1.5. To use Proposition 5.1.4 to obtain a statement about the $\mathfrak{X}_n$, we need Proposition 4.2 of Rossignol and Théret [RT10b], which we specialize to Bernoulli percolation. As is mentioned in [RT10b], this result is due to Kesten (Proposition 5.8 of [Kes86]) in dimension two. The difficulties which appear in higher dimensions are settled through Theorem 1 of Zhang [Zha07]. Zhang's argument involves an intricate construction which lies at the heart of Chapter 7 in the present paper.

For a percolation configuration ω, let $N_n(\omega)$ denote the minimum cardinality $|S|$ over all cutsets S in $\mathsf{d\text{-}cyl}(\mathsf{S}(v), 1, n)$ separating $\mathsf{d\text{-}hemi}^{\pm}(\mathsf{S}(v), 1, n)$ such that $|S|_\omega = \mathfrak{X}_n(\omega)$. The following is Proposition 4.2 of [RT10b].

Proposition 5.1.6. Let $d \geq 2$ and let $p > p_c(d)$. There are positive constants $\gamma(p,d), c_1(p,d)$ and $c_2(p,d)$ so that for all $u > 0$ and all $n \geq 1$,

$$\mathbb{P}_p(N_n \geq \gamma u \text{ and } \mathfrak{X}_n \leq u) \leq c_1 \exp(-c_2 u)\,. \tag{5.11}$$

Using Proposition 5.1.6 with Proposition 5.1.4, we deduce the following.

Corollary 5.1.7. Let $d \geq 2$, $p > p_c(d)$, $v \in \mathbb{S}^{d-1}$ and let $\epsilon > 0$. There are positive constants $c_1(p,d,\epsilon)$ and $c_2(p,d,\epsilon)$ so that for all $n \geq 1$,

$$\mathbb{P}_p\left(\frac{|\mathfrak{X}(0,v,n) - \mathbb{E}_p\mathfrak{X}(0,v,n)|}{(2n)^{d-1}} \geq \epsilon\right) \leq c_1\exp\left(-c_2 n^{(d-1)/3}\right). \tag{5.12}$$

Proof. Fix $v \in \mathbb{S}^{d-1}$ and abbreviate $\mathfrak{X}(0,v,n)$ as $\mathfrak{X}_n$. As remarked at the end of the proof of Proposition 4.2.3, we have that uniformly in $v \in \mathbb{S}^{d-1}$ and all percolation configurations

ω, we have $\mathfrak{X}_n(\omega) \le c(d)n^{d-1}$ for some positive constant $c(d)$. Let $u = c(d)n^{d-1}$ and apply Proposition 5.1.6 to obtain γ depending on p and d so that

$$\mathbb{P}_p(N_n \ge \gamma c(d)n^{d-1}) \le c_1 \exp\left(-c_2 c(d) n^{d-1}\right) . \tag{5.13}$$

We will use this bound shortly. For this γ and for $\epsilon > 0$, use Proposition 5.1.4 to obtain positive constants $c_1(p, \gamma, \epsilon)$ and $c_2(p, \gamma, \epsilon)$ so that

$$\mathbb{P}_p\left(\frac{|\mathfrak{X}_n - \mathbb{E}_p\mathfrak{X}_n|}{(2n)^{d-1}} \ge \epsilon\right) \le \mathbb{P}_p(Z_n^{(\gamma)} \ne \mathfrak{X}_n) + \mathbb{P}_p\left(\frac{|Z_n^{(\gamma)} - \mathbb{E}_p Z_n^{(\gamma)}|}{(2n)^{d-1}} \ge \epsilon\right) , \tag{5.14}$$

$$\le \mathbb{P}_p(Z_n^{(\gamma)} \ne \mathfrak{X}_n) + c_1 \exp\left(-c_2 n^{(d-1)/3}\right) . \tag{5.15}$$

Observe that $\{Z_n^{(\gamma)} \ne \mathfrak{X}_n\} \subset \{N_n \ge \gamma c(d) n^{d-1}\}$, so that by (5.13),

$$\mathbb{P}_p\left(\frac{|\mathfrak{X}_n - \mathbb{E}_p\mathfrak{X}_n|}{(2n)^{d-1}} \ge \epsilon\right) \le \mathbb{P}_p\left(N_n \ge \gamma c(d) n^{d-1}\right) + c_1 \exp\left(-c_2 n^{(d-1)/3}\right) , \tag{5.16}$$

$$\le c_1 \exp\left(-c_2 c(d) n^{d-1}\right) + c_1 \exp\left(-c_2 n^{(d-1)/3}\right) . \tag{5.17}$$

The proof is complete. □

We obtain the desired concentration estimates by combining Corollary 5.1.7 with Proposition 4.3.2, which tells us the functions $v \mapsto \mathbb{E}_p\mathfrak{X}(0, v, n)/(2n)^{d-1}$ converge uniformly to $\beta_{p,d}$. The following is the main result of the section.

Theorem 5.1.8. Let $d \ge 2$, $p > p_c(d)$, and let $\epsilon > 0$. There are positive constants $c_1(p, d, \epsilon), c_2(p, d, \epsilon)$ so that for all $x \in \mathbb{R}^d, v \in \mathbb{S}^{d-1}$, and all $r > 0$,

$$\mathbb{P}_p\left(\left|\frac{\mathfrak{X}(x, v, r)}{(2n)^{d-1}} - \beta_{p,d}(v)\right| \ge \epsilon\right) \le c_1 \exp\left(-c_2 r^{(d-1)/3}\right) . \tag{5.18}$$

Proof. We first prove (5.18) in the case that $x = 0$ and $r = n \in \mathbb{N}$. By Proposition 4.3.2, we may choose n_0 large depending on ϵ so that for all $v \in \mathbb{S}^{d-1}$, $n \ge n_0$ implies

$$\left|\frac{\mathbb{E}_p\mathfrak{X}(0, v, n)}{(2n)^{d-1}} - \beta_{p,d}(v)\right| < \epsilon/2 . \tag{5.19}$$

For $n \geq n_0$ we then have

$$\mathbb{P}_p\left(\left|\frac{\mathfrak{X}(0,v,n)}{(2n)^{d-1}} - \beta_{p,d}(v)\right| \geq \epsilon\right) \tag{5.20}$$

$$\leq \mathbb{P}_p\left(\frac{|\mathfrak{X}(0,v,n) - \mathbb{E}_p\mathfrak{X}(0,v,n)|}{(2n)^{d-1}} + \left|\frac{\mathbb{E}_p\mathfrak{X}(0,v,n)}{(2n)^{d-1}} - \beta_{p,d}(v)\right| \geq \epsilon\right), \tag{5.21}$$

$$\leq \mathbb{P}_p\left(\frac{|\mathfrak{X}(0,v,n) - \mathbb{E}_p\mathfrak{X}(0,v,n)|}{(2n)^{d-1}} \geq \epsilon/2\right), \tag{5.22}$$

and (5.18) is shown to hold by applying the concentration estimates Corollary 5.1.7 to the right-hand side.

We now consider general x and r; we claim that for any $r > 0$ and $x \in \mathbb{R}^d$, there is a positive constant $c(d)$ so that

$$\mathfrak{X}(x,v,\lceil r\rceil) - c(d)r^{d-2} \leq \mathfrak{X}(x,v,r) \leq \mathfrak{X}(x,v,\lfloor r\rfloor) + c(d)r^{d-2}. \tag{5.23}$$

To see this, let A be the collection of edges having non-empty intersection with

$$\mathcal{N}_{5d}(r\mathsf{S}(v) \setminus \lfloor r\rfloor\mathsf{S}(v)), \tag{5.24}$$

and let E be a minimal cutset in $\mathsf{d\text{-}cyl}(\mathsf{S}(v),1,\lfloor r\rfloor)$ separating $\mathsf{d\text{-}hemi}^{\pm}(\mathsf{S}(v),1,\lfloor r\rfloor)$. Our standard argument from Lemma 4.2.2 tells us that the edges of $E \cup A$ contained in $\mathsf{d\text{-}cyl}(\mathsf{S}(v),1,r)$ separate the hemispheres $\mathsf{d\text{-}hemi}^{\pm}(\mathsf{S}(v),1,r)$. That $|A| \leq c(d)r^{d-2}$ establishes the upper bound on $\mathfrak{X}(0,v,r)$ in (5.23), and we obtain the lower bound through a similar procedure.

The proof of Lemma 4.2.2 also tells us that for $x \in \mathbb{R}^d$ and $r > 0$, there exists $x' \in \mathbb{Z}^d$ so that

$$\mathfrak{X}(x',v,r+d^{1/2}) - c(d)r^{d-2} \leq \mathfrak{X}(x,v,r) \leq \mathfrak{X}(x',v,r-d^{1/2}) + c(d)r^{d-2}. \tag{5.25}$$

Apply (5.23), to conclude

$$\mathfrak{X}\left(x',v,\lceil r+d^{1/2}\rceil\right) - c(d)r^{d-2} \leq \mathfrak{X}(x,v,r) \leq \mathfrak{X}\left(x',v,\lfloor r-d^{1/2}\rfloor\right) + c(d)r^{d-2}. \tag{5.26}$$

As $x' \in \mathbb{Z}^d$, the random variables $\mathfrak{X}(x',v,\lceil r+d^{1/2}\rceil)$ and $\mathfrak{X}(x',v,\lfloor r-d^{1/2}\rfloor)$ have the same law as $\mathfrak{X}(0,v,\lceil r+d^{1/2}\rceil)$ and $\mathfrak{X}(0,v,\lfloor r-d^{1/2}\rfloor)$ respectively, so the concentration estimates

(5.18) established in the case that $x = 0$ and $r = n \in \mathbb{N}$ hold for these random variables as well. Within the high probability event

$$\left\{ \left| \frac{\mathfrak{X}\left(x', v, \lceil r + d^{1/2} \rceil\right)}{\left(2\lceil r + d^{1/2} \rceil\right)^{d-1}} - \beta_{p,d}(v) \right| < \epsilon \right\} \cap \left\{ \left| \frac{\mathfrak{X}\left(x', v, \lfloor r - d^{1/2} \rfloor\right)}{\left(2\lfloor r - d^{1/2} \rfloor\right)^{d-1}} - \beta_{p,d}(v) \right| < \epsilon \right\}, \quad (5.27)$$

and for taken r sufficiently large in a way depending on ϵ and d, we obtain

$$\beta_{p,d}(v) - 3\epsilon \le \frac{\mathfrak{X}(x, v, r)}{(2r)^{d-1}} \le \beta_{p,d}(v) + 3\epsilon\,, \quad (5.28)$$

which completes the proof. □

CHAPTER 6

Consequences of concentration estimates

We now derive important consequences of Theorem 5.1.8. In Section 6.1, we will obtain information about the random variables Ξ_{hemi} and Ξ_{face} for cylinders of small height which are based at squares and discs. This result will be crucial in Chapter 9.

In Section 6.2, we use a result powered by Theorem 5.1.8 in conjunction with Gandolfi's results (presented in the appendix) on the density of $\mathbf{C}_\infty$ in large boxes to obtain high probability upper bounds on the Cheeger constant. In fact, any polytope $P \subset [-1,1]^d$ satisfying $\mathcal{L}^d(P) \leq 2^d/d!$ gives rise to such a bound. Specializing these results to a sequence of polytopes which are progressively better approximates of the Wulff crystal, we obtain the easier half of Theorem 2.2.2.

6.1 Lower bounds for cuts in thin cylinders

We apply our concentration estimates to random variables of the form Ξ_{face} for cylinders of small height. Because of this, it is important to recall the convention established in Remark 4.1.2. Throughout this section, we adopt the notation $\mathsf{S}(x,v) := \mathsf{S}(v) + x$.

Lemma 6.1.1. Let $d \geq 2$, $p > p_c(d)$ and let $\epsilon > 0$. There exists $\eta(p,d,\epsilon) > 0$ small and positive constants $c_1(p,d,\epsilon)$, $c_2(p,d,\epsilon)$ so that for all $x \in \mathbb{R}^d$, all $v \in \mathbb{S}^{d-1}$, $h \in (0,\eta)$, and $r > 0$ taken sufficiently large depending on h, we have

$$\mathbb{P}_p\left(\Xi_{\mathrm{face}}(\mathsf{S}(x,v),h,r) \leq (1-\epsilon)\mathcal{H}^{d-1}(r\mathsf{S}(x,v))\beta_{p,d}(v)\right) \leq c_1 \exp\left(-c_2 r^{(d-1)/3}\right) . \qquad (6.1)$$

Proof. Write S for $\mathsf{S}(x,v)$ and consider a minimal cutset E in $\mathsf{d\text{-}cyl}(S,h,r)$ separating the faces $\mathsf{d\text{-}face}^{\pm}(S,h,r)$. Recall that S_h^+ and S_h^- are the top and bottom faces of the cylinder

$\mathsf{cyl}(S,h)$, and consider the collection of edges A having non-empty intersection with the neighborhood

$$\mathcal{N}_{5d}(r(\partial\mathsf{cyl}(S,h)\setminus(S_h^+\cup S_h^-)))\,. \tag{6.2}$$

The edges of $E\cup A$ contained in $\mathsf{d\text{-}cyl}(S,h,r)$ separate the hemispheres $\mathsf{d\text{-}hemi}^{\pm}(S,h,r)$ within $\mathsf{d\text{-}cyl}(S,h,r)$. It follows that these edges also separate the hemispheres $\mathsf{d\text{-}hemi}^{\pm}(S,1,r)$ in the larger cylinder $\mathsf{d\text{-}cyl}(S,1,r)$, provided that the parameters h and r are suitable for S in the sense of Remark 4.1.2. It is for this reason that, in the statement of this lemma, we must take r sufficiently large depending on h. By construction, the cardinality of A is at most $c(d)hr^{d-1}$ for some positive constant $c(d)$, so that

$$\mathfrak{X}(x,v,r)\le \Xi_{\text{face}}(S,h,r)+c(d)hr^{d-1}\,, \tag{6.3}$$

and thus,

$$\Big\{\Xi_{\text{face}}(S,h,r)\le(1-\epsilon)\mathcal{H}^{d-1}(rS)\beta_{p,d}(v)\Big\} \tag{6.4}$$

$$\subset\Big\{\mathfrak{X}(x,v,r)\le(1-\epsilon)\mathcal{H}^{d-1}(rS)\beta_{p,d}(v)+c(d)hr^{d-1}\Big\}\,, \tag{6.5}$$

$$\subset\Big\{\mathfrak{X}(x,v,r)\le(1-\epsilon/2)\mathcal{H}^{d-1}(rS)\beta_{p,d}(v)\Big\}\,, \tag{6.6}$$

where we have chosen h small depending on p,d,ϵ to obtain the second line directly above. We complete the proof by applying Theorem 5.1.8 to the event on the second line. □

We now prove the analogue of Lemma 6.1.1 for cylinders of small height based at discs instead of squares.

Proposition 6.1.2. Let $d\ge 2$, $p>p_c(d)$ and let $\epsilon>0$. Given $x\in\mathbb{R}^d$ and $v\in\mathbb{S}^{d-1}$, let $D(x,v)$ be the isometric image of the unit Euclidean ball in $\mathbb{R}^{d-1}$ centered at x and oriented so that $\mathsf{hyp}(D(x,v))$ is orthogonal to v. There exists $\eta(p,d,\epsilon)>0$ small as well as positive constants $c_1(p,d,\epsilon)$ and $c_2(p,d,\epsilon)$ so that for all $x\in\mathbb{R}^d$, $v\in\mathbb{S}^{d-1}$, $h\in(0,\eta)$ and $r>0$ taken sufficiently large depending on h, we have

$$\mathbb{P}_p\Big(\Xi_{\text{face}}(D(x,v),h,r)\le(1-\epsilon)\mathcal{H}^{d-1}(rD(x,v))\beta_{p,d}(v)\Big)\le c_1\exp\left(-c_2r^{(d-1)/3}\right)\,. \tag{6.7}$$

Proof. Let $\epsilon' > 0$, write $D = D(x, v)$ and let D' denote the closed Euclidean unit disc in $\mathbb{R}^{d-1}$ defined by $\{x \in \mathbb{R}^{d-1} : |x|_2 \leq 1\}$. Let $\varphi : D' \to \mathbb{R}^d$ be an isometry so that $\varphi(D') = D$. Let Δ^k denote the collection of closed dyadic squares in $\mathbb{R}^{d-1}$ at scale k, so that each square is a translate of $2^{-k}[-1,1]^{d-1}$. Choose $k \in \mathbb{N}$ large enough (depending on ϵ' and d) so that

$$\mathcal{L}^{d-1}\left(D' \setminus \bigcup_{S' \in \Delta^k,\, S' \subset D'} S'\right) \leq \epsilon' \mathcal{L}^{d-1}(D')\,, \tag{6.8}$$

and let us enumerate the squares $S' \in \Delta^k$ with $S' \subset D'$ as $S'_1, \ldots, S'_m$, noting that the number m of these squares depends on ϵ' and d. Shrink each square slightly to form a new disjoint collection $\{S''_i\}_{i=1}^m$ of closed squares. Specifically, S''_i shall be the $(1-\delta)$-dilate of S'_i about its center for some $\delta \in (0,1)$. For each i, define $S_i := \varphi(S''_i)$ and choose δ sufficiently small (also depending on ϵ' and d) so that

$$\mathcal{H}^{d-1}\left(D \setminus \bigcup_{i=1}^m S_i\right) \leq 2\epsilon' \mathcal{H}^{d-1}(D)\,. \tag{6.9}$$

Let $\alpha = (1-\delta)2^{-k}$. We arrange that the isometry φ is compatible with the chosen orientation, in the sense that for each i, we have $S_i = \alpha \mathsf{S}(y_i, v)$ for some $y_i \in \mathbb{R}^d$. Let $\epsilon > 0$, choose $\eta = \eta(p, d, \epsilon/2)$ as in Lemma 6.1.1 and let $h \in (0, \alpha\eta)$.

Let E be a minimal cutset separating the faces $\mathsf{d\text{-}face}^{\pm}(D, h, r)$ within $\mathsf{d\text{-}cyl}(D, h, r)$. Let E_i denote the set of all edges of E lying in edge set of $\mathsf{d\text{-}cyl}(S_i, h, r)$. Each E_i separates the faces of $\mathsf{d\text{-}face}^{\pm}(S_i, h, r)$ within $\mathsf{d\text{-}cyl}(S_i, h, r)$, so by the disjointness of the collection $\{S_i\}_{i=1}^m$, we have

$$\sum_{i=1}^m \Xi_{\text{face}}(S_i, h, r) \leq \Xi_{\text{face}}(D, h, r)\,, \tag{6.10}$$

and thus,

$$\mathbb{P}_p\Big(\Xi_{\text{face}}(D, h, r) \leq (1-\epsilon)\mathcal{H}^{d-1}(rD)\beta_{p,d}(v)\Big) \tag{6.11}$$

$$\leq \mathbb{P}_p\left(\sum_{i=1}^m \Xi_{\text{face}}(S_i, h, r) \leq (1-\epsilon)\mathcal{H}^{d-1}(rD)\beta_{p,d}(v)\right), \tag{6.12}$$

$$\leq \mathbb{P}_p\left(\sum_{i=1}^m \Xi_{\text{face}}(S_i, h, r) \leq \frac{(1-\epsilon)}{1-2\epsilon'}\sum_{i=1}^m \mathcal{H}^{d-1}(rS_i)\beta_{p,d}(v)\right), \tag{6.13}$$

with (6.13) following from our choice of δ and the squares S_i'. As $\mathcal{H}^{d-1}(rS_i)$ is the same for each i, we use a union bound to obtain

$$\mathbb{P}_p\Big(\Xi_{\text{face}}(D,h,r) \le (1-\epsilon)\mathcal{H}^{d-1}(rD)\beta_{p,d}(v)\Big) \tag{6.14}$$

$$\le \sum_{i=1}^{m} \mathbb{P}_p\left(\Xi_{\text{face}}(S_i,h,r) \le \left(\frac{1-\epsilon}{1-2\epsilon'}\right)\beta_{p,d}(v)\mathcal{H}^{d-1}(rS_i)\right), \tag{6.15}$$

$$\le \sum_{i=1}^{m} \mathbb{P}_p\left(\Xi_{\text{face}}(S_i,h,r) \le (1-\epsilon/2)\beta_{p,d}(v)\mathcal{H}^{d-1}(rS_i)\right). \tag{6.16}$$

To obtain (6.16), we have taken ϵ' small enough so that $1-\epsilon/2 > \frac{1-\epsilon}{1-2\epsilon'}$. Thus, m and α now depend on ϵ and d. We now use that φ was chosen to be compatible with the chosen orientation S, and we write each S_i as $\alpha\mathsf{S}(y_i,v)$ for some y_i. Making this switch in (6.16), we have

$$\mathbb{P}_p\Big(\Xi_{\text{face}}(D,h,r) \le (1-\epsilon)\mathcal{H}^{d-1}(rD)\beta_{p,d}(v)\Big) \tag{6.17}$$

$$\le \sum_{i=1}^{m} \mathbb{P}_p\left(\Xi_{\text{face}}(\mathsf{S}(y_i,v),h/\alpha,\alpha r) \le (1-\epsilon/2)\beta_{p,d}(v)\mathcal{H}^{d-1}(\alpha r\mathsf{S}(y_i,v))\right). \tag{6.18}$$

Having chosen h so that $h/\alpha \le \eta$, we may apply Lemma 6.1.1 to each summand directly above, using αr in place of r and $\epsilon/2$ in place of ϵ in the statement of this lemma:

$$\mathbb{P}_p\Big(\Xi_{\text{face}}(D,h,r) \le (1-\epsilon)\mathcal{H}^{d-1}(rD)\beta(v)\Big) \le mc_1\exp\left(-c_2(\alpha r)^{(d-1)/3}\right). \tag{6.19}$$

To apply Lemma 6.1.1, we must take r sufficiently large depending on h. The proof is complete upon renaming these constants and making note of their dependencies; we take $\alpha\eta$ to be the η in the statement of this proposition. □

Up until this point, we have only used half of our concentration estimates from Chapter 5, in the sense that we have only used these estimates to show that with high probability, the random variables Ξ_{face} cannot be too small. In the next section, we put the complementary estimates to good use.

6.2 Upper bounds on $\widehat{\Phi}_n$, or efficient carvings of ice

We define a *convex polytope* to be a compact subset of $\mathbb{R}^d$ which may be written as a finite intersection of closed half-spaces. We define a *polytope* to be a compact subset of $\mathbb{R}^d$ which may be written as finite union of convex polytopes. In particular, we do not require polytopes to be connected subsets of $\mathbb{R}^d$, but we say that a polytope is *connected* if its interior is a connected subset of $\mathbb{R}^d$. We say that $P \subset \mathbb{R}^d$ is a *d-polytope* if it is non-degenerate ($\mathcal{L}^d(P) > 0$).

As stated at the beginning of the section, our goal is now to use a polytope $P \subset [-1,1]^d$ to obtain upper bounds on $\widehat{\Phi}_n$. To accomplish this, we will use P to obtain a valid subgraph of $\mathbf{C}_n$, and we will control both the volume and open edge boundary of this subgraph. Equivalently, we view $\mathbf{C}_n$ as a block of ice, and we use the dilate nP as a blueprint for carving this block.

Our first task is to perform an "efficient" carving at the boundary of nP, and this is where the other side of our concentration estimates are used. It is convenient to prove a result, Proposition 6.2.1, which allows us to work on each face of the polytope P individually. We think of a $(d-1)$-polytope $\sigma \subset \mathbb{R}^d$ as one of the faces of a d-polytope P. For such σ, we let v_σ denote one of the unit vectors orthogonal to $\mathsf{hyp}(\sigma)$.

Proposition 6.2.1. Let $d \geq 2$ and let $p > p_c(d)$. Let $\sigma \subset \mathbb{R}^d$ be a connected $(d-1)$-polytope, and let $\epsilon > 0$. There is a positive constant $\eta(p,d,\epsilon,\sigma)$ and another connected $(d-1)$-polytope $\widetilde{\sigma}$ depending on σ, p, d and ϵ so that:

(i) $\widetilde{\sigma} \subset \sigma$, $\widetilde{\sigma} \cap \mathcal{N}_\eta(\partial\sigma) = \emptyset$ and $\mathcal{H}^{d-1}(\sigma \setminus \widetilde{\sigma}) \leq \epsilon \mathcal{H}^{d-1}(\sigma)$.

(ii) There are positive constants $c_1(p,d,\epsilon,\sigma)$ and $c_2(p,d,\epsilon,\sigma)$ so that when $h \in (0,\eta)$ and for all $r > 0$ taken sufficiently large in a way depending on h,

$$\mathbb{P}_p\Big(\Xi_{\mathrm{hemi}}(\widetilde{\sigma},h,r) \geq (1+\epsilon)\mathcal{H}^{d-1}(r\sigma)\beta_{p,d}(v_\sigma)\Big) \leq c_1 \exp\left(-c_2 r^{(d-1)/3}\right) . \tag{6.20}$$

Proof. Let $\epsilon' > 0$, and for the parameter $\eta > 0$, define $\widetilde{\sigma}$ to be the closure of $\sigma \setminus \mathcal{N}^{(1)}_{2\eta}(\partial\sigma)$. Then $\widetilde{\sigma}$ is a $(d-1)$ polytope, and because σ is connected (see the definition of a connected polytope given at the beginning of the present section), we may choose η sufficiently small depending on σ and ϵ so that $\widetilde{\sigma}$ is also connected, and so that

$$\mathcal{H}^{d-1}(\sigma \setminus \widetilde{\sigma}) \leq \epsilon' \mathcal{H}^{d-1}(\sigma) \,. \tag{6.21}$$

For such $\widetilde{\sigma}$, property *(i)* is already satisfied after stipulating $\epsilon' \leq \epsilon$.

To show *(ii)* holds, we employ the strategy used in the proof of Proposition 6.1.2. Let $h \leq \eta$, and let $\widetilde{\sigma}' \subset \mathbb{R}^{d-1}$ be a $(d-1)$-polytope with an isometry $\varphi : \widetilde{\sigma}' \to \widetilde{\sigma}$. Choose $k \in \mathbb{N}$ to be the smallest k such that $2^{-k} < h$, and large enough so that

$$\mathcal{L}^{d-1}\left(\widetilde{\sigma}' \setminus \bigcup_{S' \in \Delta^k,\, S' \subset \widetilde{\sigma}'} S'\right) \leq \epsilon' \mathcal{L}^{d-1}(\widetilde{\sigma}') \,, \tag{6.22}$$

where, as before, Δ^k denotes the collection of dyadic squares in $\mathbb{R}^{d-1}$ at scale k. We enumerate such squares contained in $\widetilde{\sigma}'$ as $S'_1, \dots, S'_m$. Let $\delta > 0$ and dilate each S'_i about its center by a factor of $(1-\delta)$ to produce a new, disjoint collection $\{S''_i\}_{i=1}^m$ of closed squares contained in $\widetilde{\sigma}'$. Let $S_i = \varphi(S''_i)$, and choose δ small enough so that

$$\mathcal{H}^{d-1}\left(\widetilde{\sigma} \setminus \bigcup_{i=1}^m S_i\right) < 2\epsilon' \mathcal{H}^{d-1}(\sigma) \,. \tag{6.23}$$

Let $\alpha = (1-\delta)2^{-k}$, and we lose no generality assume that $\widetilde{\sigma}'$ and φ are compatible with the chosen orientation, so that each S_i is $\alpha \mathsf{S}(y_i, v_\sigma)$ for some $y_i \in \mathbb{R}^d$.

For each i, let E_i denote a cutset in $\mathsf{d\text{-}cyl}(S_i, \alpha, r)$ separating $\mathsf{d\text{-}hemi}^{\pm}(S_i, \alpha, r)$. Let A denote the collection of edges having non-empty intersection with

$$\mathcal{N}_{5d}\left(r\left(\widetilde{\sigma} \setminus \bigcup_{i=1}^m S_i\right)\right) \tag{6.24}$$

so that $|A| \leq c(d)\epsilon' \mathcal{H}^{d-1}(r\sigma)$ for some $c(d) > 0$. The argument from the proof of Lemma 4.2.2 used throughout Chapter 4 tells us the edges of $A \cup \bigcup_{i=1}^m E_i$ which are contained in $\mathsf{d\text{-}cyl}(\widetilde{\sigma}, h, r)$ separate the hemispheres $\mathsf{d\text{-}hemi}^{\pm}(\widetilde{\sigma}, h, r)$. To ensure this argument goes through, we must make sure r is taken sufficiently large depending on h. We note that it is here we use the fact that we chose k large enough so that $\alpha \leq 2^{-k} \leq h$. Under these conditions, we conclude that

$$\Xi_{\text{hemi}}(\widetilde{\sigma}, h, r) \leq c(d)\epsilon' \mathcal{H}^{d-1}(r\sigma) + \sum_{i=1}^m \Xi_{\text{hemi}}(S_i, \alpha, r) \,, \tag{6.25}$$

and thus,

$$\mathbb{P}_p\Big(\Xi_{\text{hemi}}(\widetilde{\sigma}, h, r) \geq (1+\epsilon)\mathcal{H}^{d-1}(r\sigma)\beta_{p,d}(v_\sigma)\Big) \tag{6.26}$$

$$\leq \mathbb{P}_p\left(\sum_{i=1}^{m} \Xi_{\text{hemi}}(S_i, \alpha, r) \geq (1+\epsilon - c(p,d)\epsilon')\mathcal{H}^{d-1}(r\sigma)\beta_{p,d}(v_\sigma)\right). \tag{6.27}$$

We now choose ϵ' small enough depending on p, d, ϵ so that $1+\epsilon-c(p,d)\epsilon' \geq 1+\epsilon/2$.

$$\mathbb{P}_p\Big(\Xi_{\text{hemi}}(\widetilde{\sigma}, h, r) \geq (1+\epsilon)\mathcal{H}^{d-1}(r\sigma)\beta_{p,d}(v_\sigma)\Big) \tag{6.28}$$

$$\leq \mathbb{P}_p\left(\sum_{i=1}^{m} \Xi_{\text{hemi}}(S_i, \alpha, r) \geq (1+\epsilon/2)\sum_{i=1}^{m}\mathcal{H}^{d-1}(rS_i)\beta_{p,d}(v_\sigma)\right), \tag{6.29}$$

$$\leq \sum_{i=1}^{m}\mathbb{P}_p\left(\Xi_{\text{hemi}}(S_i, \alpha, r) \geq (1+\epsilon/2)\mathcal{H}^{d-1}(rS_i)\beta_{p,d}(v_\sigma)\right), \tag{6.30}$$

$$\leq \sum_{i=1}^{m}\mathbb{P}_p\left(\mathfrak{X}(y_i, v_\sigma, \alpha r) \geq (1+\epsilon/2)\mathcal{H}^{d-1}(\alpha r\mathsf{S}(y_i, v_\sigma))\beta_{p,d}(v_\sigma)\right). \tag{6.31}$$

where we have used a union bound and the fact that each $S_i = \alpha\mathsf{S}(y_i, v_\sigma)$ for some $y_i \in \mathbb{R}^d$. We apply our concentration estimates, Theorem 5.1.8, to each summand on the right, and we obtain

$$\mathbb{P}_p\Big(\Xi_{\text{hemi}}(\widetilde{\sigma}, h, r) \geq (1+\epsilon)\mathcal{H}^{d-1}(r\sigma)\beta_p(v_\sigma)\Big) \leq mc_1\exp\left(-c_2(\alpha r)^{(d-1)/3}\right), \tag{6.32}$$

which completes the proof. □

We now use a d-polytope P to obtain a high probability upper bound on $\widehat{\Phi}_n$ in terms of the conductance of P.

Theorem 6.2.2. Let $d \geq 2$ and let $p > p_c(d)$. Let $P \subset [-1,1]^d$ be a polytope such that $\mathcal{L}^d(P) \leq 2^d/d!$, and let $\epsilon > 0$. There exist positive constants $c_1(p,d,\epsilon,P)$ and $c_2(p,d,\epsilon,P)$ so that

$$\mathbb{P}_p\left(\widehat{\Phi}_n \geq (1+\epsilon)\left(\frac{\mathcal{I}_{p,d}(nP)}{\theta_p(d)\mathcal{L}^d(nP)}\right)\right) \leq c_1\exp\left(-c_2 n^{(d-1)/3}\right). \tag{6.33}$$

Proof. We begin by working with $P_\delta := (1-\delta)P$ for $\delta \in (0,1)$; we will choose δ carefully towards the end of the argument.

Note that the Euclidean distance from P_δ to $\partial[-1,1]^d$ is positive. Let $\epsilon, \epsilon' > 0$ and enumerate the faces of P_δ as $\sigma_1, \dots, \sigma_m$, suppressing the dependence of these faces on δ. We use Proposition 6.2.1, picking η depending on ϵ' and on each face σ_i. For each i, we produce a $\widetilde{\sigma}_i$ such that $\mathcal{H}^{d-1}(\sigma_i \setminus \widetilde{\sigma}_i) \leq \epsilon' \mathcal{H}^{d-1}(\sigma_i)$ and so that

(i) $\widetilde{\sigma}_i \subset \sigma_i$, $\widetilde{\sigma}_i \cap \mathcal{N}_\eta(\partial \sigma_i) = \emptyset$ and $\mathcal{H}^{d-1}(\sigma_i \setminus \widetilde{\sigma}_i) \leq \epsilon' \mathcal{H}^{d-1}(\sigma_i)$.

(ii) There are positive constants $c_1(p, d, \epsilon', P, \delta)$ and $c_2(p, d, \epsilon', P, \delta)$ so that if $h \in (0, \eta)$, and if $r > 0$ is taken sufficiently large depending on h,

$$\mathbb{P}_p\Big(\Xi_{\text{hemi}}(\widetilde{\sigma}_i, h, r) \geq (1+\epsilon')\mathcal{H}^{d-1}(r\sigma_i)\beta_{p,d}(v_\sigma)\Big) \leq c_1 \exp\left(-c_2 r^{(d-1)/3}\right) . \tag{6.34}$$

We assume that η is small enough so that the closed neighborhood $\mathcal{N}_\eta(P_\delta)$ is contained within $[-1,1]^d$, and we choose $h \in (0, \eta)$ so that the cylinders $\left\{\mathsf{cyl}(\widetilde{\sigma}_i, h, n)\right\}_{i=1}^m$ are disjoint. Upon choosing such an h, we treat h as fixed for the remainder of the proof. This puts us in a position to apply Proposition 6.2.1 in each cylinder $\mathsf{cyl}(\widetilde{\sigma}_i, h, n)$. This will ultimately allow us to control the open edge boundary of a subgraph of $\mathbf{C}_n$, which we will construct momentarily. Before doing so, we position ourselves to control the volume of this subgraph.

Let $Q_1, \dots, Q_\ell$ enumerate the dyadic cubes at scale k within $[-1,1]^d$. We suppose these cubes are ordered so that for $\ell_1 \leq \ell_2 \in \{1, \dots, \ell\}$, we have that $Q_1, \dots, Q_{\ell_2}$ enumerates all such dyadic cubes having non-empty intersection with $\mathcal{N}_\eta(P_\delta)$, and also that $Q_1, \dots, Q_{\ell_1}$ enumerates all such cubes contained within $P_\delta \setminus \mathcal{N}_\eta(P_\delta)$. We take k sufficiently large and take η smaller if necessary so that

$$\mathcal{L}^d\left(\bigcup_{j=\ell_1+1}^{\ell_2} Q_j\right) < \epsilon' \mathcal{L}^d(P_\delta)\,, \tag{6.35}$$

and for each $j \in \{1, \dots, \ell\}$, let $\mathcal{E}_n^{(j)}$ be the event that

$$\left\{\frac{|\mathbf{C}_\infty \cap nQ_j|}{\mathcal{L}^d(nQ_j)} \in (\theta_p(d) - \epsilon', \theta_p(d) + \epsilon')\right\} . \tag{6.36}$$

We will now construct a (random) subgraph of $\mathbf{C}_n$ from the polytope P. Fix a percolation configuration ω, and for each face σ_i, let $E_n^{(i)}(\omega)$ denote a cutset within $\mathsf{d\text{-}cyl}(\widetilde{\sigma}_i, h, n)$ separating $\mathsf{d\text{-}hemi}^\pm(\widetilde{\sigma}_i, h, n)$, and such that $|E_n^{(i)}(\omega)|_\omega$ is $\Xi_{\text{hemi}}(\widetilde{\sigma}_i, h, n)$ in the configuration ω.

We let A_n be the collection of edges having non-empty intersection with the neighborhood

$$\mathcal{N}_{5d}\left(n\left(\partial P_\delta \setminus \bigcup_{i=1}^{m} \widetilde{\sigma}_i\right)\right) . \tag{6.37}$$

Because we chose $\widetilde{\sigma}_i$ to satisfy *(i)*, we have $|A_n| \leq c(d)\epsilon' \mathcal{H}^{d-1}(\partial P_\delta) n^{d-1}$ for some positive constant $c(d)$. Define the edge set $\Gamma_n(\omega)$ as

$$\Gamma_n := \left(\bigcup_{i=1}^{m} E_n^{(i)}(\omega)\right) \cup A_n . \tag{6.38}$$

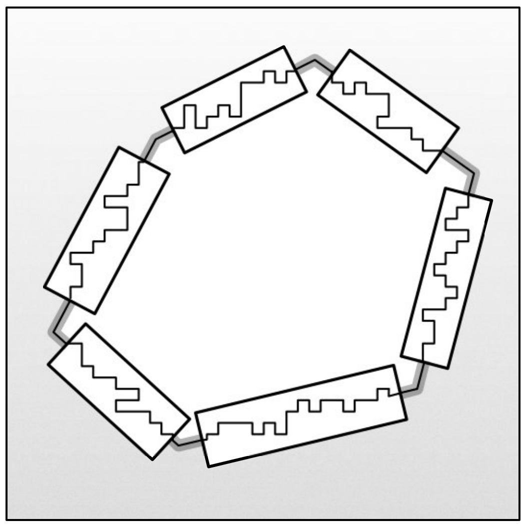

Figure 6.1: The polytope nP has six faces. Each of the boxes at the boundary of nP is one of the $\mathsf{cyl}(\widetilde{\sigma}_i, h, n)$, and within each is the corresponding cutset $E_n^{(i)}$. The set A_n is depicted as the grey outline of each corner.

We now define the vertex set $H_n(\omega)$ to be all vertices $x \in \mathbf{C}_n$ such that any path from x to ∞ in $\mathbf{C}_\infty$ must use an edge of $\Gamma_n(\omega)$. The proof of Lemma 4.2.2 tells us that not only is $H_n(\omega)$ non-empty, but that it also contains every vertex $x \in \mathbf{C}_\infty \cap Q_j$ for $j \in \{1, \ldots, \ell_1\}$. For this proof to go through, we must make sure n is large enough depending on h so that the cylinders $\mathsf{cyl}(\widetilde{\sigma}_i, h, n)$ are suitable in the sense of Remark 4.1.2. As h has been fixed and depends only on ϵ' and the polytope P, this is no issue, and $H_n(\omega)$ is well-defined for all n sufficiently large.

Let us suppress the dependence of H_n and Γ_n on the percolation configuration. We now exhibit control on the volume and open edge boundary of H_n, first working within the intersection $\mathcal{E}$ of the high probability events $\mathcal{E}_n^{(j)}$ to control $|H_n|$. For all percolation configurations within $\mathcal{E}$,

$$(\theta_p(d) - \epsilon') \left(\sum_{j=1}^{\ell_1} \mathcal{L}^d(nQ_j) \right) - \ell c(d) \left(2^{-k} n\right)^{d-1} \leq |H_n| \leq (\theta_p(d) + \epsilon') \sum_{j=1}^{\ell_2} \mathcal{L}^d(nQ_j) \,, \quad (6.39)$$

where the term we subtract on the left comes from the fact that the Q_j are not disjoint, but rather have disjoint interiors. For n sufficiently large (depending on p, d, ϵ', P), we have

$$(\theta_p(d) - 2\epsilon') \sum_{j=1}^{\ell_1} \mathcal{L}^d(nQ_j) \leq |H_n| \leq (\theta_p(d) + \epsilon') \sum_{j=1}^{\ell_2} \mathcal{L}^d(nQ_j) \,, \quad (6.40)$$

and hence that

$$(\theta_p(d) - 2\epsilon')(1 - \epsilon')\mathcal{L}^d(nP_\delta) \leq |H_n| \leq (\theta_p(d) + \epsilon')(1 + \epsilon')\mathcal{L}^d(nP_\delta) \,, \quad (6.41)$$

$$(\theta_p(d) - 2\epsilon')(1 - \epsilon')(1 - \delta)^d \mathcal{L}^d(nP) \leq |H_n| \leq (\theta_p(d) + \epsilon')(1 + \epsilon')(1 - \delta)^d \mathcal{L}^d(nP) \,. \quad (6.42)$$

We now show that H_n is a valid subgraph of $\mathbf{C}_n$ when δ is chosen appropriately. On the event $\mathcal{E}$, and due to the fact that the cubes Q_j intersect one another only at their boundaries, we have

$$|\mathbf{C}_n| \geq (\theta_p(d) - \epsilon')(2n)^d - \ell c(d)(2^{-k} n)^{d-1} \,, \quad (6.43)$$

$$\geq (\theta_p(d) - 2\epsilon')(2n)^d \,. \quad (6.44)$$

for n sufficiently large. As $\mathcal{L}^d(P) \leq 2^d/d!$, choosing δ in accordance with (6.42) so that

$$(\theta_p(d) + \epsilon')(1 + \epsilon')(1 - \delta)^d = (\theta_p(d) - 2\epsilon') \,, \quad (6.45)$$

which ensures that H_n is a valid subgraph of $\mathbf{C}_n$ on $\mathcal{E}$. Note that defining δ this way means that $\delta \to 0$ as $\epsilon' \to 0$.

Not only have we shown that H_n is valid within a high probability event, we also have exhibited a lower bound on $|H_n|$. To bound on $\widehat{\Phi}_n$ using H_n, it then suffices to find an upper bound on $\partial^\omega H_n$. From the construction of H_n, we have $\partial^\omega H_n \subset \Gamma_n$, and thanks to the disjointness of the cylinders $\mathsf{cyl}(\widetilde{\sigma}_i, h, n)$,

$$|\partial^\omega H_n| \leq \sum_{i=1}^{m} |E_n^{(i)}(\omega)|_\omega + c(d)\mathcal{H}^{d-1}(\partial P_\delta)\epsilon' n^{d-1}\,. \tag{6.46}$$

For $i \in \{1, \ldots, m\}$, let $\mathcal{F}_n^{(i)}$ be the following high probability event corresponding to *(ii)* at the beginning of the proof:

$$\Big\{ \Xi_{\text{face}}(\widetilde{\sigma}_i, \eta, n) < (1+\epsilon')\mathcal{H}^{d-1}(n\sigma_i)\beta(v_\sigma) \Big\}\,. \tag{6.47}$$

On the intersection $\mathcal{F}$ of the events $\mathcal{F}_n^{(i)}$ over all faces σ_i, we have

$$|\partial^\omega H_n| \leq (1+\epsilon')\mathcal{I}_{p,d}(nP) + c(d)\mathcal{H}^{d-1}(\partial P)\epsilon' n^{d-1}\,, \tag{6.48}$$

$$\leq (1+\epsilon' + c(p,d)\epsilon')\mathcal{I}_{p,d}(nP)\,. \tag{6.49}$$

Thus, on the intersection of $\mathcal{E}$ and $\mathcal{F}$, we have

$$\widehat{\Phi}_n \leq \left(\frac{1+\epsilon' + c(p,d)\epsilon'}{(\theta_p(d) - 2\epsilon')(1-\epsilon')(1-\delta)^d} \right) \frac{\mathcal{I}_{p,d}(nP)}{\mathcal{L}^d(nP)}\,, \tag{6.50}$$

and we take ϵ' small enough (recall $\delta = \delta(\epsilon')$ goes to zero as ϵ' does) so that

$$\widehat{\Phi}_n \leq (1+\epsilon)\frac{\mathcal{I}_{p,d}(nP)}{\theta_p(d)\mathcal{L}^d(nP)}\,. \tag{6.51}$$

We use the bounds in Corollary 10.1.5 on $\mathcal{E}$ and in Proposition 6.2.1 on $\mathcal{F}$ to conclude that

$$\mathbb{P}_p\left((\mathcal{E}\cap\mathcal{F})^c\right) \leq mc_1 \exp\left(-c_2 n^{(d-1)/3}\right) + \ell c_1 \exp\left(-c_2 (2^{-k}n)^{d-1}\right)\,, \tag{6.52}$$

which completes the proof, upon following the dependencies of ℓ and k. □

Using Proposition 10.4.2 and Borel-Cantelli, we extract from Theorem 6.2.2 a statement involving the Wulff crystal $W_{p,d}$.

Corollary 6.2.3. Let $d \geq 2$ and let $p > p_c(d)$. Consider the Wulff crystal $W_{p,d}$ corresponding to the norm $\beta_{p,d}$, and let $\epsilon > 0$. The event

$$\left\{ \limsup_{n\to\infty} n\widehat{\Phi}_n \leq (1+\epsilon)\frac{\mathcal{I}_{p,d}(W_{p,d})}{\theta_p(d)\mathcal{L}^d(W_{p,d})} \right\} \tag{6.53}$$

occurs $\mathbb{P}_p$-almost surely.

Proof. Recall that $\mathcal{L}^d(W_{p,d}) = 2^d/d!$. Let $\epsilon, \epsilon' > 0$ and apply Proposition 10.4.2 with this parameter to obtain $P_{\epsilon'} \subset W_{p,d}$ with $|\mathcal{I}_{p,d}(P_{\epsilon'}) - \mathcal{I}_{p,d}(W_{p,d})| < \epsilon'$ and with $\mathcal{L}^d(W_{p,d} \setminus P_{\epsilon'}) < \epsilon'$. Apply Theorem 6.2.2 to the polytope $P_{\epsilon'}$ to obtain positive constants $c_1(p,d,\epsilon,P_{\epsilon'})$ and $c_2(p,d,\epsilon,P_{\epsilon'})$ so that

$$\mathbb{P}_p\left(n\widehat{\Phi}_n \geq (1+\epsilon/2)\left(\frac{\mathcal{I}_{p,d}(P_{\epsilon'})}{\theta_p(d)\mathcal{L}^d(P_{\epsilon'})}\right)\right) \leq c_1 \exp\left(-c_2 n^{(d-1)/3}\right) , \tag{6.54}$$

and hence that

$$\mathbb{P}_p\left(n\widehat{\Phi}_n \geq (1+\epsilon/2)\left(\frac{1+\epsilon'}{1-\epsilon'}\right)\left(\frac{\mathcal{I}_{p,d}(W_{p,d})}{\theta_p(d)\mathcal{L}^d(W_{p,d})}\right)\right) \leq c_1 \exp\left(-c_2 n^{(d-1)/3}\right) . \tag{6.55}$$

Choosing ϵ' sufficiently small depending on ϵ and applying Borel-Cantelli completes the proof. ☐

Corollary 6.2.3 is half of Theorem 2.2.2, albeit the easier half. Before moving to the next section, let us make an observation that will facilitate the proof of Theorem 2.2.1. For $K \subset [-1,1]^d$ convex with non-empty interior, we define the *empirical measure* associated to K as

$$\overline{\nu}_K(n) := \frac{1}{n^d} \sum_{x \in \mathbf{C}_n \cap nK} \delta_{x/n} . \tag{6.56}$$

Following the proof of Theorem 6.2.2, it is not difficult to deduce the following result (recall that the metric $\mathfrak{d}$ was introduced in (3.7)).

Corollary 6.2.4. Let $d \geq 2$, $p > p_c(d)$ and let $W \subset [-1,1]^d$ be a translate of $W_{p,d}$. For $\epsilon > 0$, there are positive constants $c_1(p,d,\epsilon)$ and $c_2(p,d,\epsilon)$ so that

$$\mathbb{P}_p\left(\mathfrak{d}\left(\overline{\nu}_W(n), \nu_W\right) > \epsilon\right) \leq c_1 \exp\left(-c_2 n^{d-1}\right) . \tag{6.57}$$

Remark 6.2.5. Corollary 6.2.4 follows from the approximation result Proposition 10.4.2, from the density result Corollary 10.1.5 (used exactly as in the proof of Theorem 6.2.2) applied to a sufficiently fine mesh of dyadic cubes and finally from the definition of the metric $\mathfrak{d}$. In fact, no concentration estimates for $\beta_{p,d}$ are needed, so the proof of Corollary 6.2.4 is in some sense less involved than that of Theorem 6.2.2, and we choose to omit it.

CHAPTER 7

Coarse graining

Having spent the last section passing from continuous objects to discrete objects, we now move in the other, more difficult direction. To each Cheeger optimizer $G_n \in \mathcal{G}_n$, we would like to associate a corresponding Borel set $P_n \subset [-1,1]^d$ such that the conductance of nP_n is comparable with that of G_n. The sequence $\text{per}(P_n)$ should then be uniformly bounded in n.

Given G_n, it is natural to try to define P_n to be the following set

$$\frac{1}{n}\left(\bigcup_{x\in G_n} Q(x)\right), \tag{7.1}$$

where $Q(x)$ is the unit dual cube corresponding to $x \in \mathbb{Z}^d$. However, the perimeter of (7.1) is directly related to $|\partial G_n|$ as opposed to $|\partial^\omega G_n|$. While results like Lemma 10.2.5 allow us to control $|\partial^\omega G_n|$, we have far less control on $|\partial G_n|$ unless p is very close to one. This suggests a renormalization argument, and indeed, in the present section we will introduce a coarse graining procedure due to Zhang [Zha07].

Towards producing the desired continuum sets P_n, as an intermediate step we first augment each $G_n \in \mathcal{G}_n$ to some $F_n \subset \mathbf{C}_n$. Each F_n will be built in a way that allows us to control its boundary, and to each F_n, we may associate an empirical measure $\widetilde{\mu}_n$ as in Section 3.3 which we will show to be $\mathfrak{d}$-close to μ_n. The goal of the present section is to construct such F_n from each G_n, and to build corresponding edge sets Γ_n, which should be thought of as the boundary of F_n.

Actually, it is the Γ_n which we construct first, and from which we define the F_n. By using a modified version of Zhang's construction, we build each Γ_n from G_n, and as a consequence of this construction, each $|\Gamma_n|$ will be with high probability $O(n^{d-1})$. That each Γ_n is surface order will enable us to construct a suitable continuum set P_n from F_n, and we do this in

Chapter 8. The endgame of this construction is that, with high probability, each μ_n will be $\mathfrak{d}$-close to some ν_F representing a set F of perimeter at most some γ, where γ does not depend on n. In Chapter 9, we show that when a given μ_n is $\mathfrak{d}$-close to such a measure ν_F, we will be able to relate the conductance of G_n to that of F. Corollary 6.2.3 will tell us that unless F is the Wulff crystal, it should be impossible for μ_n to be too $\mathfrak{d}$-close to F. From this reasoning, we will obtain the main theorems of the paper.

One artifact of the construction given in this section is that we now need to restrict our attention to dimensions strictly greater than two. We will comment more on this complication in Section 7.3 (see Remark 7.3.1), but for now we simply assume $d \geq 3$ for the remainder of the paper.

Remark 7.0.6. An unfortunate defect of the formatting is that underlined text does not mesh well with captions for figures. Because underlining is used heavily in this section as a notational device, we introduce the convention that within figure captions only, the notation coarse will replace underlined text. For instance, the set $\underline{A}$ will be written instead as $\mathsf{coarse}(A)$.

7.1 Preliminary notation

Let k be a natural number which we refer to as the *renormalization parameter*. Given $x \in \mathbb{Z}^d$, we define the *k-cube* corresponding to x as:

$$\underline{B}(x) := (2k)x + [-k, k]^d \,. \tag{7.2}$$

We suppress the dependence of $\underline{B}(x)$ on k to avoid cumbersome notation. We use underscores to denote sets of k-cubes. If $\underline{G}$ is a set of k-cubes and if $x \in \mathbb{Z}^d$, we write $x \in \underline{G}$ if x is contained in one of the k-cubes of $\underline{G}$. If $e \in \mathrm{E}(\mathbb{Z}^d)$ is an edge, we write $e \in \underline{G}$ if both endpoint vertices of e lie in $\underline{G}$. We also need to introduce a type of larger cube; we define a *$3k$-cube* $\underline{B}_3(x)$ as follows:

$$\underline{B}_3(x) := (2k)x + [-3k, 3k]^d \,. \tag{7.3}$$

We emphasize that x must lie in $\mathbb{Z}^d$, so that each $3k$-cube contains exactly 3^d k-cubes. Two cubes $\underline{B}(x)$ and $\underline{B}(x')$ are *adjacent* if $x \sim x'$, or equivalently if they share a face. Two cubes $\underline{B}(x)$ and $\underline{B}(x')$ are $\mathbb{L}^d$*-adjacent* if $x \sim_{\mathbb{L}} x'$, or equivalently if either $\underline{B}(x') \subset \underline{B}_3(x)$ or $\underline{B}(x) \subset \underline{B}_3(x')$.

We now follow a construction of Zhang from Section 2 of [Zha07]. We describe Zhang's method in general first, and then apply it to G_n. The idea is to form a collection of k-cubes which contain $\partial_o G_n$, and then to discover within this collection another, more tame cutset separating G_n and ∞.

7.2 The construction of Zhang

Let $G = G(\omega) \subset \mathbf{C}_\infty$ be a finite, connected graph which is allowed to depend on the percolation configuration. From G, we define several sets of k-cubes. Firstly, we define

$$\underline{G} := \left\{ \underline{B}(x) \ : \underline{B}(x) \cap (G \cup \partial_o G) \neq \emptyset \right\}, \qquad \underline{A} := \left\{ \underline{B}(x) \ : \underline{B}(x) \cap \partial_o G \neq \emptyset \right\}. \tag{7.4}$$

Figure 7.1 depicts a possible G as well as the k-cube set $\underline{A}$. As G is finite, so too is $\underline{G}$, thus the cubes $\underline{B}(x)$ which are not in $\underline{G}$ split into a single infinite $\mathbb{L}^d$-connected component, which we label $\underline{Q}$, as well as finitely many finite $\mathbb{L}^d$-connected components $\underline{Q}'(1), \ldots, \underline{Q}'(u')$. We refer to $\underline{Q}$ as the *ocean* and following Zhang's terminology, we refer to the $\underline{Q}'(i)$ as *ponds*.

We use $\Delta\underline{Q}$ to denote the set of k-cubes $\underline{B}(x)$ which are $\mathbb{L}^d$-adjacent to a cube in the ocean $\underline{Q}$, but are not themselves contained in $\underline{Q}$. Likewise, for each pond $\underline{Q}'(i)$, we let $\Delta\underline{Q}'(i)$ denote the cubes $\mathbb{L}^d$-adjacent to $\underline{Q}'(i)$ but not contained in $\underline{Q}'(i)$.

The next step in Zhang's construction is to pass to the unique configuration ω' obtained by closing each open edge in $\partial_o G$. We do this while preserving both G and $\mathbf{C}_\infty$. In other words, when we work in the configuration ω', we still work with the graphs $G(\omega)$ and $\mathbf{C}_\infty(\omega)$, but with each modified by closing each open edge of $\partial_o G$.

Remark 7.2.1. Counterintuitively, $\mathbf{C}_\infty$ may then be a disconnected graph after passing to the configuration ω'. We hope this notation does not generate confusion, and we emphasize that $\mathbf{C}_\infty$ below is *not* $\mathbf{C}_\infty(\omega')$.

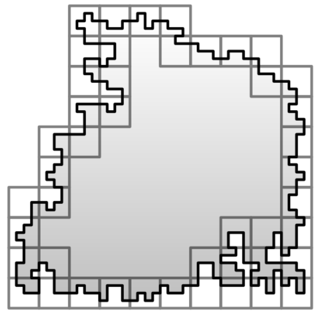

Figure 7.1: The black contour and its interior represent $\partial_o G$ and G respectively. Notice that $\mathsf{coarse}(A)$, depicted by the squares covering $\partial_o G$, is not necessarily the boundary of $\mathsf{coarse}(G)$.

We now pass to the configuration ω'. Each pond may intersect an open cluster which is connected to the ocean. Importantly, these open clusters do not need to be contained in $\mathbf{C}_\infty$. We say a pond is *live* if it intersects an open cluster which also intersects the ocean. If $\underline{Q}'(i)$ intersects an open cluster which also intersects a live pond, we say it is *almost-live*. If $\underline{Q}'(i)$ intersects an open cluster also intersecting a pond labeled as almost-live, call $\underline{Q}'(i)$ *almost-live* also. Thus, the label "almost-live" propagates through the ponds $\underline{Q}'(i)$ via open clusters, starting with the live ponds.

Finally, we say a pond is *dead* if it is neither live nor almost-live. We refine the collection of ponds $\{\underline{Q}'(i)\}_{i=1}^{u'}$ to the collection of live or almost-live ponds $\underline{Q}(1), \ldots, \underline{Q}(u)$. Figure 7.2 depicts a possible configuration of ponds. Zhang [Zha07] uses the terminology "live" and "dead," and we find it necessary to introduce the intermediate "almost-live" status.

Let C denote the collection of all open clusters which intersect $\underline{Q}$, and let C_i denote the collection of all open clusters which intersect the live or almost-live pond $\underline{Q}(i)$. We emphasize that the components of C and C_i are not necessarily in $\mathbf{C}_\infty$. To isolate the vertices of C_i within the $\underline{Q}(i)$ only, we define $Q_i := C_i \cap \underline{Q}(i)$, and we likewise define $Q := C \cap \underline{Q}$. We let bridge be the remainder of these components which lie in $\underline{G}$. Specifically,

$$\mathsf{bridge} := \left[\left(\bigcup_{\underline{B}(x) \in \underline{G}} \underline{B}(x) \right) \cap \left(C \cup \left(\bigcup_{i=1}^{u} C_i \right) \right) \right] \setminus \left(Q \cup \left(\bigcup_{i=1}^{u} Q_i \right) \right). \tag{7.5}$$

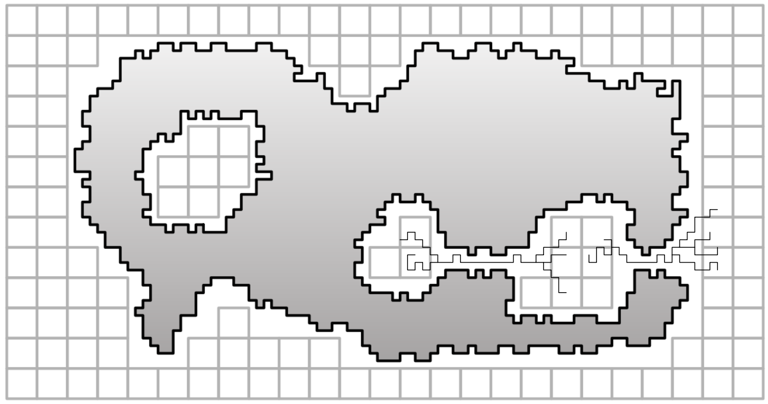

Figure 7.2: The graph G is the shaded region between closed curves. The connected components of cubes in the diagram are ponds or the ocean. The left-most pond is dead, the right-most pond is live and the middle pond is almost-live. The portions of the thin curves which do not intersect any cube represent the set bridge.

One can refer to Figure 7.2 for an example of the set bridge; we need only excise the portions of the thin black lines contained in any cubes drawn. Though bridge has been defined as a set of vertices, it is equipped with a natural graph structure which it inherits from C and the C_i. We now make one observation concerning bridge, now viewed as a graph.

Lemma 7.2.2. In the configuration ω', the vertex sets of bridge and G are disjoint, and all edges of $\partial\mathsf{bridge}$ are closed, except for those joining a vertex of bridge and a vertex in some Q_i or those joining a vertex of bridge with a vertex in Q.

Proof. To show bridge and G are disjoint, it suffices to show for each i that $C_i \cap G = \emptyset$ and that $C \cap G = \emptyset$. If the intersection of C and G were non-empty, there would then be an open path beginning from a vertex of G and ending at a vertex contained in $\underline{Q}$, which is impossible in the configuration ω'. The same reasoning shows that $C_i \cap G = \emptyset$ for each i.

Because C and the C_i are collections of open clusters, it is impossible that ∂C or any ∂C_i contain open edges. Due to the construction of bridge, the only open edges present in $\partial\mathsf{bridge}$ must either join bridge with some vertex in the ocean $\underline{Q}$, or they join bridge with a vertex in $\underline{Q}(i)$ for some i. □

We define the set of k-cubes associated to bridge as

$$\underline{\mathsf{bridge}} := \left\{ \underline{B}(x) \ : \underline{B}(x) \text{ contains a vertex of } \mathsf{bridge} \right\}, \tag{7.6}$$

and we build the central k-cube set $\underline{\Gamma}$ of Zhang's construction, depicted in Figure 7.3. Define

$$\underline{\Gamma} := \Delta\underline{Q} \cup \underline{\mathsf{bridge}} \cup \left(\bigcup_{i=1}^{u} \Delta\underline{Q}(i) \right). \tag{7.7}$$

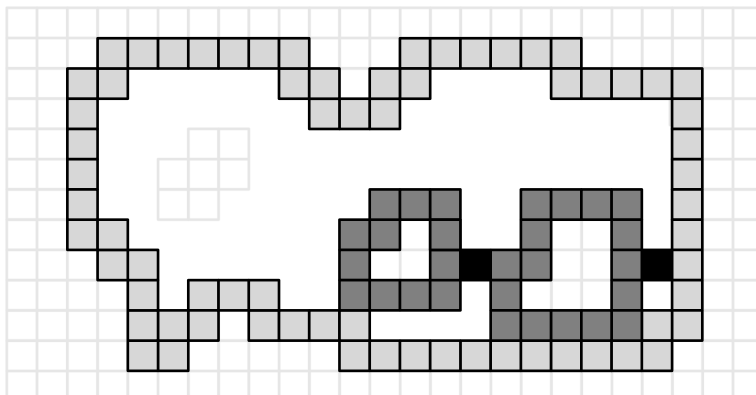

Figure 7.3: We have removed $\partial_o G$ and bridge from the diagram for the sake of clarity, but this picture is built from Figure 7.2. The light-grey cubes depict $\Delta\mathsf{coarse}(Q)$, the dark-grey cubes depict the two $\Delta\mathsf{coarse}(Q)(i)$and the black cubes depict $\mathsf{coarse}(\mathsf{bridge})$. The cubes adjacent to the black cubes are also in $\mathsf{coarse}(\mathsf{bridge})$, which illustrates that $\mathsf{coarse}(\mathsf{bridge})$ is not necessarily disjoint from the boundary of the ponds and ocean.

For each k-cube $\underline{B}(x)$, we define the corresponding *augmented* cube $\underline{B}^+(x)$ as

$$\underline{B}^+(x) := 2kx + [-2k-1, 2k+1]^d. \tag{7.8}$$

Remark 7.2.3. To any set of k-cubes, we may associate a corresponding set of augmented cubes. Our strategy for "taming" the boundary ∂G is to first form the cube-set $\underline{\Gamma}$, and then to discover within the corresponding augmented cube-set a cutset Γ which separates G from ∞, and whose size we can control.

Proposition 7.2.4. In the configuration ω', the augmented cube set $\underline{\Gamma}^+$ defined by

$$\underline{\Gamma}^+ := \left\{ \underline{B}^+(x) \ : \underline{B}(x) \in \underline{\Gamma} \right\} \tag{7.9}$$

contains a closed cutset Γ which separates G from ∞.

Proof. Let γ' be a path from G to ∞. We will show γ' uses a closed edge contained in $\underline{\Gamma}^+$. We lose no generality supposing γ' is simple. As $\underline{G}$ and each pond are finite sets, there is a first vertex $v_0 \in \underline{Q}$ used by γ'. Consider the subpath of γ' which starts at the beginning of γ' and ends at v_0. Name the reversal of this subpath γ, so that γ is a path from v_0 to G. It will suffice to show γ uses a closed edge which is contained in $\underline{\Gamma}^+$.

If the edge following v_0 in γ is closed, we are content as this edge lies in $\Delta\underline{Q}$. Thus we may suppose that the edge following v_0 in γ is open, so that γ joins the vertex set bridge. The path γ must connect with G. As bridge and G are disjoint (by Lemma 7.2.2), and because γ must eventually use a vertex of G, it must be that γ leaves bridge. If γ leaves bridge through a closed edge, this edge necessarily lies in one of the augmented cubes corresponding to the set $\underline{\mathsf{bridge}}$, in which case the proposition holds. Thus we may suppose that γ first leaves bridge through an open edge.

By Lemma 7.2.2, and thanks to the fact that γ cannot return to the ocean $\underline{Q}$, it follows that γ must pass into some Q_i. As γ is simple, and as the Q_i are disjoint from G, there is a last vertex v_1 of any Q_i used by γ. Let γ_1 denote the subpath of γ obtained by starting from this vertex v_1. If the first edge of γ_1 is closed, our claim holds as this edge lies in $\Delta\underline{Q}(i)$ for some i. Thus we may suppose the first edge of γ_1 is open, so that the path γ_1 rejoins the set bridge. However, γ_1 can no longer exit bridge through a pond or through the ocean, and thus γ_1 must exit bridge through a closed edge (again using Lemma 7.2.2).

This establishes that $\underline{\Gamma}^+$ contains a closed cutset separating G from ∞. Via some deterministic method, choose a minimal cutset within $\underline{\Gamma}^+$ and label it Γ. □

Next, we show $\underline{\Gamma}$ is contained in the coarse grained image of $\partial_o G$.

Lemma 7.2.5. The k-cube set $\underline{\Gamma}$ is contained in $\underline{A}$.

Proof. Suppose $\underline{B}(x) \in \underline{\mathsf{bridge}}$. We claim that $\underline{B}(x)$ must contain either a vertex of G or the endpoint vertex of an edge in $\partial_o G$. If not, $\underline{B}(x)$ is a member of some pond $\underline{Q}'(i)$ or of the ocean $\underline{Q}$. That $\underline{B}(x) \in \underline{\mathsf{bridge}}$ implies there is $y \in \underline{B}(x)$ which lies within an open cluster connected to a live or almost-live pond. Thus, if $\underline{B}(x)$ is a member of a pond, this pond is live or almost-live. This implies that either $y \in Q_i$ for some $i \in \{1, \dots, u\}$ or $y \in Q$, which is impossible, as we cut out such vertices in the construction of bridge.

Thus $\underline{B}(x)$ contains either a vertex of G or an endpoint vertex of $\partial_o G$. As $\underline{B}(x)$ contains some $y \in \mathsf{bridge}$, and as this y lies within an open cluster contained in C or some C_i which is disjoint from G, any path γ from y to G within the box $\underline{B}(x)$ must use some edge e of ∂G. But any $y \in C \cup \bigcup_i C_i$ is connected to ∞ via a path using no vertices of G. Thus the path γ from y to G within $\underline{B}(x)$ must actually use an edge of $\partial_o G$. This shows that $\underline{\mathsf{bridge}} \subset \underline{A}$.

We now show that $\Delta \underline{Q}(i) \subset \underline{A}$ for each live or almost-live pond $\underline{Q}(i)$. Let $\underline{B}(x) \in \Delta \underline{Q}(i)$. Then $\underline{B}(x)$ is $\mathbb{L}^d$-adjacent to a cube $\underline{B}(x') \in \underline{Q}(i)$, so $\underline{B}(x)$ must contain either a vertex of G or an endpoint vertex of an edge in $\partial_o G$, otherwise $\underline{B}(x)$ would be a member of $\underline{Q}(i)$. If $\underline{B}(x)$ contains an endpoint vertex of an edge in $\partial_0 G$, we are done, thus we suppose $\underline{B}(x)$ contains a vertex y of G.

Note that $\underline{B}(x)$ and $\underline{B}(x')$ have at least one vertex z in common, and this vertex (by virtue of lying within some $\underline{Q}(i)$) is connected to ∞ in via a $\mathbb{Z}^d$-path which does not use G. Any path joining y and z in $\underline{B}(x)$ must necessarily use an edge of $\partial_o G$. This shows $\Delta \underline{Q}(i) \subset \underline{A}$, and the final case of cubes within $\Delta \underline{Q}$ is handled identically to the case of pond boundary cubes. □

Remark 7.2.6. It is fundamental to us that $\underline{\Gamma}$ is also an $\mathbb{L}^d$-connected subset of $\underline{A}$. We know from Proposition 3.1.1 that minimal cutsets separating G from ∞ are $\mathbb{L}^d$-connected, but $\underline{\Gamma}^+$ is not necessarily the coarse grained image of the cutset Γ. Nevertheless, we can apply (analogues of) Proposition 3.1.1 to the sets $\Delta \underline{Q}(i)$ and $\Delta \underline{Q}$ to establish this result.

Lemma 7.2.7. The k-cube set $\underline{\Gamma}$ is $\mathbb{L}^d$-connected.

Proof. It follows directly from Lemma 2 of Timár [Tim13] that $\Delta \underline{Q}$ and each $\Delta \underline{Q}(i)$ are $\mathbb{L}^d$-connected cube sets.

Let D be a connected component of bridge, and let $\underline{D}$ be the collection of k-cubes containing a vertex of D, so that $\underline{D} \subset \underline{\mathsf{bridge}}$. It follows from the construction of bridge that $\underline{D}$ either intersects $\Delta \underline{Q}(i)$ for some i, or $\underline{D}$ intersects $\underline{Q}$. As D is connected in $\mathbb{Z}^d$, it is immediate that coarse grained image $\underline{D}$ is $\mathbb{L}^d$-connected. The set $\underline{\mathsf{bridge}}$ is itself the union of all such cube sets $\underline{D}$, and it follows from the defining properties of live and almost-live ponds that $\underline{\Gamma}$ is $\mathbb{L}^d$-connected. $\square$

The last step in Zhang's construction is to show each cube in $\underline{\Gamma}$ has a useful geometric property when G is sufficiently large. We introduce some more of the terminology in [Zha07]. Each k-cube $\underline{B}(x)$ has $2d$ faces $\sigma_1(x), \ldots, \sigma_{2d}(x)$, each of which an isometric image of $[-k, k]^{d-1} \subset \mathbb{R}^{d-1}$. We say that a *surface* of $\underline{B}(x)$ is a vertex set of the form $\sigma_i(x) \cap \mathbb{Z}^d$, so that each k-cube $\underline{B}(x)$ possesses $2d$ surfaces. A surface of a $3k$-cube $\underline{B}_3(x)$ is just a surface of one of the k-cubes $\underline{B}(x') \subset \underline{B}_3(x)$.

Definition 7.2.8. Say that a k-cube $\underline{B}(x)$ is *Type-I* if there is an open path γ and a surface $\sigma \cap \mathbb{Z}^d$ in $\underline{B}_3(x)$ such that γ joins a vertex in $\underline{B}^+(x)$ to a vertex of $\partial \underline{B}_3(x) \cap \mathbb{Z}^d$ and such that no vertex along γ is joined via another open path to $\sigma \cap \mathbb{Z}^d$. We require that γ uses edges which are *internal* to to $\underline{B}_3(x)$, that is, no edge of γ has both endpoints in $\partial \underline{B}_3(x)$. We also require any candidate for a path from a vertex of γ to $\sigma \cap \mathbb{Z}^d$ to use only edges internal to $\underline{B}_3(x)$.

Definition 7.2.9. We say a k-cube $\underline{B}(x)$ is *Type-II* if there are two open paths γ_1 and γ_2, each of which connects a distinct vertex of $\underline{B}^+(x)$ to distinct vertices in $\partial \underline{B}_3(x)$, and such that there is no open path in $\underline{B}_3(x)$ joining any vertex in γ_1 to a vertex in γ_2. We require that all paths in this definition use only edges internal to $\underline{B}_3(x)$.

Figure 7.4 illustrates these two geometric properties.

Remark 7.2.10. Because of the requirement that all paths in the above definitions are internal, the event that a k-cube $\underline{B}(x)$ is Type-I or Type-II does not depend on the state of any edge contained in $\partial \underline{B}_3(x)$.

Proposition 7.2.11. Suppose that no connected component of G is contained within some $3k$-cube. Then, in the configuration ω', each k-cube of $\underline{\Gamma}$ is either Type-I or Type-II.

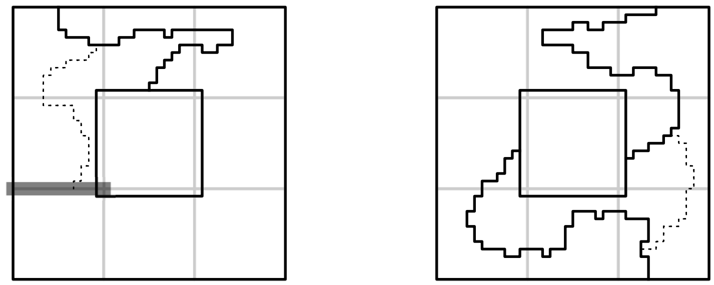

Figure 7.4: On the left, we see an illustration of what *cannot* happen in a Type-I cube. The dotted line is an open path joining the solid line (also an open path) to one of the surfaces of the $3k$-cube. Likewise, on the right, we see an illustration of what *cannot* happen in a Type-II cube.

Proof. Following Zhang, we consider two cases. In the first case, we suppose $\underline{B}(x) \in \underline{\Gamma}$ is a member of

$$\Delta\underline{Q} \cup \left(\bigcup_{i=1}^{u} \Delta\underline{Q}(i) \right) . \tag{7.10}$$

Such a $\underline{B}(x)$ is $\mathbb{L}^d$-adjacent to a cube $\underline{B}(x')$ which neither intersects G nor an endpoint vertex of $\partial_o G$. Thanks to Lemma 7.2.5, we know $\underline{B}(x) \in \underline{A}$, so that $\underline{B}(x)$ contains an endpoint vertex of $\partial_o G$. Thus, $\underline{B}^+(x)$ contains a vertex $y \in G$. There can be no open path from any surface of $\underline{B}(x')$ to y, as such a path could not use an edge of $\partial_o G$, but could be extended to a path from y to ∞ using no other vertices of G. On the other hand, because no connected component of G is contained in a $3k$-cube, there must be an open path from y to a vertex of $\partial\underline{B}_3(x)$. We may arrange this open path uses edges internal to $\underline{B}_3(x)$ by stopping it at the first vertex of $\partial\underline{B}_3(x)$ it meets. Thus in the first case, $\underline{B}(x)$ is Type-I.

In the second case, we suppose

$$\underline{B}(x) \in \underline{\mathsf{bridge}} \setminus \left(\Delta\underline{Q} \cup \left(\bigcup_{i=1}^{u} \Delta\underline{Q}(i) \right) \right) . \tag{7.11}$$

Let $y \in \underline{B}(x) \cap \mathsf{bridge}$. Then y lies in some connected component D of either C or one of the C_i. The component D cannot be contained in $\underline{B}_3(x)$, otherwise one of the k-cubes

$\mathbb{L}^d$-adjacent to $\underline{B}(x)$ would be a member of either $\underline{Q}$ or some $\underline{Q}(i)$. This is impossible as it would imply $\underline{B}(x) \in \Delta\underline{Q}$ or $\underline{B}(x) \in \Delta\underline{Q}(i)$ for some i. It follows that y is joined to the the boundary of $\underline{B}_3(x)$ by an open internal path (contained in D).

On the other hand, thanks to Lemma 7.2.5, the cube $\underline{B}(x)$ contains a vertex G or an endpoint vertex of $\partial_o G$. Thus, $\underline{B}^+(x)$ contains a vertex $z \in G$, and as no connected component of G is contained in a $3k$-cube, we have that z is connected to the boundary of $\underline{B}_3(x)$ by an open internal path (in G). But D and G are disjoint, thus the corresponding paths from y and z to the boundary of $\underline{B}_3(x)$ cannot lie in the same open cluster. We conclude that in this second case, $\underline{B}(x)$ is Type-II. □

We conclude this section with the observation that it is rare for a cube to be either Type-I or Type-II when k is large.

Proposition 7.2.12. Let $d \geq 2$ and let $p > p_c(d)$. There are positive constants $c_1(p,d)$ and $c_2(p,d)$ so that for each k-cube $\underline{B}(x)$,

$$\mathbb{P}_p(\underline{B}(x) \text{ is Type-I}) \leq c_1 \exp(-c_2 k) \tag{7.12}$$

and

$$\mathbb{P}_p(\underline{B}(x) \text{ is Type-II}) \leq c_1 \exp(-c_2 k)\,. \tag{7.13}$$

Proof. This is in part a consequence of Lemma 7.89 in Grimmett [Gri99]. The proof is also given in Section 3 of Zhang [Zha07]. □

7.3 Webbing

There are several obstacles to applying Zhang's construction to the G_n. The first is a small issue: each $G_n \in \mathcal{G}_n$ may not be connected, but we have a uniform (in n) bound with high probability on the number of connected components of each G_n thanks to Corollary 10.2.4; we may simply apply Zhang's construction to each connected component of G_n.

The second obstacle is more fundamental. Suppose that we use Zhang's construction on a Cheeger optimizer G_n to produce a cutset Γ_n. If we think of $\partial_o G_n$ as a contour, the cutset Γ_n is a deformation of this contour, and as such, the volume enclosed by Γ_n may differ from $|G_n|$. As stated in the introduction to this section, we will use Γ_n to augment G_n to some $F_n \subset \mathbf{C}_n$. We think of F_n as an approximate to G_n, so it is important that the two sets are close in volume. The set F_n will be constructed from Γ_n, and a priori, we have no control on the additional volume enclosed by the "deformed contour" Γ_n. Figure 7.1 illustrates that there may be quite a bit of room within the k-cube set $\underline{A}$, which suggests the extra volume enclosed by Γ_n may be substantial.

To get around the second obstacle, we will apply Zhang's construction a second time to surround and then excise each of the "large" components of $\mathbf{C}_n$ not in G_n which have been trapped by the initial cutsets formed. This is how we build F_n to have volume comparable with that of G_n.

This workaround creates one last complication: the boundary of F_n is now highly disconnected; it should be thought of as consisting of all contours generated by Zhang's construction so far. In order to control the total size of the boundary of F_n using a Peierls argument, it will be necessary to tie all contours together via an auxiliary edge set we call the webbing.

Remark 7.3.1. The webbing should be thought of as a one-dimensional object, and it will be used to show that the total size of all cutsets created above is still of order n^{d-1}. In particular, we need the total size of the webbing to be at most some constant times n^{d-1}. The webbing will join the large components described above to all connected components of G_n, but because the number of these large components grows with n, the size of the webbing becomes too large in dimension two. Our approach does not work in the case $d = 2$ for the same reason filaments do not exist in $d = 2$ (see Figure 2.1).

We now formalize the argument sketched at the beginning of this section. Given $G_n \in \mathcal{G}_n$, we list the connected components of G_n as $G_n^{(1)}, \dots, G_n^{(M)}$. For each connected component $G_n^{(q)}$ of G_n, let ω_q' be the unique configuration obtained from ω by closing each open edge in $\partial^\omega G_n^{(q)}$. Within the configuration ω_q', we may apply Zhang's construction to produce a closed cutset $\Gamma_n^{(q)}$ and k-cube set $\underline{\Gamma}_n^{(q)}$ with the properties discussed in Section 7.2.

Remark 7.3.2. We hope that the n in the subscript of these sets does not cause any confusion in light of our notation for $3k$-cubes; we emphasize that the sets $\underline{\Gamma}_n^{(q)}$ are collections of k-cubes.

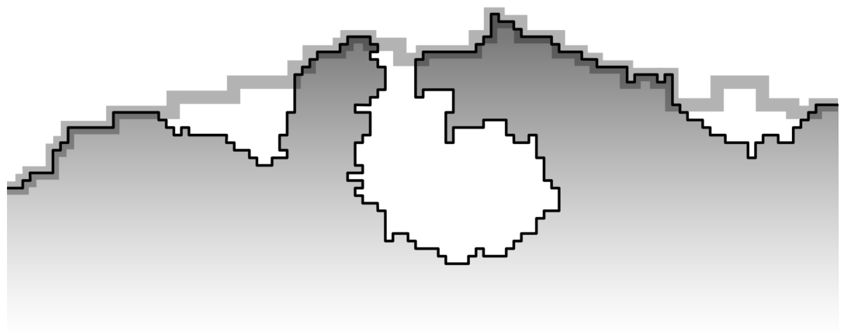

Figure 7.5: The black contour is a close-up of the boundary of some $G_n^{(q)}$. The thicker grey contour is the corresponding cutset $\mathsf{coarse}(\Gamma)_n^{(q)}$. It is possible that connected components of $\mathbf{C}_\infty$ are bounded between these two contours (see Remark 7.2.1).

Let ω' be the configuration in which all edges of $\partial^\omega G_n$ are made closed. Given a collection of edges S, we say that a connected component of Λ of $\mathbf{C}_\infty$ (in the configuration ω', see Remark 7.2.1) is *surrounded* by S if every path from Λ to ∞ must use an edge of S. We only care about the connected components of $\mathbf{C}_\infty$ surrounded by S, not other open clusters. Define

$$\epsilon(d) := 1 - \frac{d}{(d-1)^2}\,, \tag{7.14}$$

and observe that $\epsilon(d)$ is positive when $d \geq 3$. Within configuration ω', the cutsets $\Gamma_n^{(q)}$ may surround other connected components of $\mathbf{C}_\infty$ aside from the $G_n^{(q)}$ themselves. If Λ is such a component, we say that Λ is *large* if $|\Lambda| \geq n^{1-\epsilon(d)}$, and we say that Λ is *small* otherwise. We enumerate the large components $L_1, \ldots, L_m$ of $\mathbf{C}_\infty$ which are surrounded by any of the cutsets $\Gamma_n^{(q)}$; we do not include any of the $G_n^{(q)}$ in this list. We likewise enumerate the small components $S_1, \ldots, S_t$ of $\mathbf{C}_\infty$ which are surrounded by any of the cutsets $\Gamma_n^{(q)}$. Our notation suppresses the dependence of the L_i and the S_j on n, ω and G_n. Figure 7.5 depicts how these large and small components may arise. Define $F_n \subset \mathbf{C}_\infty$ from G_n as

$$F_n := G_n \cup \left(\bigcup_{j=1}^{t} S_j \right) . \tag{7.15}$$

Given a large component L_i, let ω_i' denote the configuration in which all open edges of $\partial^\omega L_i$ are made closed. Within ω_i', apply Zhang's construction to produce a closed cutset $\widehat{\Gamma}_n^{(i)}$ separating L_i from ∞, as well as a corresponding k-cube set $\widehat{\underline{\Gamma}}_n^{(i)}$. The edge sets $\Gamma_n^{(q)}$, in conjunction with the $\widehat{\Gamma}_n^{(i)}$ just constructed are thought of as representing the boundary of F_n.

From G_n, define the edge set Γ_n and k-cube set $\underline{\Gamma}_n$ as

$$\Gamma_n := \left(\bigcup_{q=1}^{M} \Gamma_n^{(q)} \right) \cup \left(\bigcup_{i=1}^{m} \widehat{\Gamma}_n^{(i)} \right) , \quad \underline{\Gamma}_n := \left(\bigcup_{q=1}^{M} \underline{\Gamma}_n^{(q)} \right) \cup \left(\bigcup_{i=1}^{m} \widehat{\underline{\Gamma}}_n^{(i)} \right) . \tag{7.16}$$

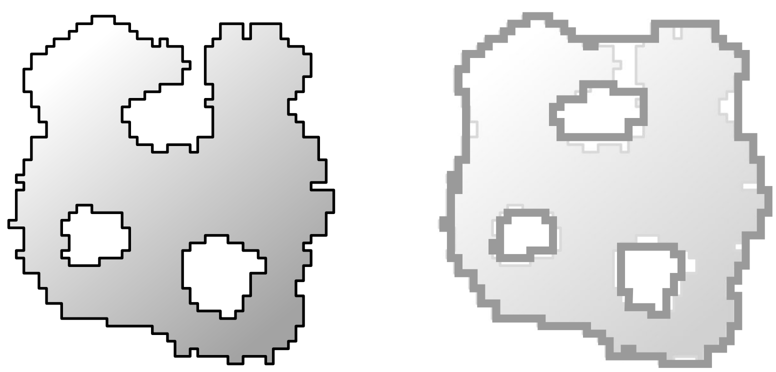

Figure 7.6: On the left is $G_n \in \mathcal{G}_n$. On the right, the thick grey contours together form the edge set Γ_n. The inner contours arise from large components and are of the form $\widehat{\Gamma}_n^{(i)}$. The outer contour corresponds to G_n itself. It is natural to wonder how these contours "interact," and we address this question at the start of Chapter 8.

We now create a web of paths between each of the cutsets used to build Γ_n. By (for instance) fixing an ordering of finite subsets of $\mathbb{Z}^d$, we may choose in a unique way an endpoint vertex ζ_q of an edge in $\Gamma_n^{(q)}$. We do the same for each $\widehat{\Gamma}_n^{(i)}$, calling these vertices z_i. Thus, once our original configuration ω is set, the optimizer G_n is determined by ω, and through Zhang's construction, so too are the ζ_q and the z_i. These will be the endpoints of the paths that make up the webbing.

Let $\alpha > 0$ be a parameter to be chosen later. Consider all cubes of the form $\lceil n^\alpha \rceil x + [-\lceil n^\alpha \rceil, \lceil n^\alpha \rceil]^d$ which intersect $[-2n, 2n]^d$, where $x \in \mathbb{Z}^d$. Let us list these cubes as $\{\overline{B}_j\}_{j=1}^{\ell}$. We may assume that the $\overline{B}_j$ are ordered so that consecutive cubes share a face. For n sufficiently large (depending on the renormalization parameter k), all vertices ζ_q and z_i will lie in the union of the $\overline{B}_j$. For each $j \in \{1, \ldots, \ell\}$, let m_j denote the number of vertices z_i which are contained in the cube $\overline{B}_j$.

Within each $\overline{B}_j$, we perform the following procedure: begin with $z_i \in \overline{B}_j$ which is least in our ordering of finite subsets of $\mathbb{Z}^d$. Pick (in some unique way) a $\mathbb{Z}^d$-path (not necessarily open) joining this "smallest" vertex to the next smallest z_j within $\overline{B}_j$, arranging that this path uses the fewest edges possible.

Continue building unique paths between the current vertex in $\overline{B}_j$ and the next smallest within $\overline{B}_j$ until all vertices z_i within $\overline{B}_j$ have been used. We refer to the union of all paths created in this process as a *tangle* and denote it T_j, which we view as a graph. Repeat this process for each $\lceil n^\alpha \rceil$-cube $\overline{B}_j$, defining the tangle T_j to be empty if $m_j = 0$.

We connect each tangle using a single long path. We construct this path by beginning with $\overline{B}_1$ and selecting the vertex $z_i \in \overline{B}_1$ minimal in our ordering. We select a minimal vertex from $\overline{B}_2$ and connect the two successive vertices by a uniquely chosen shortest path in $\mathbb{Z}^d$, or we do nothing in the case that these two vertices are identical. We continue joining vertices of consecutive cubes in $\{\overline{B}_j\}_{j=1}^{\ell}$. If at any point in our process, we find that a cube $\overline{B}_j$ contains no vertices z_i, we take instead the vertex of $\mathbb{Z}^d$ within $\overline{B}_j$ which is minimal in our ordering. When the minimal vertex of the last $\lceil n^\alpha \rceil$-cube has been used, we link this vertex to the vertex ζ_1 via a uniquely chosen shortest path, and proceed to link successive ζ_q's by uniquely chosen shortest paths until we reach ζ_M. The union of all paths created in this process shall be denoted string. Note that we have suppressed the n, ω and G_n dependence of both string and the tangles T_j. Let us define

$$\mathsf{web}_n := \mathsf{string} \cup \left(\bigcup_{j=1}^{\ell} T_j \right) , \tag{7.17}$$

viewed as a graph with the natural structure.

Remark 7.3.3. We emphasize that $\mathsf{web}_n = \mathsf{web}_n(\omega)$ and of course that web_n also depends on the Cheeger optimizer $G_n \in \mathcal{G}_n$. For each $G_n \in \mathcal{G}_n$, the edge set of web_n should have (with high probability) cardinality at most $O(n^{d-1})$ in order to execute a Peierls argument. To show this, we first compute a high probability bound on the number of large components.

Using Proposition 10.1.1 and the proof of Corollary 10.1.2, we deduce the following corollary.

Corollary 7.3.4. Let $d \geq 3$ and $p > p_c(d)$. There are positive constants $c_1(p,d), c_2(p,d)$ and $c_3(p,d)$ and an almost-surely finite random variable $R' = R'(\omega)$ so that $n \geq R'$ implies that for each ω-connected Λ satisfying $\Lambda \subset \mathcal{C}_{2n}$ and $|\Lambda| \geq n^{1-\epsilon(d)}$ we have

$$|\partial^{\omega} \Lambda| \geq c_3 |\Lambda|^{(d-1)/d} , \tag{7.18}$$

with the following tail bound on R':

$$\mathbb{P}_p(R' > n) \leq c_1 \exp(-c_2 n^{1/(d-1)}) . \tag{7.19}$$

We make direct use of Corollary 7.3.4 in the proof of the following lemma, which is why we have written it above despite its similarities to Corollary 10.1.2.

Lemma 7.3.5. Let $d \geq 3$ and $p > p_c(d)$. For each $G_n \in \mathcal{G}_n$, let $m(G_n) = m(G_n, n, \omega)$ denote the number of large components surrounded by the cutsets $\Gamma_n^{(q)}$ in Zhang's construction applied to G_n. Let $M_n = M_n(\omega)$ be the maximum $m(G_n)$ over all $G_n \in \mathcal{G}_n$. There exist positive constants $c_1(p,d,k), c_2(p,d,k)$ and $c_3(p,d)$ so that

$$\mathbb{P}_p \left(M_n > c_3 n^{d-1-1/(d-1)} \right) \leq c_1 \exp\left(-c_2 n^{1/(d-1)} \right) . \tag{7.20}$$

Proof. Fix $G_n \in \mathcal{G}_n$, let $m = m(G_n)$ and let us work within the high probability event $\{R' \leq n\}$ from Corollary 7.3.4. By Lemma 7.2.5, all large components L_i corresponding to G_n are contained in $\mathcal{C}_{2n}$. Thus, on the event $\{R' \leq n\}$, we have

$$|\partial^\omega L_i| \geq cn^{1/(d-1)}\,, \tag{7.21}$$

for each $i \in \{1, \ldots, m\}$, where we've used the defining property of large components. Let us also work within the high probability event (from Lemma 10.2.5) that there is some $\eta_3(p,d) > 0$ so that for all $G_n \in \mathcal{G}_n$, we have $\partial^\omega G_n \leq \eta_3 n^{d-1}$. Combine this upper bound with (7.21) and the fact that distinct large components must have disjoint open edge boundaries to obtain

$$m \leq \frac{\eta_3}{c} n^{d-1-1/(d-1)}\,. \tag{7.22}$$

We use the estimates from Lemma 10.2.5 and Corollary 7.3.4 to complete the proof. □

With control on the number of large components, we may now compute a high-probability bound on the number of edges in the graph web_n across all $G_n \in \mathcal{G}_n$.

Remark 7.3.6. In the following proof, we fix the parameter α (which controls the side-lengths of the cubes $\overline{B}_j$) to be $\frac{1}{(d-1)}$.

Proposition 7.3.7. Let $d \geq 3$ and $p > p_c(d)$. For each $G_n \in \mathcal{G}_n$, we construct the graph $\mathsf{web}_n = \mathsf{web}_n(G_n, \omega)$ as above. Let $W_n = W_n(\omega)$ denote the maximum cardinality of $\mathrm{E}(\mathsf{web}_n)$ taken over all $G_n \in \mathcal{G}_n$. There are positive constants $c_1(p,d,k)$, $c_2(p,d,k)$ and $c_3(p,d)$ so that

$$\mathbb{P}_p\left(W_n > c_3 n^{d-1}\right) \leq c_1 \exp\left(-c_2 n^{1/(d-1)}\right)\,. \tag{7.23}$$

Proof. We work within the high probability event from Lemma 7.3.5 that the maximal number M of large components L_i across all $G_n \in \mathcal{G}_n$ is at most $cn^{d-1-1/(d-1)}$. We also work within the high probability event from Corollary 10.2.4 that the number of connected components of any $G_n \in \mathcal{G}_n$ is at most η_4.

Fix $G_n \in \mathcal{G}_n$ for the rest of the proof. Consider the tangle T_j for G_n associated to the $\lceil n^\alpha \rceil$-cube $\overline{B}_j$. Based on our construction of each tangle, the number of edges $|\mathrm{E}(T_j)|$ is at most the ℓ^1-diameter of $\overline{B}_j$ times the number of z_i within $\overline{B}_j$. Thus, choosing the constants $c(d), c(p,d)$ below appropriately,

$$\sum_{j=1}^{\ell} |\mathrm{E}(T_j)| \leq 8dn^\alpha \sum_{j=1}^{\ell} m_j \,, \tag{7.24}$$

$$\leq c(d) m n^\alpha \,, \tag{7.25}$$

$$\leq c(p,d) n^{\alpha+n-1-1/(d-1)} \,, \tag{7.26}$$

where to obtain second line directly above, we have used that each vertex z_i is contained in at most $2d$ distinct $\lceil n^\alpha \rceil$-cubes, and to obtain the third line we have used our bound from Lemma 7.3.5. It remains to bound the size of the edge set of string. A shortest $\mathbb{Z}^d$-path between the vertices of two adjacent $\lceil n^\alpha \rceil$-cubes uses at most $16n^\alpha$ edges, and there are at most $c(d)n^{d(1-\alpha)}$ such cubes in total for some $c(d) > 0$. The final paths in the construction of string which join the vertices ζ_q each use at most $c(d)n$ edges for some $c(d) > 0$. Thus,

$$|\mathrm{E}(\mathsf{string})| \leq c(d) n^\alpha n^{d(1-\alpha)} + \eta_4 c(d) n \,, \tag{7.27}$$

so that upon choosing $\alpha = 1/(d-1)$, we have

$$|\mathrm{E}(\mathsf{web}_n)| \leq c(p,d) \left[n^{\alpha+(d-1)-1/(d-1)} + n^{\alpha+d(1-\alpha)} + n \right] \,, \tag{7.28}$$

$$\leq c(p,d) n^{d-1} \,. \tag{7.29}$$

We use the estimates from Lemma 7.3.5 and from Corollary 10.2.4 to complete the proof. □

We are finished working with $\lceil n^\alpha \rceil$-cubes. Define the coarse-grained image of each web_n as

$$\underline{\mathsf{web}}_n := \left\{ \underline{B}(x) \ : \underline{B}(x) \cap \mathsf{web}_n \neq \emptyset \right\} \tag{7.30}$$

so that each $\underline{\mathsf{web}}_n$ is a collection of k-cubes depending on n, ω and G_n. The last piece of information we need about the webbing is the following lemma.

Lemma 7.3.8. For each $G_n \in \mathcal{G}_n$, the k-cube set $\underline{\Gamma}_n \cup \underline{\mathsf{web}}_n$ corresponding to G_n is $\mathbb{L}^d$-connected.

We omit the proof of Lemma 7.3.8, as it follows directly from the construction of $\mathsf{web}_n$and from Lemma 7.2.7.

7.4 A Peierls argument

Using the webbing in conjunction with Zhang's construction, we can show that when the renormalization parameter k is taken sufficiently large, the size of each Γ_n is with high probability on the order of n^{d-1}. We make one small observation before diving into the Peierls argument.

Lemma 7.4.1. For each $G_n \in \mathcal{G}_n$, any edge of $\mathbb{Z}^d$ is contained in at most $(11k)^d$ distinct edge sets among the $\Gamma_n^{(q)}$ and the $\widehat{\Gamma}_n^{(i)}$ corresponding to G_n.

Proof. Fix $G_n \in \mathcal{G}_n$. If $\Gamma_n^{(q)}$ uses an edge $e \in \mathrm{E}(\mathbb{Z}^d)$, there is a k-cube $\underline{B}(x) \in \underline{\Gamma}_n^{(q)}$ so that $e \in \underline{B}^+(x)$. By Lemma 7.2.5, $\underline{B}^+$ also contains a vertex $y \in G_n^{(q)}$. Suppose another cutset, say $\widehat{\Gamma}_n^{(i)}$, uses the same edge e. Identical reasoning tells us that the five-fold dilate of $\underline{B}(x)$

$$\underline{B}_5(x) := 2kx + [-5k, 5k]^d \tag{7.31}$$

contains both the vertex y and a vertex $z \in L_i$. If cutsets corresponding to other connected components of G_n or other large L_i also use e, at least one vertex of each of these graphs must also lie within $\underline{B}_5(x)$. As the components of G_n and the L_i are all disjoint, and because $\underline{B}_5(x)$ contains at most $(11k)^d$ vertices, the desired claim holds. □

We now proceed with the Peierls argument, through which the renormalization parameter k is fixed once and for all.

Proposition 7.4.2. Let $d \geq 3$ and $p > p_c(d)$. There exists $\gamma = \gamma(p, d)$ and positive constants $c_1(p, d)$ and $c_2(p, d)$ so that

$$\mathbb{P}_p \left(\max_{G_n \in \mathcal{G}_n} |\Gamma_n| \geq \gamma n^{d-1} \right) \leq c_1 \exp\left(-c_2 n^{1/(d-1)}\right) . \tag{7.32}$$

Proof. Let $\mathcal{E}_{\mathsf{web}}$ be the high probability event from Proposition 7.3.7 that for all $G_n \in \mathcal{G}_n$, the corresponding graphs web_n satisfy $|\mathrm{E}(\mathsf{web}_n)| \le c(p,d)n^{d-1}$. We work with a fixed $G_n \in \mathcal{G}_n$ and corresponding Γ_n throughout the proof, and begin the proof by first using the bounds in Proposition 7.3.7 and in Lemma 10.2.5:

$$\mathbb{P}_p\Big(|\Gamma_n| \ge \gamma n^{d-1}\Big) \le c_1 \exp\Big(-c_2 n^{1/(d-1)}\Big) \tag{7.33}$$

$$+\,\mathbb{P}_p\Big(\Big\{|\Gamma_n| \ge \gamma n^{d-1}\Big\} \cap \Big\{|\partial^\omega G_n| \le \eta_3 n^{d-1}\Big\} \cap \mathcal{E}_{\mathsf{web}}\Big), \tag{7.34}$$

$$\le c_1 \exp\Big(-c_2 n^{1/(d-1)}\Big) + \sum_{j=\gamma n^{d-1}}^{\infty} \mathbb{P}_p\Big(\Big\{|\Gamma_n| = j\Big\} \cap \Big\{|\partial^\omega G_n| \le (\eta_3/\gamma)j\Big\} \cap \mathcal{E}_{\mathsf{web}}\Big). \tag{7.35}$$

Take γ large depending on k, p and d so that $\eta_3/\gamma < [2 \cdot 4^d (11k)^d (4k)^{d+1}]^{-1}$:

$$\mathbb{P}_p\Big(|\Gamma_n| \ge \gamma n^{d-1}\Big) \le c_1 \exp\Big(-c_2 n^{1/(d-1)}\Big) \tag{7.36}$$

$$+ \sum_{j=\gamma n^{d-1}}^{\infty} \mathbb{P}_p\Big(\Big\{|\Gamma_n| = j\Big\} \cap \Big\{|\partial^\omega G_n| < [2 \cdot 4^d (11k)^d (4k)^{d+1}]^{-1} j\Big\} \cap \mathcal{E}_{\mathsf{web}}\Big). \tag{7.37}$$

We equip $\underline{\Gamma}_n$ with a graph structure: the vertices of this graph will be the collection of k-cubes in $\underline{\Gamma}_n$, and we create an edge between two vertices $\underline{B}(x)$ and $\underline{B}(y)$ if $x \sim_{\mathbb{L}} y$. The maximum degree of any vertex in this graph is 3^d, so by Turán's theorem (Theorem 10.3.1 in the appendix), there exists a subcollection of cubes $\underline{\Gamma}'_n \subset \underline{\Gamma}_n$ so that whenever $\underline{B}(x), \underline{B}(y) \in \underline{\Gamma}'_n$, the corresponding $3k$-cubes $\underline{B}_3(x)$ and $\underline{B}_3(y)$ have disjoint interiors, and such that $|\underline{\Gamma}'_n| \ge |\underline{\Gamma}_n|/4^d$. Thus,

$$|\underline{\Gamma}'_n| \ge \frac{|\underline{\Gamma}_n|}{4^d} \ge \frac{|\Gamma_n|}{4^d (11k)^d (4k)^{d+1}}, \tag{7.38}$$

as when $k \ge d$, there are at most $(4k)^{d+1}$ edges of $\mathbb{Z}^d$ which have an endpoint in a given augmented k-cube, and thanks to Lemma 7.4.1, we know each such edge is contained in at most $(11k)^d$ distinct cutsets among the $\Gamma_n^{(q)}$ and $\widehat{\Gamma}_n^{(i)}$. Now, consider the following event:

$$\Big\{|\Gamma_n| = j\Big\} \cap \Big\{|\partial^\omega G_n| < [2 \cdot 4^d (11k)^d (4k)^{d+1}]^{-1} j\Big\} \cap \mathcal{E}_{\mathsf{web}}. \tag{7.39}$$

Within the event given in (7.39), we know from (7.38) that at most half of the cubes of $\underline{\Gamma}'_n$ can contain an edge of $\partial^\omega G_n$. Thus there is a further subcollection $\underline{\Gamma}''_n \subset \underline{\Gamma}'_n$ so that

$$|\underline{\Gamma}''_n| \geq \frac{|\Gamma_n|}{2 \cdot 4^d (11k)^d (4k)^{d+1}}\,, \tag{7.40}$$

and such that each k-cube of $\underline{\Gamma}''_n$ is either Type-I or Type-II by Proposition 7.2.11.

Of course, $\underline{\Gamma}''_n$ inherits from $\underline{\Gamma}'_n$ the property that any two $\underline{B}(x), \underline{B}(y) \in \underline{\Gamma}''_n$ are such that $\underline{B}_3(x)$ and $\underline{B}_3(y)$ have disjoint interiors. Thus, for distinct $\underline{B}(x)$ and $\underline{B}(y)$ in $\underline{\Gamma}''_n$, the event that $\underline{B}(x)$ is Type-I or Type-II is independent from the event that $\underline{B}(y)$ is Type-I or Type-II.

We continue to work within the event given in (7.39). Write $s = |\underline{\Gamma}_n|$; on the event $\mathcal{E}_{\mathsf{web}}$, we have $|\underline{\mathsf{web}}_n| \leq c(p,d) n^{d-1}$ for some $c(p,d) > 0$. By Proposition 7.3.8, the k-cube set $\underline{\Gamma}_n \cup \underline{\mathsf{web}}_n$ is $\mathbb{L}^d$-connected, so we use Proposition 10.3.2 to deduce that there are at most

$$(3n)^d [c(d)]^{s + cn^{d-1}} \tag{7.41}$$

distinct possibilities for the k-cube set $\underline{\Gamma}_n \cup \underline{\mathsf{web}}_n$. The factor of $(3n)^d$ is a crude upper bound on the number of vertices of $\mathbb{Z}^d$ within $[-n,n]^d$. There are at most $2^{s+cn^{d-1}}$ ways to choose $\underline{\Gamma}_n$ from $\underline{\Gamma}_n \cup \underline{\mathsf{web}}_n$, at most $2^{s+cn^{d-1}}$ ways to choose $\underline{\Gamma}'_n$ from $\underline{\Gamma}_n$ and at most $2^{s+cn^{d-1}}$ ways to choose $\underline{\Gamma}''_n$ from $\underline{\Gamma}'_n$. We use a union bound to obtain

$$\mathbb{P}_p\Big(\Big\{|\Gamma_n| = j\Big\} \cap \Big\{|\partial^\omega G_n| < [2 \cdot 4^d (11k)^d (4k)^{d+1}]^{-1} j\Big\} \cap \mathcal{E}_{\mathsf{web}}\Big) \tag{7.42}$$

$$\leq (3n)^d [c(d)]^{s+cn^{d-1}} [c_1 \exp(-c_2 k)]^{s/(2\cdot 4^d)}\,, \tag{7.43}$$

where above, we've used the lower bound on $|\underline{\Gamma}''_n|$ in terms of s following from (7.38), as well as the bounds of Proposition 7.2.12 and the independence of the events to which these bounds are applied. Note that from (7.38) we have $s \geq j/(11k)^d (4k)^{d+1}$, and as $j \geq \gamma n^{d-1}$, we may take γ larger if necessary, again in a way depending on k, p and d, so that $s \geq cn^{d-1}$, giving

$$\mathbb{P}_p\Big(\Big\{|\Gamma_n| = j\Big\} \cap \Big\{|\partial^\omega G_n| < [2 \cdot 4^d (11k)^d (4k)^{d+1}]^{-1} j\Big\} \cap \mathcal{E}_{\mathsf{web}}\Big) \tag{7.44}$$

$$\leq (3n)^d [c(d)]^{2s} [c_1 \exp(-c_2 k)]^{s/(2\cdot 4^d)} \tag{7.45}$$

Choose k large enough in a way depending on p and d so that $[c(d)]^2[c_1 \exp(-c_2 k)]^{1/2\cdot 4^d} < e^{-1}$, at which point we consider k fixed. For such k,

$$\mathbb{P}_p\Big(\Big\{|\Gamma_n| = j\Big\} \cap \Big\{|\partial^\omega G_n| < [2 \cdot 4^d \eta_4 (4k)^{d+1}]^{-1} j\Big\} \cap \mathcal{E}_{\mathsf{web}}\Big) \leq (3n)^d \exp(-s)\,, \tag{7.46}$$

and we combine this bound with (7.36) and (7.38) to obtain

$$\mathbb{P}_p\Big(|\Gamma_n| \geq \gamma n^{d-1}\Big) \leq c_1 \exp\Big(- c_2 n^{1/(d-1)}\Big) + \sum_{j=\gamma n^{d-1}}^{\infty} (3n)^d \exp(-s)\,, \tag{7.47}$$

$$\leq c_1 \exp\Big(- c_2 n^{1/(d-1)}\Big) + \sum_{j=\gamma n^{d-1}}^{\infty} (3n)^d \exp\left(-j/[(11k)^d (4k)^{d+1}]\right)\,. \tag{7.48}$$

We choose γ sufficiently large depending on p, d and $k = k(p,d)$ to complete the proof. □

Remark 7.4.3. In the proof of Proposition 7.4.2, the renormalization parameter $k = k(p,d)$ has been fixed. All constructions given in Section 7.2 and Section 7.3 depended implicitly on this parameter.

Having exhibited control on $|\Gamma_n|$, we now return to the sets F_n defined in (7.15). We asserted that Γ_n could be thought of as the boundary of F_n; the following proposition makes this rigorous. Note that, as a result of Lemma 7.2.5, we can only conclude $F_n \subset [-n-2k, n+2k]^d$ as opposed to $F_n \subset [-n,n]^d$.

Proposition 7.4.4. Let $d \geq 3$ and $p > p_c(d)$. Define $\ell(n) := \lfloor n^{(1-\epsilon(d))/2d} \rfloor$. There exist positive constants $c_1(p,d)$ and $c_2(p,d)$ so that with probability at least

$$1 - c_1 \exp\left(-c_2 n^{(1-\epsilon(d))/2d}\right)\,, \tag{7.49}$$

whenever F_n corresponds to some $G_n \in \mathcal{G}_n$, and whenever $B = [-\ell(n), \ell(n)]^d + x$, for some $x \in \mathbb{Z}^d$ is such that $B \cap F_n \neq \emptyset$ and $B \cap F_n \neq B \cap \mathbf{C}_\infty$, we have that B either contains an endpoint vertex of an edge in Γ_n , or else the three-fold dilate of B around its center contains an endpoint vertex of an edge in $\partial^\omega G_n$.

Proof. Fix $G_n \in \mathcal{G}_n$, and consider F_n corresponding to G_n. Let B be as in the statement of the proposition with $B \cap F_n \neq \emptyset$ and $B \cap F_n \neq B \cap \mathbf{C}_\infty$. Write $B = [-\ell(n), \ell(n)]^d + x$ for

$x \in \mathbb{Z}^d$, and define augmented and dilated versions of B as

$$B^+ := [-\ell(n) - 1, \ell(n) + 1]^d + x\,, \tag{7.50}$$

$$B_3 := [-3\ell(n), 3\ell(n)]^d + x\,. \tag{7.51}$$

Suppose that B contains $y \in \mathbf{C}_\infty \setminus F_n$ which is connected to ∞ by a (not necessarily open) $\mathbb{Z}^d$-path γ' using no edges of Γ_n. As $B \cap F_n \neq \emptyset$, B contains some vertex z which is either a member of some $G_n^{(q)}$ or some small component S_j. If γ is any path from z to y within B, it must be that γ uses an edge of Γ_n, otherwise we could concatenate γ with γ' to show that either some $G_n^{(q)}$ or some S_j is not surrounded by the edge set Γ_n.

Thus we may suppose every vertex $y \in (\mathbf{C}_\infty \setminus F_n) \cap B$ is surrounded by one of the cutsets $\Gamma_n^{(q)}$ or $\widehat{\Gamma}_n^{(i)}$. Any y with this property must lie in some large component L_i. Choose some $z \in F_n \cap B$, and suppose that B_3 contains no endpoint of an edge in $\partial^\omega G_n$, so that there can be no open path between between z and y in B_3.

Work within the high probability event from Lemma 10.2.3 that for each $G_n \in \mathcal{G}_n$, every connected component of G_n satisfies $|G_n^{(q)}| \geq \eta_1 n^d$ for some $\eta_1(p,d) > 0$. Then for all n sufficiently large, the connected component of F_n containing z cannot itself be contained within B_3. Likewise, by the largeness of each L_i, and due to our choice of $\ell(n)$, it cannot be that L_i is contained in B_3.

Thus there is an open path from z to $\partial(B_3 \cap \mathbb{Z}^d)$ as well as an open path from y to $\partial(B_3 \cap \mathbb{Z}^d)$, and these paths are not joined by any open path within B_3. We have shown that, if B does not contain an edge of Γ_n, and if B_3 contains no endpoint of an edge in $\partial^\omega G_n$, then B_3 has the Type-II property (see Definition 7.2.9). Use Proposition 7.2.12 in conjunction with a union bound (taken over all such B centered at $x \in \mathbb{Z}^d$ which intersect $[-n-2k, n+2k]^d$) to see that with probability at most

$$(3n)^d c_1 \exp\Big(- c_2 n^{(1-\epsilon(n))/2d} \Big)\,, \tag{7.52}$$

there is a cube B of this form with the Type-II property. Combine this with the bounds of Lemma 10.2.3 to complete the proof. □

We now demonstrate that F_n and G_n are close in some sense. Recall that in Chapter 3, we formed from each $G_n \in \mathcal{G}_n$ the empirical measure $\mu_n \in \mathcal{M}([-1,1]^d)$ defined in (3.6). For each F_n associated to $G_n \in \mathcal{G}_n$, we define the *empirical measure* $\widetilde{\mu}_n$ associated to F_n similarly as

$$\widetilde{\mu}_n := \frac{1}{n^d} \sum_{x \in F_n} \delta_{x/n} \,. \tag{7.53}$$

Note that $\widetilde{\mu}_n$ is a member of the set $\mathcal{M}([-1-2k/n, 1+2k/n]^d)$ of signed Borel measures on $[-1-2k/n, 1+2k/n]^d$. We close Chapter 7 by observing that μ_n and $\widetilde{\mu}_n$ roughly agree on Borel sets.

Lemma 7.4.5. Let $d \geq 3$ and let $p > p_c(d)$. There exist positive constants $c_1(p,d)$, $c_2(p,d)$ and $\eta_3(p,d)$ so that for each Borel $K \subset [-1-2k/n, 1+2k/n]^d$,

$$\mathbb{P}\left(\max_{G_n \in \mathcal{G}_n} |\mu_n(K) - \widetilde{\mu}_n(K)| > \eta_3 n^{-\epsilon(d)} \right) \leq c_1 \exp\left(-c_2 n^{(d-1)/2d} \right) \,. \tag{7.54}$$

Proof. Work within the high probability event from Lemma 10.2.5 that there is $\eta_3(p,d)$ so that for all $G_n \in \mathcal{G}_n$, we have $|\partial^\omega G_n| \leq \eta_3 n^{d-1}$. For each small component S_j, the open edge boundary $\partial^\omega S_j$ contains at least one edge in $\partial^\omega G_n$. The edge sets $\{\partial^\omega S_j\}_{j=1}^t$ are pairwise disjoint, thus the number t of small components S_j is at most $\eta_3 n^{d-1}$. From the definition of a small component, we observe that for each $G_n \in \mathcal{G}_n$,

$$|F_n \setminus G_n| = \sum_{j=1}^t |S_j| \leq \eta_3 n^{d-\epsilon(d)} \tag{7.55}$$

The lemma follows from the definitions of empirical measures for G_n (3.6) and F_n (7.53). $\square$

CHAPTER 8

Contiguity

In this section, we pass from each F_n (built in the previous section) to a continuum object via another coarse graining procedure. We think of each of the emperical measures $\widetilde{\mu}_n$ as becoming "flattened" into measures representing (in the sense defined in Section 3.3) sets of finite perimeter. We rely heavily on the notation introduced in the previous section.

Remark 8.0.6. We no longer use k-cubes, and the renormalization parameter k of the last section will not come up except to say that the empirical measures $\widetilde{\mu}_n$ are elements of $\mathcal{M}([-1-2k/n, 1+2k/n]^d)$. The parameter k has itself been fixed since the proof of Proposition 7.4.2.

Given a finite collection of edges S, define $\mathsf{hull}(S)$ as

$$\mathsf{hull}(S) := \left\{ x \in \mathbb{Z}^d \ : \ \text{any } \mathbb{Z}^d\text{-path from } x \text{ to } \infty \text{ must use an edge of } S \right\}, \tag{8.1}$$

and recall that for any $x \in \mathbb{Z}^d$, we defined the unit dual cube $Q(x)$ as $[-1/2, 1/2]^d + x$. Fix $G_n \in \mathcal{G}_n$, enumerate the connected components of G_n as $\{G_n^{(q)}\}_{q=1}^M$ and for each $q \in \{1, \dots, M\}$, define

$$A_q := \left\{ i \in \{1, \dots, m\} \ : L_i \text{ is surrounded by } \Gamma_n^{(q)} \right\}. \tag{8.2}$$

For each $q \in \{1, \dots, M\}$, define $H_n^{(q)}$ as

$$H_n^{(q)} := \mathsf{hull}\left(\Gamma_n^{(q)}\right) \setminus \left(\bigcup_{i \in A_q} \mathsf{hull}\left(\widehat{\Gamma}_n^{(i)}\right) \right), \tag{8.3}$$

and let $H_n = \bigcup_{q=1}^M H_n^{(q)}$. Define the polytope P_n from H_n via

$$P_n = \left(\bigcup_{x \in H_n} n^{-1} Q(x) \right) \cap [-1, 1]^d. \tag{8.4}$$

Finally, we form the representative measure $\nu_n = \nu_n(\omega, G_n)$ of P_n (in the sense of Section 3.3); for $E \subset [-1,1]^d$ Borel,

$$\nu_n(E) := \theta_p(d)\mathcal{L}^d(E \cap P_n)\,, \tag{8.5}$$

so that $\nu_n \in \mathcal{M}([-1,1]^d)$. We repeat this construction for each $G_n \in \mathcal{G}_n$, and the goal of the coarse graining argument of this section will be to show that (with high probability) for each G_n, the measures μ_n, $\widetilde{\mu}_n$ and ν_n are all $\mathfrak{d}$-close (see (3.7) for the definition of $\mathfrak{d}$).

Remark 8.0.7. Recall that we thought of the edge sets Γ_n, defined in (7.16), as collections of "contours," as in Figure 7.6. Before executing the coarse graining procedure just mentioned, we will first need to understand how these "contours" interact. Specfically, we will rule out certain pathological configurations of contours, as this will ensure the sets H_n and P_n behave in a way amenable to the coarse graining argument.

8.1 Contour control

Each large component L_i associated to G_n is surrounded by some cutset $\Gamma_n^{(q)}$, also associated to G_n. We say that a large component L_i is *bad* if L_i is surrounded by $\Gamma_n^{(q)}$ and the corresponding connected component $G_n^{(q)}$ of G_n is surrounded by the cutset $\widehat{\Gamma}_n^{(i)}$ corresponding to L_i. Such components are bad because if they exist, subtracting the hull of $\widehat{\Gamma}_n^{(i)}$ in (8.3) from the hull of $\Gamma_n^{(q)}$ removes $G_n^{(q)}$ itself. If such components exist, we cannot expect ν_n and μ_n to be close in the appropriate sense. Fortunately, the following lemma tells us these bad components do not exist.

Lemma 8.1.1. For each $G_n \in \mathcal{G}_n$, consider the large components L_i associated to G_n, and let $G_n^{(q)}$ be a connected component of G_n. It is impossible for any large component L_i to be bad: that is, L_i cannot be surrounded by some $\Gamma_n^{(q)}$ associated to $G_n^{(q)}$, with this $G_n^{(q)}$ itself surrounded by $\widehat{\Gamma}_n^{(i)}$.

Proof. Let $G_n \in \mathcal{G}_n$, and suppose that L_i is a bad large component associated to G_n. Recall that ω_q' is the configuration obtained from ω by closing each open edge in $\partial^\omega G_n^{(q)}$. Because

L_i is surrounded by $\Gamma_n^{(q)}$, it must be that $\partial_o L_i$ is a closed cutset separating L_i from ∞ in the configuration ω'_q. Thus, back in the configuration ω, it follows that

$$\partial_o L_i \cap \partial^\omega L_i \subset \partial^\omega G_n^{(q)} . \tag{8.6}$$

Let us see how this gives rise to a contradiction: let $y \in G_n^{(q)}$. Working in the original configuration ω, consider a simple path γ from y to ∞ within $\mathbf{C}_\infty$, so that γ uses only open edges. We may assume that y is the only vertex of $G_n^{(q)}$ used by γ, as γ must eventually leave $G_n^{(q)}$ and not return. Because $G_n^{(q)}$ is surrounded by $\widehat{\Gamma}_n^{(i)}$, γ must use an open edge e in $\widehat{\Gamma}_n^{(i)}$. As $\widehat{\Gamma}_n^{(i)}$ is a closed cutset in the configuration ω'_i, this open edge e must lie in $\partial^\omega L_i$. Thus we have shown γ uses a vertex of L_i, and thus γ contains an open path from L_i to ∞ which uses no vertices of $G_n^{(q)}$, which directly contradicts (8.6). □

We extract a useful observation from the proof of Lemma 8.1.1.

Lemma 8.1.2. Let $G_n \in \mathcal{G}_n$. Each large component L_i corresponding to G_n is surrounded by exactly one cutset $\Gamma_n^{(q)}$ corresponding to a connected component of G_n.

Proof. We work with a fixed $G_n \in \mathcal{G}_n$. Each large component L_i associated to G_n is surrounded by at least one of the cutsets $\Gamma_n^{(q)}$. Suppose that L_i is surrounded by both $\Gamma_n^{(q)}$ and by $\Gamma_n^{(q')}$ for $q \neq q'$. Thanks to (8.6) in the proof of Lemma 8.1.1, we know $\partial_o L_i \cap \partial^\omega L_i \subset \partial^\omega G_n^{(q)}$ and $\partial_o L_i \cap \partial^\omega L_i \subset \partial^\omega G_n^{(q')}$. As $L_i \subset \mathbf{C}_\infty$, the edge set $\partial_o L_i \cap \partial^\omega L_i$ is non-empty, thus $\partial^\omega G_n^{(q)} \cap \partial^\omega G_n^{(q')}$ is non-empty. But this is impossible, as distinct connected components of G_n must have disjoint open edge boundary. □

We use Lemma 8.1.1 to relate F_n and H_n.

Lemma 8.1.3. For each $G_n \in \mathcal{G}_n$, we have the containment $F_n \subset H_n$.

Proof. Fix $G_n \in \mathcal{G}_n$ as usual. We begin the following claim, which we refer to as $(*)$: if $y \in F_n$ and $y \notin H_n$, then there exists a large component L_i and a connected component $G_n^{(q)}$ of G_n such that y is surrounded by both $\Gamma_n^{(q)}$ and $\widehat{\Gamma}_n^{(i)}$, and such that L_i is itself surrounded by $\Gamma_n^{(q)}$.

The claim $(*)$ follows directly from the definition of H_n: as $y \in F_n$, we have $y \in \mathsf{hull}(\Gamma_n^{(q)})$ for some q. The only way we could have $y \notin H_n$ is if y were surrounded by some $\widehat{\Gamma}_n^{(i)}$ for $i \in A_q$, and we recall that A_q indexes the large components L_i which are surrounded by $\Gamma_n^{(q)}$.

Supposing that there is $y \in F_n \setminus H_n$ (for the sake of contradiction), we now consider the large component L_i given by $(*)$ and we pass to the configuration ω_i' in which each edge of $\partial^\omega L_i$ is made closed. In this configuration, $\widehat{\Gamma}_n^{(i)}$ consists only of closed edges. Let Λ be the open cluster containing y in the configuration ω_i', so that Λ is surrounded by $\widehat{\Gamma}_n^{(i)}$.

Let $F_n^{(q)}$ denote the connected component of F_n containing $G_n^{(q)}$, and let $z \in F_n^{(q)}$. Suppose for the sake of contradiction that $z \notin \Lambda$. Let γ be a path from y to z within $F_n^{(q)}$, so that in the original configuration ω, the path γ uses only open edges. From the assumption $z \notin \Lambda$, if we pass to the configuration ω_i', we see that γ must use a closed edge e, which is necessarily an element of $\partial^\omega L_i$ back in the configuration ω.

Because γ is a path in $F_n^{(q)}$, it must join two vertices which are either in G_n or in one of the small components S_j. But $e \in \partial^\omega L_i$, which means an endpoint of e must also lie in L_i. It is impossible for e to satisfy all these requirements. Thus, we conclude that $F_n^{(q)} \subset \Lambda$, and consequently $G_n^{(q)} \subset \Lambda$. In particular, this means $G_n^{(q)}$ is surrounded by $\widehat{\Gamma}_n^{(i)}$, which implies (through $(*)$) that L_i is a bad component. We apply Lemma 8.1.1 to complete the proof. $\square$

Remark 8.1.4. Lemma 8.1.3 demonstrates that $F_n \subset H_n \cap \mathbf{C}_\infty$, but the opposite containment is also immediate from the construction of H_n. Indeed, suppose $x \in H_n \cap \mathbf{C}_\infty$. Then x is surrounded by some $\Gamma_n^{(q)}$, and hence x is either in $G_n^{(q)}$ for some q, or x is an element of one of the large or small components (L_i or S_j) associated to G_n. It is impossible for x to lie in any L_i, as these components were excised in the construction (8.3) of H_n.

We finally use Proposition 7.4.2 to control the perimeter of each P_n.

Corollary 8.1.5. Let $d \geq 3$ and $p > p_c(d)$. There are positive constants $c_1(p,d), c_2(p,d)$ and $\gamma(p,d)$ so that

$$\mathbb{P}_p\left(\max_{G_n \in \mathcal{G}_n} \operatorname{per}(nP_n) \geq \gamma n^{d-1}\right) \leq c_1 \exp\left(-c_2 n^{1/(d-1)}\right) . \tag{8.7}$$

Proof. We work within the high probability event $\mathcal{E}$ corresponding to Proposition 7.4.2 that for each $G_n \in \mathcal{G}_n$, the corresponding cutset Γ_n satisfies $|\Gamma_n| < \gamma n^{d-1}$. Define the polytope $n\widetilde{P}_n$ as

$$n\widetilde{P}_n := \bigcup_{x \in H_n} Q(x)\,. \tag{8.8}$$

Every boundary face of $n\widetilde{P}_n$ has $\mathcal{H}^{d-1}$-measure one, and each boundary face of $n\widetilde{P}_n$ corresponds to a unique edge of Γ_n, thus within the event $\mathcal{E}$, the polytope $n\widetilde{P}_n$ has perimeter at most γn^{d-1}. As $nP_n = n\widetilde{P}_n \cap [-n,n]^d$, the perimeter of nP_n is at most $\gamma n^{d-1} + 2d(2n)^{d-1}$, which completes the proof upon redefining γ. □

8.2 A contiguity argument

Recall the metric $\mathfrak{d}$ introduced in (3.7) We now adapt the argument given in Section 16.2 of [Cer06] to our situation; we will show that μ_n and ν_n are $\mathfrak{d}$-close with high probability.

Remark 8.2.1. We use another renormalization argument, this time at a different scale, and we find it convenient to reuse the notation from Chapter 7. Define $\ell(n) := \lfloor n^{(1-\epsilon(d))/2d} \rfloor$, and suppress the dependence of on n by writing $\ell(n)$ as ℓ. We *redefine* $\underline{B}(x)$ to be the ℓ-cube $(2\ell)x + [-\ell,\ell]^d$. We will also work with 3ℓ-cubes, insofar as they are used in the statement of Proposition 7.4.4. These are defined as in (7.3).

Let $\delta > 0$, and introduce the $\mathbb{Z}^d$-process $\{Z_x^{(\delta)}\}_{x \in \mathbb{Z}^d}$, with each $Z_x^{(\delta)}$ the indicator function of the event

$$\left\{ \frac{|\mathbf{C}_\infty \cap \underline{B}(x)|}{\mathcal{L}^d(\underline{B}(x))} \in (\theta_p(d) - \delta,\, \theta_p(d) + \delta) \right\}\,. \tag{8.9}$$

Using Corollary 10.1.5 on each ℓ-cube intersecting $[-n,n]^d$ and via our careful examination of the contours which define F_n and nP_n, we will show $\widetilde{\mu}_n$ and ν_n are $\mathfrak{d}$-close.

Proposition 8.2.2. Let $d \geq 3$ and let $p > p_c(d)$. Let $Q \subset [-1,1]^d$ be an axis-parallel cube. For all $\delta > 0$, there are positive constants $c_1(p,d,\delta)$ and $c_2(p,d,\delta)$ so that

$$\mathbb{P}_p\left(\max_{G_n \in \mathcal{G}_n} |\widetilde{\mu}_n(Q) - \nu_n(Q)| \geq \delta \right) \leq c_1 \exp\left(-c_2 n^{(1-\epsilon(d))/2d}\right)\,. \tag{8.10}$$

Proof. Fix $G_n \in \mathcal{G}_n$, and let F_n, $\widetilde{\mu}_n$, P_n and ν_n be the objects constructed above for this G_n. Throughout the proof, we use bounds involving constants $c(d), c(p,d)$ and so on, which are understood to be positive, and which may change from line to line. Let $\underline{L}$ denote the following collection of ℓ-cubes:

$$\underline{L} := \left\{ \underline{B}(x) \ : \underline{B}(x) \cap [-n-2k, n+2k]^d \neq \emptyset \right\}, \tag{8.11}$$

For each ℓ-cube $\underline{B}(x)$, we have the bounds

$$\widetilde{\mu}_n \left(n^{-1} \underline{B}(x) \right) \leq c(d) \left(\frac{\ell}{n} \right)^d , \quad \nu_n \left(n^{-1} \underline{B}(x) \right) \leq c(d) \left(\frac{\ell}{n} \right)^d . \tag{8.12}$$

As Q is an axis-parallel cube, its boundary ∂Q intersects at most $c(d)n^{d-1}$ cubes $\underline{B}(x)$, so that by using the bounds in (8.12), we have

$$|\widetilde{\mu}_n(Q) - \nu_n(Q)| \leq c(d) \left(\frac{\ell^d}{n} \right) + \sum_{\underline{B}(x) \in \underline{L}} \left| \widetilde{\mu}_n \left(n^{-1} \underline{B}(x) \right) - \nu_n \left(n^{-1} \underline{B}(x) \right) \right| . \tag{8.13}$$

Define the following high probability events:

$$\mathcal{E}_1 := \left\{ \max_{G_n \in \mathcal{G}_n} \operatorname{per}(nP_n) < \gamma n^{d-1} \right\} , \quad \mathcal{E}_2 := \left\{ \max_{G_n \in \mathcal{G}_n} |\Gamma_n| < \gamma n^{d-1} \right\} , \tag{8.14}$$

$$\mathcal{E}_3 := \left\{ \max_{G_n \in \mathcal{G}_n} |\partial^\omega G_n| \leq \eta_3 n^{d-1} \right\} , \tag{8.15}$$

so that $\mathcal{E}_1$, $\mathcal{E}_2$ and $\mathcal{E}_3$ are respectively high probability events from Corollary 8.1.5, Proposition 7.4.2 and Lemma 10.2.5. Finally, let $\mathcal{E}_4$ be the high probability event in the statement of Proposition 7.4.4. We work within the intersection of $\mathcal{E}_1$ through $\mathcal{E}_4$; the effect of working within these events is that we may regard both nP_n and F_n as objects whose perimeters are on the order of n^{d-1}.

Motivated by the event $\mathcal{E}_4$, define $\underline{L}' \subset \underline{L}$ as

$$\underline{L}' := \left\{ \underline{B}(x) \in \underline{L} \ : \begin{array}{c} \underline{B}(x) \cap nP_n = \emptyset \ \text{ or } \ \underline{B}(x) \cap nP_n = \underline{B}(x) \\ \textbf{and} \\ \underline{B}(x) \cap F_n = \emptyset \ \text{ or } \ \underline{B}(x) \cap F_n = \underline{B}(x) \cap \mathbf{C}_\infty \end{array} \right\} . \tag{8.16}$$

As we are working within $\mathcal{E}_1$ through $\mathcal{E}_4$, it follows that are are at most $c(p,d)n^{d-1}$ ℓ-cubes $\underline{B}(x)$ which are in $\underline{L}\backslash\underline{L}'$. This is especially due to the event $\mathcal{E}_4$ from Proposition 7.4.4, which has essentially been designed for use here. Thus,

$$|\widetilde{\mu}_n(Q) - \nu_n(Q)| \le c(p,d)\left(\frac{\ell^d}{n}\right) + \sum_{\underline{B}(x)\in\underline{L}'} \left|\widetilde{\mu}_n\left(n^{-1}\underline{B}(x)\right) - \nu_n\left(n^{-1}\underline{B}(x)\right)\right| . \tag{8.17}$$

Let us define a further subcollection of boxes, $\underline{L}'' \subset \underline{L}'$:

$$\underline{L}'' := \left\{ \underline{B}(x) \in \underline{L}' \;:\; \begin{array}{c} \underline{B}(x)\cap nP_n = \emptyset \text{ and } \underline{B}(x)\cap F_n = \emptyset \\ \textbf{or} \\ \underline{B}(x)\cap nP_n = \underline{B}(x) \text{ and } \underline{B}(x)\cap F_n = \underline{B}(x)\cap \mathbf{C}_\infty \end{array} \right\} . \tag{8.18}$$

We now make the claim that $\underline{L}'' = \underline{L}'$. To prove this, we show that two of the four possible cases which define $\underline{L}'$ are impossible.

Case (i): Suppose that $\underline{B}(x)\cap F_n = \underline{B}(x)\cap\mathbf{C}_\infty$ and $\underline{B}(x)\cap nP_n = \emptyset$. This case is handled entirely by Lemma 8.1.3: as $F_n \subset H_n$, this situation is impossible unless $\mathbf{C}_\infty \cap \underline{B}(x) = \emptyset$, which is one of the two options we allow for.

Case (ii): Suppose that $\underline{B}(x)\cap F_n = \emptyset$ and $\underline{B}(x)\cap nP_n = \underline{B}(x)$. If $\underline{B}(x)\cap\mathbf{C}_\infty = \emptyset$, we are in one of the two allowed options, so we may assume there is some $y \in \underline{B}(x)\cap\mathbf{C}_\infty$. As $\underline{B}(x)\cap nP_n = \underline{B}(x)$, it follows that $y \in H_n$. Thus y is surrounded by some $\Gamma_n^{(q)}$, so that either $y \in F_n$ or $y \in L_i$ for some i. The former option is impossible by hypothesis, so we may conclude $y \in L_i$ for some i. By Lemma 8.1.2, this large component L_i containing y is surrounded by exactly one of the $\Gamma_n^{(q)}$, so that $y \in H_n$ implies $y \in H_n^{(q)}$ and $y \notin H_n^{(q')}$ whenever $q' \neq q$. But in the construction of $H_n^{(q)}$, we see that the hull of $\widehat{\Gamma}_n^{(i)}$ is removed from the hull of $\Gamma_n^{(q)}$. As the hull of $\widehat{\Gamma}_n^{(i)}$ contains L_i and hence y, it is also impossible that $y \in H_n^{(q)}$. This is a contradiction.

We thus conclude that $\underline{L}'' = \underline{L}'$, and we may then replace $\underline{L}'$ by $\underline{L}''$ in (8.17) and use the defining properties of $\underline{L}''$ in conjunction with the definitions of $\widetilde{\mu}_n$ and ν_n:

$$|\widetilde{\mu}_n(Q) - \nu_n(Q)| \le c(p,d)\left(\frac{\ell^d}{n}\right) + \sum_{\underline{B}(x)\in\underline{L}''} \left|\widetilde{\mu}_n\left(n^{-1}\underline{B}(x)\right) - \nu_n\left(n^{-1}\underline{B}(x)\right)\right| , \tag{8.19}$$

$$\le c(p,d)\left(\frac{\ell^d}{n}\right) + \sum_{\underline{B}(x)\in\underline{L}''} \left(\left|\frac{|\mathbf{C}_\infty \cap \underline{B}(x)|}{n^d} - \frac{\theta_p(d)\mathcal{L}^d(\underline{B}(x))}{n^d}\right|\right) . \tag{8.20}$$

Let us form one last high probability event $\mathcal{E}_5$ using the $\mathbb{Z}^d$-process $\{Z_x^{(\delta)}\}_{x\in\mathbb{Z}^d}$: we let $\mathcal{E}_5$ be the event that $Z_x^{(\delta)} = 1$ for all x with $\underline{B}(x) \in \underline{L}''$. As a consequence of Corollary 10.1.5, there are positive constants $c_1(p,d,\delta)$ and $c_2(p,d,\delta)$ so that

$$\mathbb{P}(\mathcal{E}_5^c) \leq c(d)n^d c_1 \exp\Big(- c_2 n^{(1-\epsilon(d))/2d}\Big) . \tag{8.21}$$

Working within the intersection of $\mathcal{E}_1$ through $\mathcal{E}_5$, we may now bound $|\widetilde{\mu}_n(Q) - \nu_n(Q)|$ as follows, continuing from (8.20).

$$\begin{aligned} |\widetilde{\mu}_n(Q) - \nu_n(Q)| &\leq c(p,d)\left(\frac{\ell^d}{n}\right) + |\underline{L}''| \max_{\underline{B}(x)\in\underline{L}''}\left(\left|\frac{|\mathbf{C}_\infty \cap \underline{B}(x)|}{n^d} - \frac{\theta_p(d)\mathcal{L}^d(\underline{B}(x))}{n^d}\right|\right), && (8.22)\\ &\leq c(p,d)\left(\frac{\ell^d}{n}\right) + \frac{|\underline{L}''|}{n^d}(2\mathcal{L}^d(\underline{B}(x))\delta), && (8.23)\\ &\leq c(p,d)\left(\frac{\ell^d}{n}\right) + c(d)\delta, && (8.24)\end{aligned}$$

where we have used the bound $|\underline{L}''| \leq |\underline{L}| \leq c(d)(n/\ell)^d$ in going from the second line to the third line directly above. Taking n sufficiently large, we have $|\widetilde{\mu}_n(Q) - \nu_n(Q)| \leq c(p,d)\delta$, and the proof is complete after using the bounds for events $\mathcal{E}_1, \dots, \mathcal{E}_5$. □

We combine the preceding result with Lemma 7.4.5 to establish $\mathfrak{d}$-closeness of μ_n and ν_n. The following is the central theorem of this section.

Theorem 8.2.3. Let $d \geq 3$, $p > p_c(d)$ and let $\delta > 0$. There are positive constants $c_1(p,d,\delta)$, $c_1(p,d,\delta)$ so that

$$\mathbb{P}_p\left(\max_{G_n\in\mathcal{G}_n} \mathfrak{d}(\mu_n,\nu_n) \geq \delta\right) \leq c_1 \exp\left(-c_2 n^{(1-\epsilon(d))/2d}\right) . \tag{8.25}$$

Proof. Let $\delta > 0$ and let Δ^k denote the collection of dyadic cubes in $[-1,1]^d$ at scale k. There should be no confusion between the integer k used for dyadic scales and the renormalization parameter from Chapter 7, as the latter has been fixed (see Remark 8.0.6). Let $Q \in \Delta^k$. Thanks to Lemma 7.4.5 and Proposition 8.2.2, there exist positive constants $c_1(p,d,\delta), c_2(p,d,\delta)$ so that when n is taken sufficiently large depending on δ,

$$\mathbb{P}_p\left(\max_{G_n\in\mathcal{G}_n} |\mu_n(Q) - \nu_n(Q)| < \delta\right) \geq 1 - c_1\exp\Big(- c_2 n^{(1-\epsilon(d))/2d}\Big). \tag{8.26}$$

Take j to be large enough so that $2^{-j} < \delta$, and we let $Q_1, \ldots, Q_m$ enumerate all dyadic cubes at scales between 0 and $j-1$. Note that the number m of these cubes depends only on δ and d. Let $\mathcal{E}_i$ be the high probability event corresponding to (8.26) for the cube Q_i, and work within the event $\mathcal{E} := \bigcap_{i=1}^m \mathcal{E}_i$, so that by definition (3.7) of the metric $\mathfrak{d}$,

$$\mathfrak{d}(\mu_n, \nu_n) \leq \sum_{k=0}^{j-1} \frac{1}{2^k} \sum_{Q \in \Delta^k} \frac{1}{|\Delta^k|} |\mu_n(Q) - \nu_n(Q)| + \sum_{k=j}^{\infty} \frac{1}{2^k} \sum_{Q \in \Delta^k} \frac{1}{|\Delta^k|} |\mu_n(Q) - \nu_n(Q)|, \quad (8.27)$$

$$\leq 2\delta + \sum_{k=j}^{\infty} \frac{1}{2^k} \sum_{Q \in \Delta^k} \frac{1}{|\Delta^k|} |\mu_n(Q) - \nu_n(Q)|. \quad (8.28)$$

We control the sum directly above via crude bounds: there is $c(d) > 0$ such that for each dyadic cube Q, we have $\mu_n(Q) \leq c(d)$ and $\nu_n(Q) \leq c(d)$. Through our choice of j, the sum in (8.28) is then bounded by $c(d)\delta$. Thus, within the event $\mathcal{E}$, we have $\mathfrak{d}(\mu_n, \nu_n) \leq c(d)\delta$. As m depends only on ϵ and d, the proof is complete upon using (8.26) in conjunction with a union bound to control the probability of $\mathcal{E}^c$. □

We explore some consequences of Theorem 8.2.3 before moving to the final chapter of this paper.

8.3 Closeness to sets of finite perimeter

Recall from Section 3.3 that $\mathcal{B}_d$ is the ball about the zero measure of radius 3^d (with respect to the total variation norm). For $\gamma, \xi > 0$, define the following collection of measures in $\mathcal{B}_d$.

$$\mathcal{P}_{\gamma,\xi} := \left\{ \nu_F \ : F \subset [-1,1]^d, \ \mathrm{per}(F) \leq \gamma, \ \mathcal{L}^d(F) \leq \mathcal{L}^d((1+\xi)W_{p,d}) \right\}, \quad (8.29)$$

where given $F \subset [-1,1]^d$ Borel, the measure ν_F representing F is defined as in Section 3.3 (as in the definition (8.5) of the ν_n).

Corollary 8.3.1. Let $d \geq 3$, $p > p_c(d)$ and let $\delta > 0$. There exist positive constants $c_1(p,d,\delta,\xi)$, $c_2(p,d,\delta,\xi)$ and $\gamma(p,d)$ so that

$$\mathbb{P}_p \left(\max_{G_n \in \mathcal{G}_n} \mathfrak{d}(\mu_n, \mathcal{P}_{\gamma,\xi}) \geq \delta \right) \leq c_1 \exp\left(-c_2 n^{(1-\epsilon(d))/2d} \right). \quad (8.30)$$

Proof. Let $\delta, \delta', \xi > 0$ and let $\gamma(p,d)$ be as in Corollary 8.1.5. We will first show that, with high probability, the measures ν_n lie in $\mathcal{P}_{\gamma,\xi}$, and then we will apply Theorem 8.2.3. Work within the intersection of the following high probability events

$$\mathcal{E}_1 := \left\{ \max_{G_n \in \mathcal{G}_n} \mathrm{per}(nP_n) < \gamma n^{d-1} \right\}, \quad \mathcal{E}_2 := \left\{ \max_{G_n \in \mathcal{G}_n} \mathfrak{d}(\mu_n, \nu_n) < \min(\delta, \delta') \right\}, \tag{8.31}$$

$$\mathcal{E}_3 := \left\{ \frac{|\mathbf{C}_n|}{(2n)^d} \in (\theta_p(d) - \delta', \theta_p(d) + \delta') \right\}, \tag{8.32}$$

respectively from Corollary 8.1.5, Theorem 8.2.3 and Corollary 10.1.5. Because we are within $\mathcal{E}_2$, for each nP_n corresponding to $G_n \in \mathcal{G}_n$ we have

$$\theta_p(d)\mathcal{L}^d(nP_n) < \delta' n^d + |G_n|, \tag{8.33}$$

$$< \delta' n^d + |\mathbf{C}_n|/d!\,. \tag{8.34}$$

From working within $\mathcal{E}_3$, we further conclude

$$\mathcal{L}^d(nP_n) < n^d \left(\frac{\delta'}{\theta_p(d)} + \frac{2^d}{d!} \left(1 + \frac{\delta'}{\theta_p(d)} \right) \right), \tag{8.35}$$

$$< n^d \left(\mathcal{L}^d((1+\xi) W_{p,d}) \right), \tag{8.36}$$

where we have taken δ' sufficiently small according to p, d and ξ. Because we are working within $\mathcal{E}_1$, we conclude that $\nu_n \in \mathcal{P}_{\gamma,\xi}$ for each $G_n \in \mathcal{G}_n$. □

Remark 8.3.2. It is important to understand how $\mathfrak{d}$ is related to the notion of weak convergence. We claim that, in fact, the metric $\mathfrak{d}$ restricted to $\mathcal{P}_{\gamma,\xi}$ encodes weak convergence of the measures in $\mathcal{P}_{\gamma,\xi}$. Let ζ, $\zeta_n \in \mathcal{P}_{\gamma,\xi}$; if $\mathfrak{d}(\zeta_n, \zeta) \to 0$, it follows from (3.7) that $\zeta_n(Q) \to \zeta(Q)$ for each dyadic cube Q. As any open subset $U \subset [-1,1]^d$ may be decomposed into a countable union of (almost disjoint) dyadic cubes, and as each measure in $\mathcal{P}_{\gamma,\xi}$ is absolutely continuous with respect to Lebesgue measure, we may conclude that $\liminf_{n\to\infty} \zeta_n(U) \geq \zeta(U)$. By the Portmanteau theorem, ζ_n converges to ζ weakly.

Conversely, if ζ, $\zeta_n \in \mathcal{P}_{\gamma,\xi}$ are such that $\zeta_n \to \zeta$ weakly, the Portmanteau theorem also tells us that $\zeta_n(A) \to \zeta(A)$ for all continuity sets $A \subset [-1,1]^d$ of ζ, in particular whenever A is a dyadic cube (using the absolute continuity of ζ). It is then immediate that $\mathfrak{d}(\zeta_n, \zeta) \to 0$.

Lemma 8.3.3. The collection of measures $\mathcal{P}_{\gamma,\xi}$ is compact subset of the metric space $(\mathcal{B}_d, \mathfrak{d})$.

Proof. It is a consequence of the Banach-Alaoglu theorem that the set $\mathcal{B}_d$ is compact when equipped with the topology of weak convergence. Because the continuous functions on $[-1,1]^d$ (equipped with the supremum norm topology) form a separable space, $\mathcal{B}_d$ is sequentially compact in the topology of weak convergence. By Remark 8.3.2, it then suffices to show $\mathcal{P}_{\gamma,\xi}$ is sequentially closed.

Let $\{\nu_{F_n}\}_{n=1}^{\infty}$ be a sequence of measures in $\mathcal{P}_{\gamma,\xi}$ which converge with respect to the metric $\mathfrak{d}$. Using the definition (3.7) of $\mathfrak{d}$, one can show (first by approximating open sets by finite unions of dyadic cubes, and then approximating measurable sets by open sets) that for any $E \subset [-1,1]^d$ measurable, the sequence $\mathcal{L}^d(E \cap F_n)$ is Cauchy. Thus the indicator functions $\mathbf{1}_{F_n}$ converge pointwise a.e. to some $\mathbf{1}_F$, and the bounded convergence theorem allows us to convert this into L^1-convergence.

As $\mathbf{1}_{F_n} \to \mathbf{1}_F$ in L^1-sense, it is immediate from (3.7) that $\mathfrak{d}(\nu_{F_n}, \nu_F) \to 0$ as $n \to \infty$, and it remains to check that $\nu_F \in \mathcal{P}_{\gamma,\xi}$. By Fatou's lemma and the lower semicontinuity of the perimeter functional (Lemma 10.4.1),

$$\mathcal{L}^d(F) \leq \liminf_{n\to\infty} \mathcal{L}^d(F_n)\,, \qquad \mathrm{per}(F) \leq \liminf_{n\to\infty} \mathrm{per}(F_n)\,, \tag{8.37}$$

so that $\nu_F \in \mathcal{P}_{\gamma,\xi}$. □

As mentioned in Section 3.3, the metric $\mathfrak{d}$ ultimately provides a way to compare discrete objects with continuous objects (both encoded as measures). We work with $\mathfrak{d}$ extensively in the next section, and in fact most of our effort will go towards proving the following precursor to Theorem 2.2.1.

Theorem 8.3.4. For $d \geq 3$ and $p > p_c(d)$, let $W_{p,d}$ be the Wulff crystal from Theorem 2.2.1. Define the following subset of $\mathcal{M}([-1,1]^d)$:

$$\mathcal{W} := \left\{ \nu_E \ : E = W_{p,d} + x, \text{ with } W_{p,d} + x \subset [-1,1]^d \right\}. \tag{8.38}$$

We have that $\mathbb{P}_p$-almost surely,

$$\max_{G_n \in \mathcal{G}_n} \mathfrak{d}(\mu_n, \mathcal{W}) \xrightarrow[n\to\infty]{} 0\,. \tag{8.39}$$

CHAPTER 9

Lower bounds and main results

We prove the main theorems of the paper in this section. The strategy is as follows: we first use Corollary 8.3.1 to anchor the empirical measures μ_n of the G_n near (in the sense of the metric $\mathfrak{d}$) measures representing sets of finite perimeter. Whenever an empirical measure μ_n is close to such a measure ν_F, we will relate the conductance of the corresponding G_n to the conductance of the continuum set F.

The challenge is to show that when $\mathfrak{d}(\mu_n, \nu_F)$ is small, $|\partial^\omega G_n|$ and $\mathcal{I}_{p,d}(nF)$ are close. This is accomplished using a covering lemma, working locally near the boundary of F. This local perspective guides a surgery we perform $\partial^\omega G_n$, which is done in order to invoke concentration estimates from Section 6.1.

This plan of attack shares much with the argument in Section 6 of [CT11]. In particular, we rely on the compactness of the space $\mathcal{P}_{\gamma,\xi}$ (Lemma 8.3.3).

9.1 Setup, the reduced boundary and a covering lemma

Recall that α_d denotes the volume of the d-dimensional Euclidean unit ball. Given a closed Euclidean ball $B(x,r)$ centered at $x \in \mathbb{R}^d$ of radius $r > 0$ and a unit vector $v \in \mathbb{S}^{d-1}$, define

$$B_-(x,r,v) := \left\{ y \in \mathbb{R}^d \, : (y-x) \cdot v \leq 0 \right\} \tag{9.1}$$

to be the *lower half-ball* of $B(x,r)$ in the direction v.

Definition 9.1.1. For $F \subset \mathbb{R}^d$ Borel, let $\nabla \mathbf{1}_F$ be the distributional derivative of the indicator function $\mathbf{1}_F$. This is a vector-valued measure whose total variation $||\nabla \mathbf{1}_F||(\mathbb{R}^d)$ is the perimeter of F. For $F \subset \mathbb{R}^d$ a set of finite perimeter, the *reduced boundary* $\partial^* F$ of F is the set of points $x \in \mathbb{R}^d$ such that (i) and (ii) hold:

(i) $||\nabla \mathbf{1}_F||(B(x,r)) > 0$ for any $r > 0$.

(ii) If we define

$$v_r(x) := -\frac{\nabla \mathbf{1}_F(B(x,r))}{||\nabla \mathbf{1}_F||(B(x,r))}\,, \tag{9.2}$$

then as $r \to 0$, the sequence $v_r(x)$ tends to a limiting unit vector $v_F(x)$, which we call the *exterior normal* to F at x.

The following "covering lemma" has been specialized to the surface energy $\mathcal{I}_{p,d}$. Its proof may be found in Section 14.3 of [Cer06].

Lemma 9.1.2. Let $F \subset \mathbb{R}^d$ be a set of finite perimeter, and let $\mathcal{I}_{p,d}$ be the surface energy defined in (3.9) for the norm $\beta_{p,d}$. Given $\delta > 0$ and $s \in (0, 1/2)$, there is a finite collection of disjoint balls $\{B(x_i, r_i)\}_{i=1}^m$ with $x_i \in \partial^* F$ and $r_i \in (0,1)$ for all $i \in \{1, \dots, m\}$. Moreover, these balls satisfy

$$\mathcal{L}^d\Big(F \cap B(x_i, r_i) \;\Delta\; B_-(x_i, r_i, v_F(x_i))\Big) \leq \delta \alpha_d r_i^d\,, \tag{9.3}$$

$$\left| \mathcal{I}_{p,d}(F) - \sum_{i=1}^m \alpha_{d-1} r_i^{d-1} \tau(v_F(x_i)) \right| \leq \widetilde{\delta}_F(s)\,, \tag{9.4}$$

where $\widetilde{\delta}_F(s) := \frac{s}{4}\mathcal{I}_{p,d}(F)$.

Remark 9.1.3. Given a set $F \subset [-1,1]^d$ of finite perimeter and a ball $B(x,r)$ with $x \in \partial^* F$ arising from Lemma 9.1.2, we abbreviate $B_-(x, r, v_F(x))$ as $B_-(x,r)$ to make our notation more concise.

We now define several global parameters which will show up throughout the section. Given a collection of balls $\{B(x_i, r_i)\}_{i=1}^m$ as in Lemma 9.1.2, we define

$$\epsilon_F := \delta \min_{i=1}^{m} \alpha_d (r_i)^d\,, \tag{9.5}$$

so that ϵ_F depends on F, δ and s. Also define

$$\lambda_F(s) := (1 - 2s)\mathcal{I}_{p,d}(F)\,. \tag{9.6}$$

Remark 9.1.4. Given $G_n \in \mathcal{G}_n$ and a ball $B(x_i, r_i)$ as in Lemma 9.1.2 for the paramters $\delta > 0$ and $s \in (0, 1/2)$, we define the event

$$\mathcal{E}(G_n, i) := \left\{ |\partial^\omega G_n \cap nB(x_i, r_i)| \leq (1-s)n^{d-1}\alpha_{d-1}(r_i)^{d-1}\beta_{p,d}(v_i) \right\} \tag{9.7}$$

for the sake of notational convenience.

The next lemma allows us to control the probability that $|\partial^\omega G_n|$ is too small when μ_n and ν_F are $\mathfrak{d}$-close using the events $\mathcal{E}(G_n, i)$.

Lemma 9.1.5. Let $d \geq 3$, let $p > p_c(d)$, and suppose that $F \subset [-1, 1]^d$ is a set of finite perimeter. Let $\{B(x_i, r_i)\}_{i=1}^m$ be a collection of balls as in Lemma 9.1.2 for $\delta > 0$ and $s \in (0, 1/2)$. Then for each $G_n \in \mathcal{G}_n$, we have

$$\left\{ |\partial^\omega G_n| \leq \lambda_F(s) n^{d-1} \right\} \subset \bigcup_{i=1}^m \mathcal{E}(G_n, i) \,, \tag{9.8}$$

where $\lambda_F(s)$ is defined in (9.6).

Proof. Because the balls $\{B(x_i, r_i)\}_{i=1}^m$ were chosen in accordance with Lemma 9.1.2, we combine (9.4) with the defintion of $\widetilde{\delta}_F(s)$ to obtain

$$\left| \mathcal{I}_{p,d}(F) - \sum_{i=1}^m \alpha_{d-1}(r_i)^{d-1}\beta_{p,d}(v_i) \right| \leq \frac{s}{2} \left(\sum_{i=1}^m \alpha_{d-1}(r_i)^{d-1}\beta_{p,d}(v_i) \right) \,, \tag{9.9}$$

so that

$$\lambda_F(s) \leq (1-s) \left(\sum_{i=1}^m \alpha_{d-1}(r_i)^{d-1}\beta_{p,d}(v_i) \right) \,. \tag{9.10}$$

Now use the disjointness of the balls in the collection $\{B(x_i, r_i)\}_{i=1}^m$, (9.10) and the definition (9.7) of the event $\mathcal{E}(G_n, i)$.

$$\left\{ |\partial^\omega G_n| \leq \lambda_F(s) n^{d-1} \right\} \subset \left\{ \sum_{i=1}^m |\partial^\omega G_n \cap nB(x_i, r_i)| \leq (1-s)n^{d-1} \sum_{i=1}^m \alpha_{d-1}(r_i)^{d-1}\beta_{p,d}(v_i) \right\} \,, \tag{9.11}$$

$$\subset \bigcup_{i=1}^m \mathcal{E}(G_n, i) \tag{9.12}$$

We complete the proof using the definition (9.5) of ϵ_F. $\square$

9.2 Local surgery on each $\partial^\omega G_n$

In this section, we think of $F \subset [-1,1]^d$ as a fixed polytope. We also work with a fixed $G_n \in \mathcal{G}_n$. Let $\{B(x_i, r_i)\}_{i=1}^m$ be a collection of balls as in Lemma 9.1.2 for F. Fix $B(x_i, r_i) \in \{B(x_i, r_i)\}_{i=1}^m$, write this ball as $B(x, r)$, and let $v := v_F(x) \in \mathbb{S}^{d-1}$ be the exterior normal vector associated to $x \in \partial^* F$.

Let $B_-(x, r)$ be the lower half-ball associated to $B(x, r)$ and the unit vector v. Let $D(x, r)$ be the closed equatorial disc of this ball, so that $\mathsf{hyp}(D(x, r))$ is orthogonal to v. For $h > 0$ small, define $r' := (1 - h^2)^{1/2} r$, and let $D(x, r') \subset D(x, r)$ be the closed disc of radius r' centered at x. Note that $D(x, r')$ is built so that

$$\mathsf{cyl}(D(x, r'), hr') \subset B(x, r)\,. \tag{9.13}$$

These geometric objects guide a surgery we perform on $\partial^\omega G_n$.

Let $J_n = J_n(\omega)$ be the collection of open edges having non-trivial intersection with $\mathcal{N}_{5d}(nD(x, r))$. If we close each edge in J_n and each edge in $\partial^\omega G_n$, we break $\mathbf{C}_\infty \cap nB(x, r)$ into a finite number of connected components. We say that one of these components Λ is *outward* if it is contained in $G_n \cap (nB(x, r) \setminus nB_-(x, r))$ and *inward* if it is contained in $nB_-(x, r) \setminus G_n$.

We are only interested in Λ which contain vertices incident to edges in J_n. We enumerate all such outward components as $\Lambda_1^+, \dots, \Lambda_{\ell^+}^+$, and all such inward components as $\Lambda_1^-, \dots, \Lambda_{\ell^-}^-$. Let us say that a component (outward or inward) is *good* if it is contained in $n\mathsf{cyl}(D(x, r), hr')$ and say that it is *bad* otherwise. Our notation suppresses the dependence of these components on G_n, $B(x, r)$ and F.

Remark 9.2.1. Every outward component is a subgraph of G_n, so "outward" should be understood as relative to the bottom half-ball $nB_-(x, r)$. See Figure 9.1 for an illustration of the objects introduced so far. In general, we regard outward and inward components $\Lambda_j^\pm$ as subgraphs of $\mathbf{C}_\infty$, so that the edge set $\mathrm{E}(\Lambda_j^\pm)$ is always a collection of open edges.

The following lemma allows us to efficiently truncate each bad component $\Lambda_j^\pm$. Let

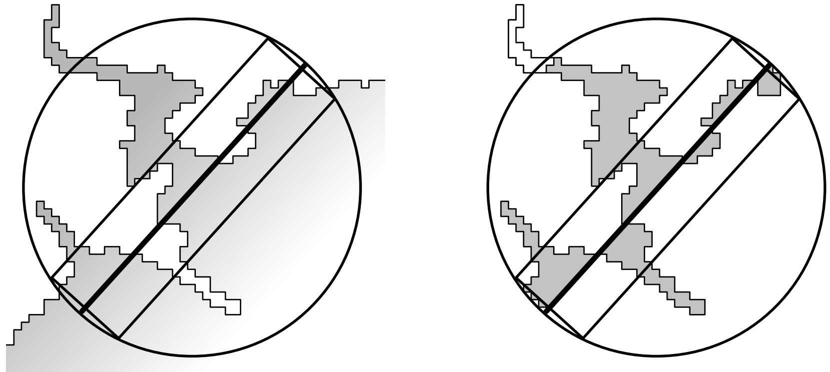

Figure 9.1: The thin cylinder $n\mathsf{cyl}(D(x,r'),hr')$ is drawn as a rectangle, the central disc $nD(x,r)$ is the bold line. On the left is G_n viewed up close. On the right, inward and outward components are in grey (outward components point up and to the left). There are three good components and three bad components, all of which are contained within $nB(x,r)$ by construction.

$\alpha \in [0,h/2]$. If Λ_j^+ is an outward component, define

$$\mathsf{slice}_j^+(\alpha) := \left\{ e \in \mathrm{E}(\Lambda_j^+) \ : e \cap [nD(x,r) + n\alpha r' v] \neq \emptyset \right\}, \tag{9.14}$$

where in the above intersection, we regard the edge e as a line-segment in $\mathbb{R}^d$. Thus $\mathsf{slice}^+(j,\alpha)$ is the set of edges in Λ_j^+ which touch a translate of $nD(x,r)$. For an inward component Λ_j^-, we likewise define

$$\mathsf{slice}_j^-(\alpha) := \left\{ e \in \ \mathrm{E}(\Lambda_j^-) \ : e \cap [nD(x,r) - n\alpha r' v] \neq \emptyset \right\}. \tag{9.15}$$

Lemma 9.2.2. Let $d \geq 3$ and $p > p_c(d)$. Let $G_n \in \mathcal{G}_n$ and $B(x,r)$ with $r \in (0,1)$ be fixed, and let $\Lambda_j^\pm$ denote the outward and inward components constructed above from G_n, $B(x,r)$ and F. Let $h > 0$. For each outward component Λ_j^+, there is $h_j^+ \in \ [0,h/2]$ so that

$$|\mathsf{slice}_j^+(h_j^+)| \leq c(d) \frac{|\Lambda_j^+|}{nhr}\,. \tag{9.16}$$

Likewise, for each inward component Λ_j^-, there is $h_j^- \in [0,h/2]$ so that

$$|\mathsf{slice}_j^-(h_j^-)| \leq c(d) \frac{|\Lambda_j^+|}{nhr}\,, \tag{9.17}$$

where in (9.16) and (9.17), $c(d) > 0$ is a universal constant depending only on the dimension.

Proof. Let Λ_j^+ be an outward component. For $k \in \{1, \dots, \lceil nh \rceil/2\}$, define $\alpha_k := k/2n$. We have

$$\bigcup_{k=1}^{\lceil nh \rceil/2} \mathsf{slice}_j^+(\alpha_k) \subset \mathrm{E}(\Lambda_j^+)\,. \tag{9.18}$$

Moreover, whenever k and k' are such that $|k - k'|_2 \geq 10d$, the edge sets $\mathsf{slice}_j^+(\alpha_k)$ and $\mathsf{slice}_j^+(\alpha_{k'})$ are disjoint. Thus,

$$\sum_{k=1}^{\lceil nh \rceil/2} |\mathsf{slice}_j^+(\alpha_k)| \leq (10d)|\mathrm{E}(\Lambda_j^+)|\,, \tag{9.19}$$

so that, for at least one $k \in \{1, \dots, \lceil nh \rceil/2\}$, we must have

$$|\mathsf{slice}_j^+(\alpha_k)| \leq c(d)\frac{|\Lambda_j^+|}{nhr}\,, \tag{9.20}$$

for some positive $c(d)$, and where we have slipped r into the denominator because $r \in (0,1)$.

By construction, any α_k satisfying (9.20) is at most $h/2$. We pick one such α_k and relabel it as h_j^+. Analogous reasoning for inward components gives $h_j^- \in [0, h/2]$ for each inward Λ_j^- so that

$$|\mathsf{slice}_j^-(h_j^-)| \leq c(d)\frac{|\Lambda_j^-|}{nhr}\,, \tag{9.21}$$

which completes the proof. ☐

Remark 9.2.3. We will continue to use the edge sets given by Lemma 9.2.2 throughout this section, but only when working with bad components. When $\Lambda_j^\pm$ is bad, we define

$$\mathsf{slice}_j^\pm := \mathsf{slice}_j^\pm(h_j^\pm)\,, \tag{9.22}$$

and if $\Lambda_j^\pm$ is good, we define $\mathsf{slice}_j^\pm$ to be empty. Figure 9.2 depicts the edge sets $\mathsf{slice}_j^\pm$. Just as we have done with the components $\Lambda_j^\pm$, we will suppress the dependence of the $\mathsf{slice}_j^\pm$ on G_n, $h > 0$, $B(x,r)$ and F.

The following is an immediate consequence of Lemma 9.2.2.

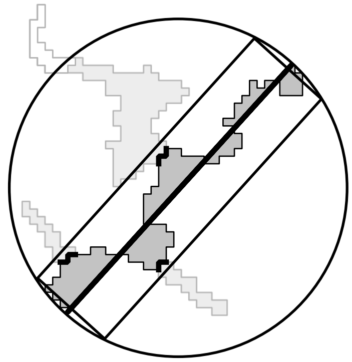

Figure 9.2: The short, bold curves are the efficiently chosen sets of open edges $\mathsf{slice}_j^\pm$ from Lemma 9.2.2. We have faded the portions of the bad components which are cut off by the $\mathsf{slice}_j^\pm$, and which stick out of the thin cylinder $n\mathsf{cyl}(D(x,r'),hr')$.

Corollary 9.2.4. Let $d \geq 3$ and $p > p_c(d)$. Let $G_n \in \mathcal{G}_n$, $h > 0$ and $B(x,r)$ be fixed. Let $\mathsf{slice}_j^\pm$ be the edge sets constructed from G_n, $h > 0$ and $B(x,r)$. There is $c(d) > 0$ so that

$$\sum_{j=1}^{\ell^+} |\mathsf{slice}_j^+| + \sum_{j=1}^{\ell^-} |\mathsf{slice}_j^-| \leq \frac{c(d)}{nhr} \left(\sum_{j=1}^{\ell^+} |\Lambda_j^+| + \sum_{j=1}^{\ell^-} |\Lambda_j^-| \right) . \tag{9.23}$$

Remark 9.2.5. Corollary 9.2.4 tells us that to control the total size of the $\mathsf{slice}_j^\pm$, it suffices to control the total volume of the $\Lambda_j^\pm$. Recall that, at the beginning of Chapter 8, we built a polytope P_n whose perimeter we could control, and whose representative measure ν_n was $\mathfrak{d}$-close to μ_n. Proposition 9.2.6 below shows that when ν_n and ν_F are $\mathfrak{d}$-close, we can control the ℓ^1 distance of the indicator functions $\mathbf{1}_{\mathbf{C}_\infty \cap nP_n}$ and $\mathbf{1}_{\mathbf{C}_\infty \cap nF}$. Proposition 9.2.7 in turn will tell us that when these indicator functions are close, we have control on the total volume of the $\Lambda_j^\pm$.

Proposition 9.2.6. Let $d \geq 3$ and $p > p_c(d)$. Let $F \subset [-1,1]^d$ be a polytope, and given $G_n \in \mathcal{G}_n$, let $P_n \subset [-1,1]^d$ be the polytope defined from G_n in (8.4) with representative measure ν_n. Let $\delta > 0$. Then there is $\epsilon(d,\delta,F) > 0$ and an event $\mathcal{E}_0$ such that for all

$G_n \in \mathcal{G}_n$,

$$\Big\{\mathfrak{d}(\nu_n, \nu_F) < \epsilon\Big\} \cap \Big\{\max_{G_n \in \mathcal{G}_n} \mathrm{per}(P_n) \leq \gamma\Big\} \cap \mathcal{E}_0 \tag{9.24}$$

$$\subset \Big\{\|\mathbf{1}_{\mathbf{C}_\infty \cap nP_n} - \mathbf{1}_{\mathbf{C}_\infty \cap nF}\|_{\ell^1} \leq \delta n^d\Big\} \cap \Big\{\max_{G_n \in \mathcal{G}_n} \mathrm{per}(P_n) \leq \gamma\Big\} \cap \mathcal{E}_0\,. \tag{9.25}$$

Moreover, there are positive constants $c_1(p, d, \delta, F), c_1(p, d, \delta, F)$ so that

$$\mathbb{P}_p(\mathcal{E}_0) \leq c_1 \exp\Big(-c_2 n^{d-1}\Big)\,. \tag{9.26}$$

Proof. Fix $G_n \in \mathcal{G}_n$ and hence also P_n and ν_n. For $k \in \mathbb{N}$, let Δ^k denote the collection of dyadic cubes in $[-1, 1]^d$ at scale k. Define:

$$\mathsf{Q}_0 := \Big\{Q \in \Delta^k\ : Q \cap (\partial F \cup \partial P_n) \neq \emptyset\Big\}\,, \tag{9.27}$$

and observe that

$$\|\mathbf{1}_{\mathbf{C}_\infty \cap nP_n} - \mathbf{1}_{\mathbf{C}_\infty \cap nF}\|_{\ell^1} \leq \sum_{Q \in \Delta^k} \|\mathbf{1}_{\mathbf{C}_\infty \cap nP_n \cap nQ} - \mathbf{1}_{\mathbf{C}_\infty \cap nF \cap nQ}\|_{\ell^1}\,, \tag{9.28}$$

$$\leq \sum_{Q \in \Delta^k \setminus \mathsf{Q}_0} \|\mathbf{1}_{\mathbf{C}_\infty \cap nP_n \cap nQ} - \mathbf{1}_{\mathbf{C}_\infty \cap nF \cap nQ}\|_{\ell^1} \tag{9.29}$$

$$+ \sum_{Q \in \mathsf{Q}_0} \|\mathbf{1}_{\mathbf{C}_\infty \cap nP_n \cap nQ} - \mathbf{1}_{\mathbf{C}_\infty \cap nF \cap nQ}\|_{\ell^1}\,. \tag{9.30}$$

Take n sufficiently large, depending on k, so that for some $c(d) > 0$, we have

$$\|\mathbf{1}_{\mathbf{C}_\infty \cap nP_n} - \mathbf{1}_{\mathbf{C}_\infty \cap nF}\|_{\ell^1} \leq \sum_{Q \in \Delta^k \setminus \mathsf{Q}_0} \|\mathbf{1}_{\mathbf{C}_\infty \cap nP_n \cap nQ} - \mathbf{1}_{\mathbf{C}_\infty \cap nF \cap nQ}\|_{\ell^1} + c(d)(\gamma + \mathrm{per}(F)) n^d 2^{-dk}\,. \tag{9.31}$$

Define the following collections of dyadic cubes:

$$\mathsf{Q}_1 := \Big\{Q \in \Delta^k\ : P_n \cap Q = Q\Big\}\,, \tag{9.32}$$

$$\mathsf{Q}_2 := \Big\{Q \in \Delta^k\ : F \cap Q = Q\Big\}\,, \tag{9.33}$$

and observe that for each $Q \in \mathsf{Q}_1$, we have $\mathbf{C}_\infty \cap nP_n \cap nQ = \mathbf{C}_\infty \cap nQ$. Likewise, for each $Q \in \mathsf{Q}_2$, we have $\mathbf{C}_\infty \cap nF \cap nQ = \mathbf{C}_\infty \cap nQ$. Using these observations, we conclude that

$$\|\mathbf{1}_{\mathbf{C}_\infty \cap nP_n} - \mathbf{1}_{\mathbf{C}_\infty \cap nF}\|_{\ell^1} \leq \sum_{Q \in \mathsf{Q}_1 \Delta \mathsf{Q}_2} \|\mathbf{1}_{\mathbf{C}_\infty \cap nP_n \cap nQ} - \mathbf{1}_{\mathbf{C}_\infty \cap nF \cap nQ}\|_{\ell^1} + c(d)(\gamma + \mathrm{per}(F))n^d 2^{-dk}\,. \tag{9.34}$$

For each dyadic cube $Q \in \Delta^k$, and for $\epsilon > 0$, introduce the event $\mathcal{E}_Q$

$$\mathcal{E}_Q := \left\{ \frac{\mathbf{C}_\infty \cap nQ}{\mathcal{L}^d(nQ)} \in (\theta_p(d)(1-\epsilon), \theta_p(d)(1+\epsilon)) \right\}\,, \tag{9.35}$$

and let $\mathcal{E}_0$ be the intersection $\bigcap_{Q \in \Delta^k} \mathcal{E}_Q$. Within the event $\mathcal{E}_0$, and by using the definition of the metric $\mathfrak{d}$, we have

$$\|\mathbf{1}_{\mathbf{C}_\infty \cap nP_n} - \mathbf{1}_{\mathbf{C}_\infty \cap nF}\|_{\ell^1} \leq n^d 2^k |\Delta^k|(1+\epsilon)\mathfrak{d}(\nu_n, \nu_F) + c(d)(\gamma + \mathrm{per}(F))n^d 2^{-dk}\,. \tag{9.36}$$

Choose k sufficiently large depending on d, δ and F so that $\delta/4 \leq c(d)(\gamma + \mathrm{per}(F))2^{-dk} \leq \delta/2$, so that within the event $\mathfrak{d}(\nu_n, \nu_F) < \epsilon$,

$$\|\mathbf{1}_{\mathbf{C}_\infty \cap nP_n} - \mathbf{1}_{\mathbf{C}_\infty \cap nF}\|_{\ell^1} \leq n^d 2^k |\Delta^k|(1+\epsilon)\mathfrak{d}(\nu_n, \nu_F) + \frac{\delta}{2} n^d\,, \tag{9.37}$$

$$\leq c(d)\delta^{-(d+1)/d}(\gamma + \mathrm{per}(F))^{(d+1)/d}(1+\epsilon)\epsilon n^d + \frac{\delta}{2} n^d\,, \tag{9.38}$$

$$\leq \delta n^d\,, \tag{9.39}$$

where to obtain the last line, we have chosen ϵ to be sufficiently small depending on d, δ and F. The proof will then be complete upon controlling the probability of $\mathcal{E}_0^c$, but this follows from a union bound in conjunction with Corollary 10.1.5 applied to each event $\mathcal{E}_Q$. □

As outlined in Remark 9.2.5, Proposition 9.2.7 below will be used in conjunction with Proposition 9.2.6 and Corollary 9.2.4 to control the total size of all $\mathsf{slice}_j^\pm$ constructed.

Proposition 9.2.7. Let $G_n \in \mathcal{G}_n$, let $F \subset [-1,1]^d$ be a polytope and let $B(x,r)$ be a ball with $r \in (0,1)$ such that

$$\mathcal{L}^d\Big((B(x,r) \cap F) \;\Delta\; B_-(x,r)\Big) \leq \delta \alpha_d r^d\,. \tag{9.40}$$

Within the event

$$\left\{\|\mathbf{1}_{\mathbf{C}_\infty\cap nP_n} - \mathbf{1}_{\mathbf{C}_\infty\cap nF}\|_{\ell^1} \leq \delta n^d \alpha_d r^d\right\} \cap \left\{\max_{G_n\in\mathcal{G}_n} |\partial^\omega G_n| \leq \eta_3 n^{d-1}\right\}, \tag{9.41}$$

there is a positive constant $c(d) > 0$ so that for n taken sufficiently large depending on d, δ, F and r, we have

$$\sum_{j=1}^{\ell^+} |\Lambda_j^+| + \sum_{j=1}^{\ell^-} |\Lambda_j^-| \leq c(d)\delta n^d \alpha_d r^d\,, \tag{9.42}$$

where the $\Lambda_j^\pm$ are the outward and inward components associated to G_n and $B(x,r)$.

Proof. Write $B(x,r)$ as B and $B_-(x,r)$ as B_- for brevity. We first handle the outward components. If Λ_j^+ is an outward component, it is contained in $G_n \cap n(B\setminus B_-)$. These components are pairwise disjoint, and hence by Remark 8.1.4,

$$\sum_{j=1}^{\ell^+} |\Lambda_j^+| \leq |G_n \cap n(B\setminus B_-)|\,, \tag{9.43}$$

$$\leq |\mathbf{C}_\infty \cap nP_n \cap n(B\setminus B_-)|\,, \tag{9.44}$$

$$\leq \|\mathbf{1}_{\mathbf{C}_\infty\cap nP_n} - \mathbf{1}_{\mathbf{C}_\infty\cap nF}\|_{\ell^1} + |\mathbf{C}_\infty \cap nF \cap n(B\setminus B_-)|\,. \tag{9.45}$$

Take n sufficiently large depending on d, F and r, to obtain

$$\sum_{j=1}^{\ell^+} |\Lambda_j^+| \leq \|\mathbf{1}_{\mathbf{C}_\infty\cap nP_n} - \mathbf{1}_{\mathbf{C}_\infty\cap nF}\|_{\ell^1} + c(d) n^d \mathcal{L}^d((B\cap F)\,\Delta\, B_-)\,. \tag{9.46}$$

Within the event (9.41) and using the bound (9.40), we find

$$\sum_{j=1}^{\ell^+} |\Lambda_j^+| \leq \delta n^d \alpha_d r^d + c(d)\delta n^d \alpha_d r^d\,. \tag{9.47}$$

As (9.47) gives the desired bound on the sum corresponding to outward components, we now turn to the inward components. These components are also pairwise disjoint, and each is contained within $\mathbf{C}_\infty \setminus G_n \cap nB_-$:

$$\sum_{j=1}^{\ell^-} |\Lambda_j^-| \leq |\mathbf{C}_\infty \setminus G_n \cap nB_-|\,, \tag{9.48}$$

$$\leq |\mathbf{C}_\infty \cap nP_n \setminus G_n| + |\mathbf{C}_\infty \setminus nP_n \cap nB_-|\,, \tag{9.49}$$

$$\leq |\mathbf{C}_\infty \cap nP_n \setminus G_n| + \|\mathbf{1}_{\mathbf{C}_\infty\cap nP_n} - \mathbf{1}_{\mathbf{C}_\infty\cap nF}\|_{\ell^1} + |\mathbf{C}_\infty \setminus nF \cap nB_-|\,. \tag{9.50}$$

where the second line above follows from Remark 8.1.4. Recall that the discrete precursor to P_n was the set $F_n \subset \mathbf{C}_\infty$, defined in (7.15). We emphasize that the polytope F introduced in the statement of this proposition is not related to this F_n (this will be the only instance the letter F is overloaded with meaning). Using the definition of F_n, in conjunction with Remark 8.1.4 which says that $\mathbf{C}_\infty \cap nP_n \subset F_n$, we find:

$$\sum_{j=1}^{\ell^-} |\Lambda_j^-| \leq |F_n \setminus G_n| + \|\mathbf{1}_{\mathbf{C}_\infty \cap nP_n} - \mathbf{1}_{\mathbf{C}_\infty \cap nF}\|_{\ell^1} + |\mathbf{C}_\infty \setminus nF \cap nB_-| \,, \tag{9.51}$$

$$\leq |F_n \setminus G_n| + \delta n^d \alpha_d r^d + c(d) n^d \mathcal{L}^d((B \cap F) \,\Delta\, B_-) \,, \tag{9.52}$$

when n is taken sufficiently large depending on r, F and d. Within the event (9.41) (using Lemma 7.4.5), we find that $|F_n \setminus G_n| \leq n^d n^{-\epsilon(d)}$, where $\epsilon(d)$ is defined in (7.14). All that matters is that $\epsilon(d) > 0$, which is the case when $d \geq 3$, and we use this in (9.52) to deduce:

$$\sum_{j=1}^{\ell^-} |\Lambda_j^-| \leq n^d n^{-\epsilon(d)} + \delta n^d \alpha_d r^d + c(d) \delta n^d \alpha_d r^d \,, \tag{9.53}$$

so that the proof is complete by taking n larger (depending on d, r) if necessary, using the above bound with (9.47). □

Remark 9.2.8. We can now control the total number of edges in the sets $\mathsf{slice}_j^\pm$ (provided we are within the correct events) by combining (9.23) of Corollary 9.2.4 and (9.42) of Proposition 9.2.7. As Figure 9.2 suggests, we may use the edge sets $\mathsf{slice}_j^\pm$ along with edges of $\partial^\omega G_n$ which lie in the thin cylinder $n\mathsf{cyl}(D(x,r'), hr')$ to form a cutset separating the faces of this cylinder. This is indeed the case, so long as we work within yet another high probability event, to be introduced in the next section.

Motivated by Remark 9.2.8, we introduce the following edge set, which depends on $G_n \in \mathcal{G}_n$, $h > 0$ and the ball $B(x,r)$.

$$E_n := \Big(\partial^\omega G_n \cap nB(x,r)\Big) \cup \left(\bigcup_{j=1}^{\ell^+} \mathsf{slice}_j^+\right) \cup \left(\bigcup_{j=1}^{\ell^-} \mathsf{slice}_j^-\right) \,. \tag{9.54}$$

Here we recall from Remark 9.2.3 that throughout this section, we have suppressed the dependence of the $\mathsf{slice}_j^\pm$ on G_n, $h > 0$ and $B(x,r)$.

9.3 Lower bounds on $|\partial^\omega G_n|$

In the previous section, we exhibited a construction which took place within a single ball, implicitly arising from Lemma 9.1.2.

Remark 9.3.1. Given a collection of balls $\{B(x_i, r_i)\}_{i=1}^m$ given by Lemma 9.1.2 for some set of finite perimeter F, and for fixed $h > 0$ and $G_n \in \mathcal{G}_n$, we may repeat the construction in the previous section within each ball $B(x_i, r_i)$. For $G_n \in \mathcal{G}_n$, $h > 0$ and each $B(x_i, r_i)$, we define the edge set $E_n^{(i)}$ as in (9.54).

For $G_n \in \mathcal{G}_n$, $h > 0$ and $B(x_i, r_i)$ arising from a collection of balls as in Lemma 9.1.2 for the parameters $\delta > 0$ and $s \in (0, 1/2)$, we define the following event:

$$\mathcal{F}(G_n, i; h) := \left\{ \left| E_n^{(i)} \right| \leq \left(1 - s + c(p,d) \frac{\delta}{h} \right) n^{d-1} \alpha_{d-1} (r_i)^{d-1} \beta_{p,d}(v_i) \right\} , \tag{9.55}$$

where $c(p, d) > 0$ is a constant we will not specify explicitly, as it arises naturally in the proof of Corollary 9.3.3. The constant $c(p, d)$ comes from the positive constants in (9.42) and (9.23) depending only on d, the ratio α_{d-1}/α_d and the extreme values of $\beta_{p,d}$ over the unit sphere.

Remark 9.3.2. The work of the preceding section allows us to bound the size of the sets $\mathsf{slice}^\pm$ within each ball, and we will use these bounds to control the size of each $E_n^{(i)}$. The $E_n^{(i)}$ will, with high probability, form a cutset separating the faces of a thin disc. Thus, we will be in a position to apply Proposition 6.1.2 to each disc, and this will show that $\mathcal{F}(G_n, i; h)$ is a rare event when δ, h and s are chosen appropriately. Corollary 9.3.3 below relates the events $\mathcal{E}(G_n, i)$ introduced at the beginning of this section to the $\mathcal{F}(G_n, i, h)$. By Lemma 9.1.5 then, knowing that each $\mathcal{F}(G_n, i; h)$ is a low-probability event will tell us it is also rare for $|\partial^\omega G_n|$ to be too small.

Corollary 9.3.3. Let $d \geq 3$ and $p > p_c(d)$. Let $G_n \in \mathcal{G}_n$, let $F \subset [-1, 1]^d$ be a polytope and let $\{B(x_i, r_i)\}_{i=1}^m$ be a collection of balls as in Lemma 9.1.2 for F and the paramters $\delta > 0, s \in (0, 1/2)$. Let $h > 0$, and let $E_n^{(i)}$ be the edge set associated to $G_n, B(x_i, r_i)$ and

$h>0$. Then, for n taken sufficiently large depending on d, ϵ_F and F,

$$\mathcal{E}(G_n,i) \cap \Big\{ \|\mathbf{1}_{\mathbf{C}_\infty \cap nP_n} - \mathbf{1}_{\mathbf{C}_\infty \cap nF}\|_{\ell^1} \leq \epsilon_F n^d \Big\} \cap \Big\{ \max_{G_n \in \mathcal{G}_n} |\partial^\omega G_n| \leq \eta_3 n^{d-1} \Big\} \subset \mathcal{F}(G_n, i; h)\,, \tag{9.56}$$

where $\mathcal{E}(G_n,i)$ and $\mathcal{F}(G_n,i;h)$ are respectively the events defined in (9.7) and (9.55), and where $\epsilon_F = \delta \min_{i=1}^m (r_i)^d \alpha_d$ was defined from the collection of balls in (9.5).

Proof. For the convenience of the reader, we recall the definition of the event $\mathcal{E}(G_n,i)$ below:

$$\mathcal{E}(G_n,i) = \Big\{ |\partial^\omega G_n \cap nB(x_i,r_i)| \leq (1-s)n^{d-1}\alpha_{d-1}(r_i)^{d-1}\beta_{p,d}(v_i) \Big\}\,. \tag{9.57}$$

Working within the event on the left-hand side of (9.56), and because each $B(x_i,r_i)$ arising from Lemma 9.1.2 satisfies

$$\mathcal{L}^d((B(x_i,r_i)\cap F)\,\Delta\, B_-(x_i,r_i)) \leq \delta\alpha_d(r_i)^d\,, \tag{9.58}$$

we may apply Proposition 9.2.7 within each ball (with n taken sufficiently large), to obtain

$$\sum_{j=1}^{\ell^+} |\Lambda_j^+(i)| + \sum_{j=1}^{\ell^-} |\Lambda_j^-(i)| \leq c(d)\delta n^d \alpha_d (r_i)^d\,, \tag{9.59}$$

where the $\Lambda_j^{\pm}(i)$ are the inward and outward components corresponding to $G_n, h>0$ and the ball $B(x_i,r_i)$. We may then immediately apply Corollary 9.2.4 to obtain:

$$\sum_{j=1}^{\ell^+} |\mathsf{slice}_j^+(i)| + \sum_{j=1}^{\ell^-} |\mathsf{slice}_j^-(i)| \leq \frac{c(d)}{nhr_i}\Big(c(d)\delta n^d \alpha_d (r_i)^d \Big)\,, \tag{9.60}$$

$$\leq c(d)\frac{\delta}{h} n^{d-1}\alpha_{d-1}(r_i)^{d-1}\,, \tag{9.61}$$

where the edge sets $\mathsf{slice}_j^{\pm}(i)$ are defined as in (9.22) for $G_n, h>0$ and each ball $B(x_i,r_i)$. It is then immediate from the definitions of $\mathcal{E}(G_n,i)$ and the edge sets $E_n^{(i)}$ that within the event on the left-hand side of (9.56), we have

$$|E_n^{(i)}| \leq (1-s)n^{d-1}\alpha_{d-1}(r_i)^{d-1}\beta_{p,d}(v_i) + c(d)\frac{\delta}{h}n^{d-1}\alpha_{d-1}(r_i)^{d-1}\,. \tag{9.62}$$

As $\beta_{p,d}$ is a continuous function on the unit sphere $\mathbb{S}^{d-1}$, the proof is complete, upon defining $\mathcal{F}(G_n,i;h)$ appropriately in (9.55). $\square$

The next lemma tells us that with high probability, each $E_n^{(i)}$ forms an open cutset.

Lemma 9.3.4. Let F be a polytope, let $h > 0$ and let $\{B(x_i, r_i)\}_{i=1}^m$ be a collection of balls as in Lemma 9.1.2 for F and the parameters $\delta > 0$ and $s \in (0, 1/2)$. For the ball $B(x_i, r_i)$, $h > 0$ and a given $G_n \in \mathcal{G}_n$, let $E_n^{(i)}$ denote the corresponding edge sets constructed above.

Let $\mathcal{E}_1$ be the event that for each $G_n \in \mathcal{G}_n$ and all $i \in \{1, \dots, m\}$, any path of open edges in $\mathsf{d\text{-}cyl}(D(x, r_i'), hr_i', n)$ joining the opposing faces $\mathsf{d\text{-}face}^{\pm}(D(x, r_i'), hr_i', n)$ must use an edge of $E_n^{(i)}$. There exist positive constants $c_1(p, d, F, \delta, s, h)$ and $c_2(p, d, F, \delta, s, h)$ so that

$$\mathbb{P}_p(\mathcal{E}_1) \geq 1 - c_1 \exp\Big(- c_2 n^{(d-1)/d} \Big) , \tag{9.63}$$

where we recall that $r_i' := (1 - h^2) r_i^2$.

Proof. Our primary tool is Theorem 10.1.6. Let us drop the indexing for the sake of clarity and work with generic objects: balls $B(x, r)$, discs $D(x, r')$ and edge sets E_n.

From the careful construction of the $\mathsf{slice}_j^{\pm}$, any open path in $\mathbf{C}_\infty$ between the opposing faces $\mathsf{d\text{-}face}^{\pm}(D(x, r'), hr', n)$ within $\mathsf{d\text{-}cyl}(D(x, r'), hr', n)$ must use an edge of E_n. As we are working within the almost sure event that there is a unique infinite cluster, the only way the faces $\mathsf{d\text{-}face}^{\pm}(D(x, r'), hr', n)$ may be joined by an open path in $\mathsf{d\text{-}cyl}(D(x, r'), hr', n)$ is if this open path lies in a finite cluster. Such a path must use at least $2r'hn$ edges, so that the cluster containing this path must have volume at least $2r'hn$. We use a union bound with Theorem 10.1.6 applied to each point in $[-n, n]^d \cap \mathbb{Z}^d$ to obtain the desired result. $\square$

Remark 9.3.5. Let $\mathcal{E}_1$ be the event from Lemma 9.3.4. For each configuration $\omega \in \mathcal{E}_1$, by completing each $E_n^{(i)}$ to a full cutset using only closed edges, we conclude that $|E_n^{(i)}| \geq \Xi_{\text{face}}(D(x_i, r_i'), hr_i', n)$ in the configuration ω.

The next proposition aggregates all of the work presented in this section so far.

Proposition 9.3.6. Let $d \geq 3$ and let $p > p_c(d)$. Let $F \subset [-1, 1]^d$ be a polytope, and for $s \in (0, 1/2)$, let $\lambda_F(s) = (1 - 2s)\mathcal{I}_{p,d}(F)$. There exist positive constants $c_1(p, d, s, F)$,

$c_2(p,d,s,F)$ and $\widetilde{\epsilon}_F(p,d,s,F)$ so that

$$\mathbb{P}_p\Big(\Big\{\exists G_n \in \mathcal{G}_n \text{ such that } |\partial^\omega G_n| \le \lambda_F(s) n^{d-1} \text{ and } \mathfrak{d}(\mu_n, \nu_F) \le \widetilde{\epsilon}_F\Big\}\Big) \le c_1 \exp\left(-c_2 n^{1/2d}\right) . \tag{9.64}$$

Proof. Let $F \subset [-1,1]^d$ be a polytope, and let $\{B(x_i, r_i)\}_{i=1}^m$ be a collection of balls as in Lemma 9.1.2 for F and the parameters $\delta > 0$ and $s \in (0, 1/2)$. The parameter δ will be fixed in terms of p, d and s later. We also introduce a height parameter $h > 0$, which will also be fixed in terms of p, d and s later in the proof.

Let $\mathcal{E}_0$ and $\mathcal{E}_1$ be the high probability events respectively from Proposition 9.2.6 and Lemma 9.3.4 for the parameter $\widetilde{\epsilon}_F$ to be determined later. Let us now define several other high probability events:

$$\mathcal{E}_2 := \Big\{ \max_{G_n \in \mathcal{G}_n} \mathfrak{d}(\mu_n, \nu_n) \le \widetilde{\epsilon}_F \Big\}, \tag{9.65}$$

$$\mathcal{E}_3 := \Big\{ \max_{G_n \in \mathcal{G}_n} |\partial^\omega G_n| \le \eta_3 n^{d-1} \Big\}, \tag{9.66}$$

$$\mathcal{E}_4 := \Big\{ \max_{G_n \in \mathcal{G}_n} \mathrm{per}(P_n) \le \gamma \Big\}, \tag{9.67}$$

where η_3 is the constant from Lemma 10.2.5, and where γ is the constant from Corollary 8.1.5. Let $\mathcal{E}^*$ be the intersection of $\mathcal{E}_0$ through $\mathcal{E}_4$. We apply Lemma 9.1.5 to conclude that

$$\Big\{\exists G_n \in \mathcal{G}_n \text{ such that } |\partial^\omega G_n| \le \lambda_F(s) n^{d-1} \text{ and } \mathfrak{d}(\mu_n, \nu_F) \le \widetilde{\epsilon}_F\Big\} \cap \mathcal{E}^* \tag{9.68}$$

$$\subset \bigcup_{G_n \in \mathcal{G}_n} \bigcup_{i=1}^m \mathcal{E}(G_n, i) \cap \Big\{\mathfrak{d}(\mu_n, \nu_F) \le \widetilde{\epsilon}_F\Big\} \cap \mathcal{E}^*, \tag{9.69}$$

$$\subset \bigcup_{G_n \in \mathcal{G}_n} \bigcup_{i=1}^m \mathcal{E}(G_n, i) \cap \Big\{\mathfrak{d}(\nu_n, \nu_F) \le 2\widetilde{\epsilon}_F\Big\} \cap \mathcal{E}^*, \tag{9.70}$$

where the last containment directly above follows from the fact that $\mathcal{E}^* \subset \mathcal{E}_2$. We now use the fact that $\mathcal{E}^*$ is contained in $\mathcal{E}_0$ and $\mathcal{E}_4$ in conjunction with Proposition 9.2.6, choosing $\widetilde{\epsilon}_F$

sufficiently small depending on ϵ_F, d and F so that

$$\Big\{\exists G_n \in \mathcal{G}_n \text{ such that } |\partial^\omega G_n| \le \lambda_F(s) n^{d-1} \text{ and } \mathfrak{d}(\mu_n, \nu_F) \le \widetilde{\epsilon}_F\Big\} \cap \mathcal{E}^* \tag{9.71}$$

$$\subset \bigcup_{G_n \in \mathcal{G}_n} \bigcup_{i=1}^m \mathcal{E}(G_n, i) \cap \Big\{\|\mathbf{1}_{\mathbf{C}_\infty \cap nP_n} - \mathbf{1}_{\mathbf{C}_\infty \cap nF}\|_{\ell^1} \le \epsilon_F n^d\Big\} \cap \mathcal{E}^*\,, \tag{9.72}$$

$$\subset \bigcup_{G_n \in \mathcal{G}_n} \bigcup_{i=1}^m \mathcal{F}(G_n, i; h) \cap \mathcal{E}^*\,, \tag{9.73}$$

where we have used Corollary 9.3.3 and taken n sufficiently large depending on d, F and ϵ_F, using that $\mathcal{E}^* \subset \mathcal{E}_3$. Finally, because $\mathcal{E}^*$ contains $\mathcal{E}_1$, we use Remark 9.3.5 to conclude

$$\Big\{\exists G_n \in \mathcal{G}_n \text{ such that } |\partial^\omega G_n| \le \lambda_F(s) n^{d-1} \text{ and } \mathfrak{d}(\mu_n, \nu_F) \le \widetilde{\epsilon}_F\Big\} \cap \mathcal{E}^* \tag{9.74}$$

$$\subset \bigcup_{G_n \in \mathcal{G}_n} \bigcup_{i=1}^m \Big\{\Xi_{\text{face}}(D(x_i, r_i'), hr_i', n) \le \Big(1 - s + c(p,d)\frac{\delta}{h}\Big) n^{d-1} \alpha_{d-1} (r_i)^{d-1} \beta_{p,d}(v_i)\Big\}\ , \tag{9.75}$$

$$\subset \bigcup_{i=1}^m \Big\{\Xi_{\text{face}}(D(x_i, r_i'), hr_i', n) \le \Big(1 - s + c(p,d)\frac{\delta}{h}\Big) n^{d-1} \alpha_{d-1} (r_i)^{d-1} \beta_{p,d}(v_i)\Big\}\,, \tag{9.76}$$

$$\subset \bigcup_{i=1}^m \Big\{\Xi_{\text{face}}(D(x_i, r_i'), hr_i', n) \le \Big(1 - s + c(p,d)\frac{\delta}{h}\Big) \frac{1}{(1-h^2)^{(d-1)/2}} n^{d-1} \alpha_{d-1} (r_i')^{d-1} \beta_{p,d}(v_i)\Big\}\,. \tag{9.77}$$

We now start fixing dependencies of various parameters. Choose h to be sufficiently small depending on p, d and $s/2$ so that the concentration estimates of Proposition 6.1.2 are applicable, when n is taken sufficiently large depending on h and ϵ_F. Next, choose δ depending on $s, c(p,d)$ and h so that

$$\Big\{\exists G_n \in \mathcal{G}_n \text{ such that } |\partial^\omega G_n| \le \lambda_F(s) n^{d-1} \text{ and } \mathfrak{d}(\mu_n, \nu_F) \le \widetilde{\epsilon}_F\Big\} \cap \mathcal{E}^* \tag{9.78}$$

$$\subset \bigcup_{i=1}^m \Big\{\Xi_{\text{face}}(D(x_i, r_i'), hr_i', n) \le \Big(1 - s/2) n^{d-1} \alpha_{d-1} (r_i')^{d-1} \beta_{p,d}(v_i)\Big\}\,, \tag{9.79}$$

and note that the number m of events in the above union now depends only on p, d, s and

F. Thus, by using Proposition 6.1.2, we find

$$\mathbb{P}_p\Big(\exists G_n \in \mathcal{G}_n \text{ such that } |\partial^\omega G_n| \leq \lambda_F(s)n^{d-1} \text{ and } \mathfrak{d}(\mu_n, \nu_F) \leq \widetilde{\epsilon}_F\Big) \tag{9.80}$$

$$\leq \sum_{i=1}^{m} c_1 \exp\Big(- c_2 n^{(d-1)/3}\Big) + \mathbb{P}_p((\mathcal{E}^*)^c)\,, \tag{9.81}$$

$$\leq c_1 \exp\Big(- c_2 n^{(d-1)/3}\Big) + \mathbb{P}_p((\mathcal{E}^*)^c)\,, \tag{9.82}$$

where c_1 and c_2 are positive constants depending on p, d, s, h, δ and F. As $\widetilde{\epsilon}_F$, δ and h all depend only on p, d, s and F, these constants have the correct dependencies, and we may also control $\mathbb{P}_p((\mathcal{E}^*)^c)$ satisfactorily. We use 9.2.6 and Lemma 9.3.4 to control $\mathbb{P}_p(\mathcal{E}_0^c)$ and $\mathbb{P}_p(\mathcal{E}_1^c)$ respectively. We further use Theorem 8.2.3, Lemma 10.2.5 and Corollary 8.1.5 to control $\mathbb{P}_p(\mathcal{E}_2^c), \mathbb{P}_p(\mathcal{E}_3^c)$ and $\mathbb{P}_p(\mathcal{E}_4^c)$ respectively, concluding finally that

$$\mathbb{P}_p\Big(\exists G_n \in \mathcal{G}_n \text{ such that } |\partial^\omega G_n| \leq \lambda_F(s)n^{d-1} \text{ and } \mathfrak{d}(\mu_n, \nu_F) \leq \widetilde{\epsilon}_F\Big) \leq \sum_{i=1}^{m} c_1 \exp\Big(- c_2 n^{(1)/2d}\Big)\,, \tag{9.83}$$

for c_1 and c_2 positive constants depending on p, d, s, F. □

Remark 9.3.7. Proposition 9.3.6 is precisely the statement we have been aiming for since the beginning of the section: we have shown that when $G_n \in \mathcal{G}_n$ is such that μ_n and ν_F are $\mathfrak{d}$-close, we have high probability lower bounds on $|\partial^\omega G_n|$. We will use Proposition 9.3.6 in conjunction with a compactness argument to prove the main results of the paper.

9.4 Proof of main results

We first give the proof of Theorem 8.3.4, from which we will deduce Theorem 2.2.2 and Theorem 2.2.1. We must introduce a quantitative version of the isoperimetric inequality for the anisotropic surface energy $\mathcal{I}_{p,d}$. Given $F \subset \mathbb{R}^d$ a set of finite perimeter, define the *asymmetry index* (or *Fraenkel asymmetry* in the Euclidean setting) of F as

$$A(F) := \inf\left\{ \frac{\mathcal{L}^d(F\Delta(x + rW_{p,d}))}{\mathcal{L}^d(F)} \ : x \in \mathbb{R}^d, \mathcal{L}^d(rW_{p,d}) = \mathcal{L}^d(F) \right\}\,. \tag{9.84}$$

For $r > 0$ chosen so that $rW_{p,d}$ and F have the same volume, also define the *isoperimetric deficit* of F as

$$D(F) := \frac{\mathcal{I}_{p,d}(F) - \mathcal{I}_{p,d}(rW_{p,d})}{\mathcal{I}_{p,d}(rW_{p,d})}\,. \tag{9.85}$$

The anisotropic isoperimetric inequality tells us that $D(F) \geq 0$ for all sets F of finite perimeter. Taylor's theorem (Theorem 3.4.1) tells us $D(F) = 0$ if and only if $A(F) = 0$. The following is a wonderful quantitative version of Taylor's theorem, which we have specialized to the surface energy $\mathcal{I}_{p,d}$. It is due to Figalli, Maggi and Pratelli [FMP10].

Theorem 9.4.1. Let $F \subset \mathbb{R}^d$ be a set of finite perimeter with $\mathcal{L}^d(F) < \infty$. There is a positive constant $c(d)$ so that

$$A(F) \leq c(d) D(F)^{1/2}\,. \tag{9.86}$$

Remark 9.4.2. As a consequence of Theorem 9.4.1, whenever $rW_{p,d}$ is a dilate of the Wulff crystal, and whenever F^r is a set of finite perimeter such that $\mathcal{L}^d(F^r) = \mathcal{L}^d(rW_{p,d})$, we have

$$\frac{\mathcal{I}_{p,d}(F^r)}{\mathcal{I}_{p,d}(rW_{p,d})} \geq 1 + c(d)(A(F^r))^2\,. \tag{9.87}$$

We will need the following result, which (essentially) generalizes Proposition 10.4.2, and which tells us that sets of finite perimeter may be realized as L^1-limits of polytopes. As usual, we specialize this result to the surface energy $\mathcal{I}_{p,d}$, and we state a version for sets of finite perimeter living in $[-1,1]^d$ as opposed to all of $\mathbb{R}^d$.

Theorem 9.4.3. Let $F \subset [-1,1]^d$ be a set of finite perimeter. There exist a sequence of polytopes $\{F_n\}_{n=1}^\infty$, each contained within $[-1,1]^d$, such that $\mathcal{L}^d(F \Delta F_n) \to 0$ as $n \to \infty$, and such that $|\mathcal{I}_{p,d}(F_n) - \mathcal{I}_{p,d}(F)| \to 0$ as $n \to \infty$.

Theorem 9.4.3 is Proposition 14.9 in [Cer06].

Proof of Theorem 8.3.4. Throughout the proof, write θ for $\theta_p(d)$. Let $\xi > 0$ and define $\eta = \eta(\xi)$ via the relation

$$(1-\eta) = \frac{1}{1+\xi}\,, \tag{9.88}$$

and use ξ and η to define the following collection of measures:

$$\mathcal{W}_\xi := \left\{ \nu_{W+x} \;:\; \begin{array}{c} x \in \mathbb{R}^d, (W+x) \subset [-1,1]^d \text{ and } W \text{ is a dilate of } W_{p,d} \\ \text{such that } \mathcal{L}^d((1-\eta)W_{p,d}) \le \mathcal{L}^d(W) \le \mathcal{L}^d((1+2\xi)W_{p,d}) \end{array} \right\}. \tag{9.89}$$

Let $\zeta > 0$, and choose $\xi = \xi(\zeta) > 0$ and $\epsilon = \epsilon(\zeta, \xi) > 0$ so that both of the following relations hold:

$$\frac{1}{1+\xi} = 1 - 2\zeta\,, \tag{9.90}$$

$$\frac{1+2\xi}{1+c(p,d)\epsilon^2} = 1 - 2\zeta\,, \tag{9.91}$$

where $c(p,d)$ shall be specified later. We remark that ϵ is not the same as the parameter $\epsilon(d)$ defined in (7.14), this latter fixed value will only show up in exponents of various upper bounds, and we will always explicate the dependence on d in this case. Our principal aim is to show that the probabilities

$$\mathbb{P}_p\Big(\exists G_n \in \mathcal{G}_n \text{ such that } \mathfrak{d}(\mu_n, \mathcal{W}_\xi) \ge \epsilon\Big) \tag{9.92}$$

decay rapidly with n.

Recall that

$$\mathcal{P}_{\gamma,\xi} = \left\{ \nu_F \;: F \subset [-1,1]^d,\ \mathrm{per}(F) \le \gamma,\ \mathcal{L}^d(F) \le \mathcal{L}^d((1+\xi)W_{p,d}) \right\}\,, \tag{9.93}$$

let $\mathcal{V}_\epsilon(\mathcal{W}_\xi)$ denote the open ϵ-neighborhood of $\mathcal{W}_\xi$ with respect to the metric $\mathfrak{d}$, and let $\mathcal{K}_{\gamma,\xi}(\epsilon)$ be the complement of this neighborhood in $\mathcal{P}_{\gamma,\xi}$. By Lemma 8.3.3, $\mathcal{K}_{\gamma,\xi}(\epsilon)$ is compact. Define the collection of measures

$$\mathcal{P}oly_{\gamma,\xi} := \left\{ \nu_F \;: F \subset [-1,1]^d \text{ is a polytope, } \mathrm{per}(F) \le 2\gamma,\ \mathcal{L}^d(F) \le \mathcal{L}^d((1+2\xi)W_{p,d}) \right\}\,, \tag{9.94}$$

so that using the definition of the metric $\mathfrak{d}$ and Theorem 9.4.3, we see that the collection of $\mathfrak{d}$-balls

$$\Big\{\mathcal{B}(\nu_F, \widetilde{\epsilon}_F/2)\Big\}_{F\in\mathcal{P}oly_{\gamma,\xi}} \tag{9.95}$$

forms an open cover of $\mathcal{K}_{\gamma,\xi}(\epsilon)$, where $\widetilde{\epsilon}_F$ is chosen as in Proposition 9.3.6 for F and $s = \zeta/2$. For a parameter $\delta' > 0$ to be used shortly, we lose no generality choosing $\widetilde{\epsilon}_F$ smaller if necessary so that

$$(1 + \widetilde{\epsilon}_F/\theta) \le (1 + \delta') \tag{9.96}$$

We also lose no generality supposing that

$$\widetilde{\epsilon}_F \le \epsilon/2 \tag{9.97}$$

holds for each F. Given $F \in \mathcal{P}oly_{\gamma,\xi}$, define

$$\lambda_F(\zeta) = (1-\zeta)\mathcal{I}_{p,d}(F)\,, \tag{9.98}$$

and use the compactness of $\mathcal{K}_{\gamma,\xi}(\epsilon)$ to extract a finite subcover from (9.95): there are polytopes $F_1, \dots, F_m$ such that

$$\Big\{\mathcal{B}(\nu_{F_j}, \widetilde{\epsilon}_{F_j}/2)\Big\}_{j=1}^{m} \tag{9.99}$$

forms an open cover of $\mathcal{K}_{\gamma,\xi}(\epsilon)$. We now begin to estimate (9.92).

$$\begin{aligned}
&\mathbb{P}_p\Big(\exists G_n \in \mathcal{G}_n \text{ such that } \mathfrak{d}(\mu_n, \mathcal{W}_\xi) \ge \epsilon\Big) && (9.100)\\
&\quad\le \mathbb{P}_p\left(\max_{G_n\in\mathcal{G}_n} \mathfrak{d}(\mu_n, \mathcal{W}_\xi) \ge \epsilon \text{ and } n\widehat{\Phi}_n \le (1+\delta')\varphi_{W_{p,d}}\right) + \mathbb{P}_p\left(n\widehat{\Phi}_n > (1+\delta')\varphi_{W_{p,d}}\right) && (9.101)\\
&\quad\le \mathbb{P}_p\left(\max_{G_n\in\mathcal{G}_n} \mathfrak{d}(\mu_n, \mathcal{W}_\xi) \ge \epsilon \text{ and } n\widehat{\Phi}_n \le (1+\delta')\varphi_{W_{p,d}}\right) + c_1\exp\left(-c_2 n^{(d-1)/3}\right) && (9.102)
\end{aligned}$$

Where we have used the bounds in the proof of Corollary 6.2.3, and where we recall that $\varphi_{W_{p,d}}$ is the continuum conductance (defined at the very end of Chapter 4) of the Wulff crystal. We now choose $\delta > 0$ so that

$$\delta \le \min_{j=1}^{m} \widetilde{\epsilon}_{F_j}/2\,, \tag{9.103}$$

and invoke Corollary 8.3.1 for δ to further deduce

$$\mathbb{P}_p\Big(\exists G_n \in \mathcal{G}_n \text{ such that } \mathfrak{d}(\mu_n, \mathcal{W}_\xi) \geq \epsilon\Big) \tag{9.104}$$

$$\leq \mathbb{P}_p\left(\max_{G_n\in\mathcal{G}_n} \mathfrak{d}(\mu_n, \mathcal{W}_\xi) \geq \epsilon \text{ and } \max_{G_n\in\mathcal{G}_n} \mathfrak{d}(\mu_n, \mathcal{P}_{\gamma,\xi}) < \delta \text{ and } n\widehat{\Phi}_n \leq (1+\delta')\varphi_{W_{p,d}}\right) \tag{9.105}$$

$$+ c_1 \exp\left(-c_2 n^{(1-\epsilon(d))/2d}\right) . \tag{9.106}$$

Using the finite open cover (9.99) and a union bound, we have

$$\mathbb{P}_p\Big(\exists G_n \in \mathcal{G}_n \text{ such that } \mathfrak{d}(\mu_n, \mathcal{W}_\xi) \geq \epsilon\Big) \tag{9.107}$$

$$\leq \sum_{i=1}^{m} \mathbb{P}_p\left(\exists G_n \in \mathcal{G}_n \text{ such that } \mathfrak{d}(\mu_n, \nu_{F_j}) \leq \epsilon_{F_j} \text{ and } n\widehat{\Phi}_n \leq (1+\delta')\varphi_{W_{p,d}}\right) \tag{9.108}$$

$$+ c_1 \exp\left(-c_2 n^{(1-\epsilon(d))/2d}\right) . \tag{9.109}$$

We focus on bounding each summand of the form $\mathbb{P}_p(\mathcal{F}_j)$ above, where the event $\mathcal{F}_j$ is defined as

$$\mathcal{F}_j := \left\{\exists G_n \in \mathcal{G}_n \text{ such that } \mathfrak{d}(\mu_n, \nu_{F_j}) \leq \epsilon_{F_j} \text{ and } n\widehat{\Phi}_n \leq (1+\delta')\varphi_{W_{p,d}}\right\} . \tag{9.110}$$

We begin by unravelling the Cheeger constant, and using (9.96).

$$\mathbb{P}_p(\mathcal{F}_j) = \mathbb{P}_p\begin{pmatrix} \exists G_n \in \mathcal{G}_n \text{ such that } \mathfrak{d}(\mu_n, \nu_{F_j}) \leq \widetilde{\epsilon}_{F_j} \\ \text{and} \\ n|\partial^\omega G_n| \leq (1+\delta')|G_n|\varphi_{W_{p,d}} \end{pmatrix} , \tag{9.111}$$

$$\leq \mathbb{P}_p\begin{pmatrix} \exists G_n \in \mathcal{G}_n \text{ such that } \mathfrak{d}(\mu_n, \nu_{F_j}) \leq \widetilde{\epsilon}_{F_j} \\ \text{and} \\ n|\partial^\omega G_n| \leq (1+\delta')n^d(\theta\mathcal{L}^d(F_j) + \widetilde{\epsilon}_{F_j})\varphi_{W_{p,d}} \end{pmatrix} , \tag{9.112}$$

$$\leq \mathbb{P}_p\begin{pmatrix} \exists G_n \in \mathcal{G}_n \text{ such that } \mathfrak{d}(\mu_n, \nu_{F_j}) \leq \widetilde{\epsilon}_{F_j} \\ \text{and} \\ |\partial^\omega G_n| \leq (1+\delta')^2 n^{d-1}\theta\mathcal{L}^d(F_j)\varphi_{W_{p,d}} \end{pmatrix} . \tag{9.113}$$

To obtain the second line directly above, we have used the definition of the metric $\mathfrak{d}$, and to obtain the third line directly above we have used (9.96). Observe that

$$\mathbb{P}_p(\mathcal{F}_j) \leq \mathbb{P}_p \begin{pmatrix} \exists G_n \in \mathcal{G}_n \text{ such that } \mathfrak{d}(\mu_n, \nu_{F_j}) \leq \widetilde{\epsilon}_{F_j} \\ \text{and} \\ |\partial^\omega G_n| \leq (1+\delta')^2 n^{d-1} \mathcal{I}_{p,d}(F_j)(\varphi_{F_j})^{-1}\varphi_{W_{p,d}} \end{pmatrix}, \tag{9.114}$$

$$\leq \mathbb{P}_p \begin{pmatrix} \exists G_n \in \mathcal{G}_n \text{ such that } \mathfrak{d}(\mu_n, \nu_{F_j}) \leq \widetilde{\epsilon}_{F_j} \\ \text{and} \\ |\partial^\omega G_n| \leq (1+\delta')^2 \frac{\mathcal{I}_{p,d}(rW_{p,d})}{\mathcal{I}_{p,d}(F_j)} r n^{d-1} \mathcal{I}_{p,d}(F_j) \end{pmatrix}, \tag{9.115}$$

where $r > 0$ is chosen so that $\mathcal{L}^d(F_j) = \mathcal{L}^d(rW_{p,d})$. We form two cases. In ***Case (i)***, $r \leq (1-\eta)$, and in ***Case (ii)***, $r \in (1-\eta, 1+2\xi]$. Focusing on the first case for now, we use Theorem 3.4.1 and the relation (9.88) between ξ and η:

$$\mathbb{P}_p(\mathcal{F}_j) \leq \mathbb{P}_p \begin{pmatrix} \exists G_n \in \mathcal{G}_n \text{ such that } \mathfrak{d}(\mu_n, \nu_{F_j}) \leq \widetilde{\epsilon}_{F_j} \\ \text{and} \\ |\partial^\omega G_n| \leq \frac{(1+\delta')^2}{1+\xi} n^{d-1} \mathcal{I}_{p,d}(F_j) \end{pmatrix} \tag{9.116}$$

As ξ was chosen as in (9.91), we may choose δ' small enough in a way depending on ξ and ζ so that

$$\mathbb{P}_p(\mathcal{F}_j) \leq \mathbb{P}_p\Big(\exists G_n \in \mathcal{G}_n \text{ such that } \mathfrak{d}(\mu_n, \nu_{F_j}) \leq \widetilde{\epsilon}_{F_j} \text{ and } |\partial^\omega G_n| \leq \lambda_{F_j}(\zeta) n^{d-1}\Big) \tag{9.117}$$

holds whenever we are in *Case (i)*.

We now maneuver into a similar position in *Case (ii)*. From (9.115), we deduce

$$\mathbb{P}_p(\mathcal{F}_j) \leq \mathbb{P}_p \begin{pmatrix} \exists G_n \in \mathcal{G}_n \text{ such that } \mathfrak{d}(\mu_n, \nu_{F_j}) \leq \widetilde{\epsilon}_{F_j} \\ \text{and} \\ |\partial^\omega G_n| \leq \frac{(1+\delta')^2}{1+c(d)(A(F_j))^2}(1+2\xi) n^{d-1} \mathcal{I}_{p,d}(F_j) \end{pmatrix}, \tag{9.118}$$

where $A(F_j)$ is the asymmetry index of F_j introduced in (9.84), and where we have used the observation in (9.87). In (9.97), we chose each $\widetilde{\epsilon}_{F_j}$ to be at most $\epsilon/2$. The finite open cover defined in (9.99) may be assumed to have no redundancies, so by construction of the compact set of measures $\mathcal{K}_{\gamma,\xi}(\epsilon)$, we have

$$\mathfrak{d}(\nu_{F_j} \mathcal{W}_\xi) \geq \epsilon/2 \tag{9.119}$$

for each F_j. Using the definition (3.7) of $\mathfrak{d}$, we have the following lower-bound on the asymmetry index of each F_j:

$$A(F_j) \geq \epsilon/4\mathcal{L}^d(F_j)\,, \tag{9.120}$$

$$\geq \epsilon/4(1-\eta)\mathcal{L}^d(W_{p,d})\,. \tag{9.121}$$

As ξ and hence η will be taken to zero, we lose no generality supposing $\eta < 1/2$. Thus, (9.121) and (9.118) together yield

$$\mathbb{P}_p(\mathcal{F}_j) \leq \mathbb{P}_p \begin{pmatrix} \exists G_n \in \mathcal{G}_n \text{ such that } \mathfrak{d}(\mu_n, \nu_{F_j}) \leq \widetilde{\epsilon}_{F_j} \\ \text{and} \\ |\partial^\omega G_n| \leq \frac{(1+\delta')^2}{1+c(p,d)\epsilon^2}(1+2\xi)n^{d-1}\mathcal{I}_{p,d}(F_j) \end{pmatrix}. \tag{9.122}$$

We now use our choice of ϵ in (9.91), and we take δ' to be sufficiently small depending on ξ and ζ so that

$$\mathbb{P}_p(\mathcal{F}_j) \leq \mathbb{P}_p\Big(\exists G_n \in \mathcal{G}_n \text{ such that } \mathfrak{d}(\mu_n, \nu_{F_j}) \leq \widetilde{\epsilon}_{F_j} \text{ and } |\partial^\omega G_n| \leq \lambda_{F_j}(\zeta)n^{d-1}\Big) \tag{9.123}$$

holds in *Case (ii)* also.

We return to (9.109) and apply the bounds (9.117) and (9.123) to each summand:

$$\mathbb{P}_p\Big(\exists G_n \in \mathcal{G}_n \text{ such that } \mathfrak{d}(\mu_n, \mathcal{W}_\xi) \geq \epsilon\Big) \tag{9.124}$$

$$\leq \sum_{i=1}^m \mathbb{P}_p\left(\exists G_n \in \mathcal{G}_n \text{ such that } \mathfrak{d}(\mu_n, \nu_{F_j}) \leq \widetilde{\epsilon}_{F_j} \text{ and } |\partial^\omega G_n| \leq \lambda_{F_j}(\zeta)n^{d-1}\right) \tag{9.125}$$

$$+ c_1 \exp\left(-c_2 n^{(1-\epsilon(d))/2d}\right)\,. \tag{9.126}$$

Thus,

$$\mathbb{P}_p\Big(\exists G_n \in \mathcal{G}_n \text{ such that } \mathfrak{d}(\mu_n, \mathcal{W}_\xi) \geq \epsilon\Big) \leq mc_1 \exp\left(-c_2 n^{1/2d}\right) + c_1 \exp\left(-c_2 n^{(1-\epsilon(d))/2d}\right) \tag{9.127}$$

Where we have used the hard-earned bounds in Proposition 9.3.6 directly above. By Borel-Cantelli,

$$\mathbb{P}_p\left(\max_{G_n \in \mathcal{G}_n} \mathfrak{d}(\mu_n, \mathcal{W}_\xi) \leq \epsilon(\zeta,\xi) \text{ for all but finitely many } n\right) = 1\,. \tag{9.128}$$

Observe that

$$\mathfrak{d}(\mathcal{W}_\xi, \mathcal{W}) \leq c(p,d) \max\left(\mathcal{L}^d(W_{p,d} \setminus (1-\eta)W_{p,d}), \mathcal{L}^d((1+2\xi)W_{p,d} \setminus W_{p,d})\right) , \tag{9.129}$$

$$\leq c(p,d,\xi) , \tag{9.130}$$

where $c(p,d,\xi)$ tends to 0 as $\xi \to 0$. From (9.91), we have that $\xi = \xi(\zeta) \to 0$ as $\zeta \to 0$ and also that $\epsilon = \epsilon(\zeta, \xi) \to 0$ as $\zeta, \xi \to 0$. Thus,

$$\mathbb{P}_p\left(\max_{G_n \in \mathcal{G}_n} \mathfrak{d}(\mu_n, \mathcal{W}) \leq c(p,d,\zeta) \text{ for all but finitely many } n \right) = 1 \tag{9.131}$$

where $c(p,d,\zeta) \to 0$ as $\zeta \to 0$. This completes the proof of Theorem 8.3.4. □

Proof of Theorem 2.2.2. Let us first show that $\mathcal{W}$ is compact with respect to the metric $\mathfrak{d}$. We appeal to the proof of Lemma 8.3.3: it suffices to show whenever $\{W_n\}_{n=1}^\infty$ is a sequence with $\nu_{W_n} \in \mathcal{W}$ such that $\mathbf{1}_{W_n}$ converges in L^1-sense to some $\mathbf{1}_F$, that F is itself a translate of $W_{p,d}$. By dominated convergence, $\mathcal{L}^d(F) = \mathcal{L}^d(W_{p,d})$. The lower semicontinuity (Lemma 10.4.1) of the surface energy $\mathcal{I}_{p,d}$ tells us that $\mathcal{I}_{p,d}(F) \leq \mathcal{I}_{p,d}(W_{p,d})$. It then follows from Theorem 3.4.1 that F must be a translate of $W_{p,d}$.

Let $\epsilon, \zeta' > 0$. For each $\nu_W \in \mathcal{W}$, choose $\widetilde{\epsilon}_W$ as in Proposition 9.3.6 for $\zeta' = 2s$. The $\mathfrak{d}$-balls $\mathcal{B}(\nu_W, \widetilde{\epsilon}_W/2)$ indexed by $\mathcal{W}$ form an open cover of $\mathcal{W}$. We may thus extract a finite collection of translates of $W_{p,d}$, enumerated as $W_1, \dots, W_m$, so that

$$\left\{ \mathcal{B}(\nu_{W_i}, \widetilde{\epsilon}_{W_i}/2) \right\}_{i=1}^m \tag{9.132}$$

covers $\mathcal{W}$. Choose $\zeta > 0$ small enough so that the term $c(p,d,\zeta)$ from (9.131) is at most $\min_{i=1}^m \widetilde{\epsilon}_{W_i}/2$, and work within the almost sure event from (9.131) that

$$\left\{ \max_{G_n \in \mathcal{G}_n} \mathfrak{d}(\mu_n, \mathcal{W}_{p,d}) \leq c(p,d,\zeta) \text{ for all but finitely many } n \right\} . \tag{9.133}$$

By (9.133), Proposition 9.3.6 and Borel-Cantelli

$$\mathbb{P}_p\left(\liminf_{n\to\infty} \min_{G_n \in \mathcal{G}_n} \frac{|\partial^\omega G_n|}{n^{d-1}} \geq (1-\zeta')\mathcal{I}_{p,d}(W_{p,d}) \right) = 1 . \tag{9.134}$$

We may take ζ smaller if necessary, in a way depending on p, d and ζ', so that from (9.133) we also have

$$\mathbb{P}_p\left(\limsup_{n\to\infty} \max_{G_n\in\mathcal{G}_n} \frac{|G_n|}{n^d} \leq (1+\zeta')\theta_p(d)\mathcal{L}^d(W_{p,d})\right) = 1\,. \tag{9.135}$$

Now choose ζ' sufficiently small depending on ϵ so that by (9.134) and (9.135), so that

$$\mathbb{P}_p\left(\liminf_{n\to\infty} n\widehat{\Phi}_n \geq (1-\epsilon)\frac{\mathcal{I}_{p,d}(W_{p,d})}{\theta_p(d)\mathcal{L}^d(W_{p,d})}\right) = 1\,. \tag{9.136}$$

The complementary upper bound on $\widehat{\Phi}_n$ was shown in Corollary 6.2.3, so the proof is complete. $\square$

Proof of Theorem 2.2.1. In (6.56), we defined the empirical measure of a translate $W \subset [-1,1]^d$ of the Wulff crystal as follows:

$$\overline{\nu}_W(n) := \frac{1}{n^d}\sum_{x\in\mathbf{C}_n\cap nW} \delta_{x/n}\,. \tag{9.137}$$

Let $\epsilon, \epsilon' > 0$. Define $M_n := n^{-1}\mathbb{Z}^d \cap [-1,1]^d$, so that $|M_n| \leq (3n)^d$. By Corollary 6.2.4, there are positive constants $c_1(p,d,\epsilon')$ and $c_2(p,d,\epsilon')$ so that

$$\mathbb{P}_p\left(\max_{x\in M_n,\,(W_{p,d}+x)\subset[-1,1]^d} \mathfrak{d}\left(\overline{\nu}_{W_{p,d}+x}(n), \nu_{W_{p,d}+x}\right) \leq \epsilon'\right) \geq 1 - c_1\exp\Big(-c_2 n^{d-1}\Big)\,, \tag{9.138}$$

and by Borel-Cantelli, the event

$$\mathcal{E}_1 := \left\{\limsup_{n\to\infty} \max_{x\in M_n,\,(W_{p,d}+x)\subset[-1,1]^d} \mathfrak{d}\left(\overline{\nu}_{W_{p,d}+x}(n),\, \nu_{W_{p,d}+x}\right) \leq \epsilon'\right\} \tag{9.139}$$

occurs almost surely.

Moreover, by choosing ζ in (9.131) to be sufficiently small depending on ϵ', the following event

$$\mathcal{E}_2 := \left\{\max_{G_n\in\mathcal{G}_n} \mathfrak{d}(\mu_n, \mathcal{W}) \leq \epsilon' \text{ for all but finitely many } n\right\} \tag{9.140}$$

also occurs almost surely. We may take n sufficiently large depending on ϵ' so that for any translate $W \subset [-1,1]^d$ of the Wulff crystal, there is $x \in M_n$ so that $\mathfrak{d}(\nu_W, \nu_{W_{p,d}+x}) \leq \epsilon'$.

Thus, for any percolation configuration $\omega \in \mathcal{E}_1 \cap \mathcal{E}_2$, there is $N = N(\omega) \in \mathbb{N}$ so that whenever $n \geq N(\omega)$,

$$\max_{G_n \in \mathcal{G}_n} \min_{x \in M_n} \mathfrak{d}(\mu_n, \overline{\nu}_{W_{p,d}+x}(n)) \leq 3\epsilon' . \tag{9.141}$$

Suppose that

$$\mathfrak{d}(\mu_n, \overline{\nu}_{W_{p,d}+x}(n)) \leq 3\epsilon' \tag{9.142}$$

for some $G_n \in \mathcal{G}_n$ and $x \in M_n$. Let $k \in \mathbb{N}$, and let Δ^k denote the set of dyadic cubes at scale k which are contained in $[-1,1]^d$. Define

$$\mathsf{Q} := \left\{ Q \in \Delta^k \ : Q \cap \partial(W_{p,d} + x) \neq \emptyset \right\} \tag{9.143}$$

Observe that

$$\left\| \mathbf{1}_{G_n} - \mathbf{1}_{n(W_{p,d}+x)\cap \mathbf{C}_n} \right\|_{\ell^1} \leq \sum_{Q \in \Delta^k} \left\| \mathbf{1}_{G_n \cap nQ} - \mathbf{1}_{n(W_{p,d}+x)\cap \mathbf{C}_n \cap nQ} \right\|_{\ell^1} , \tag{9.144}$$

$$\leq \sum_{Q \in \Delta^k \setminus \mathsf{Q}} \left\| \mathbf{1}_{G_n \cap nQ} - \mathbf{1}_{n(W_{p,d}+x)\cap \mathbf{C}_n \cap nQ} \right\|_{\ell^1} + \sum_{Q \in \mathsf{Q}} \left\| \mathbf{1}_{G_n \cap nQ} - \mathbf{1}_{n(W_{p,d}+x)\cap \mathbf{C}_n \cap nQ} \right\|_{\ell^1} , \tag{9.145}$$

$$\leq \sum_{Q \in \Delta^k \setminus \mathsf{Q}} \left\| \mathbf{1}_{G_n \cap nQ} - \mathbf{1}_{n(W_{p,d}+x)\cap \mathbf{C}_n \cap nQ} \right\|_{\ell^1} + c(p,d) 2^{-dk} n^d , \tag{9.146}$$

where $c(p,d) > 0$ accounts for the perimeter of $W_{p,d}$. For each $Q \in \Delta^k \setminus \mathsf{Q}$, we then have that $n(W_{p,d} + x) \cap \mathbf{C}_n \cap nQ = \mathbf{C}_n \cap nQ$ or $n(W_{p,d} + x) \cap \mathbf{C}_n \cap nQ = \emptyset$. It follows from the definition (3.7) of $\mathfrak{d}$ that

$$\left\| \mathbf{1}_{G_n} - \mathbf{1}_{n(W_{p,d}+x)\cap \mathbf{C}_n} \right\|_{\ell^1} \leq 2^k |\Delta^k| n^d \mathfrak{d}(\mu_n, \overline{\nu}_{W_{p,d}+x}(n)) + c(p,d) 2^{-dk} n^d , \tag{9.147}$$

and we may choose k small enough depending on ϵ, and then ϵ' small enough depending on ϵ and k (using (9.142)) to conclude that

$$n^{-d} \left\| \mathbf{1}_{G_n} - \mathbf{1}_{n(W_{p,d}+x)\cap \mathbf{C}_n} \right\|_{\ell^1} \leq \epsilon . \tag{9.148}$$

Note that the choice of k and ϵ' do not depend on $G_n \in \mathcal{G}_n$, on $x \in M_n$ or on the percolation configuration $\omega \in \mathcal{E}_1 \cap \mathcal{E}_2$. Thus, for ϵ' chosen in this way according to k and ϵ, it follows that for any $\omega \in \mathcal{E}_1 \cap \mathcal{E}_2$, and whenever $n \geq N(\omega)$, we have

$$\max_{G_n \in \mathcal{G}_n} \min_{x \in M_n} \left(n^{-d} \big\| \mathbf{1}_{G_n} - \mathbf{1}_{n(W_{p,d}+x) \cap \mathbf{C}_n} \big\|_{\ell^1} \right) \leq \epsilon \,, \tag{9.149}$$

and we thus conclude that for any $\epsilon > 0$,

$$\mathbb{P}_p \left(\limsup_{n \to \infty} \max_{G_n \in \mathcal{G}_n} \inf_{x \in \mathbb{R}^d} n^{-d} \big\| \mathbf{1}_{G_n} - \mathbf{1}_{\mathbf{C}_n \cap (x + nW_{p,d})} \big\|_{\ell^1} \leq \epsilon \right) = 1 \,, \tag{9.150}$$

which completes the proof. □

CHAPTER 10

Appendix 1: Tools from percolation, graph theory and geometry

10.1 Tools from percolation

We first present fundamental tools from percolation used both implicitly and explicitly throughout the paper. We use the notation $\Lambda(n) := [-n, n]^d \cap \mathbb{Z}^d$. The first tool was introduced by Benjamini and Mossel in [BM03], although the proof contained a gap. Mathieu and Remy filled this gap in [MR04], and alternate proofs were also given by Berger, Biskup, Hoffman and Kozma [BBH08] and by Pete [Pet08]. The version we present is Proposition A.2 of [BBH08].

Proposition 10.1.1. Let $d \geq 2$ and $p > p_c(d)$. There exist positive constants $c_1(p,d), c_2(p,d)$ and $c_3(p,d)$ so that for all $t > 0$,

$$\mathbb{P}_p(\exists \Lambda \ni 0, \omega\text{-connected}, |\Lambda| \geq t^{d/(d-1)}, |\partial^\omega \Lambda| < c_3 |\Lambda|^{(d-1)/d}) \leq c_1 \exp(-c_2 t)\,. \tag{10.1}$$

We deduce a corollary similar to Proposition A.1 in [BBH08].

Corollary 10.1.2. Let $d \geq 2$ and $p > p_c(d)$. There are positive constants $c_1(p,d)$, $c_2(p,d)$, $c_3(p,d)$ and an almost surely finite random variable $R = R(\omega)$ such that whenever $n \geq R$, we the following lower bound on $\partial^\omega \Lambda$ for each ω-connected Λ satisfying $\Lambda \subset \mathbf{C}_n$ and $|\Lambda| \geq n^{1/2}$:

$$|\partial^\omega \Lambda| \geq c_3 |\Lambda|^{(d-1)/d}\,. \tag{10.2}$$

Moreover, we have the following tail bounds on R:

$$\mathbb{P}_p(R > n) \leq c_1 n^d \exp\left(-c_2 n^{(d-1)/2d}\right) . \tag{10.3}$$

Proof. Let c_3 be as in Proposition 10.1.1, and let $\mathcal{E}_n$ denote the following event:

$$\left\{\exists \Lambda, \omega\text{-connected with } \Lambda \subset \mathbf{C}_n,\ |\Lambda| \geq n^{1/2} \text{ but } |\partial^\omega \Lambda| < c_3 |\Lambda|^{(d-1)/d}\right\}. \tag{10.4}$$

We apply Proposition 10.1.1 to every point within the box $\Lambda(n)$ with $t = n^{1/2}$ to obtain

$$\mathbb{P}_p(\mathcal{E}_n) \leq c_1 n^d \exp\left(-c_2 n^{(d-1)/2d}\right) . \tag{10.5}$$

These probabilities are summable in n. We let R be the (random) smallest natural number such that that $n \geq R$ implies $\mathcal{E}_n^c$ occurs. As $\{R > n\} \subset \mathcal{E}_n$, the proof is complete. $\square$

Proposition 10.1.1 and its corollary give us control on the open edge boundary of large subgraphs of $\mathbf{C}_n$. We now introduce a tool for controlling the density of the infinite cluster within a large box. This result was proved in two dimensions by Durrett and Schonmann [DS88] and in higher dimensions in the thesis of Gandolfi [Gan89].

Proposition 10.1.3. Let $d \geq 2$ and $p > p_c(d)$. Recall that $\theta_p(d) = \mathbb{P}_p(0 \in \mathbf{C}_\infty)$ is the density of the infinite cluster. For any $\epsilon > 0$, there exist positive constants $c_1(p, d, \epsilon)$ and $c_2(p, d, \epsilon)$ so that

$$\mathbb{P}_p\left(\frac{|\mathbf{C}_n|}{|\Lambda(n)|} \notin (\theta_p(d) - \epsilon, \theta_p(d) + \epsilon)\right) \leq c_1 \exp\Big(- c_2 n^{d-1}\Big) . \tag{10.6}$$

Remark 10.1.4. It is interesting to note that the upper deviations actually decay exponentially with n^d, whereas the lower deviations are of surface order, hence the surface term in the above proposition. Pisztora [Pis96] later refined these results, showing that with high probability $\mathbf{C}_n$ consists of a unique giant connected component whose volume is roughly $\theta_p |\Lambda(n)|$ with all other connected components of negligible size. Pisztora worked in the more general setting of the FK percolation and the Ising model, and even showed this giant component has important geometric properties, spanning each of the opposing faces of $\Lambda(n)$.

These properties are useful for renormalization arguments, but we will not need to be so delicate for our applications.

The following is an immediate corollary of Proposition 10.1.3.

Corollary 10.1.5. Let $d \geq 2$ and $p > p_c(d)$. Let $r > 0$, let $Q \subset \mathbb{R}^d$ be a translate of the cube $[-r, r]^d$ and let $\epsilon > 0$. There exist positive constants $c_1(p, d, \epsilon), c_2(p, d, \epsilon)$ so that

$$\mathbb{P}\left(\frac{|\mathbf{C}_\infty \cap Q|}{\mathcal{L}^d(Q)} \notin (\theta_p(d) - \epsilon, \theta_p(d) + \epsilon)\right) \leq c_1 \exp\left(-c_2 r^{d-1}\right) . \tag{10.7}$$

The last percolation input we need is the following fundamental result of Kesten, Zhang [KZ90], Grimmett and Marstrand [GM90], which we use in Chapter 9.

Theorem 10.1.6. *Let $d \geq 2$ and $p > p_c(d)$, and let $C(0)$ denote the open cluster containing the origin. There is a positive constant $c(p)$ so that*

$$\mathbb{P}_p(|\boldsymbol{C}(0)| = n) \leq \exp\left(-cn^{(d-1)/d}\right) . \tag{10.8}$$

10.2 Using tools from percolation

We now specialize the tools of Appendix A.1 to the Cheeger optimizers $G_n \in \mathcal{G}_n$. We begin by making a basic observation.

Lemma 10.2.1. For all n and each configuration ω, if $G_n \in \mathcal{G}_n$ and is disconnected, then G_n is a finite union of connected optimal subgraphs.

Proof. The proof follows from the identity that for $a, b, c, d > 0$, we have

$$\frac{a+b}{c+d} \geq \min\left(\frac{a}{c}, \frac{b}{d}\right) . \tag{10.9}$$

By the discussion at the end of Section 2.2, it must be that the connected components of any $G_n \in \mathcal{G}_n$ have disjoint open edge boundaries. Thus if G_n is optimal and disconnected, and if we decompose G into two disjoint subgraphs G_n' and G_n'', we must have $\varphi_{G_n} = \varphi_{G_n'} = \varphi_{G_n''}$. □

We now use Corollary 10.1.5 to obtain a high probability upper bound on $\widehat{\Phi}_n$.

Lemma 10.2.2. Let $d \geq 2$ and $p > p_c(d)$. There are positive constants $c_1(p,d), c_2(p,d)$ and $c_3'(p,d)$ so that

$$\mathbb{P}_p(\widehat{\Phi}_n > c_3' n^{-1}) \leq c_1 \exp\left(-c_2 n^{d-1}\right) . \tag{10.10}$$

Proof. We abbreviate $\theta_p(d)$ as θ. Let us work within the high probability event from Corollary 10.1.5 for the box $[-r,r]^d$, where $r = n/2(d!)^{1/d}$, and for the parameter $\epsilon > 0$. Let us also work within the corresponding high probability event for the box $[-n,n]^d$ with the same parameter ϵ. Let $\epsilon' = \epsilon/\theta$ and let H_n be $\mathbf{C}_\infty \cap [-r,r]^d$, so that

$$(1-\epsilon')\theta(2r)^d \leq |C| \leq (1+\epsilon')\theta(2r)^d , \tag{10.11}$$

$$\leq (1-\epsilon')\theta(2n)^d . \tag{10.12}$$

For ϵ chosen sufficiently small depending on d. Thus, H_n is valid with volume on the order of n^d. The open edge boundary of H_n is at most a constant depending on d times the $\mathcal{H}^{d-1}$-measure of $\partial[-r,r]^d$, which completes the proof. □

The above result may be used in conjunction with Corollary 10.1.2 to obtain a high probability lower bound on the volume of any Cheeger optimizer.

Lemma 10.2.3. Let $d \geq 2$ and $p > p_c(d)$. There exist positive constants $c_1(p,d), c_2(p,d)$ and $\eta_1(p,d)$ so that,

$$\mathbb{P}_p(\exists G_n \in \mathcal{G}_n \text{ such that } |G_n| < \eta_1 n^d) \leq c_1 \exp\left(-c_2 n^{(d-1)/2d}\right) . \tag{10.13}$$

Proof. Work within the intersection of the high probability events

$$\{\widehat{\Phi}_n \leq c_3' n^{-1}\} \cap \{R \leq n\} \tag{10.14}$$

respectively from Lemma 10.2.2 and Corollary 10.1.2. Consider $G_n \in \mathcal{G}_n$, and using Lemma 10.2.1, extract from G_n a connected subgraph $H_n \subset G_n$ with $H_n \in \mathcal{G}_n$. Suppose that $|H_n| \leq n^{1/2}$. That $H_n \subset \mathbf{C}_n$ implies $\partial^\omega H_n$ is non-empty, and hence that $\varphi_{H_n} > n^{-1/2}$. This is impossible when $\widehat{\Phi}_n \leq c_3' n^{-1}$ and n is sufficiently large.

Thus we may suppose that $H_n \geq n^{1/2}$, and using our event from Corollary 10.1.2, we have

$$|\partial^\omega H_n| \geq c_3 |H_n|^{(d-1)/d}, \tag{10.15}$$

and thus,

$$c_3 |H_n|^{-1/d} \leq \varphi_{H_n} \leq c_3' n^{-1}, \tag{10.16}$$

and the claim holds with $\eta_1 = (c_3/c_3')^d$. □

Combining Lemma 10.2.3 with Lemma 10.2.1 and the observation that $|\Lambda(n)| \leq c(d)n^d$, we immediately deduce the following.

Corollary 10.2.4. Let $d \geq 2$ and $p > p_c(d)$. There exist positive constants $c_1(p,d), c_2(p,d)$ and $\eta_4(p,d)$ so that

$$\mathbb{P}_p \begin{pmatrix} \exists G_n \in \mathcal{G}_n \text{ such that the number of} \\ \text{connected components of } G_n \text{ exceeds } \eta_4 \end{pmatrix} \leq c_1 \exp\left(-c_2 n^{(d-1)/2d}\right). \tag{10.17}$$

Having established that all Cheeger optimizers are, with high probability, volume order subgraph of $\Lambda(n)$, we now exhibit control from above and below on the open edge boundary of each Cheeger optimizer.

Lemma 10.2.5. Let $d \geq 2$ and $p > p_c(d)$. There exist positive constants $c_1(p,d), c_2(p,d)$ and $\eta_2(p,d), \eta_3(p,d)$ so that

$$\mathbb{P}_p \left(\exists G_n \in \mathcal{G}_n \text{ so that } |\partial^\omega G_n| < \eta_2 n^{d-1} \text{ or } |\partial^\omega G_n| > \eta_3 n^{d-1}\right) \leq c_1 \exp\left(-c_2 n^{(d-1)/2d}\right). \tag{10.18}$$

Proof. First work within the high probability event $\{\widehat{\Phi}_n \leq c_3' n^{-1}\}$ from Lemma 10.2.2 and consider some $G_n \in \mathcal{G}_n$. That $\varphi_{G_n} \leq c_3' n^{-1}$ and $|G_n| \leq c(d) n^d$ together imply

$$|\partial^\omega G_n| \leq c_3' c(d) n^{d-1}\,, \tag{10.19}$$

so we set $\eta_3 = c' c(d)$. To prove the second half of this Lemma, we work on the intersection of the high probability events

$$\Big\{R \leq n\Big\} \cap \Big\{\forall G_n \in \mathcal{G}_n,\ \text{we have } |G_n| \geq \eta_1 n^d\Big\} \tag{10.20}$$

from Corollary 10.1.2 and Lemma 10.2.3. Given $G_n \in \mathcal{G}_n$, we extract (through Lemma 10.2.1) to $H_n \subset G_n$ with $H_n \in \mathcal{G}_n$ and H_n connected. On the event $\{R \leq n\}$, we have

$$|\partial^\omega G_n| \geq |\partial^\omega H_n| \geq c_3 \left(\eta_1 n^d\right)^{(d-1)/d}\,, \tag{10.21}$$

and we set $\eta_2 = c_3 (\eta_1)^{(d-1)/d}$. □

10.3 Tools from graph theory

First, we state the version of Turán's theorem given directly before Lemma 6 in [SZ03]. This theorem is used in the Peierls argument of Chapter **??** (the proof of Proposition 7.4.2). Recall that for a graph (V, E), an *independent* set of vertices $A \subset \mathrm{V}$ is a collection of vertices such that no two elements of A are joined by an edge in E. For a finite graph (V, E), the *independence number* of (V, E) is

$$\alpha(\mathrm{V}, \mathrm{E}) := \max\Big\{|A|\ : A \text{ is an independent subset of V}\Big\}\,. \tag{10.22}$$

Theorem 10.3.1. Let (V, E) be a finite graph with maximal degree δ. Then

$$\alpha(\mathrm{V}, \mathrm{E}) \geq \frac{|\mathrm{V}|}{\delta + 1}\,. \tag{10.23}$$

Our next combinatorial proposition gives an exponential bound on the number of $\mathbb{L}^d$-connected subsets of $\mathbb{Z}^d$ containing the origin.

Proposition 10.3.2. There is a positive constant $c(d)$ so that the number of $\mathbb{L}^d$-connected subsets of $\mathbb{Z}^d$ containing the origin of size s is at most $[c(d)]^s$.

Proof. This is a standard estimate. See the proof of Theorem 4.20, and equation (4.24) in [Gri99]. □

10.4 Approximation and miscellany

This section contains useful information about the surface energy $\mathcal{I}_{p,d}$, as well as a proof that a nice chosen orientation (from Chapter 4) exists. We now state a fundamental property of the surface energy $\mathcal{I}_\tau$ associated to any norm τ on $\mathbb{R}^d$.

Lemma 10.4.1. Let τ be a norm on $\mathbb{R}^d$ and let $\mathcal{I}_\tau$ be the surface energy associated to τ. The surface energy $\mathcal{I}_\tau$ is lower semicontinuous, that is, if E_n is a sequence of Borel sets in $\mathbb{R}^d$ tending to E with respect to the metric d_{L^1}, we have

$$\mathcal{I}_\tau(E) \leq \liminf_{n\to\infty} \mathcal{I}_\tau(E_n)$$

The proof of Lemma 10.4.1 is immediate from the definition of the surface energy, see Section 14.2 of [Cer06]. We may use this lower semicontinuity to approximate the Wulff crystal (in both volume and surface energy) by polytopes.

Proposition 10.4.2. Consider the Wulff crystal $W_{p,d} \subset [-1,1]^d$. The Wulff crystal is a set of finite perimeter, and for all $\epsilon > 0$, there is a polytope $P_\epsilon \subset W_{p,d}$ so that

(i) $|\mathcal{I}_{p,d}(P_\epsilon) - \mathcal{I}_{p,d}(W_{p,d})| \leq \epsilon$

(ii) $\mathcal{L}^d(W_{p,d} \setminus P_\epsilon) \leq \epsilon$

Proof. The Wulff crystal is a set of finite perimeter because it is convex and bounded (see Proposition 4.3.4 and Lemma 4.3.6). For $k \in \mathbb{N}$, consider a collection of points $\{x_1^{(k)}, \dots, x_m^{(k)}\}$ (with m depending on k) in $\partial W_{p,d}$ such that for each $x \in \partial W_{p,d}$, there is some $x_i^{(k)}$ with $|x_i^{(k)} - x|_2 < 2^{-k}$.

For each $k \in \mathbb{N}$, let P_k denote the convex hull of $x_1^{(k)}, \dots, x_m^{(k)}$, and note that when k is sufficiently large, the P_k are non-degenerate polytopes contained in $W_{p,d}$, so that $\mathrm{per}(P_k) \leq \mathrm{per}(W_{p,d})$ for all k. By construction, $\mathcal{L}^d(W_{p,d} \setminus P_k)$ tends to zero as $k \to \infty$, and we use Lemma 10.4.1 to conclude

$$\lim_{k\to\infty} \mathrm{per}(P_k) \to \mathrm{per}(W_{p,d})\,. \tag{10.24}$$

As $\beta_{p,d}$ is a norm on $\mathbb{R}^d$, there are positive constants $c(\beta_{p,d}, d)$ and $C(\beta_{p,d}, d)$ so that for each $x \in \mathbb{R}^d$, we have $c|x|_2 \leq \beta_{p,d}(x) \leq C|x|_2$. This implies the following inequalities hold for all $k \in \mathbb{N}$:

$$c\mathrm{per}(P_k) \leq \mathcal{I}_{p,d}(P_k) \leq C\mathrm{per}(P_k) \quad \text{and} \quad c\mathrm{per}(W_{p,d}) \leq \mathcal{I}_{p,d}(W_{p,d}) \leq C\mathrm{per}(W_{p,d})\,. \tag{10.25}$$

which completes the proof upon using (10.24). □

Remark 10.4.3. Note that a much more general theorem of this flavor holds for all sets of finite perimeter (see Proposition 14.9 of [Cer06]). We state this result as Theorem 9.4.3 shortly before using it in the proof of Theorem 8.3.4. We have included the above approximation result because its proof is so short, and because we do not need the full power of Theorem 9.4.3 for the polyhedral approximation of the Wulff crystal used at the end of Chapter 6.

The last subject we deal with in the appendix is the so-called chosen orientation from Chapter 4, specifically the Lipschitz properties mentioned in Section 4.3. Let $\mathbb{S}_+^{d-1}$ denote the closed, upper hemisphere of the unit $(d-1)$-sphere $\mathbb{S}^{d-1}$. We denote by $T\mathbb{S}^{d-1}$ the tangent bundle of $\mathbb{S}^{d-1}$, and the restriction of this bundle to $\mathbb{S}_+^{d-1}$ shall be written as $T\mathbb{S}_+^{d-1}$. An *orthonormal k-frame* on $\mathbb{S}_+^{d-1}$ is an assignment to each point $x \in \mathbb{S}_+^{d-1}$ an ordered collection of k orthonormal vectors in $T_x\mathbb{S}_+^{d-1}$, the tangent plane to $\mathbb{S}_+^{d-1}$ at x. Each such assignment may be expressed in Euclidean coordinates thanks to the natural embedding of $\mathbb{S}^{d-1} \subset \mathbb{R}^d$ and each $T_x\mathbb{S}^{d-1}$ into $\mathbb{R}^d$, and may thus be written as a function

$$f : x \mapsto (v_1(x), \dots, v_k(x))\,, \tag{10.26}$$

with $x \in \mathbb{S}_+^{d-1} \subset \mathbb{R}^d$, and with each $v_i(x) \in \mathbb{R}^d$. The collection of all such functions f which vary smoothly has the structure of a smooth manifold, denoted $V_k(\mathbb{S}_+^{d-1})$, and called the *Stiefel manifold* for the pair $(\mathbb{S}_+^{d-1}, k)$.

Proposition 10.4.4. There exists an orthonormal $(d-1)$-frame $f \in V_{d-1}(\mathbb{S}_+^{d-1})$ and a constant $C > 0$ so that for $\epsilon > 0$, whenever $x, y \in \mathbb{S}_+^{d-1}$ satisfy $|x - y|_2 \leq \epsilon$, we have

$$|(f(x))_i - (f(y))_i|_2 \leq C\epsilon \tag{10.27}$$

for all $i \in \{1, \ldots, d-1\}$.

Proof. Let $s \in \mathbb{S}^{d-1}$ denote the south pole, with coordinate representation $(0, \ldots, 0, -1)$ in $\mathbb{R}^d$. Consider the standard stereographic projection map $P : \mathbb{S}^{d-1} \setminus \{s\} \to \mathbb{R}^{d-1}$, and note that the image of $\mathbb{S}_+^{d-1}$ under this map is a closed disc $D \subset \mathbb{R}^{d-1}$ centered at the origin. This disc D is parallelizable, that is, it is possible to construct a smooth $(d-1)$-frame g on D. Indeed, one can just take the standard basis for $\mathbb{R}^{d-1}$at each tangent space $T_y D$. Define $f \in V_{d-1}(\mathbb{S}_+^{d-1})$ as the pullback $P^* g$. As f varies smoothly over a compact domain, it follows that each of its coordinate functions is Lipschitz, which completes the proof. □

CHAPTER 11

Intrinsic isoperimetry of the giant component of supercritical bond percolation in dimension two

Isoperimetric problems, while among the oldest in mathematics, are fundamental to modern probability and PDE theory. The goal of an isoperimetric problem is to characterize sets of minimal boundary measure subject to an upper bound on the volume measure of the set. The Cheeger constant, first introduced in Cheeger's thesis [Che70] in the context of manifolds, is a way of encoding such problems. Alon [Alo86] later introduce the Cheeger constant for (finite) graphs G as the following minimum over subgraphs of G:

$$\Phi_G := \min \left\{ \frac{|\partial H|}{|H|} \;:\; H \subset G,\, 0 < |H| \leq |G|/2 \right\}, \tag{11.1}$$

Here ∂H is the edge boundary of H in G (the edges of G having exactly one endpoint vertex in H), $|\partial H|$ denotes the cardinality of this set, and $|H|$ denotes the cardinality of the vertex set of H. The Cheeger constant of a graph G measures the robustness of G; it provides information about the behavior of random walks on G and is involved in a fundamental estimate in spectral graph theory (see Chapter 2 of [Chu97]). This paper is concerned with the isoperimetric properties of random graphs arising from bond percolation in $\mathbb{Z}^2$.

Bond percolation is defined as follows: We view $\mathbb{Z}^2$ as a graph with standard nearest-neighbor graph structure and form the probability space $(\{0,1\}^{\mathrm{E}(\mathbb{Z}^2)}, \mathscr{F}, \mathbb{P}_p)$ for the *percolation parameter* $p \in [0,1]$. Here $\mathscr{F}$ denotes the product σ-algebra on $\{0,1\}^{\mathrm{E}(\mathbb{Z}^2)}$ and $\mathbb{P}_p$ is the product Bernoulli measure associated to p. Elements of this probability space are written as $\omega = (\omega_e)_{e \in \mathrm{E}(\mathbb{Z}^2)}$ and are referred to as *percolation configurations*. An edge e is *open* in the configuration ω if $\omega_e = 1$, and is *closed* otherwise. For each configuration ω, the collection of edges which are open in ω determines a subgraph of $\mathbb{Z}^2$, written as $[\mathbb{Z}^2]^\omega$. Under the probability measure $\mathbb{P}_p$, $[\mathbb{Z}^2]^\omega$ is then a random subgraph of $\mathbb{Z}^2$.

Connected components of $[\mathbb{Z}^2]^\omega$ are called *open clusters*, or just *clusters*. It is well known (Grimmett [Gri99] is a standard reference) that bond percolation on $\mathbb{Z}^2$ exhibits a phase transition: there is $p_c(2) \in (0,1)$ so that $p > p_c(2)$ implies there is a unique infinite open cluster $\mathbb{P}_p$-almost surely, and such that $p < p_c(2)$ implies there is no infinite open cluster $\mathbb{P}_p$-almost surely. Moreover, it is well known [Kes80] that $p_c(2) = 1/2$. We focus our attention on the supercritical ($p > p_c(2)$) regime, and let $\mathbf{C}_\infty = \mathbf{C}_\infty(\omega)$ denote the unique infinite cluster which exists $\mathbb{P}_p$-almost surely in this case. For $p > p_c(2)$, the quantity $\theta_p := \mathbb{P}_p(0 \in \mathbf{C}_\infty)$ is positive, and is referred to as the *density* of $\mathbf{C}_\infty$ within $\mathbb{Z}^2$.

11.1 A conjecture

It is possible to study the geometry of $\mathbf{C}_\infty$ using the Cheeger constant: define $\widetilde{\mathbf{C}}_n := \mathbf{C}_\infty \cap [-n,n]^2$, and define the *giant component* $\mathbf{C}_n$ to be the largest connected component of $\widetilde{\mathbf{C}}_n$. The random variable $\widehat{\Phi}_n := \Phi_{\mathbf{C}_n}$ is central to this paper. It is known (Benjamini and Mossel [BM03], Mathieu and Remy [MR04], Rau [Rau07], Berger, Biskup, Hoffman and Kozma [BBH08] and Pete [Pet08]) that $\widehat{\Phi}_n \asymp n^{-1}$ as $n \to \infty$, prompting the following conjecture of Benjamini, which we state in all dimensions $d \geq 2$.

Conjecture 11.1.1. *(Benjamini)* Let $d \geq 2$ and $p > p_c(d)$. The limit

$$\lim_{n\to\infty} n\Phi_{\mathbf{C}_n} \tag{11.2}$$

exists $\mathbb{P}_p$-almost surely as a deterministic constant in $(0,\infty)$.

Procaccia and Rosenthal [PR12] showed for $d \geq 2$ that $\mathrm{Var}(n\widehat{\Phi}_n) \leq cn^{2-d}$ for some positive $c(p,d)$. Biskup, Louidor, Procaccia and Rosenthal [BLP15] settled Conjecture 11.1.1 in $d = 2$ for a natural modification $\widetilde{\Phi}_n$ of the Cheeger constant. The results of [BLP15] go beyond resolving Conjecture 11.1.1 for $\widetilde{\Phi}_n$: the random variables $\widetilde{\Phi}_n$ encode a sequence of discrete, random isoperimetric problems, whose set of optimizers are the subgraphs of $\widetilde{\mathbf{C}}_n$ realizing the minimum defining $\widetilde{\Phi}_n$. The main result of [BLP15] is that these optimizers, upon rescaling, almost surely tend (with respect to Hausdorff distance) to a translate of a deterministic shape, a convex subset of $[-1,1]^2$ whose two-dimensional Lebesgue measure is half that of $[-1,1]^2$. This limit shape, known as the *Wulff shape* and denoted W_p, is the solution to a deterministic isoperimetric problem, defined in the continuum for rectifiable subsets of $[-1,1]^2$.

We settle Conjecture 11.1.1 for the original Cheeger constant $\widehat{\Phi}_n$ by employing the same overall strategy of [BLP15]. The distinction between $\widehat{\Phi}_n$ and the modified Cheeger constant $\widetilde{\Phi}_n$ is that, in the latter object, the edge boundary of a subgraph $H \subset \mathbf{C}_n$ is taken in the full infinite cluster $\mathbf{C}_\infty$ instead of just $\mathbf{C}_n$. This modification simplifies the nature of the limiting isoperimetric problem, which is the analogue of the standard Euclidean isoperimetric problem for an anisotropic perimeter functional. In our case, a *restricted perimeter* functional replaces the perimeter functional, reflecting the fact that $\widehat{\Phi}_n$ does not "see" edges outside the box $[-n,n]^2$.

11.2 The general form of the limiting variational problem

A *curve* λ in the unit square $[-1,1]^2$ is the image of a continuous function $\lambda : [0,1] \to [-1,1]^2$. A curve λ is *closed* if $\lambda(0) = \lambda(1)$ in any parametrization, *Jordan* if it is closed and one-to-one on $[0,1)$ and *rectifiable* if there is a parametrization of λ such that

$$\operatorname{length}(\lambda) := \sup_{n\in\mathbb{N}} \sup_{t_1<\cdots<t_n\in[0,1]} \sum_{j=1}^{n} |\lambda(t_j) - \lambda(t_{j-1})|_2 < \infty\,. \tag{11.3}$$

Many of the curves considered in this paper will be Jordan, and we will thus often conflate a curve λ with its image, denoted $\mathsf{image}(\lambda)$. In Chapter 13, we will study the variational problem (11.6) defined below in greater detail, and there we will be more careful. The class $\mathcal{R}$ of sets we work with is defined as follows

$$\mathcal{R} := \left\{ R \subset [-1,1]^2 \;:\; \begin{array}{c} R \text{ is compact, } R^\circ \neq \emptyset,\ \partial R \text{ is a finite union of rectifiable Jordan} \\ \text{curves, and the intersection of any two such curves is } \mathcal{H}^1\text{-null} \end{array} \right\}, \tag{11.4}$$

where $\mathcal{H}^1$ denotes the one-dimensional Hausdorff measure, and where R° denotes the interior of R. Given a norm τ on $\mathbb{R}^2$, define the restricted perimeter functional $\mathcal{I}_\tau$ on elements of $\mathcal{R}$ via

$$\mathcal{I}_\tau(\partial R) := \int_{\partial R \cap (-1,1)^2} \tau(n_x)\mathcal{H}^1(dx)\,, \tag{11.5}$$

where n_x is the normal vector to $\partial R \cap (-1,1)^2$ which exists at $\mathcal{H}^1$-almost every point on the curves $\partial R \cap (-1,1)^2$. Given the functional $\mathcal{I}_\tau$, form the following variational problem, of central interest in this paper

$$\text{minimize: } \frac{\mathcal{I}_\tau(\partial R)}{\text{Leb}(R)}\,, \quad \text{subject to: } \text{Leb}(R) \leq 2\,, \tag{11.6}$$

where $R \in \mathcal{R}$, and where Leb is the two-dimensional Lebesgue measure.

11.3 Results

Let $\mathcal{G}_n$ be the set of *Cheeger optimizers*, the subgraphs of $\mathbf{C}_n$ realizing the minimum defining $\widehat{\Phi}_n$. Recall that the Hausdorff metric on (non-empty) compact subsets of $[-1,1]^2$ is defined as follows: given $A, B \subset [-1,1]^2$ compact,

$$\mathrm{d}_H(A,B) := \max\left(\sup_{x\in A}\inf_{y\in B}\big|x-y\big|_\infty, \sup_{y\in B}\inf_{x\in A}\big|x-y\big|_\infty\right)\,, \tag{11.7}$$

where for $x, y \in \mathbb{R}^2$ and $p \in [1,\infty]$, $|x-y|_p$ denotes the ℓ^p-distance between x and y. The following shape theorem is the first of our main results.

Theorem 11.3.1. Let $d = 2$ and let $p > p_c(2)$. There is a norm β_p on $\mathbb{R}^2$ with non-empty collection of optimizers $\mathcal{R}_p$ to the associated variational problem (11.6) so that

$$\max_{G_n \in \mathcal{G}_n} \inf_{E \in \mathcal{R}_p} \mathrm{d}_H\Big(n^{-1}G_n\,,\, E\Big) \xrightarrow[n\to\infty]{} 0 \tag{11.8}$$

holds $\mathbb{P}_p$-almost surely.

The following definitions link Theorem 11.3.1 with the limit in Conjecture 11.1.1.

Definition 11.3.2. Let β_p be the norm in Theorem 11.3.1, which is the norm defined in [BLP15]. Given $R \in \mathcal{R}$, define the ratio

$$\frac{\mathcal{I}_{\beta_p}(\partial R)}{\text{Leb}(R)} \tag{11.9}$$

to be the *conductance* of R. Define the constant φ_p as

$$\varphi_p := \inf \left\{ \frac{\mathcal{I}_{\beta_p}(\partial R)}{\text{Leb}(R)} \ : R \in \mathcal{R},\, \text{Leb}(R) \leq 2 \right\} . \tag{11.10}$$

We remark that the two appearing in (11.10) and (11.6) is half the area of $[-1,1]^2$, and is an artifact of the 2 in the denominator of (11.1). Theorem 11.3.3 below is the second of our main results and settles Conjecture 11.1.1 in dimension two.

Theorem 11.3.3. Let $d = 2$ and let $p > p_c(2)$. Then $\mathbb{P}_p$-almost surely,

$$\lim_{n\to\infty} n\widehat{\Phi}_n = \frac{\varphi_p}{\theta_p}\,, \tag{11.11}$$

where $\theta_p = \mathbb{P}_p(0 \in \mathbf{C}_\infty)$, and where $\varphi_p \in (0,\infty)$ is defined in (11.10).

Definition 11.3.4. For U a subgraph of $\mathbf{C}_n$, let $\partial^n U$ denote the edge boundary of U in $\mathbf{C}_n$. We refer to this set as the *open edge boundary of U in $\boldsymbol{C}_n$*. Let $\partial^\infty U$ denote the edge boundary of U in all of $\mathbf{C}_\infty$, which we refer to as the *open edge boundary of U*. Define the *n-conductance* of U to be the ratio $|\partial^n U|/|U|$ and define the *conductance* of U to be the ratio $|\partial^\infty U|/|U|$.

Remark 11.3.5. Theorem 11.3.1 says that the optimizers to the variational problems encoded by the $\widehat{\Phi}_n$ scale to the optimizers of (11.6) for $\tau = \beta_p$. The random variable $\widehat{\Phi}_n$ is the n-conductance of any $G_n \in \mathcal{G}_n$. Theorem 11.3.3 says that these n-conductances scale to the optimal conductance (11.10) of the continuum problem (11.6) for the norm β_p.

11.4 Outline

In Chapter 12, we recall the definition of β_p from [BLP15], and we reintroduce the notion of right-most paths, which are used to define β_p. We collect the useful properties of both the norm and right-most paths.

In Chapter 13, we study the variational problem (11.6) for $\tau = \beta_p$. The main results here are existence and stability results: we first show the set $\mathcal{R}_p$ of optimizers of this problem is non-empty. We then show that if a connected set $R \in \mathcal{R}$ is d_H-far from $\mathcal{R}_p$, the conductance of R is at least φ_p plus a positive constant depending on the distance of R to $\mathcal{R}_p$.

In Chapter 14, we show the conductance of $R \in \mathcal{R}$ with $\mathrm{Leb}(R) \leq 2$ gives rise to upper bounds on $\widehat{\Phi}_n$ with high probability. Specifically, we pass from a nice object in the continuum to a subgraph of $\mathbf{C}_n$, and we relate the conductances of these two objects. We do this first for polygons, and then for more general sets, making use of the tools collected in Chapter 12. Ultimately, we show that for any $\epsilon > 0$, we have $n\widehat{\Phi}_n \leq (1+\epsilon)\varphi_p$ with high probability.

In Chapter 15, we move in the other direction, extracting from each Cheeger optimizer $G_n \in \mathcal{G}_n$ $R \in \mathcal{R}$ with $\mathrm{d}_H(G_n, nR)$ small, and relating the conductances of these objects. By controlling $\mathrm{Leb}(R)$ from above, we see that the conductance of R is at least $(1-\epsilon)\varphi_p$, which translates to a high probability lower bound on $\widehat{\Phi}_n$ of this form. This settles Theorem 11.3.3. We then use the stability result of Chapter 13 with the main result of Chapter 14 to see that it is rare for G_n to be far from $\mathcal{R}_p$, settling Theorem 11.3.1.

11.5 Discussion and context

We use many of the tools developed in [BLP15], and as such, our work can be seen as falling under the umbrella of the Wulff construction program. This was initiated in the early 1990s independently by Dobrushin, Kotecký and Shlosman [DKS92] in for the Ising model and by Alexander, Chayes and Chayes [ACC90] in percolation, both on the square lattice.

These works characterized the asymptotic shape of a large droplet of one phase of the model (for instance, a large finite open cluster in supercritical bond percolation). The probability of such an event decays rapidly in the size of the droplet, thus the theory of large deviations plays a crucial role in the analysis and is key to defining a model-dependent norm τ. Though the large droplets are not the minimizers of any isoperimetric problem, their limit shape is the minimizer of

$$\text{minimize: } \frac{\text{length}_\tau(\partial R)}{\text{Leb}(R)}, \quad \text{subject to: } \text{Leb}(R) \leq c \tag{11.12}$$

for some constant $c > 0$. The solution to (11.12), called the Wulff shape, is easily defined and was postulated by Wulff [Wul01] in 1901; it is a convex subset of $\mathbb{R}^2$ depending on τ. This solution is known to be unique up to translations and modifications on a null set thanks to the substantial work of Taylor [Tay74, Tay75, Tay78], whose results hold in all dimensions at least two. The Wulff construction has been successfully employed in dimensions strictly larger than two [Cer00, Bod99, Bod02, CP00, CP01], though with significant technical overhead due to geometric complications arising in higher dimensions. More details can be found in Section 5.5 of [Cer06] and in [BIV00].

The present work, as well as that of [BLP15], differs from the above in that we work exclusively in an event of full probability, and that we are faced with a collection of isoperimetric problems even at the discrete level. In our case, the variational problem in the continuum is a limit of these discrete problems. Because we study the unmodified Cheeger constant, our limiting variational problem (11.6) is more complicated than the variational problem given by (11.12). The shapes of droplets in the presence of a boundary, a single infinite wall, have been studied in the context of the Ising model [PV96, BIV01] using the analogue of the Wulff construction known as the Winterbottom construction [Win67]. This construction has been generalized further in a paper of Koteck$\acute{\text{y}}$ and Pfister [KP94], and related problems have been studied by Schlosman [Shl89].

11.6 Open problems

We remark on several future directions:

(1) We find it desirable to classify elements of $\mathcal{R}_p$ in terms of the Wulff shape W_p, the limit shape obtained in [BLP15] and the solution to the unrestricted isoperimetric problem (11.12) for the norm β_p. Based on work of Kotecký and Pfister [KP94] and Schlosman

[Shl89], we conjecture that the collection $\mathcal{R}_p$ consists of quarter-Wulff shapes or their complements in the square. Answering such questions may require a better understanding of the regularity of the norm. Questions regarding the regularity and strict convexity of β_p are interesting in their own right and are related to open problems in first-passage percolation (see for instance Chapter 2 of [ADH15]).

(2) Instead of studying the largest connected component of $\mathbf{C}_\infty \cap [-n, n]^2$, we can fix a Jordan domain $\Omega \subset \mathbb{R}^2$ and consider the Cheeger constant of the largest connected component of $\mathbf{C}_\infty \cap n\Omega$. The argument in this paper is likely robust enough that both Cheeger asymptotics and a shape theorem can be deduced in this case (perhaps depending on the convexity of Ω). This problem is similar in flavor to work of Cerf and Théret [CT12], in which the shapes of minimal cutsets in first passage percolation are studied for more general domains.

(3) A sharp limit and related shape theorem were recently obtained [Gol16] for the modified Cheeger constant in dimensions three and higher. It is likely that by combining the techniques of [Gol16] and the present paper, one can prove analogues of Theorem 11.3.1 and Theorem 11.3.3 for the giant component in dimensions larger than two.

11.7 Acknowledgements

I thank my advisor Marek Biskup for suggesting this problem, and for his guidance. I thank David Clyde, John Garnett, Stephen Ge, Nestor Guillen, Inwon Kim and Peter Petersen for useful conversations. This research has been partially supported by the NSF grant DMS-1407558.

CHAPTER 12

The boundary norm

The motivation for the construction of β_p goes back to the postulate of Gibbs [Gib78] that one phase of matter immersed in another will arrange itself so that the surface energy between the two phases is minimized. By regarding each $G_n \in \mathcal{G}_n$ as a droplet immersed in $\mathbf{C}_n \setminus G_n$, we can study the interface between these two "phases" and attempt to extract a surface energy.

Our tool for studying these interfaces are right-most paths, introduced in [BLP15]. Each Cheeger optimizer G_n may be expressed using finitely many right-most circuits, which together represent the boundary of G_n and hence the total interface between G_n and $\mathbf{C}_n \setminus G_n$. We assign a weight to each right-most path which depends on the percolation configuration, so that the combined weight of all right-most paths making up the boundary of G_n is exactly $|\partial^\infty G_n|$.

Given $v \in \mathbb{S}^1$, the value $\beta_p(v)$ encodes the asymptotic minimal weight of a right-most path joining two vertices $x, y \in \mathbb{Z}^d$ with $y - x$ a large multiple of v. Thus, the norm β_p encodes the surface energy minimization taking place locally at the boundary of each G_n.

12.1 Right-most paths

Consider the graph $\mathbb{Z}^2 = (\mathrm{V}(\mathbb{Z}^2), \mathrm{E}(\mathbb{Z}^2))$. Given $x, y \in \mathrm{V}(\mathbb{Z}^2)$, a *path from x to y* is an alternating sequence of vertices and edges $\gamma = (x_0, e_1, x_1, \ldots, e_m, x_m)$ such that e_i joins x_{i-1} with x_i for $i \in \{1, \ldots, m\}$, and such that $x_0 = x$ and $x_m = y$. The *length* of γ, denoted $|\gamma|$, is m. If $x_0 = x_m$, the path is said to be a *circuit*.

It is useful to regard edges in a given path γ as oriented, so that the edge e_i starting at x_{i-1} and ending at x_i, denoted $\langle x_{i-1}, x_i \rangle$, is considered distinct from the edge starting at x_i and ending at x_{i-1}, denoted $\langle x_i, x_{i-1} \rangle$. A path γ in $\mathbb{Z}^2$ is *simple* if no oriented edge is used twice. Given paths $\gamma_1 = (x_0, e_1, \ldots, e_m, x_m)$ and $\gamma_2 = (y_0, f_1, \ldots, f_k, y_k)$ with $x_m = y_0$, define the *concatenation* of γ_1 and γ_2, denoted $\gamma_1 * \gamma_2$ to be the path $(x_0, e_1, \ldots, e_m, x_m, f_1 \ldots, f_k, y_k)$.

Definition 12.1.1. Let γ be a path in $\mathbb{Z}^d$ and let x_i be a vertex in γ with x_{i-1} and x_{i+1} well-defined. The *right-boundary edges at* x_i are obtained by enumerating all oriented edges whcih start at x_i, beginning with but not including $\langle x_i, x_{i-1} \rangle$, proceeding in a counter-clockwise manner and ending with but not including $\langle x_i, x_{i+1} \rangle$. If either x_{i-1} or x_{i+1} is not well-defined, the right-most boundary edges at x_i are defined to be the empty set. The *right-boundary of* γ, denoted $\partial^+\gamma$, is the union of all right-boundary edges at each vertex of γ.

Definition 12.1.2. A path $\gamma = (x_0, e_1, x_1, \ldots, e_m, x_m)$ is said to be *right-most* if it is simple, and if no e_i is an element of $\partial^+\gamma$.

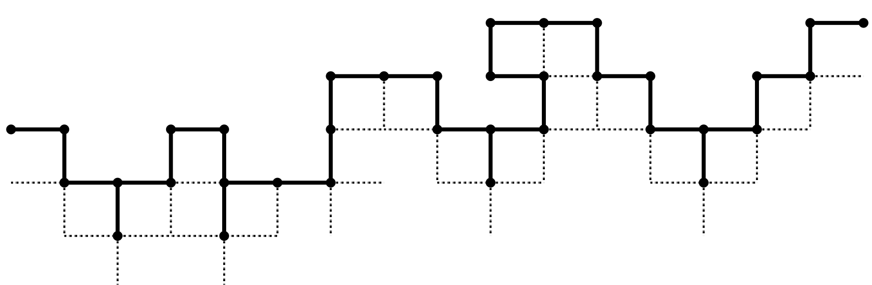

Figure 12.1: In black, a right-most path which begins on the left and ends on the right. The dotted edges are the right-most boundary of this path.

Definition 12.1.3. We assign configuration-dependent weights to right-most paths. Define the edge-sets

$$\mathfrak{b}(\gamma) := \left\{ e \in \partial^+\gamma \ : \omega(e) \text{ is open} \right\}, \tag{12.1}$$

$$\mathfrak{b}^n(\gamma) := \left\{ e \in \mathfrak{b}(\gamma) \ : e \subset [-n, n]^2 \right\}, \tag{12.2}$$

and refer to $|\mathfrak{b}(\gamma)|$ and $|\mathfrak{b}^n(\gamma)|$ respectively as the $\boldsymbol{C}_\infty$*-length* of γ and the $\boldsymbol{C}_n$*-length* of γ.

Remark 12.1.4. As we will see in Lemma 12.2.4, the boundary of a subgraph U of $\mathbf{C}_n$ may be expressed as a collection of right-most circuits. The total $\mathbf{C}_\infty$-length of these circuits will correspond to the size of $\partial^\infty U$, and the total $\mathbf{C}_n$-length of these circuits will correspond to the size of $\partial^n U$.

Following [BLP15], we let $\mathcal{R}(x, y)$ denote the collection of all right-most paths joining x to y. If vertices x and y are joined by an open path (and hence joined by an open right-most path) in the configuration ω, define the *right-boundary distance* from x to y as

$$b(x, y) := \inf \left\{ \mathfrak{b}(\gamma) \ : \gamma \in \mathcal{R}(x, y),\ \gamma \text{ uses only open edges} \right\}. \tag{12.3}$$

Remark 12.1.5. It is convenient to allow b to act on points in $\mathbb{R}^2$ by assigning to each $x \in \mathbb{R}^2$ a "nearest" point $[x]$ in $\mathbf{C}_\infty$. To do this, we augment our probability space to support a collection $\{\eta_x : x \in \mathbb{Z}^2\}$ of iid random variables uniform on $[0, 1]$ and independent of the Bernoulli random variables used to define the bond percolation. Given $x \in \mathbb{R}^2$, we let $[x]$ be the nearest (in ℓ^∞-sense) vertex in $\mathbf{C}_\infty$ to x, breaking ties using the η_x if necessary.

One can establish high-probability closeness of any $x \in \mathbb{R}^2$ with $[x]$ using a duality argument; the following is Lemma 2.7 of [BLP15].

Lemma 12.1.6. Suppose $p > p_c(2)$. There are positive constants $c_1(p), c_2(p)$ so that for all $x \in \mathbb{Z}^2$ and all $r > 0$,

$$\mathbb{P}_p\Big(|[x] - x|_2 > r \Big) \leq c_1 \exp\Big(-c_2 r \Big). \tag{12.4}$$

12.2 Properties of right-most paths

Before defining β_p, we mention some useful properties of right-most paths. In particular, we recall the correspondence between right-most paths and simple paths in the medial graph of $\mathbb{Z}^2$. Given a planar graph $G = (\mathrm{V}, \mathrm{E})$, the *medial graph* $G_\sharp = (\mathrm{V}_\sharp, \mathrm{E}_\sharp)$ is the graph with vertices $\mathrm{V}_\sharp = \mathrm{E}$, and with any two vertices in $\mathrm{V}_\sharp$ adjacent in $G_\sharp$ if the corresponding edges of G are adjacent in a face of G.

An *interface* is an edge self-avoiding oriented path in $\mathbb{Z}^2_\sharp$, which does not use its initial or terminal vertex more than once, except to close a circuit. There is a correspondence between interfaces and right-most paths: an interface $\partial = (e_1, \ldots, e_m)$, written as a sequence of vertices in $\mathbb{Z}^2_\sharp$, either reflects on a given edge e_i or cuts through a given edge.

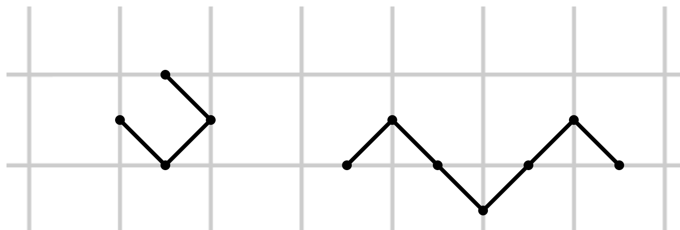

Figure 12.2: The medial path of length three on the left reflects on each edge. On the right, the medial path of length six cuts through each edge.

More rigorously, an interface $\partial = (e_1, \ldots, e_m)$ is said to *reflect* on e_i (for $i \in \{2, \ldots, m-1\}$) if e_{i-1} and e_{i+1} are on the boundary of the same face of $\mathbb{Z}^2$, and ∂ is said to *cut through* e_i otherwise. The following proposition (Proposition 2.3 of [BLP15]) provides a fundamental correspondence between interfaces and right-most paths.

Proposition 12.2.1. For each interface $\partial = (e_1, \ldots, e_m)$, the subsequence $(e_{k_1}, \ldots, e_{k_n})$ of edges which are not cut through by ∂ forms a right-most path γ. This mapping is one-to-one and onto the set of all right-most paths. In particular, γ is a right-most circuit if and only if ∂ is a circuit in the medial graph. Finally, the edges of $\partial \setminus (e_{k_1}, \ldots, e_{k_n})$ (oriented properly) form $\partial^+\gamma$.

Remark 12.2.2. Interfaces may be perturbed via "corner-rounding" to simple curves in $\mathbb{R}^2$, as illustrated at the bottom of Figure 12.3. In particular, if γ is a right-most circuit, it may be identified with a rectifiable Jordan curve λ_∂ built from the interface ∂ corresponding to γ via Proposition 12.2.1.

Definition 12.2.3. Let λ be a rectifiable curve and for $x \notin \lambda$, let $w_\lambda(x)$ denote the winding number of λ around x. Define

$$\mathsf{hull}(\lambda) := \lambda \cup \left\{x \notin \lambda \,:\, w_\lambda(x) \text{ is odd}\right\}, \tag{12.5}$$

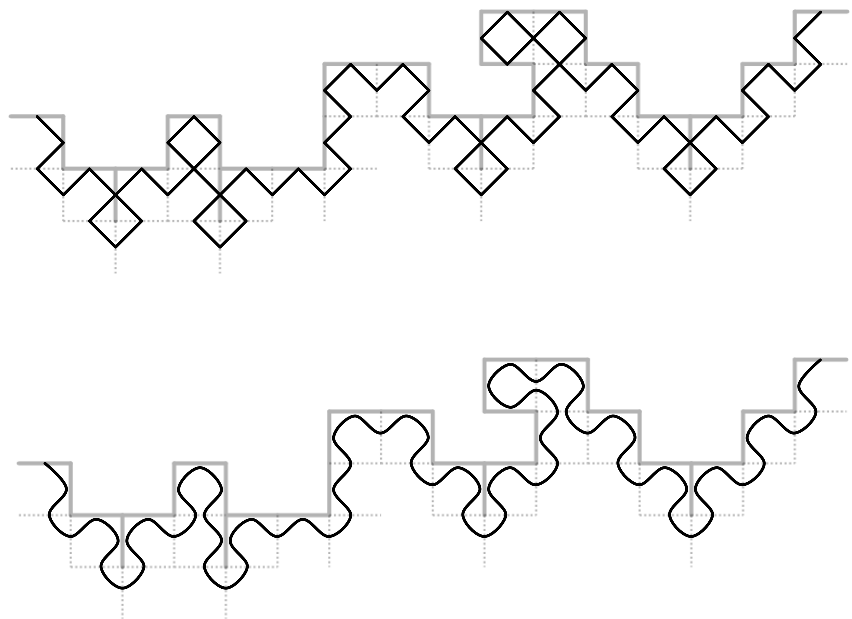

Figure 12.3: Above: the correspondence of Proposition 12.2.1, built from the right-most path in Figure 12.1. Below: the perturbed interface is a simple curve.

A fundamental property of right-most circuits is that they may be used to "carve out" subgraphs of $\mathbf{C}_n$. This is done in a way which conveniently links the total length of the circuits with the edge boundary of the subgraph, see Remark 12.1.4. Let $\mathcal{U}_n$ denote the collection of *connected* subgraphs of $\mathbf{C}_\infty \cap [-n,n]^2$ determined by their vertex set. Given an interface ∂ corresponding to a right-most circuit, let λ_∂ be the Jordan curve obtained from ∂ by rounding the corners, and write $\mathsf{hull}(\partial)$ for $\mathsf{hull}(\lambda_\partial)$. The following is proved by inducting on the size of the vertex set of U.

Lemma 12.2.4. Let $U \in \mathcal{U}_n$. The graph $\mathbf{C}_\infty \setminus U$ consists of a unique infinite connected component and finitely many finite connected components $\Lambda_1, \ldots, \Lambda_m$. There are open right-most circuits $\gamma, \gamma_1, \ldots, \gamma_m$ contained in U, where γ is oriented counter-clockwise and each γ_j is oriented clockwise so that

1. $\partial, \partial_1, \ldots, \partial_m$ are disjoint ,
2. $\mathfrak{b}(\gamma) \cup \left(\bigsqcup_{j=1}^m \mathfrak{b}(\gamma_j) \right) = \partial^\infty U$,

3. $U = \left[\mathsf{hull}(\partial) \setminus \left(\bigsqcup_{j=1}^m \mathsf{hull}(\partial_j)\right)\right] \cap \mathbf{C}_\infty$,

4. For each $j \in \{1, \dots, m\}$, we have $\Lambda_j = \mathsf{hull}(\partial_j) \cap \mathbf{C}_\infty$,

where ∂ is the counter-clockwise interface corresponding to γ, and where each ∂_j is the clockwise interface corresponding to γ_j.

The final input on right-most paths we include is Proposition 2.9 of [BLP15], which tells us $|\gamma|$ and $|\mathfrak{b}(\gamma)|$ are comparable when $|\gamma|$ is sufficiently large. This enables us to pass from discrete sets with reasonably sized open edge boundaries to rectifiable sets in the continuum.

Proposition 12.2.5. Let $p > p_c(2)$. There are positive constants α, c_1, c_2 depending only on p such that for all $n \geq 0$, we have

$$\mathbb{P}_p\Big(\exists \gamma \in \bigcup_{x \in \mathbb{Z}^2} \mathcal{R}(0,x) \,:\, |\gamma| \geq n, |\mathfrak{b}(\gamma)| \leq \alpha n\Big) \leq c_1 \exp(-c_2 n)\,. \tag{12.6}$$

12.3 The norm

We now use right-most paths to define the norm β_p on $\mathbb{R}^2$, and we aggregate several useful results from [BLP15]. The following is the main result (Theorem 2.1 and Proposition 2.2) of Section 2 in [BLP15], which we state verbatim.

Theorem 12.3.1. Let $p > p_c(2)$, and let $x \in \mathbb{R}^2$. The limit

$$\beta_p(x) := \lim_{n\to\infty} \frac{b([0],[nx])}{n} \tag{12.7}$$

exists $\mathbb{P}_p$-almost surely and is non-random, non-zero (when $x \neq 0$) and finite. The limit also exists in L^1 and the convergence is uniform on $\{x \in \mathbb{R}^2 \,:\, |x|_2 = 1\}$. Moreover,

1. β_p is homogeneous, i.e. $\beta_p(cx) = |c|\beta_p(x)$ for all $x \in \mathbb{R}^2$ and all $c \in \mathbb{R}$,

2. β_p obeys the triangle inequality

$$\beta_p(x+y) \leq \beta_p(x) + \beta_p(y)\,, \tag{12.8}$$

3. β_p inherits the symmetries of $\mathbb{Z}^2$; for all $(x_1, x_2) \in \mathbb{R}^2$, we have

$$\beta_p\big((x_1, x_2)\big) = \beta_p\big((x_2, x_1)\big) = \beta_p\big((\pm x_1, \pm x_2)\big) \tag{12.9}$$

for any choice of the signs $\pm$.

Remark 12.3.2. Theorem 12.3.1 tells us β_p defines a norm on $\mathbb{R}^2$, and that this norm inherits the symmetries of $\mathbb{Z}^2$. It is first proved by appealing to the subadditive ergodic theorem, but can also be deduced from concentration estimates developed in Section 3 of [BLP15], which we state below.

The first concentration estimate we record is measure theoretic, it is Theorem 3.1 of [BLP15].

Theorem 12.3.3. Let $p > p_c(2)$. For each $\epsilon > 0$, there are positive constants $c_1(p, \epsilon), c_2(p, \epsilon)$ so that for all $x, y \in \mathbb{Z}^2$,

$$\mathbb{P}_p\left(\left|\frac{b([x],[y])}{\beta_p(y-x)} - 1\right| > \epsilon\right) \leq c_1 \exp\Big(- c_2 \log^2 |y-x|_2\Big)\,. \tag{12.10}$$

We also require a result on the geometric concentration of right-most paths; namely that right-most paths which are almost optimal are geometrically close to the straight line joining their endpoints. Given $x, y \in \mathbf{C}_\infty$, say that $\gamma \in \mathcal{R}(x, y)$ is *ϵ-optimal* if

$$\mathfrak{b}(\gamma) - b(x, y) \leq \epsilon |y - x|_2\,, \tag{12.11}$$

and write $\Gamma_\epsilon(x, y)$ for the set of ϵ-optimal paths in $\mathcal{R}(x, y)$. The following is Proposition 3.2 of [BLP15].

Proposition 12.3.4. Let $p > p_c(2)$. There are positive constants α, c_1, c_2 so that for all $x, y \in \mathbb{Z}^2$, we have

1. For any $t > \alpha |x - y|_2$,

$$\mathbb{P}_p\Big(\exists \gamma \in \Gamma_0([x],[y]) \,:\, |\gamma| > t\Big) \leq c_1 \exp\Big(- c_2 t\Big)\,. \tag{12.12}$$

2. For all $\epsilon > 0$, once $|y - x|$ is sufficiently large depending on ϵ,

$$\mathbb{P}_p\Big(\forall \gamma \in \Gamma_\epsilon([x],[y]) \ : \mathrm{d}_H\big(\gamma, \mathsf{poly}(x,y)\big) > \epsilon |y-x|_2\Big) \leq c_1 \exp\Big(- c_2 \log^2(|y-x|_2)\Big)\,, \tag{12.13}$$

where $\mathsf{poly}(x,y)$ is the linear segment connecting x and y.

CHAPTER 13

The variational problem

Having reintroduced β_p in Chapter 12, we now discuss the variational problem (11.6) specialized to $\tau = \beta_p$. In fact, we will use nothing about β_p in this section other than the fact that it is a norm. We need two results in order to prove Theorem 11.3.1 and Theorem 11.3.3: an existence result and a stability result. We write the functional defined in (11.5) for $\tau = \beta_p$ as $\mathcal{I}_p$, and for $R \in \mathcal{R}$, we refer to $\mathcal{I}_p(\partial R)$ as the *surface energy of* R. We also introduce the β_p*-length* of a rectifiable curve $\lambda : [0,1] \to \mathbb{R}^2$:

$$\mathrm{length}_{\beta_p}(\lambda) := \sup_{n \in \mathbb{N}} \sup_{t_1 < \cdots < t_n \in [0,1]} \sum_{j=1}^{n} \beta_p(\lambda(t_j) - \lambda(t_{j-1})) \,. \tag{13.1}$$

We find it necessary to consider not just the variational problem (11.6), but a family of related problems. For $\alpha \in [-1, 1]$, define the following isoperimetric problem for sets $R \in \mathcal{R}$:

$$\text{minimize: } \frac{\mathcal{I}_p(\partial R)}{\mathrm{Leb}(R)} \,, \quad \text{subject to: } \mathrm{Leb}(R) \leq 2 + \alpha \tag{13.2}$$

The minimal value for (13.2) is

$$\varphi_p^{(2+\alpha)} := \inf \left\{ \frac{\mathcal{I}_p(\partial R)}{\mathrm{Leb}(R)} \,:\, \mathrm{Leb}(R) \leq 2 + \alpha \,, R \in \mathcal{R} \right\} , \tag{13.3}$$

and the set of optimizers for (13.2) is defined below as

$$\mathcal{R}_p^{(2+\alpha)} := \left\{ R \in \mathcal{R} \,:\, \mathrm{Leb}(R) \leq 2 + \alpha \,, \frac{\mathcal{I}_p(\partial R)}{\mathrm{Leb}(R)} = \varphi_p^{(2+\alpha)} \right\} . \tag{13.4}$$

Thus, in our new notation, the constant φ_p introduced in (11.10) is denoted $\varphi_p^{(2)}$ in this section, and the collection of optimizers $\mathcal{R}_p$ introduced in Theorem 11.3.1 is denoted $\mathcal{R}_p^{(2)}$.

13.1 Sets of finite perimeter

We extend the problem (13.2) to a larger class of sets, proving existence within this class and then recovering a representative in $\mathcal{R}$. For $E \subset [-1,1]^2$ be Borel, we define the *perimeter* of E, denoted $\mathrm{per}(\partial E)$, as

$$\mathrm{per}(\partial E) := \sup \left(\int_E \mathrm{div}(f) dx \ : f \in C_c^\infty(\mathbb{R}^2, \mathbb{R}^2) \, , \, |f|_2 \leq 1 \right) \, , \tag{13.5}$$

and say that E is a *set of finite perimeter* if $\mathrm{per}(\partial E) < \infty$. Let $\mathcal{C}$ denote the collection of all sets of finite perimeter (after Caccioppoli) contained in $[-1,1]^2$. Given $E \in \mathcal{C}$, we define the β_p*-perimeter* of E similarly:

$$\mathrm{per}_{\beta_p}(\partial E) := \sup \left(\int_E \mathrm{div}(f) dx \ : f \in C_c^\infty(\mathbb{R}^2, \mathbb{R}^2) \, , \, \beta_p^*(f) \leq 1 \right) \, , \tag{13.6}$$

where β_p^* is the dual norm to β_p. Finally, we define the surface energy of $E \in \mathcal{C}$ as:

$$\mathcal{I}_p(\partial E) := \sup \left(\int_E \mathrm{div}(f) dx \ : f \in C_c^\infty((-1,1)^2, \mathbb{R}^2) \, , \, \beta_p^*(f) \leq 1 \right) \, . \tag{13.7}$$

Remark 13.1.1. Each $R \in \mathcal{R}$ is an element of $\mathcal{C}$, and the surface energy of R defined in (11.5) agrees with the surface energy of E, defined in (13.7). This enables us to extend the variational problem (13.2) to sets of finite perimeter, and given $E \in \mathcal{C}$, we call $\mathcal{I}_p(\partial E)/\mathrm{Leb}(E)$ the *conductance* of E, which is consistent with the terminology in the introduction.

We introduce the optimal value and set of optimizers corresponding to the variational problem over this wider class of sets. Define

$$\psi_p^{(2+\alpha)} := \inf \left\{ \frac{\mathcal{I}_p(\partial E)}{\mathrm{Leb}(E)} \ : \mathrm{Leb}(E) \leq 2 + \alpha \, , E \in \mathcal{C} \right\} \, , \tag{13.8}$$

with the convention that zero divided by zero is infinity. Also define

$$\mathcal{C}_p^{(2+\alpha)} := \left\{ E \in \mathcal{C} \ : \mathrm{Leb}(E) \leq 2 + \alpha \, , \frac{\mathcal{I}_p(\partial E)}{\mathrm{Leb}(E)} = \psi_p^{(2+\alpha)} \right\} \, . \tag{13.9}$$

Lower semicontinuity is a fundamental feature of the perimeter and surface energy functionals (see for instance Section 14.2 of [Cer06]).

Lemma 13.1.2. Let $E_k \in \mathcal{C}$ be a sequence converging in L^1-sense to E. Then

1. $\mathrm{per}(\partial E) \le \liminf_{k\to\infty} \mathrm{per}(\partial E_k)$,

2. $\mathrm{per}_{\beta_p}(\partial E) \le \liminf_{k\to\infty} \mathrm{per}_{\beta_p}(\partial E_k)$,

3. $\mathcal{I}_p(\partial E) \le \liminf_{k\to\infty} \mathcal{I}_p(\partial E_k)$,

so that if $\mathrm{per}(\partial E_k)$ is uniformly bounded in k, we have $E \in \mathcal{C}$.

We now introduce some terminology in order to state a result which linking the classes $\mathcal{R}$ and $\mathcal{C}$.

Definition 13.1.3. Given $E \subset [-1,1]^2$ Borel, define the *upper density* of E at $x \in \mathbb{R}^2$ as

$$D^+(E,x) := \limsup_{r\to 0} \frac{\mathrm{Leb}(E \cap B(x,r))}{\mathrm{Leb}(B(x,r))}, \tag{13.10}$$

and define the *essential boundary* of E as

$$\partial^* E := \left\{ x \in \mathbb{R}^2 \ : D^+(E,x) > 0,\ D^+(\mathbb{R}^2 \setminus E, x) > 0 \right\} \tag{13.11}$$

Definition 13.1.4. Let $E \subset \mathbb{R}^2$ be a set of finite perimeter. Say E is *decomposable* if there is a partition of E into $A, B \subset \mathbb{R}^2$ so that $\mathrm{Leb}(A)$ and $\mathrm{Leb}(B)$ are strictly positive and so that $\mathrm{per}(\partial E) = \mathrm{per}(\partial A) + \mathrm{per}(\partial B)$. Say that E is *indecomposable* if it is not decomposable.

Recall that given a Jordan curve λ, we defined the compact set $\mathsf{hull}(\lambda)$ in (12.5). We write $\mathsf{hull}(\lambda)^\circ$ for the interior of this compact set. The following result, originally due to Fleming and Federer, allows us to think of $\partial^* E$ for $E \in \mathcal{C}$ as a countable collection of rectifiable Jordan curves. The version we state is taken roughly verbatim from Corollary 1 of [ACM01].

Proposition 13.1.5. Let $E \subset \mathbb{R}^2$ be a set of finite perimeter. There is a unique decomposition of $\partial^* E$ into rectifiable Jordan curves $\{\lambda_i^+, \lambda_j^- \ : i, j \in \mathbb{N}\}$ (modulo $\mathcal{H}^1$-null sets) so that

1. For $i \ne k \in \mathbb{N}$, $\mathsf{hull}(\lambda_i^+)^\circ$ and $\mathsf{hull}(\lambda_k^+)^\circ$ are either disjoint, or one is contained in the other. Likewise, for $i \ne k \in \mathbb{N}$, $\mathsf{hull}(\lambda_i^-)^\circ$ and $\mathsf{hull}(\lambda_k^-)^\circ$ are either disjoint, or one is contained in the other. Each $\mathsf{hull}(\lambda_j^-)^\circ$ is contained in one of the $\mathsf{hull}(\lambda_i^+)^\circ$.

2. $\mathrm{per}(\partial E) = \sum_{i=1}^{\infty} \mathcal{H}^1(\lambda_i^+) + \sum_{j=1}^{\infty} \mathcal{H}^1(\lambda_j^-)$.

3. If $\mathsf{hull}(\lambda_i^+)^\circ \subset \mathsf{hull}(\lambda_j^+)^\circ$ for $i \neq j$, then for some λ_k^-, we have $\mathsf{hull}(\lambda_i^+)^\circ \subset \mathsf{hull}(\lambda_k^-)^\circ \subset \mathsf{hull}(\lambda_j^+)^\circ$. Likewise, if $\mathsf{hull}(\lambda_i^-)^\circ \subset \mathsf{hull}(\lambda_j^-)^\circ$ for $i \neq j$, there is some λ_k^+ with $\mathsf{hull}(\lambda_i^-)^\circ \subset \mathsf{hull}(\lambda_k^+)^\circ \subset \mathsf{hull}(\lambda_j^-)^\circ$.

4. For $i \in \mathbb{N}$, let $L_i = \{j\ : \mathsf{hull}(\lambda_j^-)^\circ \subset \mathsf{hull}(\lambda_i^+)^\circ\}$, and set

$$Y_i = \mathsf{hull}(\lambda_i^+) \setminus \left(\bigcup_{j \in L_i} \mathsf{hull}(\lambda_j^-)^\circ \right) . \tag{13.12}$$

The sets Y_i are indecomposable with $\mathcal{H}^1$-null intersection, and moreover $\bigcup_{j=1}^{\infty} Y_j$ is equivalent of E modulo Lebesgue null sets.

Proposition 13.1.5 tells us that sets of finite perimeter are in some sense extensions of the class $\mathcal{R}$ to sets whose boundary consists of countably many Jordan arcs instead of finitely many. Thus, it is reasonable that the theory of such sets comes into play when discussing limits of sets in $\mathcal{R}$.

13.2 Existence

We now show that $\mathcal{R}_p^{(2+\alpha)}$ is non-empty for all $\alpha \in [-1, 1]$ by first using standard arguments to show $\mathcal{C}_p^{(2+\alpha)}$ is non-empty, and then by recovering elements of $\mathcal{R}$ from sets in $\mathcal{C}_p^{(2+\alpha)}$. We begin by making several basic observations.

The first observation implies optimal Jordan domains must have full volume.

Lemma 13.2.1. Let $\alpha \in [-1, 1]$. Let $R \in \mathcal{R}$ be such that $\mathrm{Leb}(R) < 2 + \alpha$ and such that $R = \mathsf{hull}(\lambda)$ for a rectifiable Jordan curve $\lambda \subset [-1, 1]^2$. Then there is also $R' \in \mathcal{R}$ with $\mathrm{Leb}(R) = 2 + \alpha$ and $R' = \mathsf{hull}(\lambda')$ for a rectifiable Jordan curve $\lambda' \subset [-1, 1]^2$ with

$$\frac{\mathcal{I}_p(\partial R)}{\mathrm{Leb}(R)} > \frac{\mathcal{I}_p(\partial R')}{\mathrm{Leb}(R')} . \tag{13.13}$$

Proof. Let $R \in \mathcal{R}$ be as above, and consider the open set $A = (-1,1)^2 \setminus R$. We consider three cases.

Case I: In the first, each connected component A' of A is such that $\partial A'$ intersects the interior of at most two adjacent sides of $[-1,1]^2$ non-trivially. In this first case, we can easily shrink the connected components of A to form a new open set of arbitrarily small volume, and whose surface energy is at most that of A. By complementation, we recover R' with the desired properties.

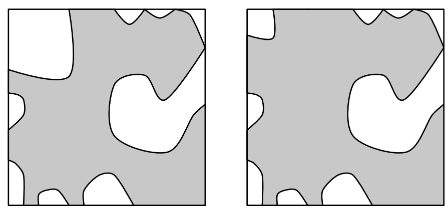

Figure 13.1: On the left, the original set $R \in \mathcal{R}$ in grey. On the right, the set $R' \in \mathcal{R}$ obtained through the procedure described in ***Case I***.

Case II: In the second case, there is a connected component A' of A such that $\partial A'$ intersects the interior of exactly three sides of $[-1,1]^2$ non-trivially. Because R is connected, $\partial A' \cap (-1,1)^2$ consists of a single arc which joins two opposing faces of the square, and this arc may be translated until it touches one of the other faces of the square. These translations naturally yield sets of the desired form and of larger measure. If the measure of these sets surpasses $2+\alpha$ before the arc reaches the boundary, we are content. Otherwise, we have obtained a set which is handled by the previous case (upon performing the same procedure on at most one other arc, perhaps).

Case III: As R must be connected, it is impossible for any connected component A' of A to have the property that $\partial A'$ intersects the interiors of two opposite sides of $[-1,1]^2$ non-trivially. Thus the last case to consider is that there is a connected component A' of A such that $\partial A'$ intersects the interior of all four sides of $[-1,1]^2$ non-trivially. In this case, ∂R intersects the interiors of at most two adjacent sides of $[-1,1]^2$ non-trivially. We may then dilate R about the corner it contains or the side it rests against until we either have a

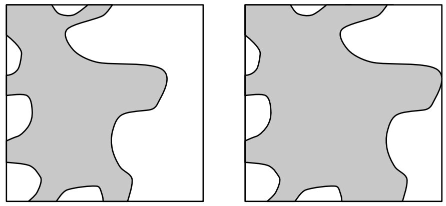

Figure 13.2: On the left, the original $R \in \mathcal{R}$ in grey. On the right, R' is obtained by "sliding" one of the contours along the boundary of the box.

set of the desired measure or we have a set falling into one of the preceding cases.

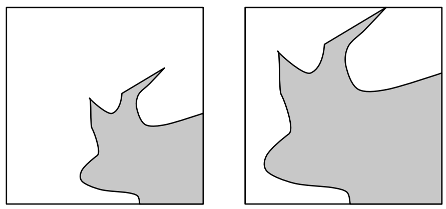

Figure 13.3: On the left, $R \in \mathcal{R}$ is in grey. On the right, $R' \in \mathcal{R}$ is obtained by dilating R.

This completes the proof. □

Lemma 13.2.1 implies that optimal sets of finite perimeter must also have full volume.

Lemma 13.2.2. Let $\alpha \in [-1, 1]$, and let $E \in \mathcal{C}$ with either $\mathrm{Leb}(E) < 2+\alpha$, or $\mathrm{Leb}(E) \leq 2+\alpha$ and E decomposable. There is $E' \in \mathcal{C}$ with $\mathrm{Leb}(E') = 2 + \alpha$ so that

$$\frac{\mathcal{I}_p(\partial E)}{\mathrm{Leb}(E)} > \frac{\mathcal{I}_p(\partial E')}{\mathrm{Leb}(E')} \,. \tag{13.14}$$

Proof. The case that $\mathrm{Leb}(E) \leq 2+\alpha$ and E is decomposable is an immediate corollary of the case $\mathrm{Leb}(E) < 2+\alpha$, so we work in the latter. Thanks to the inequality $\frac{a+b}{c+d} \geq \min(\frac{a}{c}, \frac{b}{d})$, we lose no generality supposing E is indecomposable. Using Proposition 13.1.5, it follows that E may be represented by rectifiable Jordan arcs λ and $\{\lambda_j\}_{j\geq 1}$ so that up to a Lebesgue-null set, $E = \mathsf{hull}(\lambda) \setminus \bigcup_{j\geq 1} \mathsf{hull}(\lambda_j)^\circ$. As the curves λ, λ_j have $\mathcal{H}^1$-null intersection, the sets $\mathsf{hull}(\lambda_j)^\circ$ are pairwise disjoint. Under the hypothesis that $\mathrm{Leb}(E) < 2+\alpha$, we may then shrink the curves λ_j one by one to produce a set E' having strictly smaller conductance. Thus, it suffices to consider sets E of finite perimeter which may be represented by a single rectifiable Jordan curve λ, but this is handled entirely by Lemma 13.2.1. □

We may now deduce that the collection of optimizers for (13.4) is non-empty within the class of sets of finite perimeter.

Lemma 13.2.3. The set of optimizers $\mathcal{C}_p^{(2+\alpha)}$ for the variational problem (13.4) is non-empty.

Proof. Let $E_k \in \mathcal{C}$ be a sequence of sets of finite perimeter such that

$$\frac{\mathcal{I}_p(\partial E_k)}{\mathrm{Leb}(E_k)} \to \psi_p^{(2+\alpha)} \,. \tag{13.15}$$

By Lemma 13.2.2, we lose no generality supposing $\mathrm{Leb}(E_k) = 2+\alpha$ for each k. As $\psi_p^{(2+\alpha)}$ is clearly finite, the perimeters of the E_k are uniformly bounded. We appeal to Rellich-Kondrachov and pass to a subsequence of the E_k converging to some $E \subset [-1,1]^2$ in L^1-sense. By Lemma 13.1.2, it follows that E is a set of finite perimeter with $\mathrm{Leb}(E) = 2+\alpha$ and $\mathcal{I}_p(\partial E) \leq \liminf_{k\to\infty} \mathcal{I}_p(\partial E_k)$. Thus the conductance of E is at most $\psi_p^{(2+\alpha)}$, which implies $E \in \mathcal{C}_p^{(2+\alpha)}$. □

We may now deduce that $\mathcal{R}_p^{(2+\alpha)}$ is non-empty for $\alpha \in [-1,1]$, among other things. The following is the main result of this section.

Corollary 13.2.4. Let $\alpha \in [-1,1]$.

1. If $E \in \mathcal{C}_p^{(2+\alpha)}$, then E is indecomposable and $\mathrm{Leb}(E) = 2+\alpha$.

2. $E \in \mathcal{C}_p^{(2+\alpha)}$ if and only if $E^c \in \mathcal{C}_p^{(2-\alpha)}$.

3. $\frac{2+\alpha}{2-\alpha}\psi_p^{(2+\alpha)} = \psi_p^{(2-\alpha)}$.

4. Each $E \in \mathcal{C}_p^{(2+\alpha)}$ is equivalent up to a Lebesgue-null set to some $R \in \mathcal{R}$. Thus, $\mathcal{R}_p^{(2+\alpha)}$ is non-empty and $\varphi_p^{(2+\alpha)} = \psi_p^{(2+\alpha)}$.

5. If $E \in \mathcal{C}_p^{(2+\alpha)}$, there are rectifiable Jordan curves $\lambda, \lambda' \subset [-1,1]^2$ so that up to Lebesgue-null sets, $E = \mathsf{hull}(\lambda)$ and $E^c = \mathsf{hull}(\lambda')$. Moreover, $\lambda \cap \lambda'$ is a simple rectifiable curve joining distinct points on $\partial[-1,1]^2$.

Proof. The first assertion is an immediate consequence of Lemma 13.2.2. Because each $E \in \mathcal{C}_p^{(2+\alpha)}$ satisfies $\text{Leb}(E) = 2+\alpha$, and because $\mathcal{I}_p(\partial E) = \mathcal{I}_p(\partial E^c)$, the second and third assertions follow. Thus, whenever $E \in \mathcal{C}_p^{(2+\alpha)}$, both E and E^c are indecomposable. By Proposition 13.1.5, either E or E^c is equivalent to $\mathsf{hull}(\lambda)$ for some rectifiable Jordan curve $\lambda \subset [-1,1]^2$, and the fourth assertion follows.

Turning our attention to the fifth assertion, it suffices to show that if $E \in \mathcal{C}_p^{(2+\alpha)}$ for $\alpha \in [-1,0]$, and if $E = \mathsf{hull}(\lambda)$ for a rectifiable Jordan curve $\lambda \subset [-1,1]^2$, then $\mathcal{H}^1(\lambda \cap \partial[-1,1]^2) > 0$. But this follows from the fact that if $\mathcal{H}^1(\lambda \cap \partial[-1,1]^2) = 0$, the curve λ at best can be the boundary of (a dilate of) the Wulff shape W_p (this is the limit shape of [BLP15] which is the unique solution, up to translation, of the unrestricted isoperimetric problem associated to the norm β_p). However, this shape is not optimal. For instance, a suitably dilated quarter-Wulff shape has strictly better conductance. □

Let us include one last result to be used in the proof of Theorem 11.3.1, and which guarantees the non-degeneracy of the limit in Theorem 11.3.3.

Lemma 13.2.5. For each $\alpha \in [-1,1]$, we have $\varphi_p^{(2+\alpha)} > 0$. Moreover, for each $\alpha, \alpha' \in [-1,1]$ with $\alpha > \alpha'$, we have the strict monotonicity $\varphi_p^{(2+\alpha')} > \varphi_p^{(2+\alpha)}$.

Proof. Strict monotonicity follows from Lemma 13.2.2. It suffices to show $\varphi_p^{(3)}$ is positive; let us see how this follows from the fifth assertion of Corollary 13.2.4. Given $R \in \mathcal{R}_p^{(2+\alpha)}$, we have from (5) that $\partial R \cap (-1,1)^2$ is a simple rectifiable curve η joining distinct points on the boundary of $\partial[-1,1]^2$. There are three short cases.

Case I: We suppose the endpoints of η lie on the same side of $\partial[-1,1]^2$. Thus, either R or R^c intersects at most one side of $[-1,1]^2$, and we let A denote the set among R and R^c with this property. By reflecting A about the side it borders, we produce a set A' of twice the volume, with $\mathcal{I}_p(\partial A') = 2\mathcal{I}_p(\eta) \equiv 2\text{length}_{\beta_p}(\eta)$. As $\text{Leb}(A) \geq 1$, we use the standard Euclidean isoperimetric inequality to deduce

$$\mathcal{I}_p(\eta) \equiv \text{length}_{\beta_p}(\eta) \geq \frac{c}{\sqrt{2}}\beta_p^{\min}, \tag{13.16}$$

where $c > 0$ is some absolute constant, and where $\beta_p^{\min}$ is the minimum of β_p over the unit circle.

Case II: In the second case, we suppose the endpoints of η lie on two adjacent sides of $\partial[-1,1]^2$. Either R or R^c intersects only these two sides of the square, and as before we let A denote the set among R and R^c with this property. We proceed as before, except we now reflect twice, obtaining A' with four times the volume of A, and with $\mathcal{I}_p(\partial A') = 4\mathcal{I}_p(\eta) \equiv 4\text{length}_{\beta_p}(\eta)$. Thus,

$$\mathcal{I}_p(\eta) \equiv \text{length}_{\beta_p}(\eta) \geq \frac{c}{2}\beta_p^{\min}, \tag{13.17}$$

with c and $\beta_p^{\min}$ as above.

Case III: In the final case, η joins points on two opposing sides of $\partial[-1,1]^2$. In this case, it is clear that $\mathcal{I}_p(\eta) \equiv \text{length}_{\beta_p}(\eta) \geq 2\beta_p^{\min}$, where the two arises as the Euclidean distance between two opposing sides of the square.

In each case, we conclude that $\mathcal{I}_p(\partial R) = \mathcal{I}_p(\eta) > 0$, which completes the proof. □

13.3 Stability for connected sets

Now that we have shown the set $\mathcal{R}_p^{(2+\alpha)}$ is non-empty, we show a stability result with respect to the d_H-metric. First, some preliminary results.

Lemma 13.3.1. Let $\alpha \in (-1,1)$. Suppose that $E_k \in \mathcal{C}$ are such that $\text{Leb}(E_k) \leq 2+\alpha$ and the conductances of the E_k tend to $\varphi_p^{(2+\alpha)}$. Then $\liminf_{k\to\infty}(E_k) > 0$, and if the $E_k \to E$ in L^1-sense, we have $E \in \mathcal{C}_p^{(2+\alpha)}$.

Proof. Let $\alpha' \in (-1,1)$ be strictly less than α. If $\text{Leb}(E_k) \to 0$, we would have $\varphi_p^{(2+\alpha')} \geq \varphi_p^{(2+\alpha)}$, which contradicts Lemma 13.2.5. Thus if the E_k tend to $E \subset [-1,1]^2$ in L^1-sense, it follows that $\mathrm{Leb}(E) > 0$. By Lemma 13.1.2, we have

$$\varphi_p^{(2+\alpha)} = \liminf_{k\to\infty} \frac{\mathcal{I}_p(E_k)}{\mathrm{Leb}(E_k)} \geq \frac{\mathcal{I}_p(E)}{\mathrm{Leb}(E)}\,, \tag{13.18}$$

and thus $E \in \mathcal{C}_p^{(2+\alpha)}$. □

For $E \in \mathcal{C}$ indecomposable, Proposition 13.1.5 tells us that E is equivalent (up to a Lebesgue-null set) to $\mathsf{hull}(\lambda) \setminus \left(\bigcup_{j\geq 1} \mathsf{hull}(\lambda_j)^\circ\right)$ for $\lambda, \lambda_j \subset [-1,1]^2$ rectifiable Jordan curves. Given $E \in \mathcal{C}$ indecomposable, define $\widehat{E} := \mathsf{hull}(\lambda)$, where λ corresponds to E as above.

The next result tells us that if a sequence E_k of indecomposable sets of finite perimeter tend to an optimal set, the size of the "holes" in these sets must tend to zero.

Lemma 13.3.2. Let $\alpha \in (-1,1)$. Let $E_k \in \mathcal{C}$ be indecomposable with $\mathrm{Leb}(E_k) \leq 2+\alpha$ for all $k \geq 1$. Suppose that the E_k tend to $E \in \mathcal{C}_p^{(2+\alpha)}$ in L^1-sense. Then as $k \to \infty$, we have

$$\mathrm{Leb}(\widehat{E}_k \setminus E_k) \to 0\,. \tag{13.19}$$

Proof. Suppose not, and let $\alpha' \in (-1,1)$ be strictly larger than α. We lose no generality supposing that $\mathrm{Leb}(\widehat{E}_k \setminus E_k) \geq \epsilon$ for all k. Moreover, by Lemma 13.3.1, we also lose no generality supposing that $\mathrm{Leb}(E_k) \geq 2+\alpha-\epsilon/2$ for all k (using the fact at each $E \in \mathcal{C}_p^{(2+\alpha)}$ satisfies $\mathrm{Leb}(E) = 2+\alpha$).

Note that the E_k^c also converge in L^1-sense to $E^c \in \mathcal{C}_p^{(2-\alpha)}$. The sets E_k^c however are not indecomposable by hypothesis: let A_k be the component of E_k^c of smallest conductance, so that the conductance of E_k^c serves as an upper bound for the conductance of A_k. But our hypotheses on the volumes of $\widehat{E}_k$ and E_k ensure that $\mathrm{Leb}(A_k) \leq 2-\alpha-\epsilon/2$, which implies that $\varphi_p^{(2-\alpha')} \leq \varphi_p^{(2-\alpha)}$, contradicting Lemma 13.2.5. □

Heuristically, the above lemma allows us to replace a sequence of sets in $\mathcal{R}$ by Jordan domains. The next result tells us that a sequence of Jordan domains converging in the correct sense to an element of $\mathcal{C}_p^{(2+\alpha)}$ has a limit in $\mathcal{R}$.

Lemma 13.3.3. Let $R_k \in \mathcal{R}$ be a sequence such that $R_k = \mathsf{hull}(\lambda_k)$ for rectifiable Jordan curves $\lambda_k \subset [-1,1]^2$, and suppose that the conductances of the R_k tend to $\varphi_p^{(2+\alpha)}$. Suppose also that $R_k \to K$ both in L^1-sense and in d_H-sense, where $K \subset [-1,1]^2$ is compact and $K \in \mathcal{C}_p^{(2+\alpha)}$. Then $K \in \mathcal{R}_p^{(2+\alpha)}$.

Proof. In this proof, we carefully distinguish curves (continuous functions from $[0,1]$ into $[-1,1]^2$ taking the same value at 0 and 1) from their images. Given a curve $\lambda : [0,1] \to [-1,1]^2$, let $\mathsf{image}(\lambda)$ denote the image of λ. As $K \in \mathcal{C}_p^{(2+\alpha)}$, the perimeters of the ∂R_k are uniformly bounded. By appealing to an arc length parametrization of each λ_k, we may assume each λ_k is a Lipschitz function from $[0,1]$ to $[-1,1]^2$ with a uniform bound on the Lipschitz constant across all k. Invoking Arzela-Ascoli and passing to a subsequence, we find that the λ_k tend uniformly to a rectifiable curve λ.

By appealing to the definition of the hull of a curve (using winding number), we find that $\mathsf{hull}(\lambda_k) \to \mathsf{hull}(\lambda)$ in d_H-sense, thus $K \equiv \mathsf{hull}(\lambda)$. Let $\widetilde{\lambda} : [0,1] \to (-1,1)^2$ be a reparametrization of λ of constant speed, so that $K = \mathsf{hull}(\widetilde{\lambda})$ also. Suppose that $\widetilde{\lambda}$ is not a simple curve, and moreover suppose there is $x \in (-1,1)^2$ such that $|\widetilde{\lambda}^{-1}(x)| > 1$. Let $s < t \in [0,1]$ be such that $\widetilde{\lambda}(s) = \widetilde{\lambda}(t)$. Let us write $\zeta_1 := \widetilde{\lambda}|_{[s,t)}$ and $\zeta_2 := \widetilde{\lambda}|_{[0,s]\cup(t,1]}$, so that both ζ_1 and ζ_2 are closed curves.

As $K \in \mathcal{C}_p^{(2+\alpha)}$, the set K must be indecomposable with indecomposable complement. It follows that $\mathsf{hull}(\widetilde{\lambda})^\circ$ is either $\mathsf{hull}(\zeta_1)^\circ$ or $\mathsf{hull}(\zeta_2)^\circ$. As $x \in (-1,1)^2$, we also have that $\mathcal{I}_p(\widetilde{\lambda}) > \mathcal{I}_p(\zeta_1)$ and $\mathcal{I}_p(\widetilde{\lambda}) > \mathcal{I}_p(\zeta_2)$. Without loss of generality then, we have

$$\frac{\mathcal{I}_p(\partial K)}{\mathrm{Leb}(K)} \le \frac{\mathcal{I}_p(\zeta_1)}{\mathrm{Leb}(K)} < \frac{\mathcal{I}_p(\widetilde{\lambda})}{\mathrm{Leb}(K)} \le \varphi_p^{(2+\alpha)}\,, \tag{13.20}$$

where the right-most inequality follows from lower semicontinuity of the surface energy Lemma 13.1.2 (and the hypothesis that the conductances of the R_k tend to the optimal value). This is a contradiction. Thus, if $|\widetilde{\lambda}^{-1}(x)| > 1$, it must be that $x \in \partial[-1,1]^2$, and there exists a Jordan curve $\lambda' \subset [-1,1]^2$ such that $\mathsf{hull}(\lambda') = \mathsf{hull}(\widetilde{\lambda}) = K$. We conclude that $K \in \mathcal{R}_p^{(2+\alpha)}$. □

Lemma 13.3.3 essentially allows us to recover some regularity of a suitable limit of Jordan domains. We now use this to show that the collections $\mathcal{R}_p^{(2)}$ and $\mathcal{R}_p^{(2+\alpha)}$ are close when α is small.

Lemma 13.3.4. Let $\alpha \in (0,1]$. As $\alpha \to 0$, we have $\mathrm{d}_H(\mathcal{R}_p^{(2+\alpha)}, \mathcal{R}_p^{(2)}) \to 0$.

Proof. Let $\alpha_k \in (0,1]$ be a sequence tending to zero as $k \to \infty$. Let $R_k \in \mathcal{R}_p^{(2+\alpha_k)}$. By Corollary 13.2.4 (5), there are rectifiable Jordan curves $\lambda_k \subset [-1,1]^2$ with $R_k = \mathsf{hull}(\lambda_k)$. By Corollary 13.2.4 (3), the conductances of the R_k tend to $\varphi_p^{(2)}$.

The non-empty compact subsets of $[-1,1]^2$ form a compact metric space when equipped with the d_H-metric. We pass to a subsequence (twice, using Rellich-Kondrachov) so that we lose no generality supposing $R_k \to K$ in d_H-sense and in L^1-sense, where $K \subset [-1,1]^2$ is compact. As $\mathrm{Leb}(E_k) \to 2$ as $k \to \infty$, the lower semicontinuity of the surface energy (Lemma 13.1.2) implies $K \in \mathcal{C}_p^{(2+\alpha)}$. We apply Lemma 13.3.3 to conclude that $K \in \mathcal{R}_p^{(2+\alpha)}$ to complete the proof. □

The following is the first of two stability results, and is a precursor to the main result in this section.

Proposition 13.3.5. Let $\alpha \in (-1,1)$ and let $\epsilon > 0$. There is $\delta = \delta(\alpha, \epsilon) > 0$ so that whenever $R \in \mathcal{R}$ is connected with $\mathrm{Leb}(R) \le 2 + \alpha$ and $\mathrm{d}_H(R, \mathcal{R}_p^{(2+\alpha)}) > \epsilon$, we have

$$\frac{\mathcal{I}_p(\partial R)}{\mathrm{Leb}(R)} \ge \varphi_p^{(2+\alpha)} + \delta \tag{13.21}$$

Proof. Suppose not. Then there is a sequence $R_k \in \mathcal{R}$ of connected sets with $\mathrm{Leb}(R_k) \le 2+\alpha$, and with $\mathrm{d}_H(R_k, \mathcal{R}_p^{(2+\alpha)}) > \epsilon$ and

$$\frac{\mathcal{I}_p(\partial R_k)}{\mathrm{Leb}(R_k)} \to \varphi_p^{(2+\alpha)}\,. \tag{13.22}$$

Suppose first that for each k, $R_k = \mathsf{hull}(\lambda_k)$, where $\lambda_k \subset [-1,1]^2$ is a rectifiable Jordan curve. By Rellich-Kondrachov, and by the compactness of the set of non-empty compact subsets of $[-1,1]^2$ in the metric d_H, we lose no generality (by passing to a subsequence) supposing that $R_k \to K \subset [-1,1]^2$ compact, where the convergence takes place both in L^1-sense and in d_H-sense. By Lemma 13.3.1, it follows that $K \in \mathcal{C}_p^{(2+\alpha)}$, and by Lemma 13.3.3, it then follows that $K \in \mathcal{R}_p^{(2+\alpha)}$, which is a contradiction.

Let us then suppose that none of the R_k are of the form $\mathsf{hull}(\lambda_k)$ for a sequence of rectifiable Jordan curves $\lambda_k \subset [-1,1]^2$, so that for each k, we have $\widehat{R}_k \neq R_k$. We appeal to the same compactness argument as above, and suppose that the R_k tend to $K \subset [-1,1]^2$ compact both in L^1-sense and in d_H-sense. As before, Lemma 13.3.1 tells us $K \in \mathcal{C}_p^{(2+\alpha)}$. We then use Lemma 13.3.2 to deduce that $\mathrm{Leb}(\widehat{R}_k \setminus R_k) \to 0$.

As the conductances of the R_k tend to $\varphi_p^{(2+\alpha)}$, and as $\varphi_p^{(2+\alpha+\epsilon)} \to \varphi_p^{(2+\alpha)}$ as $\epsilon \to 0$, it follows that the diameter of any connected component of $\widehat{R}_k \setminus R_k$ must also tend to zero. Thus, as $k \to \infty$, we have that $\mathrm{d}_H(\widehat{R}_k, R_k) \to 0$, and we may then realize $K \in \mathcal{C}_p^{(2+\alpha)}$ as the L^1- and d_H-limit of the $\widehat{R}_k$. As each $\widehat{R}_k$ is the hull of a rectifiable Jordan curve, we may now use Lemma 13.3.3 to deduce that $K \in \mathcal{R}_p^{(2+\alpha)}$, which is again a contradiction. □

Our second stability result upgrades Proposition 13.3.5, removing the α dependence of the constant δ. It is the main result of this section and is instrumental to the proof of Theorem 11.3.1.

Corollary 13.3.6. Let $\alpha \in (0,1]$ and let $\epsilon > 0$. There is $\delta = \delta(\epsilon) > 0$ so that whenever $R \in \mathcal{R}$ is connected with $\mathrm{Leb}(R) \leq 2 + \alpha$ and $\mathrm{d}_H(R, \mathcal{R}_p^{(2+\alpha)}) > \epsilon$, we have

$$\frac{\mathcal{I}_p(\partial R)}{\mathrm{Leb}(R)} \geq \varphi_p^{(2+\alpha)} + \delta \tag{13.23}$$

Proof. Let $\epsilon > 0$ and let α_k be a sequence in $(0,1]$ tending to zero as $k \to \infty$. Let $\tilde{\delta}(\alpha_k, \epsilon)$ be the supremum of all $\delta > 0$ for which Proposition 13.3.5 is valid for the parameters α_k and ϵ. Then, for each k, there are connected sets $R_k \in \mathcal{R}$ with $\mathrm{Leb}(R_k) \leq 2 + \alpha_k$ so that $\mathrm{d}_H(R_k, \mathcal{R}_p^{(2+\alpha_k)}) \geq \epsilon$ and

$$\frac{\mathcal{I}_p(\partial R_k)}{\mathrm{Leb}(R_k)} \leq \varphi_p^{(2+\alpha_k)} + 2\tilde{\delta}(\alpha_k, \epsilon)\,. \tag{13.24}$$

Suppose for the sake of contradiction that $\tilde{\delta}(\alpha_k, \epsilon) \to 0$ as $k \to \infty$. Then the conductances of the R_k tend to $\varphi_p^{(2)}$. Passing to a subsequence, we may assume that $R_k \to K$ compact with $K \in \mathcal{C}_p^{(2+\alpha)}$, where the convergence takes place both in L^1-sense and in d_H-sense. If each R_k is the hull of a rectifiable Jordan curve, we may invoke Lemma 13.3.3 to deduce that $K \in \mathcal{R}_p^{(2+\alpha)}$. If not, we may proceed as in the proof of Proposition 13.3.5, replacing each R_k by $\widehat{R}_k$ to deduce the same result.

Thus, the R_k get arbitrarily close in d_H-sense to $\mathcal{R}_p^{(2)}$, so that for all k sufficiently large, $\mathrm{d}_H(R_k, \mathcal{R}_p^{(2)}) \leq \epsilon/4$. Thanks to Lemma 13.3.4, we may also find k sufficiently large so that $\mathrm{d}_H(\mathcal{R}_p^{(2+\alpha_k)}, \mathcal{R}_p^{(2)}) < \epsilon/4$. This contradicts the fact that $\mathrm{d}_H(E_k, \mathcal{R}_p^{(2+\alpha_k)}) > \epsilon$ $\square$

CHAPTER 14

Continuous to discrete: upper bounds

In this chapter, we show that given $R \in \mathcal{R}$ with $\mathrm{Leb}(R) \leq 2$, there are high probability upper bounds on $n\widehat{\Phi}_n$ in terms of the conductance of R. We show first this for polygons and then use approximation to pass to more general sets.

14.1 From simple polygons to discrete sets

A *convex polygon* in $\mathbb{R}^2$ is a compact subset of $\mathbb{R}^2$ having non-empty interior which may be written as the intersection of finitely many closed half-spaces. A *polygon* is any subset of $\mathbb{R}^2$ which may be written as a finite union of convex polygons.

Recall (from the statement of Proposition 12.3.4) that given $x, y \in \mathbb{R}^2$, we use $\mathsf{poly}(x, y)$ to denote the linear segment joining x and y. Given a sequence of points $x_1, \ldots, x_m$, we define

$$\mathsf{poly}(x_1, \ldots, x_m) := \mathsf{poly}(x_1, x_2) * \cdots * \mathsf{poly}(x_{m-1}, x_m)\,, \tag{14.1}$$

where " $*$ " denotes concatenation of these curves. A *polygonal curve* is any curve of the form (14.1) for some $x_1, \ldots, x_m \in \mathbb{R}^2$ and some $m \in \mathbb{N}$ (we return to being vague about the parametrization). Polygons may be defined from polygonal curves in a natural way; we say a polygon is *simple* if it may be written as the hull of a simple polygonal circuit. The first proposition of this section associates a discrete set to any simple polygon in a convenient way.

Remark 14.1.1. In this chapter and the next we will be somewhat cavalier with notation. In particular, for $R \in \mathcal{R}$, the dilated set nR is not in general contained in $[-1, 1]^2$. The surface energy of nR, denoted $\mathcal{I}_p(n\partial R)$ is defined to be $n\mathcal{I}_p(\partial R)$. We employ a similar convention for curves.

Proposition 14.1.2. Let $p > p_c(2)$ and let $\epsilon > 0$. Let $P \subset [-1,1]^2$ be a simple non-degenerate polygon. There are positive constants $c_1(p, P, \epsilon)$ and $c_2(p, P, \epsilon)$ so that for all $n \geq 1$, with probability at least $1 - c_1 \exp(-c_2 \log^2 n)$, there is a rectifiable circuit $\lambda \equiv \lambda(P) \subset [-1,1]^2$ so that

1. $\mathrm{d}_H(n\partial P, n\lambda) \leq \epsilon n$,
2. $\mathcal{I}_p(n\partial P) \geq (1-\epsilon)|\partial^n[\mathsf{hull}(n\lambda) \cap \mathbf{C}_\infty]|$.

Proof. ***Step I: (Aggregation of high probability events)*** Let $x_1, \dots, x_m$ be the corners of nP, so that

$$nP = \mathsf{hull}(\mathsf{poly}(x_1, \dots, x_m)) \,, \tag{14.2}$$

where $x_m \equiv x_1$, and where the circuit $\mathsf{poly}(x_1, \dots, x_m)$ is oriented counter-clockwise. Let $\mathcal{E}_1$ be the high probability event from Lemma 12.1.6 that for each $i \in \{1, \dots, m\}$, we have $|[x_i] - x_i|_2 \leq \log^2 n$. Say x_i is an *interior point* if $x_i \in (-n,n)^2$, and that it is a *boundary point* otherwise. For n sufficiently large, the Euclidean ball $B_{2\log^2 n}(x_i)$ is contained in $(-n,n)^2$ for each interior point x_i. For such n and within $\mathcal{E}_1$, we have $[x_i] \in (-n,n)^2$ for each interior x_i.

For $\delta > 0$, define the high probability event $\mathcal{E}_2(\delta)$ via

$$\mathcal{E}_2(\delta) := \bigcap_{i=1}^{m-1} \Big\{ \exists \gamma \in \Gamma_\delta(x_i, x_{i+1}) \ : \mathrm{d}_H\Big(\gamma, \mathsf{poly}(x_i, x_{i+1})\Big) \leq \delta |x_{i+1} - x_i|_2 \Big\} \,, \tag{14.3}$$

so that $\mathcal{E}_2(\delta)^c$ is subject to the bounds in Proposition 12.3.4. Additionally, define

$$\mathcal{E}_3(\delta) := \bigcap_{i=1}^{m-1} \left\{ \left| \frac{b([x_i],[x_{i+1}])}{\beta_p(x_{i+1} - x_i)} - 1 \right| > \delta \right\} \,, \tag{14.4}$$

so that $\mathcal{E}_3(\delta)^c$ is subject to the bounds in Theorem 12.3.3. For the remainder of the proof, work within the intersection $\mathcal{E}_1 \cap \mathcal{E}_2(\delta) \cap \mathcal{E}_3(\delta)$.

Step II: (Constructing λ) Select $\gamma_i \in \Gamma_\delta(x_i, x_{i+1})$ with $\mathrm{d}_H(\gamma_i, \mathsf{poly}(x_i, x_{i+1})) < \delta|x_{i+1} - x_i|_2$ for each $i \in \{1, \ldots, m-1\}$. Each γ_i may be identified with an interface ∂_i via the correspondence in Proposition 12.2.1.

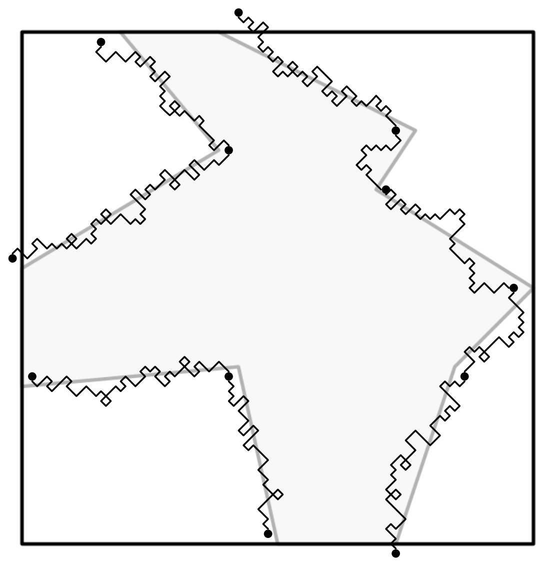

Figure 14.1: The polygon nP is in grey. The black dots are the $[x_i]$, and the contours joining these dots are the $\partial_i \equiv \lambda_i$ corresponding to the interior segments $\mathsf{poly}(x_i, x_{i+1})$.

A linear segment $\mathsf{poly}(x_i, x_{i+1})$ is an *interior segment* if at least one of x_i or x_{i+1} is an interior point, and otherwise it is a *boundary segment*. If $\mathsf{poly}(x_i, x_{i+1})$ is a boundary segment, set $\lambda_i := \mathsf{poly}(x_i, x_{i+1})$, otherwise, via "corner-rounding" (see Remark 12.2.2), regard ∂_i as a simple curve and set $\lambda_i := \partial_i$. If the endpoint of λ_i is not equal to the starting point of λ_{i+1}, let $\widetilde{\lambda}_i$ be the linear segment joining these two points. If λ_i ends at the starting point of λ_{i+1}, let $\widetilde{\lambda}_i$ be the degenerate linear segment at this endpoint. Define the circuit $n\lambda$ as the concatenation of these curves in the proper order:

$$n\lambda := \lambda_1 * \widetilde{\lambda}_1 * \lambda_2 * \widetilde{\lambda}_2 * \cdots * \lambda_m * \widetilde{\lambda}_m \,, \tag{14.5}$$

and write H_n for $\mathsf{hull}(n\lambda) \cap \mathbf{C}_n$. Let E_i be the set of all edges of $\mathbb{Z}^2$ contained in the Euclidean ball $B_{2\log^2 n}(x_i)$, so that by construction of H_n,

$$|\partial^n H_n| \leq \sum_{\substack{i\,:\,\mathsf{poly}(x_i, x_{i+1}) \\ \text{is interior}}} |\mathfrak{b}(\gamma_i)| + \sum_{i=1}^{m} |E_i| \,. \tag{14.6}$$

Step III: (Controlling $|\partial^n H_n|$***)*** We build off (14.6) and use that each γ_i is δ-optimal (see (12.11)),

$$|\partial^n H_n| \le \sum_{\substack{i\,:\,\mathsf{poly}(x_i,x_{i+1})\\ \text{is interior}}} \Big(b([x_i],[x_{i+1}]) + \delta |x_{i+1}-x_i|_2 \Big) + \sum_{i=1}^m |E_i|\,, \tag{14.7}$$

$$\le \left(\sum_{\substack{i\,:\,\mathsf{poly}(x_i,x_{i+1})\\ \text{is interior}}} b([x_i],[x_{i+1}]) \right) + 2mn\delta + C\log^4 n\,, \tag{14.8}$$

for some absolute positive constant C. As we are within $\mathcal{E}_2(\delta)$, for n sufficiently large we have

$$|\partial^n H_n| \le \left(\sum_{\substack{i\,:\,\mathsf{poly}(x_i,x_{i+1})\\ \text{is interior}}} (\beta_p(x_{i+1}-x_i) + n\delta) \right) + 4mn\delta\,, \tag{14.9}$$

$$\le \mathcal{I}_p(n\partial P) + 8mn\delta\,. \tag{14.10}$$

Step IV: (Wrapping up) Given $\epsilon > 0$, we may choose δ sufficiently small depending on P and ϵ so that from (14.10), we have

$$\mathcal{I}_p(n\partial P) \ge (1-\epsilon)|\partial^n H_n|\,. \tag{14.11}$$

Finally, the construction of λ from the γ_i ensures that

$$\mathrm{d}_H(nP, n\lambda) \le 2\delta \max_{i=1}^{m-1} |x_{i+1}-x_i|_2\,, \tag{14.12}$$

and we take δ smaller if necessary to complete the proof. □

14.2 Upper bounds on $n\widehat{\Phi}_n$ using connected polygons

We now use the output of Proposition 14.1.2 to construct a discrete approximate to more general connected polygons. We also relate the volume of the discrete approximate to the volume of this polygon.

Proposition 14.2.1. Let $p > p_c(2)$ and let $\epsilon > 0$. Let $P \subset [-1,1]^2$ be a connected polygon whose boundary consists of finitely many disjoint simple polygonal circuits. There are positive constants $c_1(p,P,\epsilon)$ and $c_2(p,P,\epsilon)$ so that for all $n \geq 1$, with probability at least $1 - c_1 \exp(-c_2 \log^2 n)$, there is subgraph $H_n \equiv H_n(P) \subset \mathbf{C}_n$ so that

1. $|\theta_p \mathrm{Leb}(nP) - |H_n|| \leq \epsilon \mathrm{Leb}(nP)$,
2. $\mathcal{I}_p(n\partial P) \geq (1-\epsilon)|\partial^n H_n|$.

Proof. ***Step I: (Using circuits to identify*** H_n***)*** Using the hypotheses on P, identify disjoint, simple polygonal circuits $\rho, \rho_1, \dots, \rho_m$ so that $\partial P = \rho \sqcup \bigsqcup_{i=1}^m \rho_m$, and so that

$$P = \mathsf{hull}(\rho) \setminus \left(\bigsqcup_{i=1}^m \mathsf{hull}(\rho_i)^\circ \right), \tag{14.13}$$

where the $\mathsf{hull}(\rho_i)$ are pairwise disjoint, and where each is a simple polygon. For $\delta > 0$, work within the high probability events given by Proposition 14.1.2 that there are circuits $\lambda, \lambda_1, \dots, \lambda_m \subset [-1,1]^2$ so that

1. $\mathrm{d}_H(n\rho_i, n\lambda_i) \leq \delta n$ for each i, and $\mathrm{d}_H(n\rho, n\lambda) \leq \delta n$.
2. $\mathcal{I}_p(n\rho_i) \geq (1-\delta)|\partial^n[\mathsf{hull}(n\lambda_i) \cap \mathbf{C}_\infty]|$ for each i, and $\mathcal{I}_p(n\rho) \geq (1-\delta)|\partial^n[\mathsf{hull}(n\lambda) \cap \mathbf{C}_\infty]|$

Define the set

$$R := \mathsf{hull}(\lambda) \setminus \left(\bigsqcup_{i=1}^m \mathsf{hull}(\lambda_i)^\circ \right), \tag{14.14}$$

and let $H_n := nR \cap \mathbf{C}_n$. By (2), the graph H_n has the second desired property:

$$\mathcal{I}_p(n\partial P) \geq (1-\delta)|\partial^n H_n|\,. \tag{14.15}$$

Step II: (Controlling the volume of H_n ***from above)*** We control the volume of H_n by appealing to Proposition 16.0.3. Let $k \in \mathbb{N}$ and let S_k denote the set of half-open dyadic squares at the scale k which are contained in $[-1,1]^2$; these are translates of $[-2^{-k}, 2^{-k})^2$. For $\delta' > 0$ and $S \in \mathsf{S}_k$, define the event

$$\mathcal{E}_S(\delta') := \left\{ \frac{|\mathbf{C}_\infty \cap nS|}{\mathrm{Leb}(nS)} \in \Big((1-\delta')\theta_p, (1+\delta')\theta_p \Big) \right\}, \tag{14.16}$$

and let $\mathcal{E}_{\mathrm{vol}}(\delta')$ be the intersection of $\mathcal{E}_S(\delta')$ over all $S \in \mathsf{S}_k$. From now on, work within the event $\mathcal{E}_{\mathrm{vol}}(\delta')$. Let $N_{2\delta}$ be the closed 2δ-neighborhood (with respect to Euclidean distance) of ∂P. Let S_k^- be the squares of S_k contained in $P \setminus N_{2\delta}$, and let S_k^+ be the squares of S_k having non-empty intersection with $P \cup N_{2\delta}$. Here we assume δ is small enough and k is large enough for S_k^- to be non-empty. Thanks to the construction of H_n, we have

$$|H_n| \leq \sum_{S \in \mathsf{S}_k^+} |nS \cap \mathbf{C}_\infty| + Cn\,, \tag{14.17}$$

where C is some absolute constant, and the term Cn directly above accounts for the vertices of $\mathbb{Z}^2$ in $\partial[-n,n]^2$, which we must be mindful of as the squares $S \in \mathsf{S}_k$ are half-open. Choose k large enough depending on δ' and P so that

$$(1-\delta')\mathrm{Leb}(P) \leq \sum_{S \in \mathsf{S}_k^-} \mathrm{Leb}(S) \leq \sum_{S \in \mathsf{S}_k^+} \mathrm{Leb}(S) \leq (1+\delta')\mathrm{Leb}(P)\,. \tag{14.18}$$

For n sufficiently large, it follows from (14.17), (14.18), the fact that we are working within $\mathcal{E}_{\mathrm{vol}}(\delta')$ that

$$|H_n| \leq (1+2\delta')^2 \theta_p \mathrm{Leb}(nP)\,. \tag{14.19}$$

Step III: (Controlling the volume of H_n from below) Work within the following high probability event from Proposition 16.0.4 for the remainder of the proof:

$$\Big\{\mathbf{C}_\infty \cap [-n+\log^2 n, n-\log^2 n] = \mathbf{C}_n \cap [-n+\log^2 n, n-\log^2 n]\Big\}\,. \tag{14.20}$$

We appeal to the construction of H_n and the disjointness of the squares in S_k, taking n sufficiently large to obtain the second line below:

$$|H_n| \geq \left(\sum_{S \in \mathsf{S}_k^-} |\mathbf{C}_\infty \cap nQ_j| \right) - |\mathbf{C}_\infty \cap [-n,n]^2 \setminus \mathbf{C}_n|\,. \tag{14.21}$$

$$\geq (1-2\delta') \sum_{S \in \mathsf{S}_k^-} \theta_p \mathrm{Leb}(nS)\,, \tag{14.22}$$

$$\geq (1-2\delta')(1-\delta')\theta_p \mathrm{Leb}(nP)\,. \tag{14.23}$$

where the last line follows from (14.18). We choose δ, δ' sufficiently small to complete the proof. □

We now use Proposition 14.2.1 to obtain high probability upper bounds on $\widehat{\Phi}_n$ in terms of the conductance of a connected, non-degenerate polygon which is not too large.

Corollary 14.2.2. Let $p > p_c(2)$ and let $\epsilon > 0$. Let $P \subset [-1,1]^2$ be a connected polygon with $\mathrm{Leb}(P) < 2$, and whose boundary is a finite disjoint union of simple polygonal curves. There are positive constants $c_1(p, P, \epsilon)$ and $c_2(p, P, \epsilon)$ so that for all $n \geq 1$, with probability at least $1 - c_1 \exp(-c_2 \log^2 n)$,

$$n\widehat{\Phi}_n \leq (1+\epsilon)\frac{\mathcal{I}_p(\partial P)}{\theta_p \mathrm{Leb}(P)}\,. \tag{14.24}$$

Proof. Define $\epsilon' := 2 - \mathrm{Leb}(P)$ and let $\delta > 0$. By combining Proposition 16.0.3 with Proposition 16.0.4, we obtain positive constants $c_1(p,\delta)$ and $c_2(p,\delta)$ so that the probability of the event

$$\left\{ \frac{|\mathbf{C}_n|}{(2n)^2} \in \Big((1-\delta)\theta_p, (1+\delta)\theta_p \Big) \right\} \tag{14.25}$$

is at least $1 - c_1 \exp(-c_2 \log^2 n)$. Work within this high probability event, and additionally work within the high probability event from Proposition 14.2.1 that there is $H_n \subset \mathbf{C}_n$ satisfying

1. $|\theta_p \mathrm{Leb}(nP) - |H_n|| \leq \delta \mathrm{Leb}(nP)$,

2. $\mathcal{I}_p(n\partial P) \geq (1-\delta)|\partial^n H_n|$.

Thus, $|H_n| \leq (\theta_p + \delta)(2-\epsilon')n^2$. Using (14.25) and choosing δ small enough depending on ϵ' so that $2(1-\delta)\theta_p \geq (\theta_p+\delta)(2-\epsilon')$, we find $|H_n| \leq |\mathbf{C}_n|/2$, and conclude that with high probability,

$$\widehat{\Phi}_n \leq \frac{|\partial^n H_n|}{|H_n|} \leq \frac{\frac{1}{1-\delta}\mathcal{I}_p(nP)}{(\theta_p - \delta)\mathrm{Leb}(nP)}\,, \tag{14.26}$$

which completes the proof, taking δ smaller if necessary. □

14.3 The optimal upper bound on $n\widehat{\Phi}_n$

We now exhibit a high probability upper bound on $n\widehat{\Phi}_n$ using the optimal conductance of φ_p defined in (11.10). We introduce results which allow us to approximate rectifiable Jordan curves by simple polygonal circuits. The following consolidates Lemma 4.3 and Lemma 4.4 of [BLP15].

Proposition 14.3.1. Let λ be a rectifiable curve in $\mathbb{R}^2$ starting at x and ending at y. Let $\epsilon > 0$. There is a simple polygonal curve ρ starting at x and ending at y such that (1) and (2) hold:

1. $\mathrm{d}_H(\lambda, \rho) \leq \epsilon\,,$
2. $\mathrm{length}_{\beta_p}(\lambda) + \epsilon \geq \mathrm{length}_{\beta_p}(\rho)\,.$

Furthermore, if λ is a closed curve (i.e. $x = y$), ρ can additionally be taken to satisfy (3):

3. $\mathrm{Leb}(\mathsf{hull}(\lambda)\,\Delta\,\mathsf{hull}(\rho)) \leq \epsilon\,.$

Remark 14.3.2. We remark that, in Proposition 14.3.1, if the curve λ is contained in $[-1,1]^2$, one can easily arrange that the polygonal approximate ρ is also contained in $[-1,1]^2$.

The following is a nearly immediate consequence Proposition 14.3.1, so we omit the proof.

Corollary 14.3.3. Let $\lambda \subset [-1,1]^2$ be a rectifiable Jordan curve such that $\lambda = \lambda_1 * \lambda_2$, where λ_1 and λ_2 are simple curves with $\lambda_1 \subset \partial[-1,1]^2$, and such that every point on the curve λ_2 except the endpoints lies in $(-1,1)^2$. Let $\epsilon > 0$. There is a simple polygonal circuit $\rho \subset [-1,1]^2$ so that

1. $\mathrm{d}_H(\lambda, \rho) \leq \epsilon\,,$
2. $\mathcal{I}_p(\lambda) + \epsilon \geq \mathcal{I}_p(\rho)\,,$
3. $\mathrm{Leb}(\mathsf{hull}(\lambda)\,\Delta\,\mathsf{hull}(\rho)) \leq \epsilon\,.$

Remark 14.3.4. If instead of a decomposition of λ into two curves as in Corollary 14.3.3, we express λ as a concatenation of finitely many curves, each having the properties of λ_1 or λ_2, the conclusion of Corollary 14.3.3 still holds. That is, for such λ, we may find a polygonal circuit ρ for which (1) – (3) hold.

We are now equipped to prove Theorem 14.3.5, which is the main theorem of the Chapter.

Theorem 14.3.5. There are positive constants $c_1(p, \epsilon)$ and $c_2(p, \epsilon)$ so that for all $n \geq 1$, with probability at least $1 - c_1 \exp(-c_2 \log^2 n)$,

$$n\widehat{\Phi}_n \leq (1 + \epsilon)\varphi_p \,, \tag{14.27}$$

where φ_p is defined in (11.10).

Proof. Let $R \in \mathcal{R}_p$. By Corollary 13.2.4, we lose no generality taking $R = \mathsf{hull}(\lambda)$, with λ as in the statement of Corollary 14.3.3. For $\delta > 0$, there is a simple polygonal circuit $\rho \subset [-1, 1]^2$ so that

1. $\mathrm{d}_H(\lambda, \rho) \leq \delta$,
2. $\mathcal{I}_p(\lambda) + \delta \geq \mathcal{I}_p(\rho)$,
3. $\mathrm{Leb}(\mathsf{hull}(\lambda) \,\Delta\, \mathsf{hull}(\rho)) \leq \delta$.

As $R = \mathsf{hull}(\lambda)$ has positive measure, there is $s > 0$ and a square of side-length S which is contained in the interior R. For δ sufficiently small, S is also contained in the interior of $\mathsf{hull}(\rho)$. Let $P_s := \mathsf{hull}(\rho) \setminus S^\circ$, and observe that P_s is a connected polygon satisfying the hypotheses of Corollary 14.2.2, as well as

4. $2 - \delta - s^2 \leq \mathrm{Leb}(P_s) \leq 2 + \delta - s^2$,
5. $\mathcal{I}_p(\partial R) + \delta + 4s\beta_p^{\max} \geq \mathcal{I}_p(\partial P_s)$,

where $\beta_p^{\max}$ is the maximum of β_p over the unit circle. By taking δ smaller if necessary so that $s^2 > 2\delta$, we find $\mathrm{Leb}(P_s) < 2$. Thus, by Corollary 14.2.2, with high probability

$$n\widehat{\Phi}_n \leq (1+\delta)\frac{\mathcal{I}_p(\partial P_s)}{\theta_p \mathrm{Leb}(P_s)}\,, \tag{14.28}$$

$$\leq (1+\delta)\frac{\mathcal{I}_p(\partial R) + \delta + 4\beta_p^{\max} s}{\theta_p(\mathrm{Leb}(R) - \delta - s^2)}\,, \tag{14.29}$$

where we have used (4) and (5). The proof is complete upon adjusting δ and s. □

CHAPTER 15

Discrete to continuous objects: lower bounds

We construct tools which allow us to pass from a subgraph of $\mathbf{C}_n$ to a connected polygon of comparable conductance. By Lemma 12.2.4, the boundary of a subgraph of $\mathbf{C}_n$ may be thought of as a finite collection of open right-most circuits. Our first goal is then to construct an approximating polygonal curve to any open right-most path.

15.1 Extracting polygonal curves from right-most paths

Our first result enables us to pass from open right-most paths of sufficient length to polygonal curves.

Lemma 15.1.1. Let $p > p_c(2)$ and let $\epsilon > 0$. There are positive constants $c_1(p, \epsilon)$ and $c_2(p, \epsilon)$ so that for all $n \geq 1$, with probability at least $1 - c_1 \exp(-c_2 \log^2 n)$, whenever $\gamma \subset [-n, n]^2$ is an open right-most path with $|\gamma| \geq n^{1/32}$, there is a simple polygonal curve $\rho = \rho(\gamma) \subset [-1, 1]^2$ with

1. $\mathrm{d}_H(\gamma, n\rho) \leq n^{1/64}$,
2. $|\mathfrak{b}(\gamma)| \geq (1 - \epsilon)\mathrm{length}_{\beta_p}(n\rho)$.

Proof. For $x, y \in [-n, n]^2 \cap \mathbb{Z}^2$ and $\epsilon > 0$, let $\mathcal{E}_{x,y}$ be the event

$$\mathcal{E}_{x,y} := \left\{ \left| \frac{b([x], [y])}{\beta_p(y - x)} - 1 \right| \leq \epsilon \right\} . \tag{15.1}$$

Let $\mathcal{E}$ be the intersetion of all $\mathcal{E}_{x,y}$ over pairs $x, y \in [-n, n]^2 \cap \mathbb{Z}^2$ satisfying $|x - y|_2 \geq n^{1/1024}$, and work within $\mathcal{E}$ for the remainder of the proof. By Theorem 12.3.3 and a union bound, there are positive constants $c_1(p,\epsilon)$ and $c_2(p,\epsilon)$ so that

$$\mathbb{P}_p(\mathcal{E}^c) \leq c_1 \exp\left(-c_2 \log^2 n\right) . \tag{15.2}$$

Step I: (Constructing a polygonal curve) Consider an open right-most path $\gamma \subset [-n,n]^2$ with $|\gamma| \geq n^{1/32}$ and express γ as an alternating sequence of vertices and edges:

$$\gamma = (x_0, e_1, x_1, \ldots, e_m, x_m) . \tag{15.3}$$

Define a subsequence of the vertices x_i as follows: let ℓ be the largest positive integer such that $(\ell-1)\lceil n^{1/256}\rceil \leq m$, and for $k \in \{0,\ldots,\ell-1\}$, set

$$y_k := x_{k\lceil n^{1/256}\rceil} \tag{15.4}$$

and set $y_\ell := x_m$. Because γ is right-most, no vertex x_j in (15.3) appears more than four times. Thus for n sufficiently large, $|y_{k+1} - y_k|_2 \geq n^{1/1024}$ for all $k \in \{0,\ldots,\ell-2\}$. Let $\rho' \subset [-1,1]^2$ be the polygonal curve defined by

$$n\rho' := \mathsf{poly}(y_0,y_1) * \mathsf{poly}(y_1,y_2) * \cdots * \mathsf{poly}(y_{\ell-1},y_\ell) . \tag{15.5}$$

We check that ρ' has the desired properties and finish the proof by perturbing ρ' to a simple polygonal curve for which these properties still hold.

Step II: (Controlling the β_p-length of ρ' from above) As $m = |\gamma| \geq n^{1/32}$, it follows that $\ell \geq \frac{1}{2} n^{(1/32)-(1/256)}$. Because $|y_{k+1} - y_k|_2 \geq n^{1/1024}$ for $k \in \{0,\ldots,\ell-2\}$, we deduce

$$\mathrm{length}_{\beta_p}(n\rho') \geq c(p) n^{29/1024} \tag{15.6}$$

for some positive constant $c(p)$. Because we are within $\mathcal{E}$, and because each γ_k is open and right-most,

$$|\mathfrak{b}(\gamma)| \geq \sum_{k=0}^{\ell-1} b(y_k, y_{k+1}) , \tag{15.7}$$

$$\geq (1-\epsilon)\mathrm{length}_{\beta_p}(n\rho') - \beta_p(y_\ell - y_{\ell-1}) . \tag{15.8}$$

As $\beta_p(y_\ell - y_{\ell-1}) \leq C(p)^{1/256}$ for some positive constant $C(p)$, by taking n sufficiently large and using (15.6), we find

$$|\mathfrak{b}(\gamma)| \geq (1-2\epsilon)\text{length}_{\beta_p}(n\rho')\,. \tag{15.9}$$

Step III: (d_H-closeness of $n\rho'$ and γ) For $k \in \{0, \dots, \ell-1\}$, let γ_k be the subpath of γ starting at y_k and ending at y_{k+1}. Observe that every vertex in γ_k has ℓ^∞-distance at most $2\lceil n^{1/256}\rceil$ from the starting point y_k. Regarding γ_k as a curve, we see $\mathrm{d}_H(\gamma_k, y_k) \leq 2\lceil n^{1/256}\rceil$. Likewise, $\mathrm{d}_H(\mathsf{poly}(y_k, y_{k+1}), y_k) \leq 2\lceil n^{1/256}\rceil$, so for each $k \in \{0, \dots, \ell-1\}$, we have

$$\mathrm{d}_H(\gamma_k, \mathsf{poly}(y_k, y_{k+1})) \leq 4\lceil n^{1/256}\rceil\,, \tag{15.10}$$

and hence, for n taken sufficiently large, we have the following desirable bound:

$$\mathrm{d}_H(\gamma, n\rho') \leq n^{1/128}\,. \tag{15.11}$$

Step IV: (Perturbation) It remains to perturb ρ' to a simple polygonal curve. For $\delta > 0$, use Proposition 14.3.1 (and Remark 14.3.2) to obtain a simple polygonal curve $\rho \subset [-1,1]^2$ so that $\mathrm{d}_H(\rho, \rho') \leq \delta$ and so that $\text{length}_{\beta_p}(\rho') + \delta \geq \text{length}_{\beta_p}(\rho)$. Using (15.11) and (15.9), we find

1. $\mathrm{d}_H(\gamma, n\rho) \leq n^{1/128} + n\delta$,
2. $|\mathfrak{b}(\gamma)| \geq (1-2\epsilon)(\text{length}_{\beta_p}(n\rho) - n\delta)$,

and the proof is complete upon setting $\delta = \min(n^{(1/128)-1}, \epsilon\text{length}_{\beta_p}(\rho'))$, adjusting ϵ and taking n larger if necessary. □

Our second result allows us to pass from right-most circuits of sufficient length to polygonal circuits. Note that the boundary of $[-1,1]^2$ now comes into play: we obtain control on the surface energy of the polygonal circuit (as opposed to simply the β_p-length) in terms of the $\mathbf{C}_n$-length of the right-most circuit (as opposed to the $\mathbf{C}_\infty$-length).

Lemma 15.1.2. Let $p > p_c(2)$ and let $\epsilon > 0$. There are positive constants $c_1(p,\epsilon)$ and $c_2(p,\epsilon)$ so that with probability at least $1 - c_1\exp(-c_2\log^2 n)$, whenever $\gamma \subset [-n,n]^2$ is an open right-most circuit with $|\gamma| \geq n^{1/4}$, there is a simple polygonal circuit $\rho = \rho(\gamma) \subset [-1,1]^2$ with

1. $\mathrm{d}_H(\gamma, n\rho) \leq n^{1/16}$,

2. $|\mathfrak{b}^n(\gamma)| \geq (1-\epsilon)\mathcal{I}_p(n\rho)$.

Moreover, if $\gamma \subset (-n,n)^2$, we may replace (2) above with

3. $|\mathfrak{b}^n(\gamma)| \geq (1-\epsilon)\mathrm{length}_{\beta_p}(n\rho)$.

Proof. Let $\gamma \subset [-n,n]^2$ be an open right-most circuit with $|\gamma| \geq n^{1/4}$, and express γ as an alternating sequence of vertices and edges

$$\gamma = (x_0, e_1, x_1, e_2, x_2, \dots, e_m, x_m)\,, \tag{15.12}$$

where $x_0 = x_m$.

Step I: (Decomposition of γ) Say that x_i is a *boundary vertex* if $x \in \partial[-n,n]^2$ and that x_i is an *interior vertex* otherwise. If no x_i in γ is a boundary vertex, our analysis is simplified, so we postpone dealing with this case. As γ is a circuit, we lose no generality supposing x_0 is a boundary vertex. Let $\widetilde{x}_0, \dots, \widetilde{x}_\ell$ enumerate the boundary vertices of γ ordered in terms of increasing index in (15.12). For $j \in \{1, \dots, \ell\}$, let γ_j be the subpath of γ starting at $\widetilde{x}_{j-1}$ and ending at $\widetilde{x}_j$. Each γ_j is right-most and has the property that only the endpoints of γ_j are boundary vertices.

Say γ_j is *long* if $|\gamma_j| \geq n^{1/32}$, and that it is *short* otherwise. For each γ_j, let γ_j' denote the unique self-avoiding path of edges contained in $\partial[-n,n]^2$ whose starting and ending points are those of γ_j.

Step II: (Polygonal approximation) Let $\epsilon > 0$ and work within the high probability event from Lemma 15.1.1 for this parameter. For each long γ_j, there is then a simple polygonal curve $\rho_j \subset [-1,1]^2$ satisfying

1. $d_H(\gamma_j, n\rho_j) \leq n^{1/64}$,

2. $|\mathfrak{b}(\gamma_j)| \geq (1-\epsilon)\text{length}_{\beta_p}(n\rho_j)$.

If γ_j is short, the path γ_j' may be regarded as a polygonal curve $n\rho_j \subset \partial[-n,n]^2$ joining $\widetilde{x}_{j-1}$ with $\widetilde{x}_j$. Thus, each γ_j gives rise to a simple polygonal curve $\rho_j \subset [-1,1]^2$ in one of two ways, according to $|\gamma_j|$. Let ρ' be the concatenation of the ρ_j in the proper order:

$$\rho' := \rho_1 * \cdots * \rho_\ell \,, \tag{15.13}$$

so that ρ' is a polygonal circuit.

We claim ρ' has the desired properties; we first check d_H-closeness of $n\rho'$ and γ. If γ_j is short, any vertex in γ_j has an ℓ^∞-distance of at most $2n^{1/32}$ to $\widetilde{x}_j$, and likewise any vertex in γ_j' has an ℓ^∞-distance of at most $2n^{1/32}$ to $\widetilde{x}_j$. It follows that $d_H(\gamma_j, n\rho_j) \leq 4n^{1/32}$ when γ_j is short. In the case that γ_j is long, (1) above provides even better control, and we conclude

$$d_H(\gamma, n\rho') \leq 4n^{1/32} + n^{1/64} \,. \tag{15.14}$$

We now turn to controlling $\mathcal{I}_p(n\rho')$. Using the decomposition $\gamma = \gamma_1 * \cdots * \gamma_\ell$ and the construction of ρ',

$$|\mathfrak{b}^n(\gamma)| \geq \sum_{j\,:\,\gamma_j \text{ long}} |\mathfrak{b}(\gamma_j)| \,, \tag{15.15}$$

$$\geq (1-\epsilon) \sum_{j\,:\,\gamma_j \text{ long}} \text{length}_{\beta_p}(n\rho_j) \,, \tag{15.16}$$

$$\geq (1-\epsilon)\mathcal{I}_p(n\rho') \,, \tag{15.17}$$

where we have used (2) to obtain the second line directly above.

Step III: (Perturbation) It remains to perturb ρ' to a simple polygonal circuit. Let $\delta > 0$, and apply Corollary 14.3.3 (and Remark 14.3.4) to ρ' with this δ, so that by (15.14) we have

$$d_H(\gamma, n\rho) \leq 4n^{1/32} + n^{1/64} + \delta n \,, \tag{15.18}$$

and by (15.17) we have

$$|\mathfrak{b}^n(\gamma)| \geq (1-\epsilon)(\mathcal{I}_p(n\rho) - \delta n). \tag{15.19}$$

The proof is complete upon setting $\delta = \min(n^{(1/32)-1}, \epsilon\mathcal{I}_p(\rho'))$, adjusting ϵ and taking n larger if necessary. In the case that γ contains no boundary vertices, we split γ into a concatenation of two long right-most paths and proceed as above. □

15.2 Interlude: optimizers are of order n^2

In the arguments to come, it will be important to know that with high probability, each Cheeger optimizer has size on the order of n^2. First, we present a self-contained argument that $\widehat{\Phi}_n$ is at most a constant times n^{-1} with high probability. This follows from results mentioned in the introduction, but the proof given here is short enough to include.

Proposition 15.2.1. Let $p > p_c(2)$. There are positive constants $c(p), c_1(p), c_2(p) > 0$ so that with probability at least $1 - c_1 \exp(-c_2 \log^2 n)$, we have

$$\widehat{\Phi}_n \leq cn^{-1} \tag{15.20}$$

Proof. We use the previous two results to provide a high-probability lower bound on $|\mathbf{C}_n|$. Fix $\delta > 0$. Using Proposition 16.0.3 and Proposition 16.0.4, we find that with probability at least $1 - c_1 \exp(-c_2 \log^2 n)$,

$$|\mathbf{C}_n| \geq |\mathbf{C}_\infty \cap [-n,n)^2| - 4n\log^2 n\,, \tag{15.21}$$

$$\geq (\theta_p - \delta)(2n)^2 - 4n\log^2 n\,, \tag{15.22}$$

$$\geq (\theta_p - 2\delta)(2n)^2\,, \tag{15.23}$$

where we have taken n sufficiently large to obtain the last line. Define $H_n := [-n/8, n/8)^2 \cap \mathbf{C}_n$. Within the above events, we have $[-n/8, n/8)^2 \cap \mathbf{C}_n = [-n/8, n/8)^2 \cap \mathbf{C}_\infty$, and thus we may also work within the high probability event that $|H_n| \in ((\theta_p - \delta)(n/4)^2, (\theta_p + \delta)(n/4)^2)$.

Thus for δ chosen well, $|H_n| \le |\mathbf{C}_n|/2$. As $|\partial^n H_n|$ is at most a constant times n, we have shown that with high probability, $\widehat{\Phi}_n \le cn^{-1}$ for some $c > 0$. □

We now deduce that with high probability, each Cheeger optimizer is large.

Proposition 15.2.2. Let $p > p_c(2)$. There are positive constants $c_1(p), c_2(p), \alpha(p)$ so that with probability at least $1 - c_1 \exp(-c_2 \log^2 n)$, we have

$$\min_{G_n \in \mathcal{G}_n} |G_n| \ge \alpha n^2 . \qquad (15.24)$$

Proof. We make two assumptions:

1. G_n is connected.
2. $|G_n| \le |\mathbf{C}_n|/2 - n^{1/8}$

Use Lemma 12.2.4 and the fact that G_n is connected to identify a right-most circuit γ as in the statement of Lemma 12.2.4. We now follow ***Step I*** in the proof of Lemma 15.1.2 and write γ as an alternating sequence of vertices and edges:

$$\gamma = (x_0, e_1, x_1, e_2, x_2, \dots, e_m, x_m) , \qquad (15.25)$$

where $x_0 = x_m$. We say x_i is a *boundary vertex* if $x_i \in \partial[-n, n]^2$ and that it is an *interior vertex* otherwise. We split the remainder of the proof into two cases.

Case I: In the first case, we suppose γ contains no boundary vertices, so that $\partial^\infty G_n = \partial^n G_n$. Thanks to Proposition 16.0.5, the following event occurs with high probability:

$$\left\{ \Lambda \subset \mathbf{C}_n, \Lambda \text{ is connected}, |\Lambda| \ge n^{1/2} \implies |\partial^\infty \Lambda| \ge \widetilde{\alpha} |\Lambda|^{1/2} \right\} . \qquad (15.26)$$

Work within this event, and also the high probability event from Proposition 15.2.1 that $\widehat{\Phi}_n \le cn^{-1}$. As $\mathbf{C}_n$ is connected, it follows that $|\partial^n G_n| \ge 1$ for each Cheeger optimizer.

Thus, within the high probability events in which we work, it follows that $|G_n| \geq c^{-1}n$, and that within this first case,

$$|\partial^n G_n| = |\partial^\infty G_n| \geq \widetilde{\alpha}|G_n|^{1/2}\,, \tag{15.27}$$

so that $|G_n| \geq (\widetilde{\alpha}/c)^2 n^2$, which is desirable.

Case II: In the second case, we suppose that γ contains at least one boundary vertex, and we continue to follow ***Step I*** in the proof of Lemma 15.1.2. Without loss of generality, x_0 is then a boundary vertex and we let $\widetilde{x}_0, \ldots, \widetilde{x}_\ell$ enumerate the boundary vertices of γ in terms of increasing order in (15.25). For $j \in \{1, \ldots, \ell\}$, we let γ_j be the subpath of γ which begins at x_{j-1} and ends at x_j. As before, we note that each γ_j is right-most and that only the endpoints of γ_j are boundary vertices. We say that γ_j is *long* if $|\gamma_j| \geq n^{1/32}$ and that γ_j is *short* otherwise.

We claim that no γ_j can be short. To see this, let $\widetilde{\gamma}_j$ be the right-most path defined by the sequence of edges, each contained in $\partial[-n,n]^2$, and which begin at $\widetilde{x}_j$ and end at $\widetilde{x}_{j-1}$. Let ∂_j be the counter-clockwise interface which corresponds to $\gamma_j * \widetilde{\gamma}_j$, and observe that

$$|\mathsf{hull}(\partial_j) \cap \mathbf{C}_n| \leq \mathrm{Leb}(\mathsf{hull}(\partial_j)) + c|\gamma_j * \widetilde{\gamma}_j|\,, \tag{15.28}$$

$$\leq c\mathrm{length}(\partial_j)^2 + c|\gamma_j * \widetilde{\gamma}_j|\,, \tag{15.29}$$

$$\leq cn^{1/16} < n^{1/8}\,. \tag{15.30}$$

Here, c is an absolute constant which is allowed to change from line to line, and we have used the standard Euclidean isoperimetric inequality to obtain the second line. The third line follows from the assumption that γ_j is short and by taking n large. Writing $G_n' := G_n \cup [\mathsf{hull}(\partial_j) \cap \mathbf{C}_n]$, and using (2), we have that $|G_n'| \leq |\mathbf{C}_n|/2$ and that the conductance of G_n' is strictly smaller than that of G_n. This is a contradiction, so our claim that no γ_j can be short holds.

By Proposition 12.2.5, it is a high-probability event that $|\mathfrak{b}^n(\gamma_j)| \geq \alpha|\gamma_j|$. Thus, writing

∂ for the interface corresponding to γ, it follows that

$$|\partial^n G_n| \geq |\mathfrak{b}^n(\gamma)| \geq c\mathcal{H}^1\big(\partial \cap (-n,n)^2\big)\,, \tag{15.31}$$

$$\geq c\mathrm{Leb}\big(\mathsf{hull}(\partial) \cap [-n,n]^2\big)^{1/2} \tag{15.32}$$

$$\geq c|G_n|^{1/2}\,, \tag{15.33}$$

where we've used the isoperimetric inequality to obtain the second line, and where the constant $c > 0$ changes from line to line.

This handles the second case, and it remains to address our assumptions (1) and (2). If $|G_n| \geq |\mathbf{C}_n|/2 - n^{1/8}$, we use (15.23) from the proof of Proposition 15.2.1 and take n large to see that $|G_n| \geq cn^2$ with high probability in this case. Finally, any G_n is a disjoint union of connected Cheeger optimizers, so the lower bounds on the connected Cheeger optimizers suffice. □

15.3 Approximating discrete sets via polygons

Now that we have tools for converting right-most circuits to polygonal circuits, we use the decomposition given by Lemma 12.2.4 to pass from subgraphs of $\mathbf{C}_n$ to connected polygons. In order to relate the conductances of these objects, we require a mild isoperimetric assumption on the subgraph of $\mathbf{C}_n$.

Recall that $\mathcal{U}_n$ denotes the collection of connected subgraphs of $\mathbf{C}_n$ which inherit their graph structure from $\mathbf{C}_n$. Given a decomposition of $U \in \mathcal{U}_n$ as in Lemma 12.2.4, define

$$\mathsf{d\text{-}per}(U) := |\gamma| + \sum_{j=1}^{m} |\gamma_j|, \tag{15.34}$$

which may be thought of as the "full" perimeter of U. We also define

$$\mathsf{vol}(U) := \mathsf{hull}(\partial) \setminus \left(\bigsqcup_{j=1}^{m} \mathsf{hull}(\partial_j) \right)\,, \tag{15.35}$$

where ∂ and the ∂_j are the interfaces corresponding to the right-most circuits γ, γ_j.

Definition 15.3.1. Say that $U \in \mathcal{U}_n$ is *well-proportioned* if

$$\mathsf{d\text{-}per}(U) \leq \mathrm{Leb}(\mathsf{vol}(U))^{2/3}\,. \tag{15.36}$$

The following coarse-graining result says that with high probability, each $U \in \mathcal{U}_n$ is d_H-close to $\mathsf{vol}(U)$. Moreover, if $U \in \mathcal{U}_n$ is well-proportioned and sufficiently large, we may deduce U has "typical" density within $\mathsf{vol}(U)$. This second statement is Lemma 5.3 of [BLP15] rephrased, and we essentially follow the proof of this lemma to deduce Lemma 15.3.2 below.

Lemma 15.3.2. Let $p > p_c(2)$ and let $\epsilon > 0$. There are positive constants $c_1(p, \epsilon)$ and $c_2(p, \epsilon)$ so that with probability at least $1 - c_1 \exp(-c_2 \log^2 n)$,

$$\mathrm{d}_H(U, \mathsf{vol}(U)) \leq \log^4 n\,. \tag{15.37}$$

Moreover, whenever $U \in \mathcal{U}_n$ satisfies

1. U is well-proportioned,

2. $\mathrm{Leb}(\mathsf{vol}(U)) \geq \log^{20} n$,

we have

$$\left| \frac{|U|}{\mathrm{Leb}(\mathsf{vol}(U))} - \theta_p \right| < \epsilon\,. \tag{15.38}$$

Proof. Let $\epsilon > 0$, and define $r := \lfloor \log^2 n \rfloor$. For $u \in \mathbb{Z}^2$, define the square $S_u := (2r)u + [-r, r)^2$, and use the density result (Proposition 16.0.3) of Durrett and Schonmann with a union bound to obtain positive constants $c_1(p, \epsilon)$ and $c_2(p, \epsilon)$ so that the event

$$\mathcal{A}_n := \left\{ u \in \mathbb{Z}^2,\, S_u \cap [-n, n]^2 \neq \emptyset \implies \left| \frac{|\mathbf{C}_\infty \cap S_u|}{\mathrm{Leb}(S_u)} - \theta_p \right| < \epsilon \right\} \tag{15.39}$$

occurs with probability at least $\mathbb{P}_p(\mathcal{A}_n) \geq 1 - c_1 \exp(-c_2 \log^2 n)$. Given $U \in \mathcal{U}_n$, let (γ, ∂) and $\{(\gamma_j, \partial_j)\}_{j=1}^m$ be pairs of corresponding right-most and interface circuits for U, as in Lemma 12.2.4. Together, these circuits allow us to form $\mathsf{vol}(U)$, defined in (15.35). Define two collections of squares:

$$\mathsf{S}_1 := \left\{ S_u \,:\, u \in \mathbb{Z}^2,\, S_u \cap \partial \mathsf{vol}(U) \neq \emptyset \right\}, \tag{15.40}$$

$$\mathsf{S}_2 := \left\{ S_u \,:\, u \in \mathbb{Z}^2,\, S_u \subset \big(\mathsf{vol}(U) \setminus \partial \mathsf{vol}(U) \big) \right\}, \tag{15.41}$$

and let $y \in \mathsf{vol}(U)$. As the S_u form a partition of $\mathbb{R}^2$, it follows that y lives in exactly one S_u, which is then either in S_1 or S_2. If $S_u \in \mathsf{S}_1$, there is $u' \in \mathbb{Z}^2$ with $|u-u'|_\infty \leq 1$ so that S_u contains a vertex in γ or some γ_j. In this case, $\mathrm{dist}_\infty(y,U) \leq 4\log^2 n$. On the other hand, if $B_u \in \mathsf{S}_2$, working within the event $\mathcal{A}_n$, we find $S_u \cap \mathbf{C}_\infty \subset U$ is non-empty and hence that $\mathrm{dist}_\infty(y,U) \leq 4\log^2 n$. As $U \subset \mathsf{vol}(U)$, it follows from the above observations that $\mathrm{d}_H(U, \mathsf{vol}(U)) \leq \log^4 n$, for n sufficiently large.

We turn to the density of U within $\mathsf{vol}(U)$, and here we follow the proof of Lemma 5.3 in [BLP15]. Let $\mathsf{vol}(U)_r$ be the union of all squares in S_1 and let $\mathsf{vol}(U)^r$ be the union of all squares in $\mathsf{S}_1 \cup \mathsf{S}_2$, so that $\mathsf{vol}(U)_r \subset \mathsf{vol}(U) \subset \mathsf{vol}(U)^r$. We continue to work within $\mathcal{A}_n$, and we now assume U is well-proportioned and satisfies $\mathrm{Leb}(\mathsf{vol}(U)) \geq \log^{20} n$. We have

$$|U| \leq |\mathsf{vol}(U)^r \cap \mathbf{C}_\infty| \leq (\theta_p+\epsilon)\mathrm{Leb}(\mathsf{vol}(U)^r) \tag{15.42}$$

$$\leq (\theta_p+\epsilon)\big(\mathrm{Leb}(\mathsf{vol}(U)) + C'\mathrm{Leb}(B_u)\mathsf{d\text{-}per}(U)\big) \tag{15.43}$$

$$\leq \theta_p \mathrm{Leb}(\mathsf{vol}(U))(1+C\epsilon) \tag{15.44}$$

for some $C', C > 0$ and where n is taken sufficiently large to obtain the last line. The lower bound $|U| \geq \theta_p\mathrm{Leb}(\mathsf{vol}(U))(1-C\epsilon)$ follows similarly, finishing the proof. $\square$

Given $U \in \mathcal{U}_n$, we we will build a polygonal approximate from a collection of simple polygonal circuits. It is convenient to introduce the following construction, used in Lemma 15.3.4 which is in turn used in the proof of Proposition 15.3.5 below.

Definition 15.3.3. Given polygonal curves $\rho, \rho_1, \dots, \rho_m \subset \mathbb{R}^2$, we define the set $\mathsf{hull}(\rho, \rho_1, \dots, \rho_m)$ to be the union of $\rho \cup \rho_1 \cup \cdots \cup \rho_m$ with

$$\left\{ x \in \mathbb{R}^2 \setminus \left(\rho \cup \bigcup_{j=1}^m \rho_j \right) : w_\rho(x) - \left(\sum_{j=1}^m w_{\rho_j}(x) \right) \text{ is odd} \right\}, \tag{15.45}$$

where we recall $w_\rho(x), w_{\rho_j}(x)$ are the winding numbers of these curves about x.

Note that, in general, $\mathsf{hull}(\rho, \rho_1, \dots, \rho_m)$ is not a polygon, though it is when the curves ρ, ρ_j are in general position.

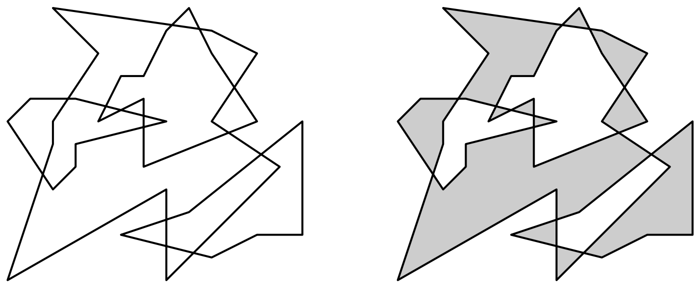

Figure 15.1: On the left, the curves $\rho, \rho_1, \rho_2, \rho_3$. On the right, $\mathsf{hull}(\rho, \rho_1, \ldots, \rho_3)$. As these curves are in general position, $\mathsf{hull}(\rho, \rho_1, \ldots, \rho_3)$ is a polygon.

Lemma 15.3.4. Let $R \in \mathcal{R}$ be connected, with ∂R consisting of the Jordan curves $\lambda, \lambda_1, \ldots, \lambda_m$. Let $\delta > 0$ and let $\rho, \rho_1, \ldots, \rho_m \subset [-1,1]^2$ be simple polygonal circuits so that $\mathrm{d}_H(\lambda, \rho) \leq \delta$ and so that $\mathrm{d}_H(\lambda_j, \rho_j) \leq \delta$ for each j. We suppose that δ is small enough so that $\mathsf{hull}(\rho_j)^\circ \cap \mathsf{hull}(\rho)^\circ$ is non-empty for each j. There are simple polygonal circuits $\rho', \rho'_1, \ldots, \rho'_m \subset [-1,1]^2$ so that

1. $\mathrm{d}_H(\rho, \rho') \leq \delta$ and $\mathrm{d}_H(\rho_j, \rho'_j) \leq \delta$ for each j,
2. $P := \mathsf{hull}(\rho', \rho'_1, \ldots, \rho'_m)$ is a connected polygon,
3. $\mathrm{d}_H(R, P) \leq 2\delta$
4. $\mathcal{I}_p(\rho) + \mathcal{I}_p(\rho_1) + \cdots + \mathcal{I}_p(\rho_m) + \delta \geq \mathcal{I}_p(\partial P)$.

Proof. Using the continuity of the norm β_p, we may perturb each $\rho, \rho_1, \ldots, \rho_m$ to a collection $\rho', \rho'_1, \ldots, \rho'_m$ of simple polygonal curves in general position with respect to each other satisfying (1) and (4). Taking δ smaller if necessary, and using the hypotheses of the lemma, we may execute this perturbation in such a way that $\mathsf{hull}(\rho'_j)^\circ \cap \mathsf{hull}(\rho')^\circ$ is non-empty for each j. Using this and the transversality of the ρ', ρ'_j, it follows that $\mathsf{hull}(\rho', \rho'_1, \ldots, \rho'_m)$ is a connected polygon, which settles (2) (connectedness can be established by inducting on the number m of polygonal curves $\rho'_1, \ldots, \rho'_m$).

We turn our attention to the Hausdorff distance between R and $P := \mathsf{hull}(\rho', \rho'_1, \ldots, \rho'_m)$. Let $x \in R$. If $x \in \partial R$, there is $y \in \partial P$ a distance of at most 2δ from x. If $x \in R$ and $x \notin P$, we appeal to the definition of hull (using winding number) to deduce that x is at most 2δ from ∂P. A symmetric argument starting with $x \in P$ settles (3). $\square$

Proposition 15.3.5 below is our first tool for passing from elements of $\mathcal{U}_n$ to connected polygons.

Proposition 15.3.5. Let $p > p_c(2)$ and let $\epsilon > 0$. There are positive constants $c_1(p, \epsilon)$ and $c_2(p, \epsilon)$ so that with probability at least $1 - c_1 \exp(-c_2 \log^2 n)$, whenever $U \in \mathcal{U}_n$ satisfies

1. U is well-proportioned,
2. $\mathrm{Leb}(\mathsf{vol}(U)) \geq n^{7/4}$,
3. $|\partial^\infty U| \leq Cn$.

there is a connected polygon $P = P(U) \in \mathcal{R}$ so that

1. $\mathrm{d}_H(U, nP) \leq n^{1/2}$,
2. $\big||U| - \theta_p \mathrm{Leb}(nP)\big| \leq \epsilon |U|$,
3. $|\partial^n U| \geq (1 - \epsilon)\mathcal{I}_p(n\partial P)$.

Proof. Let $U \in \mathcal{U}_n$. Using Lemma 12.2.4, form the pairs of right-most and interface circuits (γ, ∂) and $\{(\gamma_j, \partial_j)\}_{j=1}^m$ associated to U. We view the interfaces ∂, ∂_j as Jordan curves (via "corner-rounding," see Remark 12.2.2). Recall that we denoted $\mathsf{hull}(\partial_j) \cap \mathbf{C}_\infty$ as Λ_j, and that the Λ_j are the finite connnected components of $\mathbf{C}_\infty \setminus U$. We say Λ_j is *large* if $|\Lambda_j| \geq n^{1/2}$ and that it is *small* otherwise.

Step I: (Filling of small components) Let $(\widetilde{\gamma}_1, \widetilde{\partial}_1) \dots, (\widetilde{\gamma}_\ell, \widetilde{\partial}_\ell)$ enumerate the pairs of right-most circuits and corresponding interfaces associated to the large components Λ_j. Define

$$R := \mathsf{hull}(\partial) \setminus \left(\bigsqcup_{i=1}^{\ell} \mathsf{hull}(\widetilde{\partial}_i)^\circ \right), \tag{15.46}$$

and let $\widetilde{U} := R \cap \mathbf{C}_\infty$ (hence, $R = \mathsf{vol}(\widetilde{U})$). Observe that $\widetilde{U}$ is well-proportioned because U is. By construction, $\widetilde{U}$ is close to U both in d_H-sense and in volume. To see this, observe that the open edge boundaries of each Λ_j are disjoint and are each subsets of $\partial^\infty U$. The hypothesis $|\partial^\infty U| \leq Cn$ implies

$$|\widetilde{U} \setminus U| \leq Cn^{3/2}\,, \tag{15.47}$$

and it is immediate that

$$\mathrm{d}_H(U, \widetilde{U}) \leq n^{1/2}\,. \tag{15.48}$$

Step II: (Constructing a polygon P***)*** We use Lemma 15.1.2 and Lemma 15.3.4 to build a suitable polygon from $\widetilde{U}$. By Corollary 16.0.2, for each large $\widetilde{\gamma}_i$, we have $|\widetilde{\gamma}_i| \geq n^{1/8}$ for n sufficiently large, and likewise that $|\gamma| \geq n^{1/8}$. Work within the high probability event from Lemma 15.1.2 and find simple polygonal circuits $\rho_i \subset [-1,1]^2$ for each large $\widetilde{\gamma}_i$ so that

1. $\mathrm{d}_H(\widetilde{\partial}_i, n\rho_i) \leq 2n^{1/16}$,
2. $|\mathfrak{b}^n(\widetilde{\gamma}_i)| \geq (1-\epsilon)\mathcal{I}_p(n\rho_i)$,

as well as a polygonal circuit $\rho \subset [-1,1]^2$ corresponding to γ with

3. $\mathrm{d}_H(\partial, n\rho) \leq n^{1/16}$,
4. $|\mathfrak{b}^n(\gamma)| \geq (1-\epsilon)\mathcal{I}_p(n\rho)$.

In the case that there are no large components, we simply define $P := \mathsf{hull}(\rho)$. In the case that the collection of large components is non-empty, we define P differently below. Using Lemma 15.3.4, find polygonal circuits $\rho', \rho_1', \dots, \rho_\ell' \subset [-1,1]^2$ so that

5. $\mathrm{d}_H(n\rho, n\rho') \leq n^{1/16}$ and $\mathrm{d}_H(n\rho_i, n\rho_i') \leq n^{1/16}$ for each i,
6. $P := \mathsf{hull}(\rho', \rho_1', \dots, \rho_\ell')$ is a connected polygon,
7. $\mathrm{d}_H(R, P) \leq 2n^{1/16}$
8. $\mathcal{I}_p(\rho) + \mathcal{I}_p(\rho_1) + \cdots + \mathcal{I}_p(\rho_\ell) + n^{-15/16} \geq \mathcal{I}_p(\partial P)$.

In either case, we will show the polygon $P \subset [-1,1]^2$ has the desired properties.

Step III: (Controlling $\mathcal{I}_p(\partial P)$) Within the first case that $P = \mathsf{hull}(\rho)$, we find

$$|\partial^n U| \geq |\partial^n \widetilde{U}| = |\mathfrak{b}^n(\gamma)| \geq (1-\epsilon)\mathcal{I}_p(n\partial P)\,, \tag{15.49}$$

which is satisfactory. Thus we may suppose the set of large components is non-empty. Let $\alpha > 0$ be as in the statement of Proposition 12.2.5 and let

$$\mathcal{E}_n := \left\{ \exists \gamma \in \bigcup_{\substack{x_0 \in [-n,n]^2 \cap \mathbb{Z}^2 \\ x \in \mathbb{Z}^2}} \mathcal{R}(x_0, x) \ : n^{1/8} \leq |\gamma| \leq 100n^2\,,\ |\mathfrak{b}(\gamma)| \leq \alpha|\gamma| \right\}, \tag{15.50}$$

so that Proposition 12.2.5 with a union bound gives positive constants $c_1(p)$ and $c_2(p)$ so that $\mathbb{P}_p(\mathcal{E}_n) \leq c_1 \exp(-c_2 n)$. Work within $\mathcal{E}_n^c$ for the remainder of the proof, and use the fact that $\mathfrak{b}(\widetilde{\gamma}_i) = \mathfrak{b}^n(\widetilde{\gamma}_i)$, along with the bound $|\widetilde{\gamma}_i| \geq n^{1/8}$:

$$(1+\epsilon)|\partial^n \widetilde{U}| = (1+\epsilon)\left(|\mathfrak{b}^n(\gamma)| + \sum_{i=1}^{\ell} |\mathfrak{b}^n(\widetilde{\gamma}_i)| \right), \tag{15.51}$$

$$\geq |\mathfrak{b}^n(\gamma)| + \sum_{i=1}^{\ell} |\mathfrak{b}^n(\widetilde{\gamma}_i)| + n^{1/16}\,, \tag{15.52}$$

for n sufficiently large. Continuing from (15.52), let us use (2), (4) and (8):

$$|\partial^n U| \geq |\partial^n \widetilde{U}| \geq \frac{1}{1+\epsilon}\left(|\mathfrak{b}^n(\gamma)| + \sum_{i=1}^{\ell} |\mathfrak{b}^n(\widetilde{\gamma}_i)| + n^{1/16} \right), \tag{15.53}$$

$$\geq \frac{1-\epsilon}{1+\epsilon}\left(\mathcal{I}_p(n\rho) + \sum_{i=1}^{\ell} \mathcal{I}_p(n\rho_i) + n^{1/16} \right), \tag{15.54}$$

$$\geq \frac{1-\epsilon}{1+\epsilon}\mathcal{I}_p(n\partial P)\,, \tag{15.55}$$

so P has the desired properties as far as the surface tension in this case as well.

Step IV: (d_H-closeness of nP and $\widetilde{U}$) Let $\mathcal{A}_n$ be the high probability event from Lemma 15.3.2, and work within this event for the remainder of the proof. In the case that the collection of large components is empty, $P = \mathsf{hull}(\rho)$ implies $\mathrm{d}_H(R, nP) \leq n^{1/16}$. As

$R = \mathsf{vol}(\widetilde{U})$, it follows from working within $\mathcal{A}_n$ that

$$d_H(\widetilde{U}, nP) \le n^{1/16} + \log^4 n\,. \tag{15.56}$$

On the other hand, if the collection of large components is non-empty, (7) implies

$$d_H(\widetilde{U}, nP) \le 2n^{1/16} + \log^4 n\,, \tag{15.57}$$

as desired.

Step V: (Controlling the volume of P) Let $r = \lceil n^{1/16} \rceil$, and for $x \in \mathbb{Z}^d$ let $B_x = x + [-2r, 2r]^2$. Let $V(\widetilde{U})$ denote the vertices of $\mathbb{Z}^2$ contained in the union of paths $\gamma \cup \bigcup_{i=1}^{\ell} \widetilde{\gamma}_i$. Observe that, in either construction of P, we have

$$nP \,\Delta\, R \subset \bigcup_{x \in V(\widetilde{U})} B_x\,, \tag{15.58}$$

so that

$$\begin{aligned} \mathrm{Leb}(nP \,\Delta\, R) &\le 100 n^{1/16} \left[\mathsf{d\text{-}per}(\widetilde{U})\right]\,, && (15.59)\\ &\le 100 n^{1/16} \left[\mathrm{Leb}(\mathsf{vol}(\widetilde{U}))\right]^{2/3}\,, && (15.60)\end{aligned}$$

as $\widetilde{U}$ is well-proportioned. As $\widetilde{U}$ is also sufficiently large and we are working within the event $\mathcal{A}_n$, we also have $\big||\widetilde{U}| - \theta_p \mathrm{Leb}(R)\big| \le \epsilon \mathrm{Leb}(R)$, thus

$$\mathrm{Leb}(nP \,\Delta\, R) \le 100 n^{1/16} \left[\frac{|\widetilde{U}|}{\theta_p - \epsilon}\right]^{2/3} \le \epsilon |\widetilde{U}|\,, \tag{15.61}$$

for n sufficinently large. It follows that

$$\begin{aligned} \big||\widetilde{U}| - \theta_p \mathrm{Leb}(nP)\big| &\le \big||\widetilde{U}| - \theta_p \mathrm{Leb}(R)\big| + \epsilon |\widetilde{U}|\,, && (15.62)\\ &\le \left(\frac{\epsilon}{\theta_p - \epsilon} + \epsilon\right) |\widetilde{U}|\,. && (15.63)\end{aligned}$$

Step VI: (Wrapping up) Using (15.47), we have

$$\begin{aligned} \big||U| - \theta_p \mathrm{Leb}(nP)\big| &\le \left(\frac{\epsilon}{\theta_p - \epsilon} + \epsilon\right) (|U| + Cn^{3/2}) + Cn^{3/2}\,, && (15.64)\\ &\le C' \epsilon |U|\,, && (15.65)\end{aligned}$$

for some $C' > 0$ and when n is taken sufficiently large. By (15.48) and either (15.56) or (15.57), we also have $\mathrm{d}_H(U, nP) \leq n^{1/2}$ for n sufficiently large. Finally, recall that from either (15.49) or (15.55) we have $|\partial^n U| \geq \frac{1-\epsilon}{1+\epsilon}\mathcal{I}_p(\partial nP)$. The proof is complete upon adjusting ϵ. □

We now apply Proposition 15.3.5 to connected Cheeger optimizers. Let us define

$$\mathcal{G}_n^* := \left\{ G_n \in \mathcal{G}_n \ : G_n \text{ is connected} \right\}. \tag{15.66}$$

Proposition 15.3.6. Let $p > p_c(2)$. There are positive constants $c_1(p,\epsilon), c_2(p,\epsilon)$ so that for all $n \geq 1$, with probability at least $1 - c_1 \exp(-c_2 \log^2 n)$, for each $G_n \in \mathcal{G}_n^*$, there is a connected polygon $P_n \equiv P(G_n, \epsilon) \subset [-1,1]^2$ satisfying

1. $\mathrm{d}_H(G_n, nP_n) \leq 2n^{1/2}$,
2. $\big||G_n| - \theta_p \mathrm{Leb}(nP_n)\big| \leq \epsilon |G_n|$,
3. $|\partial^n G_n| \geq (1-\epsilon)\mathcal{I}_p(n\partial P_n)$.

Proof. Let us work within the high probability event from Proposition 15.2.2 that for some $\alpha_1 > 0$, we have $\min_{G_n \in \mathcal{G}_n} |G_n| \geq \alpha_1 n^2$. In conjunction with Proposition 15.2.1, we find that $\max_{G_n \in \mathcal{G}_n} |\partial^n G_n| \leq \alpha' n$ for some $\alpha' > 0$. As $|\partial^\infty G_n \setminus \partial^n G_n| \leq 8n$ for all $G_n \in \mathcal{G}_n$, it follows that $\max_{G_n \in \mathcal{G}_n} |\partial^\infty G_n| \leq \alpha_2 n$ for some $\alpha_2 > 0$. Fix $G \equiv G_n \in \mathcal{G}_n^*$, and observe that $G \in \mathcal{U}_n$.

We begin by following the proof of Propostion 15.3.5: consider the pairs of right-most and interface circuits (γ, ∂) and $\{(\gamma_j, \partial_j)\}_{j=1}^m$ which give rise to $\mathsf{vol}(G)$ and let Λ_j denote $\mathsf{hull}(\partial_j) \cap \mathbf{C}_\infty$. Say that Λ_j is *large* if $|\Lambda_j| \geq n^{1/2}$ and that Λ_j is *small* otherwise. Define

$$\widetilde{G} := \left[\mathsf{hull}(\partial) \setminus \left(\bigsqcup_{j\,:\,\Lambda_j \text{ large}} \mathsf{hull}(\partial_j) \right) \right] \cap \mathbf{C}_\infty, \tag{15.67}$$

As in the proof of Propostion 15.3.5, we observe $\widetilde{G}$ is close to G both in d_H-sense and in volume; as $|\partial^\infty G_n| \leq \alpha_2 n$, we find

$$|\widetilde{G} \setminus G| \leq \alpha_2 n^{3/2} \quad \text{and} \quad \mathrm{d}_H(\widetilde{G}, G) \leq n^{1/2}. \tag{15.68}$$

Step I: (Controlling $\mathsf{d\text{-}per}(\widetilde{G})$***)*** The isoperimetric inequality (Corollary 16.0.2) implies $|\gamma_j| \geq n^{1/8}$ for any Λ_j which is large. Likewise, because $|G| \geq \alpha_1 n^2$, we also have $|\gamma| \geq n^{1/8}$. Let $\alpha > 0$ be as in the statement of Proposition 12.2.5 and let $\mathcal{E}_n$ be the event defined in (15.50). Work within the high probability event $\mathcal{E}_n^c$ for the remainder of the proof, so that $|\mathfrak{b}(\gamma)| \geq \alpha|\gamma|$ and for each large $|\Lambda_j|$ we find $|\mathfrak{b}(\gamma_j)| \geq \alpha|\gamma_j|$. It follows that

$$\mathsf{d\text{-}per}(\widetilde{G}) \leq \frac{\alpha_2}{\alpha} n\,. \tag{15.69}$$

Step II: (Showing Leb$(\mathsf{vol}(\widetilde{G}))$ ***is on the order of*** n^2***)*** By construction, for some $C > 0$,

$$\begin{aligned} \mathrm{Leb}(\mathsf{vol}(\widetilde{G})) &\geq |\mathsf{vol}(\widetilde{G}) \cap \mathbb{Z}^2| - C\mathsf{d\text{-}per}(\widetilde{G})\,, &(15.70)\\ &\geq |G| - C\mathsf{d\text{-}per}(\widetilde{G}) \geq \frac{\alpha_1}{2} n^2\,, &(15.71)\end{aligned}$$

for n sufficiently large. We thus conclude that $\widetilde{G}$ is well-proportioned and satisfies $\mathrm{Leb}(\mathsf{vol}(\widetilde{G})) \geq n^{7/4}$ when n is large enough. Moreover, $\partial^\infty \widetilde{G} \subset \partial^\infty G$, so that $|\partial^\infty \widetilde{G}| \leq \alpha_2 n$, and $\widetilde{G}$ satisfies all necessary prerequisites of Proposition 15.3.5.

Step III: (Building a polygon) Work within the high probability event from Propostion 15.3.5, use (15.68) and the fact that $\partial^n \widetilde{G} \subset \partial^n G$ to obtain a polygon $P \equiv P(G, \epsilon) \subset [-1,1]^2$ with

1. $\mathrm{d}_H(G, nP) \leq 2n^{1/2}$,
2. $\big||G| - \theta_p \mathrm{Leb}(nP)\big| \leq 2\epsilon|G|$,
3. $|\partial^n G| \geq (1-\epsilon)\mathcal{I}_p(n\partial P)$,

where we have taken n sufficiently large to obtain the second item directly above. The proof is complete. □

15.4 Proofs of main theorems

We begin by proving a precursor to Theorem 11.3.1 for connected Cheeger optimizers.

Proposition 15.4.1. Let $p > p_c(2)$ and let $\epsilon > 0$. There are positive constants $c_1(p, \epsilon)$ and $c_2(p, \epsilon)$ so that for all $n \geq 1$, with probability at least $1 - c_1 \exp(-c_2 \log^2 n)$, we have

$$\max_{G_n \in \mathcal{G}_n^*} \mathrm{d}_H(n^{-1}G_n, \mathcal{R}_p) \leq \epsilon \,. \tag{15.72}$$

We emphasize that the maximum directly above runs over $\mathcal{G}_n^*$.

Proof. Let $\epsilon > 0$, and define the event

$$\mathcal{E}^{(n)}(\epsilon) := \Big\{ \exists G_n \in \mathcal{G}_n^* \ : \mathrm{d}_H(n^{-1}G_n, \mathcal{R}_p) > \epsilon \Big\} \tag{15.73}$$

Let $\epsilon' > 0$ to be determined later, and let $\mathcal{A}_1^{(n)}(\epsilon')$ be the event from Proposition 15.3.6 that for each $G_n \in \mathcal{G}_n^*$, there is a connected polygon $P_n \subset [-1, 1]^2$ so that

1. $\mathrm{d}_H(G_n, nP_n) \leq 2n^{1/2}$,
2. $\big||G_n| - \theta_p \mathrm{Leb}(nP_n)\big| \leq \epsilon' |G_n|$,
3. $|\partial^n G_n| \geq (1 - \epsilon')\mathcal{I}_p(n\partial P_n)$,

Let us first give an upper bound on $\mathrm{Leb}(P_n)$ within the event $\mathcal{A}_1^{(n)}(\epsilon')$ and another high probability event. Let

$$\mathcal{A}_2^{(n)}(\epsilon') := \left\{ \frac{|\mathbf{C}_n|}{(2n)^2} \in \Big((1-\epsilon')\theta_p, (1+\epsilon')\theta_p \Big) \right\} , \tag{15.74}$$

so that by Proposition 16.0.3 and Proposition 16.0.4, there are positive constants $c_1(p, \epsilon'), c_2(p, \epsilon')$ with $\mathbb{P}(\mathcal{A}_2^{(n)}(\epsilon')^c) \leq c_1 \exp(-c_2 \log^2 n)$. Within the intersection $\mathcal{A}_1^{(n)}(\epsilon') \cap \mathcal{A}_2^{(n)}(\epsilon')$ and using (2), we have

$$\max_{G_n \in \mathcal{G}_n^*} \mathrm{Leb}(P_n) \leq 2(1+\epsilon')^2 \,, \tag{15.75}$$

and let us choose $\alpha = \alpha(\epsilon') > 0$ so that $2 + \alpha = 2(1+\epsilon')^2$. Recall that Corollary 13.3.6 tells us there is $\delta = \delta(\epsilon) > 0$ so that when $R \in \mathcal{R}$ is connected with $\mathrm{Leb}(R) \leq 2 + \alpha$ and $\mathrm{d}_H(R, \mathcal{R}_p^{(2+\alpha)}) > \epsilon/100$, we have

$$\frac{\mathcal{I}_p(\partial R)}{\mathrm{Leb}(R)} > \varphi_p^{(2+\alpha)} + \delta \,. \tag{15.76}$$

We now take ϵ' small enough so that (using Lemma 13.3.4), we have $\mathrm{d}_H(\mathcal{R}_p^{(2+\alpha)}, \mathcal{R}_p) \leq \epsilon/4$. Thus, for this ϵ', within $\mathcal{E}_n(\epsilon) \cap \mathcal{A}_1^{(n)}(\epsilon') \cap \mathcal{A}_2^{(n)}(\epsilon')$ and for n sufficiently large (using (1)), the following event occurs

$$\left\{ \mathrm{d}_H(P_n, \mathcal{R}_p^{(2+\alpha)}) > \epsilon/4 \right\} , \tag{15.77}$$

so that by (15.76), (2) and (3), we have

$$n\widehat{\Phi}_n \geq (1-\epsilon')^2 \theta_p^{-1} \frac{\mathcal{I}_p(\partial P_n)}{\mathrm{Leb}(P_n)} , \tag{15.78}$$

$$\geq (1-\epsilon')^2 \theta_p^{-1} \left[\varphi_p^{(2+\alpha)} + \delta \right] , \tag{15.79}$$

within $\mathcal{E}_n(\epsilon) \cap \mathcal{A}_1^{(n)}(\epsilon') \cap \mathcal{A}_2^{(n)}(\epsilon')$. Working within this intersection, we use Corollary 13.2.4 to deduce

$$n\widehat{\Phi}_n \geq (1-\epsilon')^2 \theta_p^{-1} \left(\frac{2-\alpha}{2+\alpha} \varphi_p^{(2-\alpha)} + \delta \right) , \tag{15.80}$$

$$\geq (1-\epsilon')^2 \theta_p^{-1} \left(\frac{2-\alpha}{2+\alpha} \varphi_p + \delta \right) , \tag{15.81}$$

$$\geq \theta_p^{-1} \left(\varphi_p + \delta/2 \right) , \tag{15.82}$$

where we have taken ϵ' sufficiently small (depending on δ and hence ϵ) to obtain the last line, and where we emphasize the cruciality that δ does not depend on ϵ'. Thus,

$$\mathbb{P}_p(\mathcal{E}_n(\epsilon)) \leq \mathbb{P}_p(\mathcal{A}_1^{(n)}(\epsilon')^c) + \mathbb{P}_p(\mathcal{A}_2^{(n)}(\epsilon')^c) + \mathbb{P}_p\left(n\widehat{\Phi}_n \geq \theta_p^{-1}(\varphi_p + \delta/2) \right) \tag{15.83}$$

We have established that $\mathcal{A}_1^{(n)}(\epsilon')^c$ and $\mathcal{A}_2^{(n)}(\epsilon')^c$ are low-probability events; we bound the last term in (15.83) using Theorem 14.3.5 to complete the proof. □

Proof of Theorem 11.3.3. Let $\delta > 0$, and let $\mathcal{A}_1^{(n)}(\delta)$ and $\mathcal{A}_2^{(n)}(\delta)$ be the high-probability events from the proof of Proposition 15.4.1 for the parameter δ in place of ϵ'. Within the intersection $\mathcal{A}_1^{(n)}(\delta) \cap \mathcal{A}_2^{(n)}(\delta)$, we have for each $G_n \in \mathcal{G}_n^*$ a connected polygon $P_n \subset [-1,1]^2$ satisfying

1. $\mathrm{Leb}(P_n) \leq 2(1+\delta)^2$,

2. $\big| |G_n| - \theta_p \mathrm{Leb}(nP_n) \big| \leq \delta |G_n|$,

3. $|\partial^n G_n| \geq (1-\delta)\mathcal{I}_p(n\partial P_n)$,

and as before we define $\alpha = \alpha(\delta) > 0$ so that $2(1+\delta)^2 = 2+\alpha$. Thus, within $\mathcal{A}_1^{(n)}(\delta) \cap \mathcal{A}_2^{(n)}(\delta)$ we have

$$n\widehat{\Phi}_n \geq (1-\delta)^2 \frac{\mathcal{I}_p(\partial P_n)}{\theta_p \mathrm{Leb}(P_n)}, \tag{15.84}$$

$$\geq (1-\delta)^2 \frac{\varphi_p^{(2+\alpha)}}{\theta_p}, \tag{15.85}$$

$$\geq \frac{(1-\delta)^2(2-\alpha)}{2+\alpha} \frac{\varphi_p}{\theta_p}, \tag{15.86}$$

where we have used Corollary 13.2.4 and the fact that $\varphi_p^{(2-\alpha)} \geq \varphi_p$ to obtain the last line. Thus, for $\epsilon > 0$, we may take δ and hence α sufficiently small so that within $\mathcal{A}_1^{(n)}(\delta) \cap \mathcal{A}_2^{(n)}(\delta)$, we have $n\widehat{\Phi}_n \geq (1-\epsilon)(\varphi_p/\theta_p)$. Using Theorem 14.3.5, we then conclude that for all $n \geq 1$, there are positive constants $c_1(p,\epsilon)$ and $c_2(p,\epsilon)$ so that with probability at least $1 - c_1 \exp(-c_2 \log^2 n)$, we have

$$(1+\epsilon)\varphi_p \geq n\widehat{\Phi}_n \geq (1-\epsilon)\varphi_p. \tag{15.87}$$

We apply Borel-Cantelli to complete the proof. □

Proof of Theorem 11.3.1. Our strategy is to show that each $G_n \in \mathcal{G}_n^*$ is large. By Lemma 13.2.5, we have $\varphi_p^{(7/4)} > \varphi_p$. Let $\epsilon > 0$ be small enough so that $\varphi^{(7/4)} > (1+\epsilon)\varphi_p$, and choose δ depending on this ϵ so that

$$(1-\delta)^2 \varphi_p^{(7/4)} \geq (1+\epsilon)\varphi_p\,. \tag{15.88}$$

For this δ, work within the intersection $\mathcal{A}_1^{(n)}(\delta) \cap \mathcal{A}_2^{(n)}(\delta)$, the events introduced in the proof of Proposition 15.4.1, so that for each $G_n \in \mathcal{G}_n^*$, there is a connected polygon $P_n \subset [-1,1]^2$ with

1. $\mathrm{d}_H(G_n, nP_n) \leq 2n^{1/2}$,

2. $\big||G_n| - \theta_p \mathrm{Leb}(nP_n)\big| \leq \delta |G_n|$,

3. $|\partial^n G_n| \geq (1-\delta)\mathcal{I}_p(n\partial P_n)$,

Thus by (2), (3) and (15.88)

$$\mathcal{A}_1^{(n)}(\delta) \cap \mathcal{A}_2^{(n)}(\delta) \cap \{\exists G_n \in \mathcal{G}_n^* \ : \mathrm{Leb}(P_n) \leq 7/4\} \subset \left\{ n\widehat{\Phi}_n \geq (1-\delta)^2 \varphi_p^{(7/4)} \right\} , \tag{15.89}$$

$$\subset \left\{ n\widehat{\Phi}_n \geq (1+\epsilon)\varphi_p \right\} . \tag{15.90}$$

Let us write $\mathcal{F}_n(\epsilon)$ for the complement of the event in (15.90). Theorem 14.3.5 tells us $\mathcal{F}_n(\epsilon)$ occurs with high probability, so that on the intersection $\mathcal{A}_1^{(n)}(\delta) \cap \mathcal{A}_2^{(n)}(\delta) \cap \mathcal{F}_n(\epsilon)$, we have

$$\min_{G_n \in \mathcal{G}_n^*} \mathrm{Leb}(P_n) > 7/4 \,, \tag{15.91}$$

and hence by (2),

$$\min_{G_n \in \mathcal{G}_n^*} |G_n| \geq \frac{1}{1+\delta} \theta_p \left(\frac{7}{4} \right) n^2 \,. \tag{15.92}$$

As we are working within $\mathcal{A}_2^{(n)}(\delta)$, we also have $|\mathbf{C}_n| \leq 4n^2 \theta_p (1+\delta)$, so that from (15.92) and by taking δ smaller if necessary, we find

$$\min_{G_n \in \mathcal{G}_n^*} |G_n| \geq \left(\frac{5}{16} \right) |\mathbf{C}_n| \,. \tag{15.93}$$

The inequality $\frac{a+b}{c+d} \geq \min\left(\frac{a}{c}, \frac{b}{d}\right)$ tells us that each $G_n \in \mathcal{G}_n$ is a disjoint union of elements of $\mathcal{G}_n^*$. The constraint $|G_n| \leq |\mathbf{C}_n|/2$ and (15.93) tell us that

$$\mathcal{A}_1^{(n)}(\delta) \cap \mathcal{A}_2^{(n)}(\delta) \cap \mathcal{F}_n(\epsilon) \subset \left\{ \mathcal{G}_n^* \equiv \mathcal{G}_n \right\} . \tag{15.94}$$

Thus, on the intersection of $\mathcal{A}_1^{(n)}(\delta) \cap \mathcal{A}_2^{(n)}(\delta) \cap \mathcal{F}_n(\epsilon)$ and the high-probability event from Proposition 15.4.1, we find that there are positive constants $c_1(p,\epsilon)$ and $c_2(p,\epsilon)$ so that for each $n \geq 1$, with probability at least $1 - c_1 \exp(-c_2 \log^2 n)$, we have

$$\max_{G_n \in \mathcal{G}_n} \mathrm{d}_H (n^{-1} G_n, \mathcal{R}_p) \leq \epsilon \,, \tag{15.95}$$

where we emphasize the above maximum now runs over all of $\mathcal{G}_n$. The proof is complete upon applying Borel-Cantelli. □

CHAPTER 16

Appendix 2: Percolation inputs and miscellany

Recall that $\mathcal{U}_n$ denotes the connected subgraphs of $\mathbf{C}_\infty \cap [-1,1]^2$ which are defined by their vertex set. For $U \in \mathcal{U}_n$, Lemma 12.2.4 furnishes pairs of right-most circuits and corresponding interfaces $(\gamma, \partial), (\gamma_1, \partial_1), \ldots, (\gamma_m, \partial_m)$ which "carve" U out of $\mathbf{C}_\infty$. Recall that we used these pairs to define the value $\mathsf{d\text{-}per}(U)$ in (15.34) and the set $\mathsf{vol}(U)$ in (15.35). Recall that we identify the interfaces $\partial, \partial_1, \ldots, \partial_m$ with simple closed curves, see Remark 12.2.2.

Lemma 16.0.2. There is $c > 0$ so that for all $n \geq 1$ and for all $U \in \mathcal{U}_n$,

$$\mathsf{d\text{-}per}(U) \geq c\text{Leb}(\mathsf{vol}(U))^{1/2} . \tag{16.1}$$

Proof. Using the correspondence of Proposition 12.2.1, we find constants $c_1, c_2 > 0$ so that whenever γ' is a right-most circuit with corresponding interface ∂', we have

$$c_1|\gamma'| \leq \text{length}(\partial') \leq c_2|\gamma'| , \tag{16.2}$$

where we view ∂' as a simple circuit in $\mathbb{R}^2$. As the circuits $\partial, \partial_1, \ldots, \partial_m$ make up the boundary of the set $\mathsf{vol}(U)$, the standard Euclidean isoperimetric inequality gives $c > 0$ so that

$$\text{length}(\partial) + \sum_{i=1}^{m} \text{length}(\partial_i) \geq c\text{Leb}(\mathsf{vol}(U))^{1/2} . \tag{16.3}$$

The proof is complete upon combining (16.2) with (16.3). $\square$

The next three results are more general percolation inputs. The following result of Durrett and Schonmann ([DS88] Theorems 2 and 3) allows us to control the density of the infinite cluster within large boxes.

Proposition 16.0.3. Let $p > p_c(2)$, let $\epsilon > 0$ and let $r > 0$, and let $B_r \subset \mathbb{R}^2$ be a translate of $[-r, r)^2$. There are positive constants c_1, c_2 depending on p and ϵ so that

$$\mathbb{P}_p\left(\frac{|\mathbf{C}_\infty \cap B_r|}{(2r)^2} \notin (\theta_p - \epsilon, \theta_p + \epsilon)\right) \leq c_1 \exp\Big(-c_2 n\Big). \tag{16.4}$$

The next result, due to Benjamini and Mossel, allows us to pass from $\widetilde{\mathbf{C}}_n = \mathbf{C}_\infty \cap [-n, n]^2$ to $\mathbf{C}_n$ (see Proposition 1.2 of [BM03] and Lemma 5.2 of [BLP15]).

Proposition 16.0.4. Let $p > p_c(2)$. There is a positive constant $c(p)$ such that for all $n \geq 1$, with probability at least $1 - \exp(-C \log^2 n)$, and for any $n' \leq n - \log^2 n$, we have

$$\mathbf{C}_\infty \cap [-n', n']^2 = \mathbf{C}_n \cap [-n', n']^2. \tag{16.5}$$

Finally we need Proposition A.2 of [BBH08], which we state in dimension two only.

Proposition 16.0.5. Let $p > p_c(2)$. There are positive constants $c_1(p), c_2(p)$ and $\widetilde{\alpha}(p)$ so that for all $t > 0$,

$$\mathbb{P}_p(\exists \Lambda \subset \mathbf{C}_\infty, \omega\text{-connected}, 0 \in \Lambda, |\Lambda| \geq t^2, |\partial^\infty \Lambda| < \widetilde{\alpha} |\Lambda|^{1/2}) \leq c_1 \exp(-c_2 t). \tag{16.6}$$

[ABC96] G. Alberti, G. Bellettini, M. Cassandro, and E. Presutti. "Surface tension in Ising systems with Kac potentials." *J. Statist. Phys.*, **82**(3-4):743–796, 1996.

[ACC90] K Alexander, JT Chayes, and L Chayes. "The Wulff construction and asymptotics of the finite cluster distribution for two-dimensional Bernoulli percolation." *Communications in mathematical physics*, **131**(1):1–50, 1990.

[ACM01] L. Ambrosio, V. Caselles, S. Masnou, and J.-M. Morel. "Connected components of sets of finite perimeter and applications to image processing." *J. Eur. Math. Soc. (JEMS)*, **3**(1):39–92, 2001.

[ADH15] A. Auffinger, M. Damron, and J. Hanson. "50 years of first passage percolation." 11 2015. arXiv:1511.03262.

[Alo86] N. Alon. "Eigenvalues and expanders." *Combinatorica*, **6**(2):83–96, 1986. Theory of computing (Singer Island, Fla., 1984).

[AP96] Peter Antal and Agoston Pisztora. "On the chemical distance for supercritical Bernoulli percolation." *Ann. Probab.*, **24**(2):1036–1048, 1996.

[BBH08] Noam Berger, Marek Biskup, Christopher E Hoffman, and Gady Kozma. "Anomalous heat-kernel decay for random walk among bounded random conductances." In *Annales de l'IHP Probabilités et statistiques*, volume 44, pp. 374–392, 2008.

[BIV00] Thierry Bodineau, Dmitri Ioffe, and Yvan Velenik. "Rigorous probabilistic analysis of equilibrium crystal shapes." *Journal of Mathematical Physics*, **41**(3):1033–1098, 2000.

[BIV01] T. Bodineau, D. Ioffe, and Y. Velenik. "Winterbottom construction for finite range ferromagnetic models: an $\mathbb{L}_1$-approach." *J. Statist. Phys.*, **105**(1-2):93–131, 2001.

[Bl05] Viktor Blåsjö. "The Isoperimetric Problem." *The American Mathematical Monthly*, **112**(6):526–566, 2005.

[BLP15] Marek Biskup, Oren Louidor, Eviatar B. Procaccia, and Ron Rosenthal. "Isoperimetry in Two-Dimensional Percolation." *Communications on Pure and Applied Mathematics*, **68**(9):1483–1531, 2015.

[BM03] Itai Benjamini and Elchanan Mossel. "On the mixing time of a simple random walk on the super critical percolation cluster." *Probability Theory and Related Fields*, **125**(3):408–420, 2003.

[Bod99] T Bodineau. "The Wulff construction in three and more dimensions." *Communications in mathematical physics*, **207**(1):197–229, 1999.

[Bod02] T. Bodineau. "On the van der Waals theory of surface tension." *Markov Process. Related Fields*, **8**(2):319–338, 2002. Inhomogeneous random systems (Cergy-Pontoise, 2001).

[BP12] Andrea Braides and Andrey Piatnitski. "Variational problems with percolation: dilute spin systems at zero temperature." *Journal of Statistical Physics*, **149**(5):846–864, 2012.

[BP13] Andrea Braides and Andrey Piatnitski. "Homogenization of surface and length energies for spin systems." *Journal of Functional Analysis*, **264**(6):1296–1328, 2013.

[CD81] J. Theodore Cox and Richard Durrett. "Some limit theorems for percolation processes with necessary and sufficient conditions." *Ann. Probab.*, **9**(4):583–603, 1981.

[Cer00] Raphaël Cerf. "Large deviations for three dimensional supercritical percolation." *Astérisque*, (267):vi+177, 2000.

[Cer06] Raphaël Cerf. *The Wulff Crystal in Ising and Percolation Models: Ecole d'Eté de Probabilités de Saint-Flour XXXIV-2004*, volume 34. Springer, 2006.

[Che70] Jeff Cheeger. "A lower bound for the smallest eigenvalue of the Laplacian." In *Problems in analysis (Papers dedicated to Salomon Bochner, 1969)*, pp. 195–199. Princeton Univ. Press, Princeton, N. J., 1970.

[Chu97] Fan R. K. Chung. *Spectral graph theory*, volume 92 of *CBMS Regional Conference Series in Mathematics*. Published for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 1997.

[CP00] Raphaël Cerf and Ágoston Pisztora. "On the Wulff crystal in the Ising model." *Annals of probability*, pp. 947–1017, 2000.

[CP01] Raphaël Cerf and Agoston Pisztora. "Phase coexistence in Ising, Potts and percolation models." In *Annales de l'IHP Probabilités et statistiques*, volume 37, pp. 643–724, 2001.

[CT11] Raphaël Cerf and Marie Théret. "Lower large deviations for the maximal flow through a domain of $\mathbb{R}^d$ in first passage percolation." *Probab. Theory Related Fields*, **150**(3-4):635–661, 2011.

[CT12] Raphaël Cerf and Marie Théret. "Maximal stream and minimal cutset for first passage percolation through a domain of $\mathbb{R}^d$." *Annals of Probability 2014, Vol. 42, No. 3, 1054-1120*, 01 2012.

[DKS92] Roland Lvovǐc Dobrushin, Roman Kotecky, and S Shlosman. *Wulff construction: a global shape from local interaction*, volume 104. American Mathematical Society Providence, Rhode Island, 1992.

[DP96] Jean-Dominique Deuschel and Agoston Pisztora. "Surface order large deviations for high-density percolation." *Probab. Theory Related Fields*, **104**(4):467–482, 1996.

[DS88] R. Durrett and R. H. Schonmann. "Large deviations for the contact process and two-dimensional percolation." *Probab. Theory Related Fields*, **77**(4):583–603, 1988.

[FMP10] Alessio Figalli, Francesco Maggi, and Aldo Pratelli. "A mass transportation approach to quantitative isoperimetric inequalities." *Inventiones mathematicae*, **182**(1):167–211, 2010.

[Gan89] Alberto Gandolfi. *Clustering and uniqueness in mathematical models of percolation phenomena*. PhD thesis, Technische Universiteit Delft, 1989.

[Gar09] Olivier Garet. "Capacitive flows on a 2D random net." *The Annals of Applied Probability*, pp. 641–660, 2009.

[Gib78] Josiah Willard Gibbs. "On the equilibrium of heterogeneous substances." *American Journal of Science*, (96):441–458, 1878.

[GK84] Geoffrey Grimmett and Harry Kesten. "First-passage percolation, network flows and electrical resistances." *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, **66**(3):335–366, 1984.

[GM90] G. R. Grimmett and J. M. Marstrand. "The supercritical phase of percolation is well behaved." *Proc. Roy. Soc. London Ser. A*, **430**(1879):439–457, 1990.

[GMP15] Olivier Garet, Régine Marchand, Eviatar B Procaccia, and Marie Théret. "Continuity of the time and isoperimetric constants in supercritical percolation." *arXiv preprint arXiv:1512.00742*, 2015.

[Gol16] J. Gold. "Isoperimetry in supercritical bond percolation in dimensions three and higher." *preprint*, 02 2016. arXiv:1602.05598.

[Gri99] Geoffrey Grimmett. *Percolation*, volume 321 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, second edition, 1999.

[Iof95] Dmitry Ioffe. "Exact large deviation bounds up to T_c for the Ising model in two dimensions." *Probability theory and related fields*, **102**(3):313–330, 1995.

[IS98] Dmitry Ioffe and Roberto H Schonmann. "Dobrushin–Koteckỳ–Shlosman theorem up to the critical temperature." *Communications in mathematical physics*, **199**(1):117–167, 1998.

[Kes80] H. Kesten. "The critical probability of bond percolation on the square lattice equals $\frac{1}{2}$." *Comm. Math. Phys.*, **74**(1):41–59, 1980.

[Kes86] Harry Kesten. "Aspects of first passage percolation." In *École d'été de probabilités de Saint-Flour, XIV—1984*, volume 1180 of *Lecture Notes in Math.*, pp. 125–264. Springer, Berlin, 1986.

[Kes87] Harry Kesten. "Surfaces with minimal random weights and maximal flows: a higher-dimensional version of first-passage percolation." *Illinois J. Math.*, **31**(1):99–166, 1987.

[KP94] R. Kotecký and C. E. Pfister. "Equilibrium shapes of crystals attached to walls." *Journal of statistical physics*, **76**(1-2):419–445, 1994.

[KZ90] Harry Kesten and Yu Zhang. "The probability of a large finite cluster in supercritical Bernoulli percolation." *Ann. Probab.*, **18**(2):537–555, 1990.

[MP05] Ben Morris and Yuval Peres. "Evolving sets, mixing and heat kernel bounds." *Probability Theory and Related Fields*, **133**(2):245–266, 2005.

[MR04] Pierre Mathieu and Elisabeth Remy. "Isoperimetry and heat kernel decay on percolation clusters." *Ann. Probab.*, **32**(1A):100–128, 2004.

[MS67] R. A. Minlos and Ja. G. Sinaĭ. "The phenomenon of "separation of phases" at low temperatures in certain lattice models of a gas. I." *Mat. Sb. (N.S.)*, **73 (115)**:375–448, 1967.

[MS68] R. A. Minlos and Ja. G. Sinaĭ. "The phenomenon of "separation of phases" at low temperatures in certain lattice models of a gas. II." *Trudy Moskov. Mat. Obšč.*, **19**:113–178, 1968.

[Pet08] Gábor Pete. "A note on percolation on $\mathbb{Z}^d$: isoperimetric profile via exponential cluster repulsion." *Electron. Commun. Probab.*, **13**:377–392, 2008.

[Pis96] Agoston Pisztora. "Surface order large deviations for Ising, Potts and percolation models." *Probab. Theory Related Fields*, **104**(4):427–466, 1996.

[PR12] Eviatar B. Procaccia and Ron Rosenthal. "Concentration estimates for the isoperimetric constant of the supercritical percolation cluster." *Electron. Commun. Probab.*, **17**:no. 30, 11, 2012.

[PV96] C-E Pfister and Yvan Alain Velenik. "Mathematical theory of the wetting phenomenon in the 2D Ising model." *Helvetica Physica Acta*, **69**:949–973, 1996.

[PV97] C-E Pfister and Yvan Velenik. "Large deviations and continuum limit in the 2D Ising model." *Probability theory and related fields*, **109**(4):435–506, 1997.

[Rau07] Clément Rau. "Sur le nombre de points visités par une marche aléatoire sur un amas infini de percolation." *Bull. Soc. Math. France*, **135**(1):135–169, 2007.

[RT10a] Raphaël Rossignol and Marie Théret. "Law of large numbers for the maximal flow through tilted cylinders in two-dimensional first passage percolation." *Stochastic Processes and their Applications*, **120**(6):873–900, 2010.

[RT10b] Raphaël Rossignol and Marie Théret. "Lower large deviations and laws of large numbers for maximal flows through a box in first passage percolation." *Ann. Inst. Henri Poincaré Probab. Stat.*, **46**(4):1093–1131, 2010.

[Shl89] S. B. Shlosman. "The droplet in the tube: a case of phase transition in the canonical ensemble." *Comm. Math. Phys.*, **125**(1):81–90, 1989.

[SZ03] J. Michael Steele and Yu Zhang. "Nondifferentiability of the time constants of first-passage percolation." *Ann. Probab.*, **31**(2):1028–1051, 2003.

[Tal95] Michel Talagrand. "Concentration of measure and isoperimetric inequalities in product spaces." *Publications Mathématiques de l'Institut des Hautes Etudes Scientifiques*, **81**(1):73–205, 1995.

[Tay74] Jean E. Taylor. "Existence and structure of solutions to a class of nonelliptic variational problems." In *Symposia Mathematica, Vol. XIV (Convegno di Teoria Geometrica dell'Integrazione e Varietà Minimali, INDAM, Roma, Maggio 1973)*, pp. 499–508. Academic Press, London, 1974.

[Tay75] Jean E. Taylor. "Unique structure of solutions to a class of nonelliptic variational problems." In *Differential geometry (Proc. Sympos. Pure. Math., Vol. XXVII, Stanford Univ., Stanford, Calif., 1973), Part 1*, pp. 419–427. Amer. Math. Soc., Providence, R.I., 1975.

[Tay78] Jean E. Taylor. "Crystalline variational problems." *Bull. Amer. Math. Soc.*, **84**(4):568–588, 1978.

[The14] Marie Theret. "Upper large deviations for maximal flows through a tilted cylinder." *ESAIM: Probability and Statistics*, **18**:117–129, 2014.

[Tim13] Ádám Timár. "Boundary-connectivity via graph theory." *Proc. Amer. Math. Soc.*, **141**(2):475–480, 2013.

[Win67] WL Winterbottom. "Equilibrium shape of a small particle in contact with a foreign substrate." *Acta Metallurgica*, **15**(2):303–310, 1967.

[Wul01] G Wulff. "Zur Frage der Geschwindigkeit des Wachstums und der Auflosung der Kristallflachen." *Z. Kryst. Miner*, **34**:449, 1901.

[Zha07] Yu Zhang. "Limit theorems for maximum flows on a lattice." 10 2007.